



GigaVUE Cloud Suite Deployment Guide - VMware (NSX-T)

GigaVUE Cloud Suite

Product Version: 6.14.xx

Document Version: 2.0

(See Change Notes for document updates.)

Copyright © 2026 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
6.14.01/6.14.02	2.0	09/07/2026	This release incorporates changes in the following topics for the listed features: • Port Optimization ◦ Network Firewall Requirements
6.14.00	1.0	07/10/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Cloud Suite Deployment Guide - VMware (NSX-T)	1
Change Notes	3
Contents	4
Get Started	9
Overview of GigaVUE Cloud Suite for VMware	9
Architecture for GigaVUE Cloud Suite for VMware NSX-T	10
GigaVUE-FM Orchestrated Fabric Deployment	11
Key Concepts	12
Secure Communication between GigaVUE Fabric Components	12
GigaVUE-FM acts as the PKI	14
Bring Your Own CA	14
Secure Communication in FMHA Mode	14
Supported Platforms	14
Supported Components	14
Rules and Notes	15
Increase or Decrease GigaVUE V Series Node	15
Increase or Decrease GigaVUE V Series Nodes using GigaVUE-FM	15
Increase or Decrease GigaVUE V Series Nodes using VMware NSX-T Manager	16
Sharing the Same Host across Different Monitoring Domains	18
Deployment Overview	18
Deployment Planning	19
Identify the Deployment Method	19
Identify the Supported Platform Versions	20
Identify the Traffic Visibility Scope	20
Identify the Host, Cluster, and Transport Zone Requirements	20
Identify Access, Roles, and Connectivity	20
Identify the Image Source and Certificate Strategy	21
Identify the Node Deployment Model and Scale	21
Identify Monitoring Session and Traffic Forwarding Requirements	21
Prerequisites for Integrating GigaVUE V Series Nodes with NSX-T	22
Licensing for GigaVUE Cloud Suite for VMware NSX-T	23
Default Trial Licenses	24
Volume-Based License	25

Network Firewall Requirements	28
Recommended Form Factor (Instance Types)	30
Required VMware Virtual Center Privilege	30
Required Roles in VMware NSX-T	30
Disable Certification Validation in VMware NSX-T	31
Default Login Credentials	31
Supported Hypervisors for VMware	32
Points to Note (VMware NSX-T)	33
Install and Upgrade GigaVUE-FM	33
Before You Install	34
Prerequisites for GigaVUE-FM	34
VMware ESXi Hardware Requirements	34
Recommended Resource Requirements for Scaled Environments	34
Rules, Notes, and Limitations	36
Reference Values for Disk Space Estimation	37
Recommended Form Factor for VMware vCenter (Instance Types)	37
Deploy OVA file on VMware vCenter	38
Troubleshooting	41
Install GigaVUE-FM using ISO Image File on VMware ESXi Host	41
Pre-requisites	42
Installation Steps	42
What to do Next	43
Install GigaVUE-FM using ISO Image File on VMware vCenter	44
Initial GigaVUE-FM Configuration	45
Perform Initial Configuration	46
Perform Network Configurations	46
Deploy	48
Deploy GigaVUE Cloud Suite for VMware (NSX-T)	48
Deployment Options for GigaVUE Cloud Suite for VMware (NSX-T)	49
Deploy GigaVUE V Series Nodes using GigaVUE-FM	49
What to do next:	55
Deploy GigaVUE V Series Nodes using VMware NSX-T Manager	60
What to do next:	66
Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform	68
License	68
Supportability	69
Rules and notes	70
Prerequisites	70
Get Started	70
Deployment Using Terraform	71
Install the Gigamon Terraform Provider Without Internet Access	73
Single-file configuration	75
Module-based configuration	76

Run the Terraform Deployment File	77
Best Practices	81
Verification and Troubleshooting	82
Configure	84
Configure Monitoring Session	84
Create a Monitoring Session (VMware NSX-T)	85
Monitoring Session Page	86
Configure Monitoring Session Options	87
Create Ingress and Egress Tunnel (VMware NSX-T)	88
Create Raw Endpoint (VMware NSX-T)	97
Rules and Notes	97
Configure Raw Endpoint in Monitoring Session	98
Configuration Support for Non IP-Addressable (I2 supported) Tools	98
Create a New Map	99
Example- Create a New Map using Inclusion and Exclusion Maps ..	103
Map Library	104
Deploy Monitoring Session	104
V Series Applications	107
Masking	112
Slicing	115
Application Intelligence	118
Points to Note for Application Intelligence	118
Application Visualization	120
Configure Application Visualization for Virtual Environment	120
Configure Filtering and Metadata Export for Selected Applications in Application Visualization	122
View Application Statistics for Application Visualization	123
User Defined Application	124
Create Rules for User Defined Application	130
Application Filtering Intelligence	131
Configure Application Filtering Intelligence for Virtual Environment	132
View Application Statistics for Application Filtering	137
Application Metadata Intelligence	137
Configure Application Metadata Intelligence for Virtual Environment	138
View Application Statistics for Application Metadata	147
Application Metadata Exporter	147
Application Metadata Exporter Deployment Options	148
Application Metadata Exporter - Export Scenarios	151
Rules and Notes	157
Configure Metadata Export Filter for AMX	158

Deploy Application Metadata Exporter in VMware NSX-T using Third Party Orchestration	159
Configure Cloud Tool Exporter	170
Deploy the Monitoring Session	170
Configure Kafka Exporter	172
Deploy the Monitoring Session	172
Extend Your AMI and AMX Deployment	173
Deploy Application Metadata Exporter (AMX) on KVM Hypervisor ..	174
Prerequisite:	175
View Application Metadata Exporter Statistics	176
Attributes for GigaVUE Enriched Metadata for Mobile Networks	180
Attributes for GigaVUE Enriched Metadata for Cloud Workloads ...	185
Supported Element Types for NetFlow Integrator Element Mapping	192
Troubleshoot AMI/AMX Deployment	196
Scenario 1: AMI not exporting IPFIX/CEF records (exporter side)	196
Scenario 2: AMX not receiving AMI/NetFlow records (ingestor side)	196
Scenario 3: AMX not exporting JSON records (exporter side)	197
GigaSMART NetFlow Generation	198
Create NetFlow Session for Virtual Environment	198
Examples- Configuring Application Intelligence Solution with Other Applications	203
Slicing and Masking with Application Filtering Intelligence	204
De-duplication with Application Metadata Intelligence	204
Migrate Application Intelligence Session to Monitoring Session	205
Post Migration Notes for Application Intelligence	207
De-duplication	209
Feature Overview	209
Configure De-duplication Application	210
Distributed De-duplication	211
Limitations	212
GENEVE Decapsulation	213
What's Next	213
Header Stripping	214
Limitation	215
Configure Header Stripping Application	215
Create Custom Port Template	217
Load Balancing	218
Enhanced Load Balancing	220
What's Next	224
5G Service Based Interface Application	224
Ericsson vTAP	224
How Packet Ordering Works	225
DTLS Support	225

Ericsson vTAP Domain Tagging	226
How SBI Application works	227
Supported Platforms:	228
Rules and Notes	228
End-to End Solution Overview and Configuration:	229
Ingress: Ericsson vTAP UDP-GRE Feed	229
PCAPNG Processing and Flow Tracking	230
Configure 5G-SBI Application	231
Egress Tunnels and VLAN Tag Insertion	234
Deploy the Monitoring Session	235
Adding CSV file for IP Mapping	235
Recommended VM Specifications	237
FHA Dashboards for 5G Applications	237
Interface Mapping	238
Monitor and Troubleshoot	238
Monitor Cloud Health	238
Configuration Health Monitoring	239
Traffic Health Monitoring	239
Supported Resources and Metrics	240
Create Threshold Templates	243
Apply Threshold Template	244
Clear Thresholds	244
View Health Status	245
View Health Status of an Application	245
View Operational Health Status of an Application	246
View Health Status for Individual GigaVUE V Series Nodes	247
View Application Health Status for Individual V Series Nodes	248
View Monitoring Session Statistics	248
Visualize the Network Topology (VMware NSX-T)	249
Administer GigaVUE Cloud Suite for VMware NSXT	250
Create Service Chain and NSX-T Group	250
Create Service Chain	250
Create Group	251
Create and Publish a Policy	251
Configure VMware Settings	252
Configure Certificate Settings	253
High Availability for Cloud Infrastructure	254
About GigaVUE-FM High Availability	254
Supportability Information	257
Recommendations, Rules, Notes, and Limitations	257
Licensing Information	260
Hostname Setups	261
Configure GigaVUE-FM HA	261

Configure Secure Communication between Fabric Components in FMHA	269
GigaVUE-FM HA Landing Page	270
Remove Standby GigaVUE-FM Instance	270
Disassemble GigaVUE-FM High Availability Group	271
GigaVUE-FM High Availability States	271
Fail over Mechanism	273
GigaVUE-FM High Availability Scenarios	274
Troubleshoot GigaVUE-FM High Availability Issues	274
Analytics for Virtual Resources	275
Virtual Inventory Statistics and Cloud Applications Dashboard	275
Remove Gigamon Service from NSX-T and GigaVUE-FM	280
Step 1: Remove the Rule and Policy	280
Step 2: Remove the Service Chain	280
Step 3: Delete the Monitoring Session	280
Step 4: Delete the Monitoring Domain	281
GigaVUE V Series Deployment Clean up	281
Step 1: Delete Rule and Policy in NSX-T Manager	281
Remove Service Deployments	282
Remove Service Reference	284
Remove Service Manager	284
Remove Vendor Template and Service Definition	285
Debuggability and Troubleshooting	285
Sysdumps	286
Sysdumps—Rules and Notes	286
Generate a Sysdump File	286
FAQs - Secure Communication between GigaVUE Fabric Components (VMware-VSeries-NSX-T)	287

Get Started

Overview of GigaVUE Cloud Suite for VMware

GigaVUE Cloud Suite for VMware enables you to manage and monitor network traffic in virtual environments. It captures, improves, and sends selected network traffic to your security and monitoring tools.

This solution works closely with VMware tools to give you clear visibility into traffic from virtual machines. It helps you understand what's happening in your private cloud.

GigaVUE-FM, a key part of the Cloud Suite, works with VMware vCenter to automatically set up GigaVUE V Series Node to support a growing private cloud infrastructure. It also helps track changes in workloads and keeps traffic policies working properly.

Benefits:


- **Flexible Traffic Acquisition:** Collects traffic using port mirroring in VMware ESXi.
- **Automated Visibility Provisioning:** Automatically sets up and applies traffic rules as new users or groups are added.
- **Improved Tool Efficiency:** Filters and balances traffic to reduce the load on your monitoring tools.
- **Application Intelligence Solution:** Detects thousands of applications and accesses over 7,000 application metadata elements to understand your network better.

Architecture for GigaVUE Cloud Suite for VMware NSX-T

This section provides an overview of the GigaVUE V Series Node deployment on the VMware NSX-T platform and describes the procedure for setting up the traffic monitoring sessions using the GigaVUE V SeriesNodes. The GigaVUE V Series Nodes support traffic visibility on the NSX-T NVDS switch.

GigaVUE-FM creates, manages and deletes the GigaVUE V Series Nodes in the VMware NSX-T based on the configuration information provided by the user. GigaVUE-FM can communicate directly with the GigaVUE V Series Nodes.

The following diagram provides a high-level overview of the deployment:



GigaVUE-FM Orchestrated Fabric Deployment

GigaVUE Cloud Suite for VMware comprises multiple elements that enable traffic acquisition, aggregation, intelligence and distribution, along with centralized, single-pane-of-glass orchestration and management.

Main Components

The following table lists the fabric components for GigaVUE Cloud Suite for VMware NSX-T.

Component	Description
GigaVUE-FM fabric manager (GigaVUE-FM)	Represents a web-based tool that helps you manage physical and virtual network traffic that forms GigaVUE Cloud Suite for VMware. It gives you complete visibility and control of your entire VMware cloud suite from one dashboard. GigaVUE-FM generates a complete network map to easily see which cloud systems are connected to the deep observability pipeline. It can manage hundreds of visibility nodes across on-premises and cloud environments. It also handles the setup for all other components in your platform.
GigaVUE® V Series Node	Represents a node that collects mirrored traffic, applies filters, and processes data using GigaSMART applications. It then sends the

Component	Description
	optimized traffic to your cloud-based tools or back to on-premises tools.

Key Concepts


GigaVUE Cloud Suite for VMware (NSX-T) supports the following features:

- [Secure Communication between GigaVUE Fabric Components](#)
- [Increase or Decrease GigaVUE V Series Node](#)
- [Sharing the Same Host across Different Monitoring Domains](#)

Secure Communication between GigaVUE Fabric Components

The Secure Communication feature in GigaVUE-VM uses mutual TLS (mTLS) authentication to improve network security. It ensures all GigaVUE Fabric Components communicate over encrypted, verified connections using certificates issued by a Certificate Authority (CA), without relying on static credentials.

How it Works!



In this setup:

- GigaVUE-FM establishes an mTLS connection and checks for GigaVUE V Series Proxy availability.
- If GigaVUE V Series Proxy is unavailable, it directly connects to the GigaVUE V Series Node through mTLS.
- If a GigaVUE V Series is available, GigaVUE-FM first connects to the GigaVUE V Series Proxy and establishes an mTLS connection with the GigaVUE V Series Node.
- GigaVUE-FM also initiates an mTLS connection to the UCT-V Controller, establishing an mTLS connection with UCT-V.

This structured flow ensures secure communication using mTLS-based authentication across all the fabric components.

GigaVUE-FM acts as the PKI

GigaVUE-FM manages all certificates for fabric components. It acts as a private PKI and uses Step-CA with the ACME protocol to issue and renew certificates. This automated process reduces the need for manual certificate handling and avoids external dependencies.

Bring Your Own CA

If your organization already uses a corporate CA, you can import those certificates into GigaVUE-FM. This allows your existing PKI infrastructure to work with Gigamon's secure communication system.

For more details on how to integrate your PKI infrastructure with GigaVUE-FM, refer to [Integrate Private CA](#)

Secure Communication in FMHA Mode

In FMHA (Fabric Manager High Availability) mode:

- The active GigaVUE-FM instance shares intermediate CA files with all standby nodes.
- Only the active instance handles certificate requests. In case of a failover, a standby node takes over.
- The root and intermediate CAs are copied to all nodes to ensure continuity.
- If an instance is removed, it generates a new self-signed CA on restart.

Supported Platforms

- AWS
- Azure
- OpenStack
- Nutanix
- GCP
- OCI
- Openshift(UCT-C)
- VMware ESXi
- VMware NSX-T

Supported Components

- GigaVUE V Series Node
- GigaVUE V Series Proxy

- UCT-V
- UCT-V Controller
- UCT-C Controller and Tap

Rules and Notes

- If a public IP is revoked in public cloud platforms, you can issue a new certificate to remove the old IP.
- This feature is optional.
- Ensure NTP (Network Time Protocol) runs if GigaVUE-FM and components are on different hosts.
- Applying a certificate may temporarily cause a component to show as Down, but it recovers automatically.
- In AWS, disable the Source/Destination Check on network interfaces for GigaVUE V Series Proxy.

Note: Enabling this check may block traffic if the IP address does not match the associated interface.

Increase or Decrease GigaVUE V Series Node

You can add more nodes or remove nodes from an existing monitoring domain using GigaVUE-FM or VMware NSX-T manager, based on method you have deployed the GigaVUE V Series Nodes.

For more information, refer to the following topics:

- [Increase or Decrease GigaVUE V Series Nodes using GigaVUE-FM](#)
- [Increase or Decrease GigaVUE V Series Nodes using VMware NSX-T Manager](#)

Increase or Decrease GigaVUE V Series Nodes using GigaVUE-FM

You can add or remove nodes in an existing Monitoring Domain directly from GigaVUE-FM if you deployed GigaVUE V Series Nodes using GigaVUE-FM.

NOTE: Increasing or Decreasing the number of nodes in a cluster is only applicable when using Clustered deployment.

Refer to the following topics for instructions:

- [Add GigaVUE V Series Nodes to Existing Monitoring Domain](#)
- [Decrease GigaVUE V Series Nodes from Existing Monitoring Domain](#)

Add GigaVUE V Series Nodes to Existing Monitoring Domain

To increase the number of nodes in a Monitoring Domain,

1. On the Monitoring domain page, select the desired monitoring domain to add more GigaVUE V Series Nodes.
2. Select **Action>Deploy Fabric**.
3. In the **VMware Fabric Deployment** page, perform the following:
 - a. Enter the details as mentioned in [Increase or Decrease GigaVUE V Series Node](#)



- Ensure the Deployment type is set to Clustered to have multiple deployment on the same cluster.
- A cluster can have only one Host Based Deployment. However, you can deploy multiple clustered instances on the same cluster.

- b. In the **Deployment Count** field, enter the number of GigaVUE V Series Nodes you wish to add.
4. Select **Deploy**.

GigaVUE-FM adds the new nodes to the selected Monitoring Domain. The nodes appear under a new Deployment Name within the domain.

Decrease GigaVUE V Series Nodes from Existing Monitoring Domain

To decrease the number of nodes,

1. On the Monitoring domain page, perform one of the following:
 - Select the desired **Deployment** to remove
 - Select the entire Monitoring Domain to remove all the related deployments.

NOTE: To select the deployment, you can use the check-box on the left side or select the deployment name.

2. Select **Action>Delete Deployment**.

GigaVUE-FM deletes all nodes associated with the selected deployment. The node count in the Monitoring Domain decreases accordingly.

Example use-case for Increase or Decrease GigaVUE V Series Nodes using GigaVUE-FM

This feature is helpful when migrating from a GigaVUE-VM visibility solution to a GigaVUE V Series solution. Instead of undeploying and redeploying the Monitoring Domain each time you need to scale, simply add or remove GigaVUE V Series Nodes within the existing domain.

Increase or Decrease GigaVUE V Series Nodes using VMware NSX-T Manager

You can add or remove nodes in an existing Monitoring Domain directly from VMware NSX-T Manager if you deployed GigaVUE V Series Nodes using VMware NSX-T Manager.

Refer to the following topics for instructions:

- [Add V Series Nodes to Existing Monitoring Domain](#)
- [Decrease V Series Nodes from Existing Monitoring Domain](#)

Add V Series Nodes to Existing Monitoring Domain

To increase the number of nodes in a Monitoring Domain,

1. In the VMware NSX-T manager, go to the **Service Deployment** page and select **Deployment**.
This page lists the service deployments that are already deployed.
2. Select **Deploy Service**.
For more details on how to deploy a service, refer to [Deploy a Partner Service](#).
3. Enter the same details as given for the service mapped to the existing monitoring domain in GigaVUE-FM to which you wish to add more nodes.
4. In the **Clustered Deployment Count**, enter the number of nodes you wish to add to the existing monitoring domain.
5. Select **Save**.

When the Service deployment is successful and the nodes are deployed, you can view the nodes on the monitoring domain page of GigaVUE-FM.

Example: Consider a scenario where the monitoring domain in GigaVUE-FM has two V Series Nodes.

To increase the number of nodes in this monitoring domain, go to VMware NSX-T Manager and create a new service using the steps mentioned above. Then, the number of V Series Nodes in the monitoring domain in GigaVUE-FM goes up by the number you have mentioned in **Clustered Deployment Count** column in the VMware NSX-T.

Decrease V Series Nodes from Existing Monitoring Domain

To decrease the number of nodes in an existing monitoring domain,

1. In the VMware NSX-T manager, go to the **Service Deployment** page and select **Deployment**.
The service deployment page lists the service deployments that are already deployed. .
2. Select the service deployment that you want to delete.

The GigaVUE V Series Nodes in that service deployment are deleted from the host. These GigaVUE V Series Nodes are also removed from the monitoring domain in the GigaVUE-FM.

Example: Consider a scenario where the monitoring domain in GigaVUE-FM has five V Series Nodes.

To reduce the number of nodes in this monitoring domain, go to VMware NSX-T Manager and delete a service deployment using the steps mentioned above. Then, the number of V Series Nodes in the monitoring domain in GigaVUE-FM goes down by the number you have mentioned in **Clustered Deployment Count** column of the service you have deleted.

Sharing the Same Host across Different Monitoring Domains

GigaVUE-FM enables you to share a host between VMware ESXi and VMware NSX-T monitoring domains. You can deploy multiple V Series nodes from VMware NSX-T monitoring domain and one V Series Node from VMware ESXi monitoring domain on the same host.


As a result, you can monitor the workload of virtual machines using the following two options:

- Connected to NSX segments using the V Series nodes deployed in NSX-T monitoring domain.
- Connected to regular VSS / VDS networks using the V Series node deployed in the ESXi monitoring domain.

NOTE: GigaVUE-FM cannot provide visibility in the ESXi platform to a Virtual Machine with NICs attached to both VMware NSX-T segments and ESXi VDS or VSS port groups.

Deployment Overview

GigaVUE Cloud Suite for VMware NSX-T deployment follows a set of steps that include planning and preparation before the actual solution deployment. This section provides an overview of the activities associated with GigaVUE Cloud Suite for VMware NSX-T deployment. The top-level deployment activities are:



Deployment Planning

This section focuses on the key decisions and information you must gather before deploying GigaVUE Cloud Suite for VMware NSX-T. It assumes you understand the Cloud Suite components, VMware NSX-T architecture, and the supported deployment workflows for V Series nodes.

Refer to the following sections for more detailed information:

- [Identify the Deployment Method](#)
- [Identify the Supported Platform Versions](#)
- [Identify the Traffic Visibility Scope](#)
- [Identify the Host, Cluster, and Transport Zone Requirements](#)
- [Identify Access, Roles, and Connectivity](#)
- [Identify the Image Source and Certificate Strategy](#)
- [Identify the Node Deployment Model and Scale](#)
- [Identify Monitoring Session and Traffic Forwarding Requirements](#)

Identify the Deployment Method

GigaVUE Cloud Suite for VMware NSX-T supports two deployment methods for GigaVUE V Series nodes: deployment through GigaVUE-FM or deployment through VMware NSX-T Manager.

Choose GigaVUE-FM when you want Fabric Manager to establish the NSX-T connection and launch the fabric directly from the VMware Fabric Launch Configuration page.

Choose VMware NSX-T Manager when you prefer to create the monitoring domain in GigaVUE-FM first and then deploy the V Series nodes from VMware NSX-T Manager by selecting the partner service that matches the monitoring domain.

Identify the Supported Platform Versions

Before deployment, confirm that your vCenter, ESXi, NSX-T, and GigaVUE-FM versions match the supported combinations listed in the guide. Use only supported VMware NSX-T environments, and verify the target release against the tested NSX-T, vCenter, and ESXi matrix before creating the monitoring domain or deploying nodes. For more information on supported hypervisors, refer to [Supported Hypervisors for VMware NSX-T](#).

Identify the Traffic Visibility Scope

Define which workload virtual machines require monitoring and whether the visibility use case is based on NSX-T service insertion or customer-orchestrated traffic delivery.

GigaVUE-FM supports service insertion only on overlay transport zones associated with east-west traffic. It does not support service insertion on VLAN transport zones for north-south traffic or on NSX-T Managers running in federation mode.

If your environment uses NSX-T federation or other restrictions that prevent service insertion, plan to acquire traffic by tunneling it directly to the GigaVUE V Series node by using Customer Orchestrated Source.

Identify the Host, Cluster, and Transport Zone Requirements

Prepare the ESXi hosts as NSX-T transport nodes and connect the hosts that run monitored workload VMs to an overlay transport zone.

Create the service segment in NSX-T on the overlay transport zone before deploying GigaVUE V Series nodes through GigaVUE-FM.

Also decide which vCenter data center, cluster, datastore, management network, and tunnel network will be used for the V Series deployment, because these values are required on the VMware Fabric Launch Configuration page.

Identify Access, Roles, and Connectivity

Plan the accounts that GigaVUE-FM uses to communicate with VMware vCenter, VMware NSX-T, and GigaVUE-FM itself during monitoring domain creation.

For vCenter, create a user with the minimum required VMware Virtual Center privileges.

For NSX-T, assign roles based on the deployment method and NSX-T version. For GigaVUE-FM-driven deployment on NSX-T 3.2.x and 4.x.x, use NETX Partner Admin and Security Admin. For deployment directly through NSX-T Manager, the minimum required role is NETX Partner Admin, and Security Admin is required only if you want GigaVUE-FM to orchestrate deployment.

Open the required firewall paths between GigaVUE-FM, NSX-T Manager, vCenter, GigaVUE V Series nodes, administrators, and any external image server used for OVA access.

Identify the Image Source and Certificate Strategy

Decide whether the monitoring domain uses an internal image or an external image source.

If you use an external image, extract the OVA contents into OVF and VMDK files and host them on an HTTP or HTTPS server in the directory referenced by the image URL.

If you use GigaVUE-FM as the image server, upload the OVA file to GigaVUE-FM first and then select that image when creating the monitoring domain or upgrading nodes.

If your deployment requires enterprise PKI integration, plan whether to use the default GigaVUE-FM PKI model or upload your own CA-signed certificates through the private CA workflow.

Identify the Node Deployment Model and Scale

Select whether the deployment will be Clustered or Host-Based, because the deployment type affects how nodes are scaled later.

Clustered deployment is the correct choice when you want to increase or decrease the number of nodes in a cluster later by using GigaVUE-FM.

Also identify the required form factor for each V Series node. The available node sizes are Small, Medium, and Large, with different memory and vCPU allocations. For more information on form factor, refer to [Recommended Form Factor](#).

Identify Monitoring Session and Traffic Forwarding Requirements

Plan how traffic is selected, processed, and forwarded after the V Series nodes are deployed.

A monitoring session can include or exclude target instances, monitor ingress, egress, or all traffic, and automatically update when matching virtual machines are added or removed.

If NSX-T service insertion is used, also plan the NSX-T group, service chain, and policy that will redirect selected traffic to the GigaVUE Cloud Suite service.

Prerequisites for Integrating GigaVUE V Series Nodes with NSX-T

The following are the prerequisites for integrating GigaVUE V Series Nodes with VMware NSX-T:

Host and Network Requirements

- Prepare ESXi hosts as NSX-T Data Center transport nodes using transport node profiles.
- Connect ESXi hosts running the workload VMs (to be monitored) to an overlay transport zone.
- Create a service segment in the NSX-T Manager on the overlay transport zone before deploying GigaVUE V Series Nodes through GigaVUE-FM. For instructions, refer to [Create a Service Segment in VMware NSX-T](#).
- Ensure that your deployment meets the supported versions of VMware vCenter, VMware ESXi, and VMware NSX-T. For version compatibility, refer to [Supported Hypervisors for VMware](#).
- Only IPv4 traffic is supported.
- If a guest VM running on an ESXi host is connected to a VLAN segment, and that ESXi host is not configured to an Overlay Transport zone, then the traffic destined to a service VM is disrupted. Such a configuration can also route traffic to a black hole.
- For more detailed VMware requirements on East-West traffic monitoring, refer to the below links and select the appropriate NSX-T version.
 - For versions 4.x.x: [NSX Requirements for East-West Traffic](#).

Service Insertion Support

- GigaVUE-FM supports service insertion only on overlay transport zones associated with east-west (E-W) traffic.
- GigaVUE-FM does not support service insertion on:
 - VLAN transport zones associated with north-south (N-S) traffic.
 - NSX-T Managers running in federation mode.

Software and Image Requirements

- Ensure your ESXi hosts meet the resource requirements for deploying GigaVUE V Series Nodes. For details, refer to [Prerequisites for Integrating GigaVUE V Series Nodes with NSX-T](#).
- Download the GigaVUE V Series Node OVA image file from the [Gigamon Customer Portal](#).

NOTE: If you use the **External Image** option in the monitoring domain, host the OVF and VMDK files (extracted from the OVA) on an external HTTP(S) server. For information, refer to [Create Monitoring Domain for VMware NSX-T](#).

Unsupported Configurations when using VMware NSX-T:

- Service Insertion:
 - Not supported on Global NSX-T managers in federation mode. Use Local NSX-T Managers for deploying our solution.
 - Not supported on Multi tenancy environments.
- Multiple monitoring domains: Cannot configure multiple monitoring domains with the same NSX-T manager.

Refer to the following topics:

- [Network Firewall Requirements](#)
- [Recommended Form Factor \(Instance Types\)](#)
- [Required VMware Virtual Center Privilege](#)
- [Required Roles in VMware NSX-T](#)
- [Disable Certification Validation in VMware NSX-T](#)
- [Default Login Credentials](#)

Licensing for GigaVUE Cloud Suite for VMware NSX-T

GigaVUE Cloud Suite for VMware (NSX-T) requires a valid license to deploy and operate the GigaVUE V Series Nodes and to run applications on them. Licenses are managed centrally through GigaVUE-FM, which tracks usage across all connected V Series Nodes and the applications running on them.

GigaVUE Cloud Suite uses Volume-Based Licensing (VBL), available as a monthly subscription. In the VBL scheme, specific applications on the V Series Nodes are entitled to a specified amount of total data volume over the term of the license.

Refer to the following sections:

- [Default Trial License](#)
- [Volume-Based License](#)

Default Trial Licenses

After installing GigaVUE-FM, you receive a one-time, free 1TB SecureVUE Plus trial Volume-Based License (VBL) for 60 days, starting from the installation date.

This license includes the following applications:

- ERSPAN
- GENEVE
- Slicing
- Masking
- Trailer
- Tunneling
- Load Balancing
- Enhanced Load Balancing
- Flow map
- Header Stripping
- Header Addition
- De-duplication
- NetFlow
- Application Packet Filtering
- Application Filtering Intelligence
- Application Metadata Intelligence
- Application Metadata Exporter
- Inline SSL
- SSL Decrypt
- Precryption

NOTE: If you do not have any other volume-based licenses installed, the deployed monitoring sessions are undeployed from the existing GigaVUE V Series Nodes after 60 days at the expiration of the trial license.

When you install a new Volume-Based License (VBL), the existing trial license remains active alongside the new VBL. When the trial license period expires, it is automatically deactivated. After deactivation, the trial license moves to the Inactive tab on the VBL page.

License Requirements for Monitoring Sessions

You must use one of the following GigaVUE Cloud Suite™ licenses to deploy any supported application in a Monitoring Session:

- CoreVUE Base Bundle (minimum required)

- NetVUE
- SecureVUE Plus

The default GigaVUE-FM licenses do not provide entitlement for deploying Monitoring Session applications.

The default trial VBL license (SecureVUE Plus) will support deploying applications in Monitoring Sessions only for the duration of the trial license.

Volume-Based License

All the GigaVUE V Series Nodes connected to GigaVUE-FM periodically report statistics on the amount of traffic that flows through the V Series Nodes. The statistics reflect the data volume flowing through the V Series Nodes, with the usage statistics of all licensed applications that run on these nodes.

GigaVUE Cloud Suite uses volume-based licensing (VBL), available as monthly subscription licenses. In the Volume-based Licensing (VBL) scheme, specific applications on the V Series Nodes are entitled to a specified amount of total data volume over the term of the license.

Distributing the license to individual nodes becomes irrelevant for Gigamon accounting purposes. GigaVUE-FM monitors overall consumption across all nodes and tracks individual application usage and overages.

Related Information

- [Contact Sales](#): For purchasing licenses with the Volume-Based License (VBL) option.
- For more information, refer to the Data Sheet for the required GigaVUE Cloud Suite.

Base Bundles

In volume-based licensing scheme, licenses are offered as bundles. The following three base bundle types are available:

- CoreVUE
- NetVUE
- SecureVUE Plus

The bundles are available as SKUs¹. The SKUs are named such that the number indicates the total volume allowance of the SKU for that base bundle. For example, VBL-250T-BN-CORE indicates a daily volume allowance of 250 Terabytes (250T) for the CoreVUE bundle.

The features supported for base bundle licenses are given in the following table:

¹Stock Keeping Unit. Refer to the [What is a License SKU?](#) section in the FAQs for Licenses chapter.

Feature	CoreVUE	NetVUE	SecureVUE Plus
Slicing	√	√	√
Masking	√	√	√
Advanced Load Balancing	√	√	√
Advanced Tunneling	√	√	√
Deduplication	-	√	√
NetFlow Generation	-	√	√
Adaptive Packet Filtering (APF)	-	-	√
Application Filtering Intelligence (AFI)	-	-	√
Application Metadata Intelligence (AMI)	-	-	√
Application Metadata Exporter (AMX)	-	-	√
Inline Decryption	-	-	√
OOB Decryption	-	-	√
Preencryption	-	-	√
Gigamon Enriched Metadata	-	-	√

Bundle Replacement Policy

Refer to the following notes:

- You can only upgrade to a higher bundle.
- You cannot have two different base bundles at the same time. However, you can have multiple base bundles of the same type.
- As soon as you upgrade to a higher bundle, the existing lower bundles are automatically deactivated.

Add-on Packages

GigaVUE-FM allows you to add add-on packages to the base bundles. These add-on packages allow you to add additional applications to your base bundles. Add-on packages have their own start/end date and volume specifications.

The following add-on SKUs are available:

- VBL-50T-ADD-5GC
- VBL-250T-ADD-5GC
- VBL-2500T-ADD-5GC
- VBL-25KT-ADD-5GC

Rules for add-on packages:

- An active base bundle is required to use an Add-on package.
- Your base bundle limits the total volume usage of the add-on package in the following ways:
 - If the volume allowance of your add-on package is less than the base bundle, then your add-on package can only handle the volume allocated for the add-on package.
 - When the life term of an add-on package extends beyond the base bundle, and the base bundle expires, the add-on package's volume allowance is reduced to zero until you add a new base bundle.
 - The total volume is cumulative when multiple base bundles of the same type are active within the same time interval.

For more information about SKUs, refer to the respective Data Sheets as follows:

- [GigaVUE Cloud Suite for VMware Data Sheet](#)
- [GigaVUE Cloud Suite for AWS Data Sheet](#)
- [GigaVUE Cloud Suite for Azure Data Sheet](#)
- [GigaVUE Cloud Suite for OpenStack](#)
- [GigaVUE Cloud Suite for Nutanix](#)

How GigaVUE-FM Tracks Volume-Based License Usage

GigaVUE-FM applies the following methods to track the license usage for each GigaVUE V Series Node:

- When you create and deploy a monitoring session, GigaVUE-FM allows you to use only applications with active licenses.
- When a license expires, you are notified with an audit log. For more information, refer to the *About Audit Logs* section in the respective GigaVUE Cloud Suite Deployment Guide.
- When a license expires (and has not been renewed yet), the monitoring sessions using the corresponding license are not undeployed.
- For releases prior to 6.4:
 - The Monitoring Sessions using the corresponding license are undeployed, but not deleted from the database.
 - Any undeployed monitoring sessions are redeployed when you renew a license or newly import the same.

NOTE: GigaVUE-FM displays a notification on the screen when the license expires.

Network Firewall Requirements

Following are the Network Firewall Requirements for GigaVUE V Series Node deployment.

Source	Destination	Source Port	Destination Port	Protocol	Service	Purpose
GigaVUE-FM	NSX-T Manager	Any (1024-65535)	443	TCP	https	Allows GigaVUE-FM to communicate with vCenter and NSX-T.
	vCenter					
GigaVUE-FM	GigaVUE V Series Node	Any (1024-65535)	8889	TCP	Custom API	Allows GigaVUE-FM to communicate with GigaVUE V Series Node
GigaVUE-FM	GigaVUE V Series Nodes	Any (1024-65535)	80	TCP	Custom TCP	Allows GigaVUE-FM to send ACME challenge requests to GigaVUE V Series Node.
Administrator	GigaVUE-FM	Any (1024-65535)	443	TCP	https	Management connection to GigaVUE-FM
			22		ssh	
Administrator	GigaVUE V Series Nodes	Not Applicable	22		ssh	Troubleshooting GigaVUE V Series Nodes.
GigaVUE-FM	GigaVUE V Series Node	Any (1024-65535)	5671	TCP	Custom TCP	Allows GigaVUE-FM to receive the traffic health updates with GigaVUE V Series Node

Remote Source	GigaVUE V Series Node	Custom Port (VXLAN and UDPGRE),N/A for GRE	4789	UDP	VXLAN	Allows to UDPGRE Tunnel to communicate and tunnel traffic to GigaVUE V Series Nodes (Applicable for Tunnel Ingress option only)
			N/A	IP 47	GRE	
			4754	UDP	UDPGRE	
GigaVUE V Series Node	Tool/ GigaVUE HC Series instance	Custom Port (VXLAN),N/A for GRE	4789	UDP	VXLAN	AllowsGigaVUE V Series Node to communicate and tunnel traffic to the Tool
			N/A	IP 47	GRE	
GigaVUE V Series Node	Tool/ GigaVUE HC Series instance	N/A	N/A	ICMP	echo Request	Allows V Series node to health check tunnel destination traffic (Optional)
					echo Response	
GigaVUE V Series Node	GigaVUE-FM	Any (1024-65535)	5671	TCP	Custom TCP	Allows GigaVUE V Series Nodes to communicate the traffic health updates with GigaVUE-FM
GigaVUE V Series Nodes	GigaVUE-FM	Any (1024-65535)	9600	TCP	Custom TCP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.
GigaVUE-FM	External Image Server URL	Any (1024-65535)	Custom port on web Server	TCP	http	Access to image server to image lookup and checks, and downloading the image
NSX-T Manager						

vCenter						
NSX-T Manager	GigaVUE-FM	Any (1024-65535)	443	TCP	http	When using GigaVUE-FM as the image server for uploading the GigaVUE V Series Image.
vCenter						

Recommended Form Factor (Instance Types)

The form factor (instance type) of the GigaVUE V Series Node is configured on the OVF file and packaged as part of the OVA image file. Instance types can differ for GigaVUE V Series Nodes in different ESXi hosts. Small is the default type.

The table below lists the available form factors (instance types) based on memory and the number of vCPUs for a single GigaVUE V Series Node.

Type	Memory	vCPU	Disk space
Small	4GB	2vCPU	8GB
Medium	8GB	4 vCPU	8GB
Large	16GB	8 vCPU	8GB

Required VMware Virtual Center Privilege

This section lists the minimum privileges required for the GigaVUE-FM user in vCenter.

Category	Required Privilege	Purpose
vApp	vApp application configuration	V Series Node Deployment
Virtual machine	Interaction - Power on - Power Off	- V Series Node Deployment - Used to power on and power off GigaVUE V Series Node.

Required Roles in VMware NSX-T

This section lists the minimum roles required for the GigaVUE-FM user in VMware NSX-T.

Deploying GigaVUE V Series Node using GigaVUE-FM

When deploying GigaVUE V Series Node using GigaVUE-FM, the required user roles vary depending on the NSX-T version:

For **NSX-T version 3.2.x** and **NSX-T version 4.x.x**, select the following Roles:

- NETX Partner Admin
- Security Admin

For **NSX-T version 3.1.x**, select LDAP with any one of the following Role combinations:

- NETX Partner Admin and Security Operator
- NETX Partner Admin and Network Operator

For details, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#) section.

Deploying GigaVUE V Series Nodes using VMware NSX-T

When deploying GigaVUE V Series Node using VMware NSX-T manager, the minimum required role is NETX Partner Admin.

NOTE: Assign the NSX Security Admin role to enable GigaVUE-FM to orchestrate V Series deployment. If you skip this role, you must deploy the V Series manually in NSX Manager.

For details, refer to [Deploy GigaVUE V Series Nodes using VMware NSX-T Manager](#).

Disable Certification Validation in VMware NSX-T

If you use an uncertified GigaVUE V Series Node image, the deployment may fail due to VMware NSX-T's certificate validation requirements. To avoid this issue, disable certificate validation before deploying the node.

To disable Certificate Validation,

1. Using CLI with root credentials, log in to NSX-T manager (all three NSX-T manager when it is in a cluster).
2. In a text editor, open the file: **/config/vmware/auth/ovf_validation.properties**
3. Set a value for **THIRD_PARTY_OVFS_VALIDATION_FLAG** as **2**. The definition of the legends are as follows:
 - 0: only VMware-signed OVF's are allowed for deployment
 - 1: only VMware-signed and well-known CA-signed OVF's are allowed for deployment
 - 2: no validation
4. Save the file and exit the editor.

Default Login Credentials

Using the default credentials, you can log in to the GigaVUE V Series Node, GigaVUE V Series proxy, and UCT-V Controller.

Product	Login credentials
GigaVUE V Series Node	You can log in to the GigaVUE V Series Node by using ssh. The default username and password are: - Username: gigamon - Password: Gigamon123!

Supported Hypervisors for VMware

The table lists the supported hypervisor versions for vCenter, VMware ESXi and VMware NSX-T.

Tested Platforms- vSphere ESXi			
	vCenter Server	ESXi	GigaVUE-FM
v6.7	v6.7U3	v6.7U3	v5.10.02, v5.11.01, v5.12.00, v5.13.00, v5.13.01
v7.0	v7.0	v7.0	v5.10.02, v5.11.01, v5.12.00, v5.13.00, v5.13.01, v5.14.00, v5.15.00, v5.16.00, v6.0.00, v6.1.00
v7.0	v7.0U3	v7.0U3	v5.15.00, v5.16.00, v6.0.00, v6.1.00, v6.2.00, v6.3.00, v6.4.00, v6.5.00, v6.6.00, v6.7.00, v6.8.00, v6.9.00, v6.10.00, v6.11.00, v6.12.00
v8.0	v7.0U3	v8.0U2	v6.9.00
v8.0	v8.0	v8.0	v6.3.00, v6.4.00, v6.5.00, v6.6.00, v6.7.00, v6.8.00, v6.9.00
v8.0	v8.0U2, v8.0U3	v8.0U2, v8.0U3	v6.8.00, v6.9.00, v6.10.00, v6.11.00, v6.12.00, v6.13.00, v6.14.00

Tested Platforms- vSphere NSX-T			
	vCenter Server	ESXi	GigaVUE-FM
v3.1.0	v7.0	v7.0	v5.11.01, v5.12.00
v3.1.2	v7.0	v6.7U3, v7.0U1	v5.12.00, v5.13.00, v5.13.01
v3.1.3	v7.0	v6.7U3, v7.0U1	v5.13.01, v5.14.00, v6.0.00
v3.2.0	v7.0, v7.0U3	v6.7U3, v7.0U1, v7.0U3	v5.14.01, v5.15.00, v5.16.00, v6.0.00
v3.2.1	v7.0U3	v6.7U3, v7.0U1,	v6.0.00, v6.1.00, v6.2.00

Tested Platforms- vSphere NSX-T			
	vCenter Server	ESXi	GigaVUE-FM
		v7.0U3	
v3.2.2	v7.0U3	v7.0U3	v6.3.00, v6.4.00
v3.2.3	v7.0U3	v7.0U3	v6.5.00, v6.6.00, v6.7.00, v6.8.00, v6.9.00, v6.10.00, v6.11.00, v6.12.00, v6.12.00
v4.0.0	v7.0U3	v7.0U3	v6.0.00, v6.1.00, v6.2.00, v6.3.00
v4.1.0	v7.0U3	v7.0U3	v6.3.00, v6.4.00, v6.5.00
	v8.0U2	v8.0U2	v6.5.00, v6.6.00, v6.7.00
v4.1.2	v8.0U2, v8.0U3	v8.0U2, v8.0U3	v6.8.00, v6.9.00
v4.2	v8.0U2	v8.0U2, v8.0U3	v6.9.00, 6.10.00, v6.11.00, v6.12.00, v6.13.00, 6.14.00

Points to Note (VMware NSX-T)

- **VMware NSX-T Setup:** This guide assumes that you install and configure VMware NSX-T. For configuration details, see the [VMware Documentation](#).
- **Service Insertion Support:**
 - GigaVUE-FM supports service insertion only for overlay transport zone associated with the E-W traffic.
 - GigaVUE-FM does not support Service insertion for VLAN transport zone associated with the N-S traffic or when the VMware NSX-T manager runs in federation mode.
- **Acquiring Traffic:** You can use UCT-V to acquire traffic from workload virtual machines in NSX-T federated environments.

For deployment steps, see

Install and Upgrade GigaVUE-FM

You can install GigaVUE-FM on many supported platforms. These include both cloud and on-premises environments.

GigaVUE-FM supports flexible deployment. You can easily install or upgrade it on your chosen platform.

For details, see the following:

- **Installation:** Refer to GigaVUE-FM Installation and Upgrade Guide available in the [Gigamon Documentation Library](#).
- **Upgrade:** Refer to Upgrade GigaVUE-FM topic in GigaVUE-FM Installation and Upgrade Guide.

Before You Install

This section describes the hardware and virtual computing requirements for GigaVUE-FM. Ensure that the time on the VMware ESXi host is set correctly to ensure accuracy of the trending data that is captured.

Prerequisites for GigaVUE-FM

Before Installing GigaVUE-FM, ensure that VMware vSphere Standard, Enterprise, or Enterprise Plus is installed on hardware that meets minimum requirements. Refer [VMware ESXi Hardware Requirements](#) for minimum hardware requirements.

VMware ESXi Hardware Requirements

The following table describes the hardware requirements on which VMware ESXi runs GigaVUE-FM.

Table 1: Hardware Requirements for VMware Hypervisor

Hardware Requirements	
VMware Hypervisor	vSphere ESXi: v8.xx and above. Refer to Supported Hypervisors for VMware for more detailed information.
CPU	One or more 64-bit x86 CPUs with virtualization assist (Intel-VT or AMD-V) enabled. Note: To run GigaVUE-FM, hardware support for virtualization must be enabled on the VMware ESXi host. Make sure that the BIOS option for virtualization support is not disabled. For more information, see your BIOS documentation
RAM	At least 16GB
Disk Space	At least 84GB shared (FC, iSCSI, NFS, or FCoE) or locally attached storage (PATA, SATA, SCSI) • System Disk - 40GB • Swap Disk - 4GB • Data Disk - 40GB Note: Increase the disk space depending on the configuration load requirements for the given GigaVUE-FM instance.
Network	At least one 1Gb NIC

Recommended Resource Requirements for Scaled Environments

Refer to the following sections for large deployment scenarios of GigaVUE-FM.

For Standalone GigaVUE-FM

The following table displays the minimum recommended requirements for a standalone GigaVUE-FM:

Requirements	Basic Configuration	Medium Configuration	Large Configuration	Large Configuration (GigaVUE-FM Hardware Appliance 3)
CPU				
CPU count (minimum)	2	4	12	32
CPU Min speed (per CPU)	2.30GHz	2.30GHz	2.30GHz	2.0GHz
Memory				
Memory Size	16GB	32GB	128GB	128GB
Disk Space				
OS Disk Size (min)	40GB	40GB	40GB	40GB
Config Disk Size - 15 Days Data	80GB	620GB	1100GB	1100GB
Config Disk Size - 35 Days Data	160GB	1240GB	2200GB	2200GB
Config Disk Size - Rollup 35 Days + 120 Days Rollup	200GB	1600GB	2500GB	2500GB
** The disk usage is estimated based on the configuration load defined in the table below (Reference Values for Disk Space Estimation). You must periodically monitor the disk utilization and increase the disk size on need.				
NIC				
Number of NICs	1	1	1	1
Management/Data	1GB	1GB	1GB	1GB
Appliance Scalability				
Number of GigaVUE HC Series/GigaVUE TA Series Devices	Up to 50	Up to 500	Up to 1000	Up to 1000
Number of GigaVUE V Series Nodes	Up to 50	Up to 500	Up to 1000	Up to 1000
Number of Worker Pods monitored	1 Cluster with 100 Nodes and 10000 pods with 2 rules	3 Clusters, each with 100 Nodes and 10000 pods with 2 rules.	–	–

For GigaVUE-FM High Availability Instances

The following table displays the minimum recommended requirements for GigaVUE-FM High Availability group:

Configuration	Basic Configuration	Medium Configuration	Large Configuration	Large Configuration (GigaVUE-FM Hardware Appliance 3)
CPU				
CPU count (minimum)	2	4	12	32
CPU Min speed (per CPU)	2.30GHz	2.30GHz	2.30GHz	2.0GHz
Memory				
Memory Size	16GB	32GB	128GB	128GB
Disk Space				
OS Disk Size	40GB	40GB	40GB	40GB
Config Disk Size - 15 Days Data	80GB	620GB	1600GB	1100GB
Config Disk Size - 35 Days Data	160GB	1240GB	3000GB	3000GB
Config Disk Size - Rollup 35 Days + 120 Days Rollup	200GB	1600GB	4300GB	4300GB
** The disk usage is estimated based on the configuration load defined in the table below (Reference Values for Disk Space Estimation). You must periodically monitor the disk utilization and increase the disk size on need.				
NIC				
Number of NICs	1	1	1	1
Management/Data	1GB	1GB	1GB	1GB
Appliance Scalability				
Number of GigaVUE HC Series/GigaVUE TA Series Devices	Up to 50	Up to 500	Up to 3020	Up to 3020

Rules, Notes, and Limitations

- When deploying GigaVUE-FM in VMware ESXi, use the 'reservation' option and allocate fixed resource for CPU and memory.
- Values may change depending on the hardware used. You must ensure that the hardware you use supports the data rate specified in the table.
- Syslogs must be disabled if more than 1000 nodes are configured in GigaVUE-FM.

- Config sync time for 3000 node system on an average takes 10 minutes and on a maximum can go up to 12 minutes.

Reference Values for Disk Space Estimation

The following table defines the load in system for various system configurations that is used for qualification. Use this as a reference to estimate the disk usage and I/O operation in the system.

	Default (Basic)	Medium	Large	GigaVUE-FM High Availability Large 3000 Nodes
Devices	50	500	1000	3020
Ports	4550	45500	91000	273000
Maps	380	4750	13300	39900
Rules	1900	23750	66500	199500
GigaSMART Operations	250	1250	3500	10500
GigaSMART Groups	40	500	1400	4200
Virtual Ports	250	1250	3500	10500
Syslog Rate (per minute)	40	500	1400	4200
SNMP Traps Rate (per minute)	80	1000	2800	7200
Devices	50	500	1000	3020

Recommended Form Factor for VMware vCenter (Instance Types)

The form factor (instance type) of the GigaVUE V Series Node is configured on the OVF file and packaged as part of the OVA image file.

Instance types can differ for GigaVUE V Series Nodes in different ESXi hosts. Small is the default type.

The table below lists the available form factors (instance types) based on memory and the number of vCPUs for a single GigaVUE V Series Node.

Type	Memory	vCPU	Disk space	vNIC
Small	4GB	2 vCPU	10GB	1 Management interface, 1 Tunnel interface, and 8 vTAP interfaces
Medium	8GB	4 vCPU		
Large	16GB	8 vCPU		

NOTE: You can contact your account manager or Gigamon Technical Support to identify a suitable form factor.

Deploy OVA file on VMware vCenter

The GigaVUE-FM software package is distributed as an OVA file.

Use the vSphere Client to install the GigaVUE-FM OVA file. Starting from software version 5.3, you cannot deploy GigaVUE-FM directly from the ESXi host. You must login to the vCenter on the vSphere client to deploy a GigaVUE-FM instance.

IMPORTANT NOTE:



- The OVA file must be stored in a location that is accessible to the vSphere Client. This location cannot be a datastore accessible to the ESXi host which will be the target of the deployment.
- GigaVUE-FM v5.8 is used for illustration purposes only throughout this example. In your deployment, select the version that you are deploying.

The following the steps given below to deploy a GigaVUE-FM instance:

1. Log in to the VMware vCenter web interface.
2. In the vSphere Client, select an inventory object that is a valid parent object of a virtual machine such as data center, cluster, or ESXi host.
3. Right-click the ESXi Host, Cluster, or data center on which you want to deploy GigaVUE-FM 6.14, and then select **Deploy OVF Template**. The Deploy OVF Template wizard opens.
4. In the **Select an OVF template** page of the **Deploy OVF Template** wizard, choose one of the following options:
 - a. **URL**—Enter the URL from where you want to download and install the OVF package.
 - b. **Local file**—Click **Browse** to navigate to the OVA file available on your local machine, and then select the OVA file.
5. Click the **Next**. The **Select a name and folder** page of the **Deploy OVF Template** wizard appears.
6. Specify a unique name for the GigaVUE-FM instance, and then select a location and host to which you want to deploy the GigaVUE-FM instance.
7. Click **Next**. The **Select a compute resource** page of the **Deploy OVF Template** wizard appears.
8. Select a destination compute host for the OVF deployment. The Deploy OVF Template wizard performs a validation to ensure that the selected host has all the resources required for the GigaVUE-FM deployment.
9. Click **Next**. The **Review details** page of the **Deploy OVF Template** wizard appears.

10. Verify the OVF template details, and then click **Next**. The **Select storage** page of the **Deploy OVF Template** wizard appears.
11. From the **Select virtual disk format** drop-down list, select **Thick Provisioning** as the format for the virtual disks and provisioning.

NOTE: You must deploy GigaVUE-FM using **Thick Provisioning**. Any other choice results in GigaVUE-FM not working correctly.

12. Select the datastore where the virtual machine's files will be stored.
13. Click **Next**. The **Select networks** page of the **Deploy OVF Template** wizard appears.
14. Set network mapping by doing either of the following, depending on how you are deploying, and then click **Next**:
 - **If you are not deploying on a standalone ESXi host**, the **Network Mapping** displays under **Source Networks**. Use the drop-down lists to assign the correct **Destination Network** to the source network.
 - **If you are deploying GigaVUE-FM on a standalone ESXi host**, the network mapping is set automatically by assigning the destination network to the any destination Network. In case of multiple port groups, you need to manually assign the destination network to the any destination network VM Network.
15. For the **IP protocol** drop-down list, ensure to maintain the default settings intact. Do not make any changes to the **IP protocol** drop-down list, as any changes may cause problems for vCenter.
16. Click **Next**. The **Customize template** page of the **Deploy OVF Template** wizard appears.
17. In the **System Hostname** field, enter the hostname of the GigaVUE-FM instance.
18. Set the Administrative Login Password. The password must be a minimum of 11 characters and a maximum of 30 characters. The characters must include at least:
 - One numerical character
 - One upper case character
 - One lower case character
 - One special character (Example- !@#\$%^&*)

NOTE: The password that you set here is applicable only to the GigaVUE-FM Shell. To access the GigaVUE-FM UI, the default username is admin and password is admin123A!@. You can change the password when you first login to the GigaVUE-FM UI.

If the password does not meet the complexity requirements:

- The system displays an error message in the GigaVUE-FM UI.
- The parameters configured above are not passed on to the GigaVUE-FM except the host name parameter.

- o The static IP addresses that were configured originally will be unavailable after GigaVUE-FM is deployed.
19. (Optional) Set the Administrative Login Public Key. Provide the SSH public key associated with the UCT-V controller instance to configure administrative access.
 20. (Optional) Complete the IP Networking configurations such as the default gateway for the management port, IP address of the DNS server, and the domain name for this GigaVUE-FM instance.
 21. (Optional) Complete the NTP configurations such as the IP address and version of the NTP server, and then click **Next**.

Deploy OVF Template

- Select an OVF template
- Select a name and folder
- Select a compute resource
- Review details
- Select storage
- Select networks
- Customize template**
- Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values

System	3 settings
Hostname	This is the Hostname of this UCT-V Controller instance. Hostnames may only contain letters, numbers and hyphens, but may not begin with a hyphen. uctv-ctrlr
Administrative Login Password	This is the Password for this UCT-V Controller instance. Password Confirm Password
Administrative Login Public Key	This is the SSH user public Key for this UCT-V Controller instance.

CANCEL
BACK
NEXT

Figure 1 vSphere Client: Customize Template

NOTE: You can also choose to perform the IP Networking and NTP configurations by running the **fmctl jump-start** command after you power on the GigaVUE-FM instance at the successful completion of the OVF template deployment. For details, refer to [Perform Network Configurations](#).

The **Ready to Complete** page of the **Deploy OVF Template** wizard appears.

22. Verify that all of the settings are correct, and then click **Finish**. The Recent Tasks pane of the vSphere Client Home page shows the progress of the deployment operation. When the operation completes, you have successfully deployed a GigaVUE-FM instance.

Important: Clear the browser cache before logging in to GigaVUE-FM.

Troubleshooting

1. What should be done when GigaVUE-FM does not load on ESXi 7.0?

When GigaVUE-FM 6.8 does not load due to the missing of AVX CPU flags on ESXi 7.0, perform the following steps to bring up GigaVUE-FM:

1. Deploy GigaVUE-FM 6.8.00 on VMware vCenter 7.0 using OVA.

NOTE: After deployment, don't power on the GigaVUE-FM until the following steps are verified.

2. Check the **Compatibility** version in the **Summary** tab of installed GigaVUE-FM; it should be greater than 13. If it is 13 or less than 13, do the following:
 - a. Click **Actions > Compatibility > Upgrade VM Compatibility** to upgrade the VM compatibility.
The **VM Compatibility Upgrade** message box appears.
 - b. Click **Yes** to view the **Configure VM Compatibility** dialog box.
 - c. Select the compatible version as **ESXi 7.0 and later** from the drop-down list box and click **OK**.
3. Enable EVC mode. To enable EVC mode for GigaVUE-FM, do the following:
 - a. Go to **Configure > VMware EVC > EDIT**. You can view the **Change EVC Mode** screen.
 - b. Click the radio button to **Enable EVC for Intel Hosts**.
 - c. Select **Intel Ivy Bridge Generation** from the drop-down list in the **CPU Mode** field.
 - d. Press **OK**.
4. Power on GigaVUE-FM and verify whether GigaVUE-FM loads without any issues.

NOTE: Alternatively, instead of using ESXi 7.0, you can upgrade using ESXi 7.2 or 7.3.

Install GigaVUE-FM using ISO Image File on VMware ESXi Host

You can install GigaVUE-FM on VMware ESXi host using the ISO image file.

Pre-requisites

- Download the GigaVUE-FM ISO image file from the Software and Release Notes page on the [VUE Gigamon Community portal](#) to your local.
- Ensure you have an ESXi host environment that meets CPU, memory, and storage requirements. Refer to *Before you Install* in GigaVUE-FM Installation and Upgrade Guide for detailed information.

Installation Steps

Follow the steps given below to install GigaVUE-FM using ISO image on ESXi host:

1. Log in to the VMware ESXi host web interface.
2. In the ESXi Host Client, go to **Storage > Datastore browser** for a datastore that is accessible to the host.
3. Click **Upload** and select the GigaVUE-FM ISO image from your local system to a suitable folder on the datastore.
4. After uploading the ISO image file, go to the **Virtual Machines** and select **Create/Register VM**. The New Virtual Machine wizard appears.
5. In the **Select creation type** page, select **Create a new virtual machine** and click **Next**.
6. In the **Select a name and guest OS** page,
 - a. Enter a name for the virtual machine.
 - b. Select the **Guest OS family** as Linux.
 - c. Select the **Guest OS version** as Other Linux (64-bit).
 - d. Click **Next**.
7. In the **Select storage** page, choose the datastore where the VM files will reside and click **Next**.

8. In the **Customize settings** page,
 - a. On the **Virtual Hardware** tab, enter the following details:
 - CPU: 2
 - Memory: 16GB
 - Hard disk 1: 40GB
 - Click **Add hard disk** to add two more disks with 4GB and 40GB.

NOTE: Increase the disk space depending upon the configuration load requirements for the given GigaVUE-FM instance.

- b. In the **CD/DVD Drive**, select **Datastore ISO file**. The Datastore browser appears.
 - c. Select the ISO image uploaded in Step 3.

NOTE: By default, the **Connect** checkbox for both **Network Adapter** and **CD/DVD Drive** is enabled.

- d. On the **VM Option** tab, expand **Boot Options**, and under Firmware, select **EFI**.
 - e. Click **Next**.

NOTE: If you leave the firmware as BIOS, GigaVUE-FM attempts to boot in BIOS mode, and the boot process fails .

9. Review all settings on the **Ready to complete** page and click **Finish** to create the virtual machine.
10. Select the newly created VM and click **Power on**.
11. Open the VM console to monitor the boot process. It may take approximately 15–20 minutes for the VM to boot.

What to do Next

1. **Log in with default credentials:** In the VM console, enter the default credentials:
 - Username: admin
 - Password: admin123A!!
2. **Change the default password:** On first login, you will be asked to change the password. Enter the current password (admin123A!!), then set and confirm a new one.
3. **Accessing the Shell Bash in GigaVUE-FM :**After the password is successfully changed, no reboot is required. You are logged in to the fmctl command prompt instead of a standard bash shell. If you need to run other Linux commands, you can switch to the (admin) bash shell by running the shell command in the fmctl command prompt. For the list of fmctl commands, refer to *GigaVUE-FM CLI Commands* in GigaVUE-OS CLI Reference Guide.

4. **(Optional):** If you want to change any of the network configurations or increase memory refer to the instructions in *Perform Initial Configuration* in GigaVUE-FM Installation and Upgrade Guide.

Install GigaVUE-FM using ISO Image File on VMware vCenter

You can install GigaVUE-FM on VMware vCenter using the ISO image file. Follow the steps given below to install GigaVUE-FM using ISO image on vCenter:

1. Log in to the VMware vCenter web interface.
2. In the vSphere Client, select an inventory object that is a valid parent object of a virtual machine such as data center, cluster, or ESXi host.
3. Right-click the ESXi Host, Cluster, or data center on which you want to deploy GigaVUE-FM6.14, and then select **Files** in the top navigation bar. Then click on **Upload Files** and then upload the ISO image file.
4. Click on the **Actions** button on the top and select **New Virtual Machine**. The New Virtual Machine wizard opens.
5. In the **Select a creation type** tab, select **Create a new virtual machine**. Then, click **Next**.
6. In the **Select a name and folder** tab, specify a unique name for the GigaVUE-FM instance, and then select a location and host to which you want to deploy the GigaVUE-FM instance.
7. Click **Next**. The **Select a compute resource** page of the **New Virtual Machine** wizard appears.
8. Select a destination compute host for the virtual machine deployment. The New Virtual Machine wizard performs a validation to ensure that the selected host has all the resources required for the GigaVUE-FM deployment.
9. Click **Next**. The **Select storage** page of the **New Virtual Machine** wizard appears. Select the datastore where the virtual machine's files will be stored. Then, Click **Next**.
10. The **Select compatibility** page of the **New Virtual Machine** wizard appears. Select compatibility for this virtual machine depending on the ESXi hosts used in your environment. Then, Click **Next**.

11. The **Select a guest OS** page of the **New Virtual Machine** wizard appears.
 - a. Select the **Guest OS family** as Linux.
 - b. Select the **Guest OS version** as Rocky Linux 8.8 (64-bit).
 - c. Then, click **Next**.

NOTE: GuestOS version is displayed as RockyLinux only if the vSphere version is either 6.7 U3p or 7.0 U3c. For other versions of vSphere, CentOS 4/5 is only displayed. Use the drop-down to select RockyLinux.

12. The **Customize hardware** page of the **New Virtual Machine** wizard appears, enter the CPU, Memory and Hard disk details. Refer to Minimum Requirements for more detailed information on the minimum requirements required to install GigaVUE-FM.
 - a. Go to **Add New Device > Hard disk** to more than one hard disk. Add a secondary hard disk with at least 40GB memory.
 - b. From the **Disk Provisioning** drop-down list, select **Thick Provisioning** as the format for the virtual disks and provisioning.

NOTE: You must deploy GigaVUE-FM using **Thick Provisioning**. Any other choice results in GigaVUE-FM not working correctly.

- c. Select a **Network Adapter**.
 - d. In the **CD/DVD Drive**, select **Datastore ISO file** and the **CD/DVD Media** select the ISO image uploaded in Step 3.
 - e. Click Next.
13. The **Ready to Complete** page of the **New Virtual Machine** wizard appears.
14. Verify that all of the settings are correct, and then click **Finish**. The Recent Tasks pane of the vSphere Client Home page shows the progress of the deployment operation.
15. When the operation is complete, you have successfully deployed a GigaVUE-FM instance.

Important: Clear the browser canche before logging in to GigaVUE-FM.

Initial GigaVUE-FM Configuration

After you have deployed a new GigaVUE-FM instance, you can choose to perform some initial configuration before you start using GigaVUE-FM. This procedure only needs to be performed once for each GigaVUE-FM instance deployed.

NOTE: Use care when Shutting Down or Rebooting a GigaVUE-FM. **Never** directly Power-Off the virtual machine. In VMware ESXi environment when using vSphere client, ALWAYS use **Shut Down Guest OS** or **Restart Guest** functions from VMware. Access is available from either the FILE menu or from the appropriate buttons on the GigaVUE-FM console. Failure to follow these steps may lead to database corruption issues.

Perform Initial Configuration

You can optionally perform the following configurations:

- [Perform Network Configurations](#)
- [Add Additional vNIC, vCPUs, and Increase Memory](#)

Perform Network Configurations

To perform the initial configuration:

1. Log in to GigaVUE-FM using one of the following methods:
 - If you have a IP address set up for the GigaVUE-FM, log into GigaVUE-FM directly using the IP address.
 - a. Log into GigaVUE-FM directly using the IP address.
 - b. Enter the following default login Credentials:
 - Username: admin
 - Password: admin123A!!
 - c. On first login, you will be asked to change the password. Enter the current password, then set and confirm a new one.
 - If you do not have a IP address set up for the GigaVUE-FM, complete the following steps:
 - a. Log in to the VMware vCenter web interface.
 - b. Select the required GigaVUE-FM virtual machine, and click **Console**.
 - c. After the password is successfully changed, no reboot is required. You are logged in to the fmctl command prompt instead of a standard bash shell. If you need to run other Linux commands, you can switch to the (admin) bash shell by running the **shell** command in the fmctl command prompt. For the list of fmctl commands, refer to [GigaVUE-FM CLI Commands](#).
2. Perform the following steps to complete the initial configuration:

NOTE: In the VM console, if you are in (admin) bash shell command prompt, you can run both fmctl and bash shell commands. To run fmctl commands, prefix every command given below with fmctl.

- a. Provide a unique hostname for GigaVUE-FM. Note that the hostname may contain letters, numbers, periods (.), and hyphens (-), but may not begin with a hyphen. No other special characters are permitted.

Example: set hostname <host-name>

- b. Enter the required domain name.

Example: set domain <domain-name>

- c. Decide whether to use DHCP or STATIC for the management interface. Choose one of the following options:

- If you want to set DHCP the **IPv4 Address** or **IPv6 Address, CIDR, Gateway, DNS Servers**, and **Search Domains** are dynamically set up for the management interface.

Example: set ip dhcp

- If you want to set static ip for GigaVUE-FM, you must provide the following:
 - i. **IPv4 Address** or **IPv6 Address**, and the respective **CIDR** and **Gateway**.

Example: set ip static<IP address>/<mask> <gateway>

Example: set ip6 static<IPv6 address>/<mask> <gateway>

- ii. In the **DNS Servers** field, provide the address of any additional name servers required. Enter the set of IP addresses with spaces in between.

Example: set nameservers <IP address>

- iii. In the **Search Domains** field, provide the required DNS domains with spaces in between.

Example: set searchdomain <domain-name>

- d. You can choose to either enable or disable NTP using fmctl command.

Example: set ntp enable 0.pool.ntp.org, 1.pool.ntp.org, 2.pool.ntp.org

- e. Reboot your GigaVUE-FM virtual machine.

You can now access GigaVUE-FM by opening a browser and entering its IP address (the IP address you specified).

NOTE: You must reboot the GigaVUE-FM, when there is a change in the IP address of the GigaVUE-FM.

Add Additional vNIC, vCPUs, and Increase Memory

Gigamon allows you to configure GigaVUE-FM with two network interfaces—eth0 and eth1. The network interface eth0 can be configured to connect to a network used to manage Gigamon devices. The other network interface eth1 can be configured to connect to a network hosting different servers like SMTP server, Archive server, and so on.

Follow the steps given below to add Additional vNIC, vCPUs, and increase memory:

1. Right-click the GigaVUE-FM instance and select **Edit Settings...** . The **Edit Settings** dialog box opens.
2. Click on the **Virtual Hardware** tab, edit the memory and CPU details to increase the memory and add vCPUs.

Edit Settings | FM-6.0.00-Latest

Virtual Hardware | VM Options

ADD NEW DEVICE ▾

> CPU	2 ▾					
> Memory	16	▾	GB	▾		
> Hard disk 1	40		GB	▾		
> Hard disk 2	4		GB	▾		
> Hard disk 3	40		GB	▾		
> SCSI controller 0	VMware Paravirtual					
> Network adapter 1	VM Network ▾				☒ Connected	
> New Network *	VM Network ▾				☒ Connected	
> CD/DVD drive 1	Datastore ISO File ▾				☒ Connected	

CANCEL OK

3. Click on the Add New Device drop-down in the top right corner.
4. Select Network Adapter under Network. A New device (additional vNIC) is added to GigaVUE-FM.

Deploy

Deploy GigaVUE Cloud Suite for VMware (NSX-T)

This section introduces the workflow to deploy GigaVUE Cloud Suite for VMware NSX-T using GigaVUE-FM or VMware NSX-T Manager. It describes the key deployment stages required to set up the fabric components, configure traffic monitoring, and integrate the solution with the NSX-T environment for traffic visibility. Refer to the deployment options topic for more information.

- [Deployment Options for GigaVUE Cloud Suite for VMware \(NSX-T\)](#)

Deployment Options for GigaVUE Cloud Suite for VMware (NSX-T)

This section describes the available methods to configure GigaVUE Cloud Suite (NSX-T) to provide visibility into both physical and virtual traffic.

Depending on how you plan to deploy the GigaVUE V Series Nodes, you can configure the suite using two methods.

Before you begin, make sure you review the [Prerequisites for Integrating GigaVUE V Series Nodes with NSX-T](#) section. For details and the work flow, refer the following topics:

- [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#)
- [Deploy GigaVUE V Series Nodes using VMware NSX-T Manager](#)

Deploy GigaVUE V Series Nodes using GigaVUE-FM

This section provides an overview of the workflow to deploy GigaVUE V Series nodes for VMware NSX-T using GigaVUE-FM.

- [Create Users in VMware vCenter, VMware NSX-T, and GigaVUE-FM](#)
- [Upload GigaVUE V Series Node Image into GigaVUE-FM](#)
- [Create a Service Segment in VMware NSX-T](#)
- [Create Monitoring Domain for VMware NSX-T](#)
- [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#)
- [Configure and Deploy Monitoring Session](#)

Create Users in VMware vCenter, VMware NSX-T, and GigaVUE-FM

For VMware NSX-T and GigaVUE-FM to communicate, you need to create a user in VMware NSX-T Manager, VMware vCenter, and GigaVUE-FM.

NOTE: GigaVUE-FM connects to NSX-T Manager that supports TLSv1.0, TLSv1.1, and TLSv1.2.

Create User in VMware vCenter

For GigaVUE-FM to communicate with vCenter, you need to create a user with the minimum required privileges in VMware vCenter.

For details, refer to [Required VMware Virtual Center Privilege](#).

Create User in NSX-T manager

For GigaVUE-FM to communicate with NSX-T, you need to create a user with the minimum required role in NSX-T manager.

To create a user in VMware NSX-T,

1. In NSX-T, navigate to **System > Settings > User Management** and select the **User Role Assignment** tab.
2. On the **User Role Assignment** tab, select **ADD**.
3. Select the Roles based on the GigaVUE V Series Node deployment type as mentioned in [Required Roles in VMware NSX-T](#)
4. Select **Save**

A GigaVUE-FM user is created in NSX-T.

Create user in GigaVUE-FM

For VMware NSX-T Manager to communicate with GigaVUE-FM, you need to create a user admin role in GigaVUE-FM.

For details, refer to *Add Users* section in *GigaVUE Administration Guide*.

Tips: You can follow these tips to easily identify the user created for VMware NSX-T.

- In the **Name** field, enter the name of the call back user. For example, you can use NSX-T Manager Callback as the user name to help you associate this user with the NSX-T Manager.
- In the **Username** field, enter a username for the user. For example, you can use nsxv to help you remember that the user is associated with NSX-T.


The username and password created for vCenter, NSX-T Manager, and GigaVUE-FM in this section is used when creating Monitoring Domain in GigaVUE-FM. For details, refer to [Create Monitoring Domain for VMware NSX-T](#) section.

Upload GigaVUE V Series Node Image into GigaVUE-FM

You can upload your V Series Node image to GigaVUE-FM to use it as an internal image server.. This step is optional.

To upload the V Series image into GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > VMware NSX-T (V Series)**, and then click **Settings > OVA Files**.
2. The OVA Files page appears.



3. On the OVA Files page, select **Browse** to select the *gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova* file.
4. Select **Upload to Server**.

The selected OVA image file is uploaded to the GigaVUE-FM server.

NOTE: You can upload a maximum of three OVA files to GigaVUE-FM for VMware NSX-T.

Create a Service Segment in VMware NSX-T

Before you create a service segment, you must register your NSX-T details in GigaVUE-FM.

To create a service segment in VMware NSX-T:

1. On the NSX manager, go to **Security**.
2. From the left pane, select **Network Introspection**.
3. The **Network Introspection Settings** page opens.
4. From the top navigation bar, select **Service Segment**.
5. On the **Service Segment** page, select **ADD SERVICE SEGMENT**.
6. A new row appears to create a service segment.
7. Enter the name and map it to the overlay transport zone created for the VMs.
8. Select **Save**.

The service segment you create here is used as the service attachment when deploying GigaVUE V Series Nodes using GigaVUE-FM. For details, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#).

Create Monitoring Domain for VMware NSX-T

This topic describes how to create a Monitoring Domain for deploying GigaVUE V Series Nodes in the VMware NSX-T environment through GigaVUE-FM. Creating a Monitoring Domain in GigaVUE-FM allows you to establish the mandatory connection between your VMware NSX-T environment and GigaVUE-FM.



Before you begin:

- You can create multiple Monitoring Domains using a single VMware NSX-T Manager. However, each Monitoring Domain must have unique VMware vCenters associated with it.
- If a Monitoring Domain has deployed GigaVUE V Series Nodes, the **Use External Image** and **Use FM to Launch Fabric** toggle buttons are disabled.
- If the NSX-T user password changes in VMware NSX-T manager, update in GigaVUE-FM. Otherwise, the NSX-T connection fails with an authentication error.

Prerequisites:

- If you prefer to use the **Use External Image** option, extract all the contents of the OVA file into VMDK and OVF files and place them in the directory that matches the Image URL.
- If you want to use GigaVUE-FM as your image server, save the OVA files in the dedicated directory. For details, refer to [Upload GigaVUE V Series Node Image into GigaVUE-FM](#).

To create a Monitoring Domain in GigaVUE-FM for VMware NSX-T,

1. Go to **Inventory > VIRTUAL > VMware NSX-T (V Series)**.
2. Select **Monitoring Domain**. The Monitoring Domain page appears. The Monitoring Domain page has three tabs, namely Monitoring Domains, Service Deployments and Fabric.
3. On the **Monitoring Domain** tab, select **New**.
4. The **Create Monitoring Domain Configuration** page appears.

The screenshot shows the 'Create Monitoring Domain Configuration' page in the VMware NSX-T interface. The page has a dark header with navigation tabs: 'Monitoring Domains', 'Service Deployments', 'Fabric', and 'Settings'. Below the header, there's a sidebar with icons for various functions. The main area contains a form with the following fields:

- Monitoring Domain***: Text input field with placeholder 'Enter a monitoring domain name'.
- Connection Alias***: Text input field with placeholder 'Alias'.
- Virtual Center***: Text input field with placeholder 'Virtual Center'.
- Username***: Text input field with placeholder 'Username'.
- Password***: Password input field with placeholder 'Password' and an eye icon to toggle visibility.
- NSX-T Manager***: Text input field with placeholder 'IP address or hostname'.
- NSX-T Username***: Text input field with placeholder 'NSX-T Manager username'.
- NSX-T Password***: Password input field with placeholder 'NSX-T Manager password' and an eye icon to toggle visibility.
- FM Username***: Text input field with placeholder 'FM username'.
- FM Password***: Password input field with placeholder 'FM password' and an eye icon to toggle visibility.
- Use External Image Server**: Toggle switch (currently on) and an 'Image URL' text input field below it.
- Use FM to Launch Fabric**: Toggle switch (currently on).
- Use Public IP**: Toggle switch (currently off) with a 'No' label.

At the top right of the form area, there are 'Save' and 'Cancel' buttons.

5. On the **Create Monitoring Domain Configuration** page, enter or select the following:

Field	Description
Monitoring Domain	Name of the monitoring domain
Connection Alias	Name of the connection
Virtual Center	IP address or Hostname of the vCenter. NOTE: To ensure the validity of Nutanix Prism central certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. For details, refer to Trust Store.
Username	Username of the vCenter user.
Password	vCenter password to connect to the vCenter
NSX-T Manager	IP address or Hostname of your VMware NSX-T
NSX-T Username	Username of your NSX-T account.
NSX-T Password	Password of your NSX-T account.

Field	Description
FM Username	Username of your GigaVUE-FM account.
FM Password	Password of your GigaVUE-FM account.
Use External Image	Yes - Enter the web server URL of the directory where VMDK and OVF files are available, or enter the Web Server URL in the following format: <i>http(s)://<server-IP:port>/<path to where the OVF files are saved></i> or enter a valid number for the port. The default port number is 80. No - Select this option to use an internal image. To use an internal image, select the uploaded OVA files from the Select an image dropdown list.
Use FM to Launch Fabric	Enable this option to deploy GigaVUE V Series Nodes using GigaVUE-FM.
Use Public IP	Enable this to use the public IP address of the FM instance for communication from fabric nodes. This could be Public IP, Elastic IP, or Floating IP depending on the platform.

6. Select **Save**. The **VMware NSX-T Fabric Deployment** page appears. For details on how to deploy GigaVUE V Series Nodes in the **VMware NSX-T Fabric Deployment** page, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#). When creating multiple monitoring domains with the same NSX-T Manager, each monitoring domain is associated with a unique service name. You can view the service name of each monitoring domain on the **Monitoring Domains** tab.

**Notes:**

- All V Series Nodes within a single Monitoring Domain must run the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.
- Similarly, when upgrading a V Series Node, ensure that the GigaVUE-FM version is the same or higher than the V Series Node version.

You can perform the following actions in the Monitoring Domains, Service Deployments and Fabric tabs:

Actions	Description
Monitoring Domains tab	
Manage Certificates	You can use this button to perform the following actions: • Re-issue: Required to address security compromises, key changes, or configuration updates, like validity period adjustments. • Renew: Extends the expiration date and usually happens

[illegible]

What to do next:

- **Use FM to Launch Fabric** is enabled: You are navigated to the **VMware NSX-T Fabric Deployment** page.
- **Use FM to Launch Fabric** is disabled: You must deploy GigaVUE V Series Nodes using VMware NSX-T Manager.

Deploy GigaVUE V Series Nodes using GigaVUE-FM

After you create a monitoring domain in GigaVUE-FM for VMware NSX-T, GigaVUE-FM establishes a connection with the VMware NSX-T manager and launches the **VMware NSX-T Fabric Deployment** page.

For details, refer to [Create Monitoring Domain for VMware NSX-T](#).

Deploy GigaVUE V Series Node from GigaVUE-FM

Launch the **VMware NSX-T Fabric Deployment** page and deploy.

To deploy,

1. Launch the **VMware NSX-T Fabric Deployment** page using one of the following options:
 - Create a Monitoring domain and you are immediately navigated to the **VMware NSX-T Fabric Deployment** page.
 - On the **Monitoring Domain** page, go to **Inventory > VIRTUAL > VMware NSX-T**

(V Series) and select **Actions > Deploy Fabric**.

2. On the **VMware NSX-T Fabric Deployment** page, select or enter the following details:

Field	Description
Deployment Name	Name of the deployment (NSX-T service deployment)
Datacenter	vCenter Data Center with the NSX-T hosts provisioned with V Series nodes
Cluster	Cluster where you want to deploy GigaVUE V Series Nodes
Enable Custom Certificates	Enable this option to validate the custom certificate during SSL Communication. GigaVUE-FM validates the Custom certificate with the trust store. If the certificate is not available in Trust Store, communication does not happen, and a handshake error occurs. Note: If the certificate expires after the successful deployment of the fabric components, then the fabric components move to failed state.
Custom SSL Certificate Note: This field appears only when Enable Custom Certificates is enabled.	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes. For details, refer to Install Custom Certificate .
Datastore	Network datastore shared among all NSX-T hosts in a cluster.
SSL Key	Reserved for future use.
Name Server	The server that stores the mapping between the domain names and the IP address. You can enter a maximum of three name servers. Enter the valid IPv4 address, separated by comma.
Management	
Network	Management network for GigaVUE V Series Nodes
IP Type	Select the management network IP type as Static or DHCP
IP Pool Note: This field appears only when the Management IP type is Static.	Select the IP Pool
MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that you can transfer as a single entity in a network connection. Enter value between 1280 to 9000.
Tunnel	
Network	Tunnel Network for the V Series nodes
IP Type	Select the tunnel network IP address type as Static or DHCP
Gateway IP (optional)	Gateway IP address of the Tunnel Network
IP Pool Note: This field appears	Select the IP Pool

Field	Description
only when the Tunnel IP type is Static.	
MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that you can transfer as a single entity in a network connection. Enter value between 1280 to 9000.
Use IPv6	Enable to use IPv6.
User Password: (gigamon)	SSH Password for the built-in user, ' gigamon ' on the V Series node
Confirm Password	Confirm the SSH Password of the GigaVUE V Series Node
Form Factor	Instance size of the GigaVUE V Series Node. (For example, Small, Medium or Large)
Service Attachment	Service segment created in VMware NSX-T Manager. For details, refer to Create a Service Segment in VMware NSX-T .
Deployment Type	Type of GigaVUE V Series Node deployment. The options available are Clustered or Host-Based deployment type. Note: Select the deployment type as Clustered if you wish to increase or decrease the number of nodes in a cluster using GigaVUE-FM. For details, refer Deploy GigaVUE V Series Nodes using GigaVUE-FM .
Deployment Count (for Clustered deployment type)	Number of GigaVUE V Series Nodes (Service Instances) to deploy

3. Select **Deploy**.

After deployment, the GigaVUE V Series Node appears on the Monitoring Domain page under the **Deployment Name** on the selected Monitoring Domain. You can click on the deployment name on the Monitoring Domain page to select the corresponding service deployment.

<

To view the fabric launch configuration specification of a fabric component, select a GigaVUE V Series Node. A quick view of the **Fabric Launch Configuration** appears.



Points to Note:

- VMware NSX-T automatically generates the deployed GigaVUE V Series Node name. Do not change the name of the GigaVUE V Series Node in the vCenter.



- When rebooting a GigaVUE V Series Node, the existing traffic flows stops redirecting to it. However, new flows initiated after the reboot are redirected to the GigaVUE V Series Nodes.
- In a cluster-based deployment, do not use Storage vMotion to migrate a GigaVUE V Series Node to a different host or datastore. GigaVUE-FM does not automatically update the datastore information. To move a node, delete and redeploy it on the desired host and datastore.

Configure and Deploy Monitoring Session

In GigaVUE-FM, you must do the following to configure and deploy the Monitoring Session. Refer to [Configure Monitoring Session](#) section for more details.

1. **Access the Monitoring Session Page:**
 - a. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > VMware (NSX-T)**.
 - b. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
2. **Add components to the canvas:**
 - a. Drag and drop the required components to the canvas.
 - b. From the **Applications** expand menu, drag and drop the required applications to the graphical workspace.
3. **Deploy the Monitoring Session:**
 - a. From the **Actions** menu, select **Deploy**.
 - b. After successful deployment on all the V Series Nodes, the status appears as **Success** on the Monitoring Sessions **Sources** tab.

What to do Next

You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab. For more details, refer to [View Monitoring Session Statistics \(VMware NSX-T\)](#).

Deploy GigaVUE V Series Nodes using VMware NSX-T Manager

This section provides an overview of the workflow to deploy GigaVUE V Series nodes for VMware NSX-T using GigaVUE-FM.

- [Create Users in VMware vCenter, VMware NSX-T, and GigaVUE-FM](#)
- [Upload GigaVUE V Series Node Image into GigaVUE-FM](#)
- [Create a Service Segment in VMware NSX-T](#)

- [Create Monitoring Domain for VMware NSX-T](#)
- [Deploy GigaVUE V Series Nodes using VMware NSX-T Manager](#)
- [Configure and Deploy Monitoring Session](#)

Create Users in VMware vCenter, VMware NSX-T, and GigaVUE-FM

For VMware NSX-T and GigaVUE-FM to communicate, you need to create a user in VMware NSX-T Manager, VMware vCenter, and GigaVUE-FM.

NOTE: GigaVUE-FM connects to NSX-T Manager that supports TLSv1.0, TLSv1.1, and TLSv1.2.

Create User in VMware vCenter

For GigaVUE-FM to communicate with vCenter, you need to create a user with the minimum required privileges in VMware vCenter.

For details, refer to [Required VMware Virtual Center Privilege](#).

Create User in NSX-T manager

For GigaVUE-FM to communicate with NSX-T, you need to create a user with the minimum required role in NSX-T manager.

To create a user in VMware NSX-T,

1. In NSX-T, navigate to **System > Settings > User Management** and select the **User Role Assignment** tab.
2. On the **User Role Assignment** tab, select **ADD**.
3. Select the Roles based on the GigaVUE V Series Node deployment type as mentioned in [Required Roles in VMware NSX-T](#)
4. Select **Save**

A GigaVUE-FM user is created in NSX-T.

Create user in GigaVUE-FM

For VMware NSX-T Manager to communicate with GigaVUE-FM, you need to create a user admin role in GigaVUE-FM.

For details, refer to [Add Users](#) section in *GigaVUE Administration Guide*.

Tips: You can follow these tips to easily identify the user created for VMware NSX-T.

- In the **Name** field, enter the name of the call back user. For example, you can use NSX-T Manager Callback as the user name to help you associate this user with the NSX-T Manager.

- In the **Username** field, enter a username for the user. For example, you can use nsxv to help you remember that the user is associated with NSX-T.


The username and password created for vCenter, NSX-T Manager, and GigaVUE-FM in this section is used when creating Monitoring Domain in GigaVUE-FM. For details, refer to [Create Monitoring Domain for VMware NSX-T section](#).

Upload GigaVUE V Series Node Image into GigaVUE-FM

You can upload your V Series Node image to GigaVUE-FM to use it as an internal image server.. This step is optional.

To upload the V Series image into GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > VMware NSX-T (V Series)**, and then click **Settings > OVA Files**.
2. The OVA Files page appears.



3. On the OVA Files page, select **Browse** to select the *gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova* file.
4. Select **Upload to Server**.

The selected OVA image file is uploaded to the GigaVUE-FM server.

NOTE: You can upload a maximum of three OVA files to GigaVUE-FM for VMware NSX-T.

Create a Service Segment in VMware NSX-T

Before you create a service segment, you must register your NSX-T details in GigaVUE-FM.

To create a service segment in VMware NSX-T:

1. On the NSX manager, go to **Security**.
2. From the left pane, select **Network Introspection**.
3. The **Network Introspection Settings** page opens.

4. From the top navigation bar, select **Service Segment**.
5. On the **Service Segment** page, select **ADD SERVICE SEGMENT**.
6. A new row appears to create a service segment.
7. Enter the name and map it to the overlay transport zone created for the VMs.
8. Select **Save**.

The service segment you create here is used as the service attachment when deploying GigaVUE V Series Nodes using GigaVUE-FM. For details, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#).

Create Monitoring Domain for VMware NSX-T

This topic describes how to create a Monitoring Domain for deploying GigaVUE V Series Nodes in the VMware NSX-T environment through GigaVUE-FM. Creating a Monitoring Domain in GigaVUE-FM allows you to establish the mandatory connection between your VMware NSX-T environment and GigaVUE-FM.



Before you begin:

- You can create multiple Monitoring Domains using a single VMware NSX-T Manager. However, each Monitoring Domain must have unique VMware vCenters associated with it.
- If a Monitoring Domain has deployed GigaVUE V Series Nodes, the **Use External Image** and **Use FM to Launch Fabric** toggle buttons are disabled.
- If the NSX-T user password changes in VMware NSX-T manager, update in GigaVUE-FM. Otherwise, the NSX-T connection fails with an authentication error.

Prerequisites:

- If you prefer to use the **Use External Image** option, extract all the contents of the OVA file into VMDK and OVF files and place them in the directory that matches the Image URL.
- If you want to use GigaVUE-FM as your image server, save the OVA files in the dedicated directory. For details, refer to [Upload GigaVUE V Series Node Image into GigaVUE-FM](#).

To create a Monitoring Domain in GigaVUE-FM for VMware NSX-T,

1. Go to **Inventory > VIRTUAL > VMware NSX-T (V Series)**.
2. Select **Monitoring Domain**. The Monitoring Domain page appears. The Monitoring Domain page has three tabs, namely Monitoring Domains, Service Deployments and Fabric.
3. On the **Monitoring Domain** tab, select **New**.

4. The **Create Monitoring Domain Configuration** page appears.

5. On the **Create Monitoring Domain Configuration** page, enter or select the following:

Field	Description
Monitoring Domain	Name of the monitoring domain
Connection Alias	Name of the connection
Virtual Center	IP address or Hostname of the vCenter. NOTE: To ensure the validity of Nutanix Prism central certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. For details, refer to Trust Store.
Username	Username of the vCenter user.
Password	vCenter password to connect to the vCenter
NSX-T Manager	IP address or Hostname of your VMware NSX-T
NSX-T Username	Username of your NSX-T account.
NSX-T Password	Password of your NSX-T account.

Field	Description
FM Username	Username of your GigaVUE-FM account.
FM Password	Password of your GigaVUE-FM account.
Use External Image	Yes - Enter the web server URL of the directory where VMDK and OVF files are available, or enter the Web Server URL in the following format: <i>http(s)://<server-IP:port>/<path to where the OVF files are saved></i> or enter a valid number for the port. The default port number is 80. No - Select this option to use an internal image. To use an internal image, select the uploaded OVA files from the Select an image dropdown list.
Use FM to Launch Fabric	Disable this option to deploy GigaVUE V Series Nodes using VMware NSX-T manager.
Use Public IP	Enable this to use the public IP address of the FM instance for communication from fabric nodes. This could be Public IP, Elastic IP, or Floating IP depending on the platform.

6. Select **Save**. The **VMware NSX-T Fabric Deployment** page appears. For details on how to deploy GigaVUE V Series Nodes in the **VMware NSX-T Fabric Deployment** page, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#). When creating multiple monitoring domains with the same NSX-T Manager, each monitoring domain is associated with a unique service name. You can view the service name of each monitoring domain on the **Monitoring Domains** tab.

**Notes:**

- All V Series Nodes within a single Monitoring Domain must run the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.
- Similarly, when upgrading a V Series Node, ensure that the GigaVUE-FM version is the same or higher than the V Series Node version.

You can perform the following actions in the Monitoring Domains, Service Deployments and Fabric tabs:

Actions	Description
Monitoring Domains tab	
Manage Certificates	You can use this button to perform the following actions: Re-issue: Required to address security compromises, key changes, or configuration updates, like validity period adjustments. Renew: Extends the expiration date and usually happens automatically unless you decide to do it during scheduled downtime. Auto-renewal is performed based on the duration specified in the Certificate Settings page. For details, refer to Configure Certificate Settings.
Refresh Inventory	Used to refresh the monitoring domain inventory.

Actions	Description
Edit Monitoring Domain	Used to edit a monitoring domain. You cannot edit fabric nodes of different monitoring domains at the same time.
Deploy Fabric	Used to deploy GigaVUE V Series Nodes.. This option is enabled only when no existing fabric launch configuration is present. and orchestration is enabled.
Upgrade Fabric	Used to upgrade GigaVUE V Series Nodes. For details, refer to Upgrade GigaVUE V Series Nodes for VMware NSX-T
Delete Monitoring Domain	Used to delete a monitoring domain.
Edit SSL Configuration	Used to add Certificate Authority and the SSL Keys when using the secure tunnels.
Service Deployments tab	
Delete Deployment	Used to delete the service deployments for any service.
Fabric tab	
Generate Sysdump	

What to do next:

- **Use FM to Launch Fabric** is enabled: You are navigated to the **VMware NSX-T Fabric Deployment** page.
- **Use FM to Launch Fabric** is disabled: You must deploy GigaVUE V Series Nodes using VMware NSX-T Manager.

Deploy GigaVUE V Series Nodes using VMware NSX-T Manager

You can deploy your GigaVUE V Series Nodes using VMware NSX-T Manager. The GigaVUE V Series Nodes register themselves with GigaVUE-FM using the information provided by the user in the NSX-T Manager. When the nodes are registered with GigaVUE-FM, you can configure Monitoring Session and related services in GigaVUE-FM.

After creating a Monitoring Domain, use a VMware NSX-T user account to deploy the GigaVUE V Series Nodes with at least the NETX Partner Admin role assigned. When creating multiple domains with the same NSX-T Manager, GigaVUE-FM assigns a unique service name to every Monitoring Domain. You can view these names on the Monitoring Domain page.

To deploy the V Series Node:

1. In **VMware NSX-T Manager**, go to the **Service Deployment** page and select **Deployment**.
2. From the **Partner Service** drop-down, select the service name corresponding to your Monitoring Domain. For detailed information, refer to [Deploy a Partner Service](#) topic in the VMware documentation.

3. Select the **Deployment template** and **Deployment Specification**, and click **Configure Attributes**.
4. In the **Configure Attributes** page, enter the
 - Service VM Host Name
 - Admin user password
 - (Optional) **SSL Private Key** and the **SSL Certificate** to use a custom certificate.

For details, refer to the Secure Communication topic.

After deployment,

- NSX-T Manager displays the runtime status of node as UP.
- GigaVUE-FM updates the node status to Launching. After successful registration, the status changes to **OK**.
- The deployed node appears on the Monitoring Domain page in GigaVUE-FM.

NOTE: VMware NSX-T automatically generates the V Series Node name. Do not rename the node in vCenter.

Compatibility and Limitations

- All GigaVUE V Series Nodes in a Monitoring Domain must run the same version. GigaVUE-FM does not support mixed versions.
- IPv6 is not supported for the gateway of the tunnel interface when deploying nodes through NSX-T Manager.
- Avoid renaming or manually managing the V Series Node from vCenter to prevent deployment issues.

Configure and Deploy Monitoring Session

In GigaVUE-FM, you must do the following to configure and deploy the Monitoring Session. Refer to [Configure Monitoring Session](#) section for more details.

1. **Access the Monitoring Session Page:**
 - a. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > VMware (NSX-T)**.

- b. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
2. **Add components to the canvas:**
 - a. Drag and drop the required components to the canvas.
 - b. From the **Applications** expand menu, drag and drop the required applications to the graphical workspace.
3. **Deploy the Monitoring Session:**
 - a. From the **Actions** menu, select **Deploy**.
 - b. After successful deployment on all the V Series Nodes, the status appears as **Success** on the Monitoring Sessions **Sources** tab.

What to do Next

You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab. For more details, refer to View Monitoring Session Statistics (VMware NSX-T).

Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform

Terraform automates the deployment and management of supported GigaVUE Cloud Suite configurations. It uses reusable configuration files to define the required infrastructure. This approach enables consistent deployment, updates, and maintenance without manual work in the GigaVUE-FM UI.

Terraform applies these configurations to keep the system in the desired state. It updates only the parts that change, which helps reduce errors and maintain consistency across environments. Version-controlled files make it easier to track and manage configuration changes.

Terraform supports repeatable deployments and helps integrate GigaVUE Cloud Suite into automated workflows.

License

The GigaVUE-FM Terraform Collection does not require any additional licensing. However, you must procure a valid GigaVUE-FM license and the corresponding feature licenses to enable the GigaVUE-FM features required for your deployment.

For detailed information on license types, acquisition, and usage, refer to the Gigamon Licensing Guide.

Supportability

Use the Terraform version 1.14 or later for the GigaVUE-FM deployment. If multiple provider builds are available, select the build validated for the target deployment.

Platform support depends on the deployment model and traffic acquisition method.

Supported Deployment Models

GigaVUE-FM Orchestration

- VMware ESXi-Traffic acquisition using Platform Mirroring.

Third-Party Orchestration (Generic Mode)

Supported platforms

- AWS
- Azure
- Google Cloud Platform (GCP)
- Nutanix
- OpenStack
- VMware ESXi
- NSX-T Federation
- Oracle Cloud Infrastructure (OCI)

Supported traffic acquisition methods

- UCT-V
- Customer-Orchestrated Source

Supported GigaSMART Applications

The Gigamon Terraform provider supports the following GigaSMART applications:

- AMX
- Dedup
- Header Stripping
- Load Balancing

- Masking
- Slicing

Rules and notes

- [Install](#) a supported version of Terraform.
- Ensure GigaVUE-FM is installed and reachable from the system running Terraform.
- Always review `terraform plan` output before running `terraform apply` Or `terraform destroy`.

Prerequisites

Before starting a Terraform deployment, ensure the following:

- Install a supported version of Terraform (1.14 or later).
- Ensure GigaVUE-FM is installed and reachable from the system running Terraform. GigaVUE-FM token authentication is required for the Terraform provider to connect to GigaVUE-FM. For details on generating a token, refer to [Configure Tokens](#).
- Refer the sample configurations from the [Terraform Registry](#).
- Prepare the required deployment values, including connection details, authentication credentials, and environment-specific inputs.

Get Started

Install Terraform

Install Terraform 1.14 or later on the system from which you will run the deployment.

Terraform runs the provider and the deployment profile that you create for GigaVUE Cloud Suite. If Terraform is not installed, or if the installed version is not supported, the deployment can fail during initialization or when Terraform reads the deployment files.

- Download and install Terraform version 1.14 or later on the deployment system:
 - [Install Terraform on Windows](#)
 - [Install Terraform on Linux](#)
- Open a terminal or command prompt on the deployment system.
- To verify the installation, run the following command:

```
terraform version
```

Confirm that the command runs successfully and that the displayed version is 1.14 or later.

If the command returns an error such as **command not found**, verify that Terraform is installed correctly

The deployment system is ready to initialize the provider and run the Terraform workflow.

Create a working directory

1. Create a folder on the deployment system where you will store the Terraform files and run the Terraform commands.
2. This folder is the working directory. Terraform reads all .tf files in the working directory together as one configuration.
3. Use one working directory for each deployment profile so that the provider settings, resource definitions, variables, values, and outputs for that deployment stay together.
4. Use the published Gigamon Terraform documentation as the primary reference when you create the deployment profile.
5. The provider documentation tells you:
 - which resources are available
 - which arguments each resource requires
 - which values each argument accepts
 - how each resource is used in Terraform

Refer to: [Gigamon Terraform provider documentation](#)

Deployment Using Terraform

This section describes the supported Terraform-based deployment paths for GigaVUE Cloud Suite.

Third-Party Orchestration

- [Acquire traffic using UCT-V](#)
- [Acquire traffic using Customer-Orchestrated Source](#)

Deploy Third Party Orchestration by using Terraform

Third Party Orchestration (3PO) deployments (Terraform-based) use Generic Mode. In this mode of deployment, the third-party orchestrator is responsible for deploying and managing the GigaVUE V Series nodes and the underlying infrastructure. Terraform communicates with GigaVUE-FM only to create and manage the monitoring configuration. i.e the monitoring domain, the connection, and the monitoring session objects.

Supported platforms (Generic Mode):

- AWS
- Azure
- Google Cloud Platform (GCP)
- Nutanix
- OpenStack
- VMware ESXi
- NSX-T Federation
- Oracle Cloud Infrastructure (OCI)

This deployment supports the following traffic acquisition methods:

- UCT-V
- Customer-Orchestrated Source

Before you begin

Complete the following setup before starting the Terraform deployment.

1. Install GigaVUE-FM.

- Terraform uses GigaVUE-FM as the management system for this deployment. For installation instructions, see the [GigaVUE-FM installation topic](#)
- GigaVUE-FM token authentication is required for the Terraform provider to connect to GigaVUE-FM. For details on generating a token, refer to [Configure Tokens](#).
- If GigaVUE-FM is not installed or is not reachable from the deployment system, Terraform cannot create or update the 3PO deployment objects.

2. Ensure that the third-party orchestration environment is available.

- Confirm that the third-party orchestrator is running and that the target platform environment is reachable. The third-party orchestrator is responsible for deploying and managing the GigaVUE V Series nodes.
- Terraform does not deploy V Series nodes in a 3PO deployment.

3. Verify third-party networking configuration

Ensure that the required management, data, and tool-side networks are available and correctly configured in the target environment. The third-party orchestrator uses these values when deploying the GigaVUE V Series nodes. For more details, refer to [Prerequisites](#).

4. **Install and verify Terraform 1.14 or later**

Run the `terraform version` command and confirm that respective version 1.14 or later is displayed.

5. **Run the `terraform init` command**

Terraform downloads the Gigamon provider automatically from the Terraform Registry and prepares the working directory for deployment. You do not need to download or extract the provider package manually.

The Gigamon provider is available at the [Terraform Registry](#).

Install the Gigamon Terraform Provider Without Internet Access

If the deployment system does not have internet access or cannot reach the Terraform Registry, you can download the Gigamon Terraform provider binary directly from the [Terraform provider GitHub Releases page](#) and install it locally

Step 1: Download and install the provider binary

1. **Download the provider zip file for your platform from the GitHub releases page.**

The following example uses the Linux AMD64 platform. Replace `linux_amd64` with the directory name that matches your operating system and architecture.

2. **Create the local plugin directory:**

```
mkdir -p ~/.terraform.d/plugins/local/gigamon/gigamon/6.14.0/linux_
amd64
```

3. **Extract the zip file:**

```
unzip terraform-provider-gigamon_6.14.0_linux_amd64.zip
```

4. **Move the binary to the local plugin directory:**

```
mv terraform-provider-gigamon
~/.terraform.d/plugins/local/gigamon/gigamon/6.14.0/linux_amd64/
```

5. **Make the binary executable:**

```
chmod +x ~/.terraform.d/plugins/local/gigamon/gigamon/6.14.0/linux_
amd64/terraform-provider-gigamon
```

Step 2: Update the provider block

In your .tf file, update the provider source to point to the local plugin directory instead of the Terraform Registry:

```
terraform {
  required_providers {
    gigamon = {
      source = "local/gigamon/gigamon"
      version = ">= 6.14"
    }
  }
}
```

NOTE: The source value `local/gigamon/gigamon` tells Terraform to look for the provider in the local plugin directory. Do not use the registry source address when using the local installation method.

Step 3: Continue with the standard deployment flow

After the binary is in place and the provider block is updated, the remaining deployment steps are the same as the registry-based installation. Run `terraform init`, `terraform plan`, and `terraform apply` from the working directory as described in [Run the Terraform Deployment File](#).

Prepare the Third Party Orchestration Deployment Files

Sample configuration files for the UCT-V and Customer-Orchestrated Source deployment paths are available in the [Terraform Registry](#).

These sample files provide working examples that demonstrate the resource structure and required configuration arguments for supported traffic acquisition methods. Use them as a reference when creating deployment files, or building custom deployment files based on requirements. Review and update all placeholder values before using the configuration in an environment.

NOTE: Terraform reads all .tf files in the working directory as a single configuration and processes them together during plan and apply operations.

Before you create the deployment files, decide how to handle the monitoring domain:

- If the monitoring domain does not exist in GigaVUE-FM, include a resource block to create it as part of the deployment.
- If the monitoring domain already exists in GigaVUE-FM, supply the monitoring domain name and connection name in the values file. Do not include a resource block to create it.

You can use either of the following approaches depending on the complexity of your deployment.

1. Single-file configurationModule-based configuration

Single-file configuration

Write all provider settings and resource definitions in a single main.tf file. Use this approach for straightforward deployments where you want to keep all configuration in one place.

Your **main.tf** (Main deployment configuration file) must include the following blocks:

1. Provider block

Declare the Gigamon provider and specify the GigaVUE-FM connection details.

Example:

```
terraform {
  required_providers {
    gigamon = {
      source  = "gigamon-engg/gigamon"
      version = ">= 6.14"
    }
  }
}

provider "gigamon" {
  fm_ip = "<GigaVUE-FM IP address or hostname>"
  token = "<Token>"
}
```

2. Resource blocks

Add a resource block for each third-party orchestration object that Terraform must create or manage. The resources you include depend on your traffic acquisition method.

For UCT-V deployments, define the resources in the following sample order:

- [Monitoring domain](#) (if not pre-existing in GigaVUE-FM)
- [Connection](#)
- [Monitoring session](#)
- [Applications](#)
- [Maps](#)
- [Links](#)
- [Tunnels](#)

For Customer-Orchestrated Source deployments, define the resources in the following sample order:

- [Monitoring domain](#) (if not pre-existing in GigaVUE-FM)
- [Connection](#)
- [Monitoring session](#)
- [Applications, maps, links, and tunnels](#)

Refer to the [Terraform Registry](#) documentation for the resource type names, required arguments, and argument descriptions for each resource.

Module-based configuration

Organize the configuration across multiple files or modules. Use this approach when you want to reuse configuration across environments, manage a large deployment, or separate concerns across teams.

In a module-based setup, you create a root configuration that calls one or more child modules. Each module groups related resources and accepts input variables. The module structure varies depending on the traffic acquisition method you use.

Example main.tf - calling modules

```
module "monitoring_domain" {
  source      = "./modules/monitoring-domain"
  fm_ip       = var.fm_ip
  domain_alias = var.domain_alias
  # Omit this module block if using an existing monitoring domain.
}

module "connection" {
  source              = "./modules/connection"
  monitoring_domain_id = module.monitoring_domain.id
  connection_alias    = var.connection_alias
  # Add required variables for your environment.
}

module "monitoring_session" {
  source              = "./modules/monitoring-session"
  monitoring_domain_id = module.monitoring_domain.id
  connection_id       = module.connection.id
  # Add required variables for your environment.
}
```

Each module's **main.tf** contains the resource blocks for that group. Pass values between modules using module outputs.

Refer to the [Terraform provider](#) documentation for the full list of resources and arguments for each module.

Run the Terraform Deployment File

After you finish the sample-file updates, run the deployment from the working directory.

1. Initialize the working directory. Run the following command:

```
terraform init
```

Terraform initializes the working directory, downloads Gigamon provider automatically from the Terraform Registry and prepares the working directory for deployment.

2. Optional: Validate the configuration.

- Run:

```
terraform validate
```

- Run this command when you want Terraform to check the configuration for syntax and structural issues before it generates the deployment plan.

3. Review the deployment plan.

- Run:

```
terraform plan
```

- Review the output carefully. Confirm that Terraform will create or update the expected third-party objects.
- Do not continue until the plan matches the deployment that you expect.

4. Apply the configuration.

- Run:

```
terraform apply
```

- When prompted, confirm the deployment.
- Terraform connects to GigaVUE-FM, creates or updates the third-party monitoring domain and connection, and applies the monitoring session configuration defined in the deployment files.

NOTE: Manage monitoring domain, connection, CA, and SSL key/certificate resources separately from the monitoring session configuration. Create them in GigaVUE-FM and reference them as data sources, or manage them through a separate Terraform configuration. This approach ensures the resources exist before deployment and reduces the risk of unintended modifications during update or destroy operations.

Update an existing Third Party Orchestration deployment

To update an existing deployment, change only the required values in the deployment files and rerun:

```
terraform plan
```

```
terraform apply
```

Make small, controlled changes and verify the deployment after each update. This approach simplifies review and troubleshooting.

Destroy a Third-Party Orchestration Deployment

Use terraform destroy to remove the Third-Party Orchestration objects that Terraform manages in GigaVUE-FM.

NOTE: In a Third-Party Orchestration deployment, terraform destroy cannot complete in a single run if the configuration includes a monitoring domain and connection. GigaVUE-FM does not allow deletion of a monitoring domain or connection while UCT-V agents or V Series nodes remain registered. Terraform

will destroy the monitoring session, certificate application, CA, and SSL key or certificate in the first run, but will fail when it attempts to delete the connection and monitoring domain.

To destroy the deployment completely, follow the two-step process below.

1. Run the first destroy

- Run the following command from the working directory:

```
terraform destroy
```

- Terraform removes the monitoring session in this run:

Terraform cannot delete the connection and monitoring domain during this run. This occurs because GigaVUE-FM does not allow deletion of a monitoring domain or connection while UCT-V agents or V Series nodes remain registered.

2. Unregister agents and nodes, then re-run destroy

- In GigaVUE-FM, manually unregister the following from the monitoring domain and connection before the CA and SSL keys and certificates are deleted:
 - **UCT-V deployment** - Unregister all UCT-V agents, the UCT-V controller, and V Series nodes.
 - **Customer-Orchestrated Source deployment** - Unregister all V Series nodes.
- Confirm that no agents or nodes remain registered under the monitoring domain or connection.
- Re-run the destroy command:

```
terraform destroy
```

Terraform deletes the CA and SSL key or certificate, connection and monitoring domain.

NOTE: If you configure the monitoring domain and connection as data sources instead of managed resources in your Terraform configuration, terraform destroy completes in a single run. In this case, Terraform does not attempt to delete the monitoring domain or connection, and manual unregistration is not required.

Verify the deployment

After the deployment completes, verify the deployment in GigaVUE-FM.

1. Verify the monitoring domain and connection

- Go to the monitoring domain page.
- Confirm that:
 - monitoring domain is present
 - connection is present
 - connection is associated with the expected monitoring domain

If the monitoring domain or connection does not appear in the expected location, review the monitoring domain name and connection name in the values file and rerun the deployment.

2. **Verify node and tunnel health**

- Confirm that the deployed components appear under the correct monitoring domain and connection.
- Confirm that node health is normal. GigaVUE-FM places components in the correct monitoring domain and connection by using the supplied registration data and tracks health through heartbeat messages.

3. **Verify the monitoring session configuration**

- Go to the monitoring session view for the deployment.
- Confirm that:
 - monitoring session is present
 - expected applications are attached
 - expected maps are attached
 - expected links are attached
 - expected tunnels are attached
- If the deployment uses secure tunnel settings, confirm that the secure tunnel configuration is present and healthy.

If any required object is missing, review **main.tf** and the monitoring session values in the values file, then rerun terraform plan and terraform apply.

4. **Verify traffic flow**

- Confirm that traffic reaches the configured tool destination.
- If traffic does not reach the tool destination, review the traffic acquisition method, monitoring session values, application values, and tunnel or link configuration in the sample files.

After you complete this verification, GigaVUE-FM contains the third-party orchestration monitoring domain, the connection, and the monitoring session configuration defined in the Terraform deployment files

OpenTofu Support

OpenTofu is an open-source alternative to Terraform that is compatible with the Gigamon Terraform provider. All configuration files, scripts, and deployment steps described in this guide work with OpenTofu without modification.

OpenTofu Registry

To use OpenTofu instead of Terraform:

- Install OpenTofu v1.12.3 or later on the deployment system:
 - [Install OpenTofu on Windows](#)
 - [Install OpenTofu on Linux](#)
- Replace all terraform CLI commands with the equivalent tofu commands.
- Follow the same deployment steps described in this guide.

For the Gigamon provider, refer to the [OpenTofu Registry](#) to download the provider package compatible with OpenTofu.

Best Practices

Use the following best practices when you deploy by using Terraform:

1. Store Terraform configuration files in version control.
2. Run `terraform validate` before `terraform plan` to catch syntax and configuration errors early.
3. Review `terraform plan` before you run `terraform apply`.
4. Pin the provider version in `provider.tf` to avoid unexpected behavior from automatic provider upgrades.
5. Keep sensitive values out of configuration files and protect access to the state file.
6. Make small, controlled changes and verify the deployment after each update.

7. Use a shared remote backend for production deployments.

- A shared backend ensures that all users working on the same deployment share the same Terraform state file. Without a remote backend, concurrent changes from multiple users can corrupt the state, and accidental deletion of a local state file results in loss of all tracked resources.
- Configure the remote backend in the .tf file using the http backend block. Store the backend connection values separately in a .hcl file and pass it at initialization, do not store credentials directly in the .tf file.

Example: main.tf - backend and provider configuration

```
terraform {
  required_providers {
    gigamon = {
      source = "gigamon-engg/gigamon"
      version = ">= 6.14"
    }
  }
}
```

Example: backend.hcl - backend connection values

```
address          = "https://{fm_address}/terraform-state/{project_name}"
lock_address      = "https://{fm_address}/terraform-state/{project_name}/lock"
unlock_address    = "https://{fm_address}/terraform-state/{project_name}/lock"
skip_cert_verification = true
username          = "{project_name}"
password          = "{fm_token}"
```

- Pass the .hcl file when you initialize the working directory:

```
terraform init -backend-config=<hcl file location>
```

NOTE:

- Save the **.hcl** file locally and do not commit it to version control. The file contains credentials and backend connection details that must remain secure.
- Keep credentials and sensitive values out of configuration files and protect access to the state file.

Verification and Troubleshooting

Verify the deployment

After you run Terraform, verify the deployment in GigaVUE-FM.

- Log in to GigaVUE-FM.
- Verify that the monitoring domain and connection are present.
- Verify that the monitoring session is present.
- Verify that the required applications, maps, links, and tunnels are attached to the monitoring session.
- Verify that the deployed nodes are healthy.
- Verify that traffic is flowing as expected.

If any required object is missing or unhealthy, review the Terraform output, correct the configuration, and rerun the deployment.

Troubleshoot the deployment

If the deployment does not complete successfully, use the following checks.

Provider issues

- Verify that the provider source is correct.
- Verify that the provider version is supported for your release.
- Rerun `terraform init`.

Authentication or connectivity issues

- Verify the GigaVUE-FM IP address or hostname.
- Verify the authentication values.
- Verify network connectivity between the system running Terraform and GigaVUE-FM.

Configuration issues

- Run `terraform validate`.
- Run `terraform plan`.
- Correct any configuration errors before you rerun `terraform apply`.

Deployment timeout issues

- Verify whether the deployment includes operations that require additional time to complete.
- Increase the timeout value only if your deployment requires it.

State file issues

- Verify that all users managing the same production deployment use the same shared backend.
- Do not use separate local state files for the same production deployment.

Recover after a failed deployment

- Identify and correct the root cause.
- Rerun the Terraform workflow.
- Verify that the deployment reaches the intended state.

Configure

Configure Monitoring Session

GigaVUE-FM collects inventory data on all V Series nodes deployed in your environment through vCenter connections. 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 you add a new target VM to your environment, GigaVUE-FM automatically detects it. If the VM matches your session's criteria, GigaVUE-FM adds it to the monitoring session. Likewise, when you remove a target VM, GigaVUE-FM updates the monitoring session to reflect the change.

Before deploying a monitoring session, you need to deploy a V Series node in each host where you want to monitor the traffic.



Notes:

- Link transformation and multiple links between two entities are not supported in V Series nodes of ESXi.
- VMware ESXi running with V Series nodes does not support Pre-filtering.

Refer to the following topics for details:

- [Create a Monitoring Session \(VMware NSX-T\)](#)
- [Interface Mapping](#)

- [Create Ingress and Egress Tunnel \(VMware NSX-T\)](#)
- [Create a New Map](#)
- [Add Applications to Monitoring Session](#)
- [Deploy Monitoring Session](#)
- [View Monitoring Session Statistics](#)

Create a Monitoring Session (VMware NSX-T)

GigaVUE-FM automatically collects inventory data on all target instances in your cloud environment. You can design your Monitoring Session to:

- Include or exclude the instances that you want to monitor.
- Monitor egress, ingress, or all traffic.

Target Instance

- When a new target instance is added to your cloud environment, GigaVUE-FM automatically detects and adds it to your Monitoring Session based on your selection criteria. Similarly, when an instance is removed, it updates the Monitoring Sessions.

You can create multiple Monitoring Sessions within one Monitoring Domain.

To create a new Monitoring Session:

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Session** page appears.
2. Select **New Monitoring Session** to open the New Monitoring Session configuration page.
3. In the configuration page, perform the following:
 - In the **Alias** field, enter the name of the Monitoring Session.
 - From the **Monitoring Domain** drop-down list, select the desired Monitoring Domain or select **Create New** to create a Monitoring Domain. For details, refer to the Create a Monitoring Domain section in the respective cloud guides.
 - From the **Connections** drop-down list, select the required connections to include as part of the Monitoring Domain.
 - From the **VPC** drop-down list, select the required VPCs to include as part of the Monitoring Domain.
 - Enable the **Distributed Deduplication** option to identify duplicate packets across different GigaVUE V Series Nodes when traffic from various targets is routed to these instances for monitoring. Distributed Deduplication is only supported on GigaVUE V Series Node version 6.5.00 and later.
4. Select **Save**.
The Monitoring Session Overview page appears.

**Points to Note:**

- In a Monitoring Session, if a selected VM is connected to VSS and VDS, then the GigaVUE-FM can create tapping for both VSS and VDS network.
- When you undeploy or edit a Monitoring Session and redeploy it, the existing traffic flows redirected to the GigaVUE V Series Node will stop being redirected to the GigaVUE V Series Nodes. However, new flows initiated after the redeployment will be redirected to the GigaVUE V Series Nodes.

Monitoring Session Page

The table outlines the functions of each tab in the Monitoring Session Page for managing and analyzing network traffic.

Tab	Description
Overview	You can view the high level information of the selected Monitoring Session such as, connections, tunnel details, health status, deployment status, and information related to Application Intelligence statistics. You can also view the statistics of the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis. You can filter the statistics based on the elements associated with the Monitoring Session. For more information, refer to View Monitoring Session Statistics
Sources	Displays the sources and target details monitored by the Monitoring Session. You can view and edit the connection details of the Monitoring Session. You can view the deployment status, number of targets, and targets source health. In the Selection Status section, you can view the VM status. The status indicates whether the VM is supported, not supported, selected, or not selected. When you hover over the status, a tooltip displays the reason for that status. In the case of OVS Mirroring, the Sources tab also displays the Hypervisor details along with the Instances.
Traffic Processing	You can view, add, and configure applications, tunnel endpoints, raw endpoints, and maps. You can view the statistical data for individual applications and also apply threshold template, enable user defined applications, and enable or disable distributed De-duplication. Refer to Configure Monitoring Session Options for more detailed information.
V Series Nodes	You can view the V Series nodes associated with the Monitoring Session. In the split view, you can view details such as name of the V Series Node, health status, deployment status, Host VPC, version, and Management IP. You can also change the interfaces mapped to an individual GigaVUE V Series Node. Refer to Interface Mapping section for details.
Topology	Displays the fabric and monitored instances based on the connections configured in your network. You can select a specific connection to explore its associated subnets and instances in the topology view, offering a clear visualization of the monitored network elements. Refer to Visualize the Network Topology (VMware NSX-T) .

NOTE: Ensure that the GigaVUE V Series Node and GigaVUE-FM are time synchronized or configure NTP time synchronization.

The Monitoring Session page **Actions** button has the following options. The Actions menu is placed common in all the tabs explained above.

Button	Description
Delete	Deletes the selected Monitoring Session.
Clone	Duplicates the selected Monitoring Session.
Deploy	Deploys the selected Monitoring Session.
Undeploy	Undeploys the selected Monitoring Session.

You can use the  icon on the left side of the Monitoring Session page to view the Monitoring Sessions list. Click  to filter the Monitoring Sessions list. In the side bar, you can:

- Create a new Monitoring Session
- Rename a Monitoring Session
- Hover over, click the check box of the required Monitoring Session(s) and perform bulk actions (Delete, Deploy, or Undeploy).

Configure Monitoring Session Options

In the Monitoring Session page, you can perform the following actions in the **TRAFFIC PROCESSING** tab.

- [Apply Threshold Template](#)
- [Enable User-Defined Applications](#)
- [Enable Distributed De-duplication](#)

Access the TRAFFIC PROCESSING tab

To navigate to **TRAFFIC PROCESSING** tab, follow these steps:

1. Go to **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. On the left pane with the Monitoring Sessions list view, select a Monitoring Session.
3. Select the **TRAFFIC PROCESSING** tab.

Apply Threshold Template

You can apply the Threshold configuration to a Monitoring Session before deployment.

To apply a threshold,

1. In the **TRAFFIC PROCESSING** page, select **Options > Thresholds**.
2. Select an existing threshold template from the **Select Template** drop-down list. You can create a template using **New Threshold Template** option and apply it.
3. Select **Apply**.

For more information, refer to the [Traffic Health Monitoring](#) section.

The template is added to the Monitoring Session.



Notes:

- Undeploying the Monitoring Session does not remove the applied Thresholds.
- You can also view the details related to the applied thresholds, such as traffic element, metrics, type, trigger values, and time intervals, in the threshold window.

Select **Clear Thresholds** to clear the applied thresholds across the selected Monitoring Session.

Enable User-Defined Applications

To enable a defined application,

1. In the Monitoring Session **TRAFFIC PROCESSING** page, select **Options > User Defined Applications**.
2. Enable the **User-defined Applications** toggle button.
3. From the **Actions** drop-down, add one of the existing applications or create a User-Defined Application.

For more information, refer to [User Defined Application](#).

Enable Distributed De-duplication

Enabling the Distributed De-duplication option identifies duplicate packets across different GigaVUE V Series Nodes when traffic from various targets is routed to these instances for monitoring. For more information, refer to [Distributed De-duplication](#).

To enable,

1. In the TRAFFIC PROCESSING page, select **Options > Distributed De-duplication**.
2. Enable the toggle.



Notes:

- Supported only on V Series version 6.5.00 and later.
- From version 6.9, the Traffic Distribution option is renamed to Distributed De-duplication.

Create Ingress and Egress Tunnel (VMware NSX-T)

GigaVUE V Series Node sends traffic to tunnel endpoints in a monitoring session. You can create a tunnel endpoint using L2GRE, VXLAN, UDPGRE, TLS-PCAPNG, UDP, or ERSPAN.

Create a new tunnel endpoint

To create,

1. Perform one of the following and navigate to the **TRAFFIC PROCESSING** tab:
 - Create a new monitoring session
 - Select **Actions > Edit** on an existing monitoring session.

The GigaVUE-FM Monitoring Session canvas page appears.


2. On the left pane of the canvas, select the  icon to view the traffic processing elements.
3. Select **New > New Tunnel**, drag and drop a new tunnel template to the workspace.

The **Add Tunnel Spec** quick view appears.

4. Enter the **Alias**, **Description**, and **Type** details.

For details, refer to [Details - Add Tunnel Specifications](#) table.

5. Select **Save**.



To delete a tunnel, select the  menu button of the required tunnel and select **Delete**.

Apply a threshold template to Tunnel End Points

1. Select the  menu button of the required tunnel endpoint on the canvas and click **Details**.
2. In the quick view, go to the **Threshold** tab.

For details on creating or applying a threshold template, refer to the Monitor Cloud Health topic in the respective Cloud guides.

You can use the configured Tunnel End Points to send or receive traffic from GigaVUE HC Series and GigaVUE TA Series. Provide the IP address of the GigaVUE HC Series and GigaVUE TA Series as the Source or the Destination IP address as required when configuring Tunnel End Points.

After configuring the tunnels and deploying the Monitoring Session, you can view the number of ingress and egress tunnels configured for a Monitoring Session. Select the numbers of tunnels displayed in the **OVERVIEW** tab to view the tunnel names and their respective **ADMIN STATUS** and **HEALTH STATUS**.

Table 2: Details - Add Tunnel Specifications

Field	Description
Alias	The name of the tunnel endpoint.
Description	The description of the tunnel endpoint.
Admin State Note: This option appears only after the Monitoring session deployment.	Use this option to send or stop the traffic from GigaVUE-FM to the egress tunnel endpoint. Admin State is enabled by default. You can use this option to stop sending traffic to unreachable or down tools. Each egress tunnel configured on the GigaVUE V SeriesNode has an administrative state that enables GigaVUE-FM to halt the tunnel's traffic flow. GigaVUE-FM only disable the tunnels when it receives a notification via REST API indicating that a tool or group of tools is down. Note: This option is not supported for TLS-PCAPNG tunnels.
Type	The type of the tunnel. Select from the options below to create a tunnel. ERSPAN, L2GRE, VXLAN, TLS-PCAPNG, UDP, or UDPGRE.
VXLAN	
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. Note: In the scenario where secure tunnels need to be established between a GigaVUE V Series Node and a GigaVUE HC Series, you can utilize the Configure Physical Tunnel option provided in the GigaVUE V Series Secure Tunnel page. This allows you to configure secure tunnels on your physical device conveniently. For details, refer to Secure Tunnels .	
In	Choose In (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.

Field	Description	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP/FQDN	For egress tunnels, specify the destination endpoint by using an IPv4 address, IPv6 address, or fully qualified domain name (FQDN). FQDN support applies only to VXLAN and L2GRE egress tunnels. By using an FQDN, you can keep the same tunnel configuration even when the destination IP address changes. If the FQDN resolves to both IPv4 and IPv6 addresses, the system selects the address based on the preferred IP stack configured for the monitoring domain.
	VXLAN Network Identifier	Unique value that is used to identify the VXLAN. The value ranges from 1 to 16777215.
	Source L4 Port	The port used to establish the connection to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port used to establish the connection will be established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
Out	Choose Out (Encapsulation) for creating an egress tunnel from the GigaVUE V Series Node to the destination endpoint.	
	Remote Tunnel IP	For egress tunnel, the Remote Tunnel IP is the IP address of the tunnel destination endpoint.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.
	Time to Live	Enter the value of the time interval for which the session needs to be available. The value ranges from 1 to 255. The default value is 64.
	DSCP	Differentiated Services Code Point (DSCP) is a value that network devices use to identify traffic to be handled with higher or lower priority. The values ranges from 0 to 63 with 0 being the highest priority and 63 being the lowest priority.
	Flow Label	Unique value, which is used to identify packets that belong to the same flow. A flow is a sequence of packets that need to be treated as a single entity that may require special handling. The accepted value is between 0 and 1048575.
	VXLAN Network Identifier	Unique value which is used to identify the VXLAN. The value ranges from 1 to 16777215.
	Multi Tunnel	Enable the multi-tunnel flag to create multiple tunnels for flow distribution to the 5G-Cloud application. Refer to 5G-Cloud Ericson SCP Support . Applicable Platforms: OpenStack, Third Party Orchestration,

Field	Description	
		VMware ESXi  Notes: - You can configure either a single-tep or multi-tep setup for the egress tunnel. Switching between these configurations is not allowed; to make changes, you must undeploy and redeploy the Monitoring Session. - When you enable Multi-Tunnel on a VXLAN tunnel and set the number of tunnels, GigaVUE-FM automatically creates the additional VXLAN tunnel endpoints. Any later changes to the original VXLAN tunnel, such as disabling Multi-Tunnel or modifying Domain Tagging do not update these auto created endpoints. They continue to retain the configuration that existed at the time they were created. To apply updated settings, you must delete the VXLAN TEP and the associated LB application, then recreate the LB and VXLAN TEP with the new configuration, and re-establish the link between them.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
	Domain Tagging	Enable this option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. NOTE: This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
UDPGRE		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	Choose In (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of

Field	Description	
		the tunnel source.
	Key	Identifier used to differentiate different UPDGRE/L2GRE tunnels. It routes the encapsulated frames to the appropriate tunnel on the remote endpoint. Enter a value between 0 and 4294967295.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
L2GRE		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. Note: In the scenario where secure tunnels need to be established between a GigaVUE V Series and a GigaVUE HC Series, you can utilize the Configure Physical Tunnel option provided in the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels on your physical device . For details, refer to the Secure Tunnels .		
In	Choose In (decapsulation) to create an ingress tunnel, which will carry traffic from the source to the GigaVUE V Series Node.	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Key	Identifier used to differentiate different UPDGRE/L2GRE tunnels. It is used to route the encapsulated frames to the appropriate tunnel on the remote endpoint. Enter a value between 0 and 4294967295.
Out	Choose Out (Encapsulation) for creating an egress tunnel from the V Series Node to the destination endpoint.	
	Remote Tunnel IP/FQDN	For egress tunnels, specify the destination endpoint by using an IPv4 address, IPv6 address, or fully qualified domain name (FQDN). FQDN support applies only to VXLAN and L2GRE egress tunnels. By using an FQDN, you can keep the same tunnel configuration even when the destination IP address changes. If the FQDN resolves to both IPv4 and IPv6 addresses, the system selects the address based on the preferred IP stack configured for the monitoring domain.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.

Field	Description	
	Time to Live	Enter the value of the time interval for which the session needs to be available. The value ranges from 1 to 255. The default value is 64.
	DSCP	Differentiated Services Code Point (DSCP) is a value that network devices use to identify traffic to be handled with higher or lower priority. The values ranges from 0 to 63 with 0 being the highest priority and 63 being the lowest priority.
	Flow Label	Unique value, which is used to identify packets that belong to the same flow. A flow is a sequence of packets that need to be treated as a single entity that may require special handling. The accepted value is between 0 and 1048575.
	Key	Identifier used to differentiate different UPDGRE/L2GRE tunnels. It is used to route the encapsulated frames to the appropriate tunnel on the remote endpoint. Enter a value between 0 and 4294967295.
	Domain Tagging	Enable this option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. NOTE: This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
ERSPAN		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Flow ID	The ERSPAN flow ID is a numerical identifier that distinguishes different ERSPAN sessions or flows. The value ranges from 1 to 1023.
TLS-PCAPNG		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. Note: In the scenario where secure tunnels need to be established between a GigaVUE V Series and a GigaVUE HC Series, you can utilize the Configure Physical Tunnel option provided in the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels on your physical device . For details, refer to Secure Tunnels section.		

Field	Description	
In	IP Version	The version of the Internet Protocol. Only IPv4 is supported.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
	Key Alias	Select the Key Alias from the drop-down.
	Cipher	Only SHA 256 is supported.
	TLS Version	Only TLS Version 1.3.
	Selective Acknowledgments	Enable to receive the acknowledgments.
	Sync Retries	Enter the number of times the sync has to be tried. The value ranges from 1 to 6.
	Delay Acknowledgments	Enable to receive the acknowledgments for a delay.

Field	Description	
Out	IP Version	The version of the Internet Protocol. Only IPv4 is supported.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.
	Time to Live	Enter the value of the time interval for which the session needs to be available. The value ranges from 1 to 255. The default value is 64.
	DSCP	Differentiated Services Code Point (DSCP) is a value that helps network devices identify the higher or lower priority to handle traffic. The values ranges from 0 to 63 with 0 being the highest priority and 63 being the lowest priority.
	Flow Label	Unique value which is used to identify packets that belong to the same flow. A flow is a sequence of packets that need to be treated as a single entity that may require special handling. The accepted value is between 0 and 1048575.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
	Cipher	Only SHA 256 is supported.
	TLS Version	Only TLS Version 1.3.
	Selective Acknowledgments	Enable the receipt of acknowledgments.
	Sync Retries	Enter the number of times you can try the sync. The value ranges from 1 to 6.
	Delay Acknowledgments	Enable the receipt of acknowledgments when there is a delay.
UDP:		

Field	Description	
Out	L4 Destination IP Address	Enter the IP address of the tool port or when using Application Metadata Exporter (AMX), enter the IP address of the AMX application. For details, refer to Application Metadata Exporter .
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.

Tunnel End Points created will be listed in the **Tunnel Specifications** page. You can create, edit, and delete tunnel end point from this page. Refer to [Create Tunnel Specifications](#) for more detailed information on how to create tunnel end points.

Create Raw Endpoint (VMware NSX-T)

This topic explains how to configure a RAW Endpoint (REP) in the Monitoring Session using VMware NSX-T.

Refer to the below sections for more detailed information:

- [Rules and Notes](#)
- [Create Raw Endpoint \(VMware NSX-T\)](#)
- [Configure Raw Endpoint in Monitoring Session](#)
- [Configuration Support for Non IP-Addressable \(I2 supported\) Tools](#)

Rules and Notes

- You must use VSS port groups. REPs are not supported on NSX Segments.
- GigaVUE-FM requires a valid IP address on the interface used for the RAW Endpoint.
- For interface roles and details, refer to the GigaVUE V Series Node interface configuration table in the Monitoring Domain documentation.
- The maximum number of links that can egress from any endpoint in V Series is four.

Points to Note:

Refer to the following table for more detailed information on the number of interfaces and their roles in the GigaVUE V Series Nodes deployed in the Monitoring Domain.

Display Name	Interface Name	Role	Comments
Network Adapter 1	ens160	Management	-
Network Adapter 2	ens192	Tunnel	Supports Tunnel and Egress RAW endpoint.
Network Adapter 3	ens224	Data	Reserved and used by NSX-T for Service Insertion (Traffic Acquisition)

Configure Raw Endpoint in Monitoring Session

A Raw Endpoint (REP) passes traffic from a GigaVUE V Series Node interface. You can use it to ingest traffic from physical sources and optionally forward it to monitoring tools.

To add Raw Endpoint to the Monitoring Session:

1. Go to the graphical workspace.
2. From the **New** menu, drag **New Raw Endpoint** into the workspace.
3. Select the menu icon on the new endpoint and choose **Details**.
4. In the **Raw** quick view, enter an Alias and Description, then select **Save**.
5. To deploy the Monitoring Session after adding the Raw Endpoint:
 - a. On the **Traffic Processing** page, go to the **Actions** menu and select **Deploy**.
 - b. In the **Deploy Monitoring Session** dialog, select the appropriate V Series Nodes.
 - c. For each node, map the correct interface (REP and TEP) using the drop-down menus.
 - d. Select **Deploy**.
6. Select **Export** to download all or selected V Series Nodes in CSV and XLSX formats.

Configuration Support for Non IP-Addressable (I2 supported) Tools

You can configure to deliver RAW traffic to Layer 2 tools or sensors that do not support IP addressing.

To configure,

1. Create a Virtual Standard Switch (VSS).
No uplink or physical NIC is required.
2. Create two port groups on the VSS:
 - EgressPortgroup: For the GigaVUE V Series Node to send RAW traffic.
 - IngressPortgroup: Enable promiscuous mode to mirror traffic to the sensor/tool.

NOTE: For details on creating a VSS and port groups, refer to [Port Group Configuration for Virtual Machines](#) section in the VMware Documentation.

3. Connect the tool or sensor interface to IngressPortgroup.



NOTE:

- Deploy the tool or sensor on the same ESXi host as the GigaVUE V Series Node.
- Use Host-Based Deployment for the GigaVUE V Series Node. For details, refer to [Deploy GigaVUE V Series Nodes using GigaVUE-FM](#).

4. Select EgressPortgroup as the tunnel interface (Network Adapter 2 or ens192) when deploying the node.
5. In the Monitoring Session, select ens192 as the RAW interface during interface mapping as in the [Create Raw Endpoint \(VMware NSX-T\)](#).

NOTE: Promiscuous mode allows Port Group to mirror the entire switch traffic.

Create a New Map

For new users, the free trial bundle expires after 30 days. GigaVUE-FM prompts you to buy a new license. For licensing information, refer to GigaVUE Licensing Guide.

A map is used to filter the traffic flowing through the GigaVUE V Series Nodes. 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.

Parameters to create a map

Parameter	Description
Rules	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.
Priority	Priority determines the order in which the rules are executed. The priority value can range from 1 to 5, with 1 being the highest and 5 is the lowest priority.

Pass	The traffic from the virtual machine is passed to the destination.
Drop	The traffic from the virtual machine is dropped when passing through the map.
Traffic Filter Maps	A set of maps to match traffic and perform various actions on the matched traffic.
Inclusion Map	An inclusion map determines the instances to include for monitoring. This map is used only for target selection.

Exclusion Map	An exclusion map determines the instances to exclude from monitoring. This map is used only for target selection.
Automatic Target Selection (ATS)	A built-in feature that automatically selects the cloud instances based on the rules defined in the traffic filter maps, inclusion maps, and exclusion maps in the Monitoring Session. The below formula describes how ATS works: Selected Targets = Traffic Filter Maps $\cap$ Inclusion Maps - Exclusion Maps Below are the filter rule types that work in ATS: • mac Source • mac Destination • ipv4 Source • ipv4 Destination • ipv6 Source • ipv6 Destination • VM Name Destination • VM Name Source • VM Tag Destination - Not applicable to Nutanix. • VM Tag Source - Not applicable to Nutanix. • VM Category Source - Applicable only to Nutanix. • VM Category Destination - Applicable only to Nutanix. • Host Name -Applicable only to Nutanix and VMware. The traffic direction is as follows: • For any rule type as Source - the traffic direction is egress. • For Destination rule type - the traffic direction is ingress. • For Hostname - As it doesn't have Source or Destination rule type, the traffic direction is Ingress and Egress. Note: • For OpenStack environment, Subnet Name Source and Subnet Name Destination are the exclusion filters available as part of Exclusion Maps with Traffic Acquisition method as OVS Mirroring in the Monitoring Domain. • If no ATS rule filters listed above are used, all VMs and vNICs are selected as targets. When any ATS rule results in a null set, no target is selected and V Series Node does not receive traffic from any VM or vNIC. • When using VM Name filter for selecting the Virtual Machines using Inclusion and Exclusion Maps, wild- cards in VM names are not supported. You must use the prefix of the Virtual Machine name.
Group	A group is a collection of pre-defined maps saved in the map library for reuse.

Rules and Notes:

- Directional rules do not work on single NIC VMs that are running a Windows UCT-V.
- Loopback captures bidirectional traffic from both ingress and egress. To prevent duplicate tapping, only egress tapping is permitted.
- If you are running GigaVUE Cloud Suite on OpenStack, you can add a subnet to the exclusion map. To do this, create an exclusion map and select the Subnet name in the ruleset.
- If a packet is fragmented then all the fragments are destined to the same application end point. You can find the stats of mapped fragmented traffic in GigaVUE-FM. For details, refer to "Review Map Statistics with Map Rule Counters" section in *GigaVUE Fabric Management Guide*.

To create a new map,

1. Drag and drop **New Map** from the **New** expand menu to the graphical workspace. The **Map** quick view appears.
2. On the new Map quick view, select the **General** tab and enter the following information:
 - a. Enter the **Name** and **Description** of the new map.
 - b. Enable the **Application Filtering** option if you wish to use Application Filtering Intelligence. Enabling this option allows you to filter traffic based on Application name or family. Refer to [Application Filtering Intelligence](#).

NOTE: Pass and Drop rule selection with Automatic Target Selection (ATS) differ with the Map type as follows: Traffic Map—Only Pass rules for ATS Inclusion Map—Only Pass rules for ATS Exclusion Map—Only Drop rules for ATS

3. Select the **Rule Sets** tab.
 - a. Perform the following to create a new rule set:
 - i. Select **Actions > New Ruleset**.
 - ii. Enter a **Priority** value from 1 to 5 for the rule with 1 being the highest and 5 is the lowest priority.
 - iii. In the **Application EndPoint ID** field, enter the Application Endpoint.
 - iv. Select a required condition from the drop-down list.
 - v. Select the rule to **Pass** or **Drop** through the map.
 - b. Perform the following to create a new rule:
 - i. Select **Actions > New Rule**.
 - ii. Select a required condition from the drop-down list.
 - iii. Click  and select **Add Condition** to add more conditions.
 - iv. Select the rule to **Pass** or **Drop** through the map.
4. Select **Save**.

Through the map, packets is dropped or passed based on the highest to lowest rule priority. You can add 5 rule sets on a map. Use the + and - buttons to add or remove a rule set in the map. Each rule set can have only 25 rules per map and each rule can have a maximum of 4 conditions.

You must select at least one rule condition to add ATS rules for an Inclusion/Exclusion map. For details, refer to [Create a New Map](#).

You can also perform the following action in the Monitoring session canvas.

- To edit a map, click the  menu button of the required map on the canvas and select **Details**, or select **Delete** to delete the map.
- To apply threshold template to maps, select the required map on the canvas and select **Details**. The quick view appears. Select the Thresholds tab. For more details on how to create or apply threshold templates, refer to [Monitor Cloud Health](#).
- Hover over the rules and apps buttons on the map to view the rule and applications configured for the selected map. Select the rules and apps buttons to open the quick view menu for RULESETS.

Example- Create a New Map using Inclusion and Exclusion Maps

Consider a Monitoring Session with 5 cloud instances. Namely target-1-1, target-1-2, target-1-3, target-2-1, target-2-2.

1. Drag and drop a new map template to the workspace.
The New map quick view appears.
2. In the **GENERAL** tab, enter the name as Map 1 and enter the description.
3. In the **RULESETS** tab, enter the priority and Application Endpoint ID.
4. Select the condition as VM Name and enter the **target**.
This includes the instances target-1-1, target-1-2, target-1-3, target-2-1, and target-2-2.
5. At the bottom of the Monitoring session canvas, select the Expand icon.
The Inclusion Maps and Exclusion Maps section appears.
6. Drag and drop a new map template to the Inclusion Maps region.
The New Map quick view appears.
7. Enter following for the map.
 - a. In the **GENERAL** tab, enter the name as Inclusionmap1 and enter the description.
 - b. In the **RULESETS**, enter the priority and Application Endpoint ID.
 - c. Select the condition as VM Name and enter the VM Name as **target-1**.
Then the instance with VM names **target-1-1, target-1-2, and target-1-3** is included.
8. Drag and drop a new map template to the Exclusion Maps region.
The New Map quick view appears.

9. Enter the details as mentioned in the above section.
 - a. In the **GENERAL** tab, enter the name as Exclusionmap1 and enter the description. In the **RULESETS** tab, enter the priority and Application Endpoint ID.
 - b. Select the condition as VM Name and enter the VM Name as **target-1-3**. Then, the instance **target-1-3** is excluded.

Based on this configuration, the Automatic Target Selection selects the instances target-1-1 and target-1-2 as target.

Map Library

Map Library is available in the **TRAFFIC PROCESSING** canvas page. You can add and use the maps from the Monitoring Session.

To add a map,

1. From the **Monitoring Session** screen, select **TRAFFIC PROCESSING**.

The GigaVUE-FM Canvas page appears.

2. From the page, select the desired map and save it as a template.
3. Select **Details**.

The Application quick view appears.

4. Select **Add to Library** and perform one of the following:

- From the **Select Group** list, select an existing group.
- Select **New Group** to create a new one.

5. In the **Description** field, add details, and select **Save**.

The map is added to Map Library. You can use the added map for all the monitoring sessions.

Reusing a map

From the **Map Library**, drag and drop the saved map.

Deploy Monitoring Session

You can deploy the Monitoring Session on all the nodes and view the report.

To deploy the Monitoring Session,

1. Add components to the canvas

Drag and drop the following items to the canvas as required:

- **Ingress tunnel** (as a source): From the **New** section.
- **Maps**: From the **Map Library** section.
- **Inclusion and Exclusion maps**: From the Map Library to their respective section at the bottom of the workspace.
- GigaSMART **apps**: From the **Applications** section.
- **Egress tunnels**: From the **Tunnels** section.

2. Connect components

Perform the following steps after placing the required items in the canvas.

- a. Hover your mouse on the map
- b. Select the dotted lines
- c. Drag the arrow over to another item (map, application, or tunnel).

You can drag multiple arrows from a single map and connect them to different maps.

3. (Optional) Review Sources

Select the SOURCES tab to view details about the subnets and monitored instances.

The monitored instances and the subnets are visible in orange.

Not applicable for NSX-T solution and Customer Orchestrated Source as Traffic Acquisition Method.

4. Deploy the Monitoring Session

From the **Actions** menu, select **Deploy**.

After successful deployment on all the V Series Nodes, the status appears as **Success** on the **Monitoring Sessions** page.

View the Deployment Report

You can view the Monitoring Session Deployment Report in the **SOURCES** and **V SERIES NODES** tab.

- When you select the **Status** link, the Deployment Report is displayed.
- When the deployment is incorrect, the Status column displays one of the following errors:
 - **Success**: Not deployed on one or more instances due to V Series Node failure.
 - **Failure**: Not deployed on all V Series Nodes or Instances.

The **Monitoring Session Deployment Report** displays the errors that appeared during deployment.

The Monitoring Session Deployment includes two key configuration:

- [Interface Mapping](#)
- [Tool Exclusion](#)

Interface Mapping

It allows to associate specific network interfaces (from monitored instances) with monitoring tools. This ensures that traffic from selected sources is accurately mirrored and routed for analysis. You can:

- Select interfaces from available instances.
- Map each interface to one or more monitoring tools.
- Apply filters or conditions to refine traffic selection.

Tool Exclusion

It excludes specific monitoring tools from receiving mirrored traffic during a monitoring session. This option is available only when the Traffic Acquisition method is set to **VPC Traffic Mirroring**.

Deploy Monitoring Session

INTERFACE MAPPING | **TOOL EXCLUSION**

Tool instances should be excluded to avoid traffic looping. Review the instances with the same IP address below and select the tool instance to exclude.

IP ADDRESS	TOOL EXCLUSION
10.10.10.100	Excluded
10.10.10.200	--
10.10.10.300	Excluded

VM NAME	ID
VM100	i-0cae6ab7c57a9d237
☒ Tool	i-0cae6ab7c57a9d328
VM200	i-0cae6ab7c57a9f395

Cancel Deploy

- Review the list of available monitoring tools.
- Select the tools to exclude from traffic flow.
- Confirm the exclusion before deploying the session.

V Series Applications

GigaVUE V Series Node is a virtual machine running in the customer's infrastructure which processes and distributes network traffic. It plays the same role as an HC Series appliance in a physical deployment, running many of the same GigaSMART applications and feeding data to tools in a similar manner. Because GigaVUE V Series nodes reside in a virtual environment, inbound and outbound traffic is tunneled (because there are no physical device ports).

GigaVUE V Series Applications are hosted on GigaVUE V Series Nodes and are licensed through a Volume-Based License (VBL). These applications help optimize the traffic forwarded from your instances to the monitoring tools. To use them, configure the applications on the Monitoring Session canvas, then deploy the Monitoring Session to route the traffic accordingly.

The following table summarizes GigaSMART feature support across GigaVUE Cloud Suite platforms, covering packet operations, tunnel handling, application intelligence, and 5G capabilities.

GigaSMART Operation	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
Application Filtering Intelligence (APF without RegEx+ASF)	✓	✓	✓	✓	✓	✓	✓
Application Metadata Exporter	✓	✓	✓	✓	✓ (Only when deploying GigaVUE V Series Node using Third party)	✓	✓

GigaSMA RT Operatio n	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestrati on	GigaVUE Cloud Suite for Nutanix
					Orchestrati on)		
Application Metadata Intelligence	✓	✓	✓	✓	✓	✓	✓
Application Visualization	✓	✓	✓	✓	✓	✓	✓
De- duplication	✓	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Encapsulati on Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospectiv e GigaVUE Cloud Suite Deployment Guide.	✓	x	✓	✓	✓	✓	✓
VXLAN Tunnel Encapsulati on Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospectiv	✓	✓	✓	✓	✓	✓	✓

GigaSMA RT Operatio n	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenSta ck	GigaVUE Cloud Suite for VMwar e (VMwa re vCente r)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestrati on	GigaVUE Cloud Suite for Nutani x
See GigaVUE Cloud Suite Deployment Guide.							
L2GRE Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✓
VXLAN Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓	✓
ERSPAN Tunnel	✓	✗	✓	✓	✓	✓	✓

GigaSMART Operation	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestration	GigaVUE Cloud Suite for Nutanix
Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.							
UDP GRE Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✗
GigaSMART NetFlow Generation	✓	✓	✓	✓	✓	✓	✓
GRE-In-UDP Tunnel Decapsulation	✗	✗	✓	✓	✗	✓	✗

GigaSMA RT Operatio n	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenSta ck	GigaVUE Cloud Suite for VMwar e (VMwa re vCent er)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestrati on	GigaVUE Cloud Suite for Nutani x
GENEVE Decapsulati on	✓	✗	✗	✗	✓	✗	✗
Header Stripping	✓	✓	✓	✓	✓	✓	✓
Load Balancing	✓	✓	✓	✓	✓	✓	✓
Masking	✓	✓	✓	✓	✓	✓	✓
SSL Decrypt	✓	✓	✓	✓	✗	✓	✓
Secure Tunnels Encapsulati on	✗	✓	✗	✓	✓	✓	✓
Secure Tunnels Decapsulati on	✗	✓	✗	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓	✓	✓
5G-Cloud Application	✗	✗	✓	✓	✗	✓	✗
5G-Service Based Interface Application	✗	✗	✓	✓	✓	✓	✗

For more details about supported V Series applications, refer to [Supported V Series Applications](#).

- Application Visualization
- Application Filtering Intelligence
- Application Metadata Intelligence
- Slicing
- Masking
- De-duplication

- Load Balancing
- PCAPng Application
- Header Stripping
- Application Metadata Exporter
- SSL Decrypt
- GigaSMART NetFlow Generation
- 5G-Service Based Interface Application
- 5G-Cloud Application

For details on configuring these applications, refer to *GigaVUE V Series Applications Guide*.

Masking

Masking lets you overwrite specific packet fields with a specified pattern so that sensitive information is protected during network analysis. Masking operations consist of an **offset**, **length**, and **pattern**.

To add a Masking application:

1. Drag and drop **Masking** from **APPLICATIONS** to the graphical workspace.
2. Click the Masking application and select **Details**. The Application quick view appears.

masking



X

Application

Save

Application	Masking
Alias	masking
Protocol	none
Offset	0
Pattern	0x
Length	8

3. In the Application quick view, enter the information as follows:

Component	Description
Alias	Enter a name for the application.
Protocol	The following are the protocols that you can select from the protocol drop-down list: - o None – Mask starting a specified number of bytes from the start of the packet. - o IPV4 – Mask starting a specified number of bytes after the IPv4 header. - o IPV6 – Mask starting a specified number of bytes after the IPv6 header. - o UDP – Mask starting a specified number of bytes after the UDP header. - o TCP – Mask starting a specified number of bytes after the TCP header. - o ftp-data – Identify using TCP port 20. Mask payloads using offset from the TCP header. - o HTTPS – Identify using TCP port 443. Mask payloads using offset from the TCP header. - o SSH – Identify using TCP port 22. Mask payloads using offset from the TCP header. - o GTP – Mask starting a specified number of bytes after the outer GTP header. - o GTP-IPV4 – Mask starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. - o GTP-UDP – Mask starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. - o GTP-TCP – Mask starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Offset	Specifies where the application should start masking data with the supplied pattern. You can specify this in terms of either a static offset from the start of the packet or a relative offset from a particular protocol layer. This lets you automatically compensate for variable length headers, specifying a mask target in terms of a particular packet header.
Length	Specifies how much of the packet should be masked. The specified one-byte pattern can be repeated to mask from 1-9600 bytes.
Pattern	Specifies what pattern the application should use to mask the specified portion of the packet. You can specify a one-byte hex pattern (for example, 0xFF).

4. Click **Save**.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on Traffic health monitoring and how to create threshold template, refer to Traffic Health Monitoring section in the respective cloud deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to Monitor Cloud Health section in the respective cloud deployment guides.

You can view the statistics of the application in the **STATISTICS** tab.

Slicing


Packet Slicing allows you to truncate packets after a specified header and slice length, preserving the portion of the packet required for monitoring purposes. Slicing operations are typically configured to preserve specific packet header information, allowing effective network analysis without the overhead of storing full packet data.

Packets can have multiple variable-length headers, depending on where they are captured, the different devices that have attached their own headers along the way, and the protocols in use (for example, IPv4 versus IPv6). Because of this, slicing operations with a hard-coded offset will not typically provide consistent results.

To address this, the Slicing application lets you configure Packet Slicing using protocols that allow you to start slicing from a particular number of bytes after a specific packet header (IPv4, IPv6, UDP, and so on). The Slicing application parses through Layer 4 (TCP/UDP) to identify the headers in use, slicing based on the variable offset identified for a particular header instead of a hard-coded number of bytes.

To add a Slicing application:

1. Drag and drop **Slicing** from **APPLICATIONS** to the graphical workspace.
2. Click the Slicing application and select **Details**. The Application quick view appears.



The screenshot shows the 'Application' quick view for the 'Slicing' application. The interface includes a sidebar with a 'slicing' icon, a header with a close button and a 'Save' button, and a main area with configuration fields.

Field	Value
Application	Slicing
Alias	slicing
Protocol	none
Offset	64
Enhanced Name	none

3. In the Application quick view, enter the information as follows:

Component	Description
Alias	Enter a name for the application.
Protocol	The following are the protocols that you can select for from the protocol drop-down list: - o None – Slice starting a specific number of bytes from the start of the packet. - o IPV4 – Slice starting a specified number of bytes after the IPv4 header. - o IPV6 – Slice starting a specified number of bytes after the IPv6 header. - o UDP – Slice starting a specified number of bytes after the UDP header. - o TCP – Slice starting a specified number of bytes after the TCP header. - o FTP – Identify using TCP port 20 and slice payloads using offset from the TCP header. - o HTTPS – Identify using TCP port 443. Slice encrypted payloads using offset from the TCP header. - o SSH – Identify using TCP port 22. Slice encrypted payloads using offset from the TCP header. The Slicing application can provide slicing for GTP tunnels, provided the user payloads are unencrypted. Both GTPv1 and GTPv2 are supported – GTP' (GTP prime) is not supported. Keep in mind that only GTP-u (user plane packets) are sliced. Control plane packets (GTP-c) are left unmodified because of their importance for analysis. - o GTP – Slice starting a specified number of bytes after the outer GTP header. - o GTP-IPV4 – Slice starting a specified number of bytes after the IPv4 header inside the encapsulating GTP packet. - o GTP-UDP – Slice starting a specified number of bytes after the UDP header inside the encapsulating GTP packet. - o GTP-TCP – Slice starting a specified number of bytes after the TCP header inside the encapsulating GTP packet.
Offset	Specify the length of the packet that must be sliced.

4. Click **Save**.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on Traffic health monitoring and how to create threshold template, refer to Traffic Health Monitoring section in the respective cloud deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to Monitor Cloud Health section in the respective cloud deployment guides.

You can view the statistics of the application in the **STATISTICS** tab.

Application Intelligence

Application Intelligence provides a comprehensive solution that:

- identifies the applications contributing to the network traffic.
- isolates preferred application-specific traffic and directs it to the appropriate tools.
- exports relevant application metadata for further analytics and analysis.

Application Intelligence provides the following capabilities for virtual nodes:

- [Application Visualization](#)
- [Application Filtering Intelligence](#)
- [Application Metadata Intelligence](#)
- [Application Metadata Exporter](#)

Points to Note for Application Intelligence

Point to note when configuring Application Intelligence:

1. For a monitoring domain, the following application can be configured only once, and all these applications must be configured in a single monitoring session.
 - a. Application Visualization
 - b. Application Filtering
 - c. Application Metadata
2. For GigaVUE V Series Node version lesser than 6.3.00, Application Visualization, Application Filtering, Application Metadata, and Application Metadata Exporter (AMX) applications are not supported in the Monitoring Session.

3. When undeploying and redeploying the Monitoring session which has the Application Intelligence application, ensure to follow the steps given below:
 - a. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Sessions** page appears. Select the Monitoring Session for which you enabled Secure Tunnels. Click **Actions > Undeploy**. The monitoring session is undeployed.
 - b. Select the Monitoring Session for which you enabled Secure Tunnels. Click **Actions > Edit**. The **Edit Monitoring Session** Canvas page appears.
 - c. Add the Application Intelligence applications.
 - d. Modify the Number of Flows as per the below table:

Cloud Platform	Instance Size	Maximum Number of Flows
VMware	Large (8 vCPU and 16GB RAM)	200k
AWS	AMD - Large (c5n.2xlarge)	300k
	AMD - Medium (t3a.xlarge)	100k
	ARM - Large (c7gn.2xlarge)	100k
	ARM - Medium (m7g.xlarge)	200k
Azure	Large (Standard_D8s_V4)	500k
	Medium (Standard_D4s_v4)	100k
Nutanix	Large (8 vCPU and 16GB RAM)	200k

- Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
 - If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.
- e. Click **Deploy**.
4. After adding the above-listed applications and deploying the Monitoring Session, you cannot edit the Number of flows and Fast Mode. For more detailed information on Number of flows and Fast Mode, refer to [Number of Flows](#) and [Fast Mode](#).
 5. Once the Number of flows is added in any of the above-listed applications, the same value is applied to all the above-listed applications configured in that Monitoring Session. You cannot change it.
 6. Once Fast Mode is enabled in any of the above-listed applications, then it is enabled for all the above-listed applications configured in that Monitoring Session. You cannot change it.

7. You can also configure Application Intelligence with Precryption, prefiltering, and secure tunnels. Refer to Precryption, Prefiltering, and Secure Tunnels topics in the respective cloud deployment guides for more detailed information on how to configure these features.
8. Small Form Factor for VMware ESXi is not supported when using applications like Application Visualization, Application Metadata, Application Filtering. Refer to *Configure GigaVUE V Series Nodes for VMware ESXi* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more detailed information on how to deploy GigaVUE V Series Node, where you select the Form Factor.

Application Visualization

Application Visualization identifies and monitors all applications contributing to the network traffic and reports on the total applications and the total bandwidth they consume over a select period. Application Visualization allows you to identify more than 3,200 applications. It displays the traffic statistics in bytes and packets.

Refer to the following topics for more detailed information on how to configure the application and view the statistics:

- [Configure Application Visualization for Virtual Environment](#)
- [Configure Filtering and Metadata Export for Selected Applications in Application Visualization](#)
- [View Application Statistics for Application Visualization](#)

Configure Application Visualization for Virtual Environment

Application Visualization can be configured in the **Edit Monitoring Session** Canvas Page. To add an Application Visualization application to the canvas, follow the steps given below:

1. Drag and drop **Application Visualization** from **APPLICATIONS** to the graphical workspace.
2. Click the Application Visualization application and select **Details**. The Application

quick view appears.

3. In the Application quick view, enter or select the following details:

Parameter	Description																							
Name	Enter a name for the application.																							
Description	Enter the description.																							
Export Interval	The time interval in seconds at which the export must be done. The export interval is set to 300 seconds. It cannot be modified.																							
Advanced Settings																								
Number of Flows	The number of flows supported by the application. Refer to the following table for the maximum number of flows supported for VMware, AWS, Nutanix, and Azure platforms. <table border="1"> <thead> <tr> <th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr> </thead> <tbody> <tr> <td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> <tr> <td rowspan="4">AWS</td><td>AMD - Large (c5n.2xlarge)</td><td>300k</td></tr> <tr> <td>AMD - Medium (t3a.xlarge)</td><td>100k</td></tr> <tr> <td>ARM - Large (c7gn.2xlarge)</td><td>100k</td></tr> <tr> <td>ARM - Medium (m7g.xlarge)</td><td>200k</td></tr> <tr> <td rowspan="2">Azure</td><td>Large (Standard_D8s_v4)</td><td>500k</td></tr> <tr> <td>Medium (Standard_D4s_v4)</td><td>100k</td></tr> <tr> <td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> </tbody> </table> ▪ Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k.. ▪ If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.	Cloud Platform	Instance Size	Maximum Number of Flows	VMware	Large (8 vCPU and 16GB RAM)	200k	AWS	AMD - Large (c5n.2xlarge)	300k	AMD - Medium (t3a.xlarge)	100k	ARM - Large (c7gn.2xlarge)	100k	ARM - Medium (m7g.xlarge)	200k	Azure	Large (Standard_D8s_v4)	500k	Medium (Standard_D4s_v4)	100k	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows																						
VMware	Large (8 vCPU and 16GB RAM)	200k																						
AWS	AMD - Large (c5n.2xlarge)	300k																						
	AMD - Medium (t3a.xlarge)	100k																						
	ARM - Large (c7gn.2xlarge)	100k																						
	ARM - Medium (m7g.xlarge)	200k																						
Azure	Large (Standard_D8s_v4)	500k																						
	Medium (Standard_D4s_v4)	100k																						
Nutanix	Large (8 vCPU and 16GB RAM)	200k																						
Monitoring	You can use this option, enable or disable the Application Visualization application functionality.																							
Fast Mode	Enable the Fast Mode option for performance (less CPU cycles and less memory utilization) improvement. When the Fast Mode is enabled, some or all of the attributes of the applications will be disabled. If all the attributes of the application are disabled then the application itself is disabled. Note: This option is disabled for NetVUE Base Bundle License.																							

4. Click **Save**.

Configure Filtering and Metadata Export for Selected Applications in Application Visualization

This section describes how to perform filtering and exporting metadata for selected applications when configuring Application Visualization. Refer to the following steps for more detailed information:

- [Filter Traffic for Selected Applications](#)
- [Export Metadata for Selected Applications](#)

Filter Traffic for Selected Applications

1. Click the Application Visualization application in the Monitoring Session Canvas page and select **Details**. The Application quick view appears.
2. Click on **STATISTICS** tab.
3. Click on the drop-down menu that displays **Top 10** and select the **All**.
4. Select the applications for which you want to filter traffic.

NOTE: Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each application.

5. Click **Actions > Filter Selected Applications**. The **Filter Selected Applications** dialog box opens.
6. In the **Filter Selected Applications** dialog box, Select the existing Application Filtering map or New Map from the **Send to Map** drop-down menu.
 - **New Map:** Select this option if you wish to create a new application filtering map, to filter the traffic from the applications.
 - **Existing map:** Select this option if you have already configured an application filtering map and you wish to send the traffic for filtering to that map.
7. Under the **Applications** section, choose the traffic as pass or drop for the selected applications.
8. Click **Send to Map**.

NOTE: If the **New Map** option is selected from the **Send to Map** drop-down menu. Then, the map quick view appears. Refer to step 3, 4, 5, 6, and 7 in [Configure Application Filtering Intelligence for Virtual Environment](#) section for more detailed instructions on how to configure Application Filtering.

Export Metadata for Selected Applications

1. Click the Application Visualization application and select **Details**. The Application quick view appears.
2. Click on **STATISTICS** tab.
3. Click on the drop-down menu that displays **Top 10** and select the **All**.

4. Select the applications for which you want to export metadata.

NOTE: Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each application.

5. Click **Actions > Export Metadata for Selected Applications**. The **Export Metadata for Selected Applications** dialog box opens.
6. In the **Export Metadata for Selected Applications** dialog box, Select the existing Application Metadata application or New Exporter from the **Send to Exporter** drop-down menu.
 - **New Application Metadata:** Select this option if you wish to create a new application metadata, to export metadata from the applications.
 - **Existing Application Metadata Application:** Select this option if you have already configured an Application Metadata Intelligence application and you wish to send the traffic for exporting metadata from the applications.
7. Click **Export**.

NOTE: If the **New Application Metadata** option is selected from the **Send to Exporters** drop-down menu. Then, the application quick view appears. Refer to steps 3 and 4 in [Configure Application Metadata Intelligence for Virtual Environment](#) section for more detailed instructions on how to configure Application Metadata application.

View Application Statistics for Application Visualization

To view the application Statistics for the Application Visualization application, follow the steps given below:

1. Click **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. Select a monitoring session from the list view, click **Actions > Edit**. The Edit Monitoring Session page appears.
3. Click the application and select **Details**. The Application quick view appears.
4. Click on **STATISTICS** tab.

5. You can view the following in the statistics page:
 - a. **Total Traffic:** Displays the total traffic of the network. Use the drop-down menus to change the parameters. You can use the **Select Tags** filter option to filter the traffic related to the selected application tags.
 - b. **Total Applications :** You can view all the applications and their bandwidth in the network. You can also perform filtering and exporting metadata for selected applications. Refer to [Configure Filtering and Metadata Export for Selected Applications in Application Visualization](#) for more detailed information on how to perform filtering and exporting metadata for selected applications.
 - c. **Top 10 Applications:** Click on the drop-down menu that displays **All** and select the **Top 10** . Displays the Top 10 applications running in the network based on the metrics.
 - d. **Top 10 Application Families:** You can view a graphical representation of top 10 applications running in the network based on the metrics. When you hover over the Pie-chart, GigaVUE-FM shows the application families in the network.

You can view the statistics for past hour, past 24 hours, or past 7 days. GigaVUE-FM also allows you to view statistics for a particular period by selecting the date and time. The selected data and time must be within the past 7 days.

GigaVUE-FM takes more than five minutes to display the application statistics since the export interval is fixed at five minutes. For the first fifteen minutes after creating the solution, if GigaVUE-FM receives traffic, it will show real-time data. If there is no traffic during this time, it will take at least eleven minutes to display the statistics once traffic is received.

User Defined Application

This feature allows you to identify unclassified TCP, UDP, HTTP, and HTTPS applications and extract their application name and ID.

Refer to the following topic for more detailed information:

- [Supported Protocols and Attributes](#)
- [Metadata](#)
- [Supported RegExp Syntax](#)
- [Limitations](#)
- [Create Rules for User Defined Application](#)

Supported Protocols and Attributes

The DPI engine will match the rules defined based on the following protocols and attributes within the first 500 bytes of a packet payload.

For supported Regex patterns, refer [Supported RegExp Syntax](#)

Protocol	Attributes	Attribute Labels	Description	Direction	Supported Data Type	Example Value
http	cts-uri	Request URI	Partially Normalized URL (path + request)	Client to Server Only	REGEX P	\fupload\(create_file new_slice upload_slice)\?.*upload_token=.*
	cts-server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEX P	(.*\.)?gigamon\.com
	mime_type	MIME Type	Content type of Request or the Web page	Both, Client to Server or Server to Client	REGEX P	http
	cts-user_agent	User Agent	Software / Browser used for request	Client to Server Only	REGEX P	mozilla
	cts-referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEX P	http:\Vgigamon.com\
	stc-server_agent	Server Agent	Software used for the server	Server to Client Only	REGEX P	NWS_TCloud_PX
	stc-location	Redirect Location	Destination address where	Server to Client Only	REGEX P	.*\Vfootball\.*

			the client is redirected to			
	cts-cookie	Cookie (Raw)	Raw value of the HTTP Cookie header line	Client to Server Only	REGEX P	.*tEstCoOkie.*
	content	Content	Message body content	Both, Client to Server or Server to Client	REGEX P	.*GIGAMON.* mindata = 206 Refer Mindata
http2	cts-uri	Request URI	Partially Normalized URL (path + request)	Client to Server Only	REGEX P	\fupload\(create_file new_slice upload_slice)\?.*upload_token=.*
	cts-server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEX P	(.*\.)?gigamon\.com
	cts-user_agent	User Agent	Software / Browser used for request	Client to Server Only	REGEX P	mozilla
	cts-referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEX P	http:\Vgigamon.com\

ssl	common_name	Domain Name	Domain name from Client Hello message or the certificate		REGEX P	(.*\.)?gigamon\.com
	stc-subject_alt_name	Subject Alt Name (s)	List of host names which belong to the same certificate	Server to Client Only	REGEX P	(.*\.)?gigamon\.com
rtmp	cts-page_url	Page URL	URL of the webpage where the audio/video content is streamed	Client to Server Only	REGEX P	http://www.music.tv/recorded/1234567
tcp	stream	Payload Data	Data payload for a packet, excluding the header.		REGEX P	.*GIGAMON.* mindata = 70 Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEX P String	80-4350

udp	stream	Payload Data	Data payload for a packet, excluding the header		REGEX P	.*GIGAMON.* mindata = 100 Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEX P String	80-4350
sip	user_agent	User Agent	Software used	Both, Client to Server or Server to Client	REGEX P	GVUE-release 6.2.0
icmp	code	Message Code	Code of the ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEX P String	200
	typeval	Message Type	Type of ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEX P String	10
ip	address	Server IP Address	IP address of the server		IPV4 as REGEX P String	62.132.12.30/24

	dscp	DSCP Value	DSCP from Differentiated Service (DS) Field in IP header		UINT8 as REGEX P String	33
	resolv_name	DNS Name	Server's DNS name		REGEX P	gigamon.com
ipv6	address	Server IP Address	IP address of the server		IPV6 as REGEX P String	2001:0:9d38:6ab8:307b:16a4:9c66:5f4 2001:0:9d38::9c66:5f4/64
	dscp	DSCP Value	DSCP from Differentiated Service (DS) Field in IP header		UINT8 as REGEX P String	43

Mindata

The mindata value is the number of payload bytes to buffer and match a given pattern. You can configure mindata value for HTTP content, TCP stream, and UDP stream. The buffer size is calculated from the start of the payload and the default buffer size is different for each protocol (HTTP - 206, TCP - 67, and UDP - 48.)

For example, for pattern `".*TEST.*"` that may be present within the first 67 bytes of TCP payload, you can specify the mindata value as 4 (which is the length of the input string) or as 67 (which is the default buffer size of TCP payload). In case, the pattern is present in between 65 to 68 bytes of the payload and the mindata is specified as 4 or 67, it will not match. For this case, you must specify the mindata value as 68.

Supported RegExp Syntax

Pattern	Description
.	Matches any symbol

*	Searches for 0 or more occurrences of the symbol or character set that precedes it
+	Searches for 1 or more occurrences of the symbol or character set that precedes it
?	Searches for 0 or 1 occurrence of the symbol or character set that precedes it
()	Groups a series of expressions together
[]	Matches any value included within the bracket at its current position Example: [Dd]ay matches Day and day
 [<start>-<end>]	Separates values contained in (). Searches for any one of the values that it separates. Example: The following expression matches dog or cat: (dog cat). Matches any value contained within the defined range (a hyphen indicates the range). You can mix character class and a hexadecimal range Example: [AaBbCcDdEeFf0-9]
\0 <octal_ number>	Matches for a direct binary with octal input
\x<hexadecimal- number>\x	Matches for a direct binary with hexadecimal input
\[<character- set>\]	Matches a character set while ignoring case. WARNING: Not performance friendly

Limitations

- The maximum number of user defined application that can be configured is 120 per GigaVUE-FM. These applications can be spread across one or more application intelligence sessions.
- The maximum number of rules that can be created per application is 8.
- The maximum number of protocols that can be configured per rule is 3.

Create Rules for User Defined Application

To create a new application:

You can create rules for User Defined Application in two ways:

1. Go to **Inventory > Resources > User-defined Applications**. Refer to User Defined Applications section in GigaVUEGigaVUE Fabric Management Guide for details.
2. Go to **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**. Follow the steps listed below:
 - a. Select a Monitoring Session from the list view. Navigate to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas page appears.
 - b. In the canvas, click on the  icon on the left side of the page to view the traffic processing elements.
 - c. Select **User Defined Applications** under **Options** menu.
 - d. Enable the **User-defined Applications** toggle button.
 - e. Click **New Application**. The New Application page appears.
 - f. Enter the **User-Defined Application Name**.
 - g. Enter **Priority**. The value must be between 1 and 120.

NOTE: The lowest value has the highest priority.

- h. In the Rules dialog box, select the following details:
 - Choose the **Protocol** from the list of protocols.
 - Choose the **Attributes** from the list of attributes.
 - Choose the **Values** from the list of values.

Using the **Actions** Button, you can perform the following actions:

- i. Click **Save**.

To add the created applications to the Monitoring Session:

1. In the **User Defined Applications** tab, click Add Application button.
2. Select the applications that must be added.
3. Click **Done**.

After creating rules for User defined Applications, you can add it to Application Filtering when configuring the applications. Refer to [Add Application](#) section for more detailed information on how to add User defined Application when configuring Application Filtering.

Application Filtering Intelligence

Application Filtering Intelligence allows filtering of traffic based on the application (such as YouTube, Netflix, Sophos, or Facebook) or application family (such as antivirus, web, erp, or instant-messaging) or application tags. Enables traffic filtering by layer 7

applications, which means you can filter out high-volume, low-risk traffic from reaching the tools and distribute high-risk network traffic of interest to the right tool at the right time.

Refer to the following topics for more detailed information and step-by-step instructions on how to configure Application Filtering Intelligence application and view the statistics:

- [Configure Application Filtering Intelligence for Virtual Environment](#)
- [View Application Statistics for Application Filtering](#)

Configure Application Filtering Intelligence for Virtual Environment

Application Filtering Intelligence (AFI) can be configured in the Monitoring Session Canvas. To add Application Filtering application to the canvas, follow the steps given below:

1. Drag and drop **New Map** from **New** to the graphical workspace.
2. Click the application and select **Details**. The Application quick view appears.
3. Enable **Application Filtering** in the **GENERAL** tab.

4. In the Application quick view, enter or select the following details in the **GENERAL** tab:

Parameter	Description
Name	Enter a name for the application.
Description	Enter the description.
Application Filtering Settings	
Bidirectional	Enable or Disable Bi-Directional Flow behavior. Bi-Directional is enabled by default. Disable this option for Uni-Directional Flow behavior.
Timeout	Specify the traffic flow inactivity timeout, in seconds. The session will be removed due to inactivity when no packets match.
Buffer	This option is enabled by default.
Buffer Count Before	Number of packets that should be buffered until the flow is identified. If the flow is not identified even after reaching the maximum number of packets buffered, then all the subsequent packets of this session will be dropped.
Protocol	Select the Protocol. The packet matching the selected protocol will be filtered. The default value is TCP-UDP.
Packet Count	Enable or Disable Packet Count. Packet Count is disabled by default.
Number of packets Note: This field appears only when Packet Count field is enabled.	Specifies the number of packets to forward to the tool port for each session match. After the packet count is reached, subsequent packets for the session are dropped. The packet count includes the packet that triggered the creation of the session. The default is disable, which means that all packets will be forwarded to the tool port. The range is from 2 to 100.
Session Fields	
Session Field	The Packet fields to be considered for creating the Session / traffic flow (Session key fields)
Action	Add or Remove 'VlanId' Packet field for creating the session / traffic flow.

Parameter	Description
Advanced Settings	
Number of Flows	The number of flows supported by the application. Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms.

- Click the **RULESETS** tab. Through the map, packets can be dropped or passed based on the highest to lowest rule priority. You can add 5 rule sets on a map. Each rule set can have only 25 rules per map and each rule can have a maximum of 4 conditions.

Enter the following details for each of the Rule Set created:

Parameter	Description
Priority	A priority determines the order in which the rules are executed. The priority value can range from 1 to 5, with 1 being the highest and 5 is the lowest priority.
AE ID	Application Endpoint ID will be used as source or destination object for creating or connecting links
Actions	Using this option, you can perform the following functions: • New Ruleset- Use to add a new Rule Set. Note: A maximum of 5 Rule Sets can be created. • New Rule- Use to add a New Rule Note: A maximum of 25 Rules can be created per rule set. • Delete this Ruleset- Use to delete the Ruleset
RULES	
Rule	Use the toggle button to Pass or Drop the traffic through the map.

Parameter	Description
Condition	Select any one of the conditions from the drop-down menu and search or select the attributes. Use the + and - buttons to add or remove a condition with a Rule. Click  and select Add Condition to add more conditions. Note: A maximum of 4 conditions can be created per Rule.
APPLICATION FILTERING Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each application.	
Add Application	Click on the Add Application button. The Add Application dialog box opens. Select a Type. The available options are: • Application Family: Each application is mapped only mapped to one Application Family ■ Select an Application Family and the Applications that needs to be filtered from the traffic. ■ In the Traffic Action column, select Pass or Drop to pass or drop the traffic. You can also use Pass All or Drop All to allow or drop the traffic for all the applications. • Application Tag: Each application can be mapped to one or more Application Tags. ■ Select an Application Tag and the Applications that needs to be filtered from the traffic. ■ In the Traffic Action column, select Pass or Drop to pass or drop the traffic. You can also use Pass All or Drop All to allow or drop the traffic for all the applications. User Defined Applications: To configure User Defined Applications for AFI, follow the steps given below. - Enable User Defined Applications toggle button in the Options page. Refer to User Defined Application topic for more detailed information on what is user defined applications and how to configure it. - In this Add Application dialog box, select User Defined Applications from the Application Family list.

- To pass or drop any remaining traffic in the network, enter the priority and AE ID in the default rule set available. Select **Pass** or **Drop** option for **Any Remaining Traffic** field.
- Click the **THRESHOLDS** tab. For more details on how to create and apply threshold template, refer to Traffic Health Monitoring section in the respective cloud deployment guides.

8. To reuse the configuration, click **Add to Library**. Save the application filtering configurations using one of the following ways:
 - a. Select an existing group from the **Select Group** list or create a **New Group** with a name.
 - b. Enter a description in the **Description** field, and click **Save**.
 - c. The saved map can be found in the **Map Library** in the **Edit Monitoring Session** Canvas Page.
9. Click **Save**.

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

When using the Application Filtering application, you can either use a single tunnel, to tunnel all the filtered traffic from the application or use a separate tunnel for each rule configured.

View Application Statistics for Application Filtering

To view the application Statistics for the Application filtering application, follow the steps given below:

1. Click **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. Select a monitoring session from the list view, click **Actions > Edit**. The Edit Monitoring Session page appears.
3. Click the application and select **Details**. The Application quick view appears.
4. Click on **STATISTICS** tab.
5. You can view the following in the Application Filtering application statistics page:
 - a. Rules - Displays the rules created in this application.
 - b. Pass App (Bytes) - Displays the packet that pass through the applications selected.
 - c. Drop App (Bytes) - Displays the packets that are dropped by the applications selected.
 - d. Pass Rule (Bytes) - Displays the packet that pass through the rule sets configured.
 - e. Drop Rule (Bytes) - Displays the packets that are dropped by the rule sets configured.

Application Metadata Intelligence

Application Metadata Intelligence allows you to export metadata from applications that are detected in the network traffic. The records can be exported to a collector either in IPFIX or CEF format through the IP interface or the management interface. You can also use the application metadata attributes for purposes other than security, such as to determine the network or application health, to track the long-lived sessions seen in the network, and so on.

Application Metadata Intelligence generates more than 6000 attributes for more than 4000 applications without impacting the users, devices, applications, or the network appliances. The feature identifies applications even when the traffic is encrypted.

Application Metadata Intelligence (AMI) is enabled to multi-collect protocols with more than one metadata attribute of the same type. The multi-collect feature supports additional protocols such as DNS, GTP, GTPV2, DHCP, HTTP, HTTPS, SSL, HTTP_PROXY, HTTP2, KERBEROS5, and DHCP6.

The generated metadata is exported in IPFIX (IP Flow Information Export) format and CEF (Common Event Format) to security analytics and forensics tools thereby providing greater visibility to enforce corporate compliance.

The output from the Application Metadata Intelligence in CEF format can also be converted to JSON format using Application Metadata Exporter (AMX) application. To learn more about AMX application refer to Application Intelligence—Application Metadata Exporter

Refer to following topics for more detailed information and step-by-step instructions on how to configure Application Metadata Intelligence and view the statistics:

- [Configure Application Metadata Intelligence for Virtual Environment](#)
- [View Application Statistics for Application Metadata](#)

You can convert the output from the Application Metadata Intelligence (AMI) which is in CEF format into JSON format and send it to the cloud tools and Kafka. Refer to [Application Metadata Exporter](#) for detailed information on AMX and how to configure it.

Configure Application Metadata Intelligence for Virtual Environment

Application Metadata Intelligence (AMI) can be configured in the Monitoring Session Canvas. To add Application Metadata Intelligence application to the canvas, follow the steps given below:



Notes:

- Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
- Flow Direction/Behaviour under Application Metadata Settings must be configured as **Bidirectional** for exporting application Metadata. Each exporter can be configured to export application metadata for up to 256 applications, and up to 64 attributes can be exported for each application. GigaVUE-FM allows selecting metadata for applications that support exporting metadata. Applications which do not support exporting metadata are not listed in the GigaVUE-FM. By default, application name is exported for all such applications.

1. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows**. Select the required cloud platform.
2. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
3. Drag and drop **Application Metadata** from **APPLICATIONS** to the graphical workspace.
4. Click the Application Metadata application and select **Details**. The Application quick view appears.

5. In the Application quick view, enter or select the following details in the **General** tab:

Parameter	Description
Name	Enter a name for the application.
Description	Enter the description.
Application Metadata Settings	
Flow Direction	Enable or Disable Bi-Directional Flow behavior. Bi-Directional is enabled by default. Disable this option for Uni-Directional Flow behavior.
Timeout	Specify the traffic flow inactivity timeout, in seconds. The session will be removed due to inactivity when no packets match.
Multi Collect	Enable: Enables the multi-collect of attributes within a given Metadata Store cache which means that if a configured attributes is seen in multiple packets within the same flow, each of these information is collected. Multi Collect is enabled by default, when a new cache is created. Multi Collect is enabled, when upgraded from an older release. Disable: Disables the multi-collect of attributes within a given Metadata Store cache. This option is disabled by default. Note: Do not enable this option if you are going to export the Application Metadata using the AMX application. There can be only one attribute in a JSON object, therefore Multi-collect is not supported when configuring the AMX application.
Data Link	If you want to include the VLAN ID along with the 5-tuple to identify the traffic flow, select the Data Link and enable the VLAN option.
Observation ID	Enter a value to identify the source from where the metadata is collected. The range is from 0 to 255. The calculated value of Observation Domain Id in Hexadecimal is 00 01 02 05 , and in Decimal is 66053 .
Enable DPI Packet limit	This field is used to restrict the number of packets in a particular session to be sent to the DPI engine instead of sending all the packets in order to improve the AMI performance. The value must range between 20 - 50 as the first 20 to 50 packets contains the most significant attributes.
Recurring attributes (SNI)	When enabled, it keeps the TLS SNI (SSL server name) for the full lifetime of each SSL/TLS flow. This means every active or inactive timeout export record will include the hostname, giving consistent visibility for long-lived flows. When disabled, SNI is exported only in the first record of a session. Later records for the same flow may not contain the hostname. This feature is supported only on GigaVUE HC Series Gen3 and GigaVUE V Series. It requires a bidirectional cache and exporter templates that already export SNI. The setting applies only to new flows after it is enabled.
Advanced Settings	
Number of Flows	The number of flows supported by the application. Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms.

Parameter	Description		
	Cloud Platform	Instance Size	Maximum Number of Flows
	VMware	Large (8 vCPU and 16GB RAM)	200k
	AWS	AMD - Large (c5n.2xlarge)	300k
		AMD - Medium (t3a.xlarge)	100k
		ARM - Large (c7gn.2xlarge)	100k
		ARM - Medium (m7g.xlarge)	200k
	Azure	Large (Standard_D8s_V4)	500k
		Medium (Standard_D4s_v4)	100k
	Nutanix	Large (8 vCPU and 16GB RAM)	200k
	- Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k.. - If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.		
Fast Mode	Enable the Fast Mode option for performance (less CPU cycles and less memory utilization) improvement. When the Fast Mode is enabled, some or all of the attributes of the applications will be disabled. If all the attributes of the application are disabled then the application itself is disabled.		
Aggregate Round-trip Time	Enable this option to export the minimum, maximum, and mean of RTT values for the following list of supported protocols and attributes and also the aggregate of TCP Lost byte values collected per export time interval.		
	Protocol	Attribute	
	http	rtt	
	icmp	rtt	
	icmp6	rtt	
	ssh	rtt	
	tcp	rtt	
	tcp	rtt_app	

Parameter	Description	
	telnet	rtt
	wsp	connect_rtt
	wsp	query_rtt

6. In the Application quick view, enter or select the following details in the **Exporters** tab:

Parameter	Description
Exporter Name	Enter a name for the Exporter.
Actions	Using this option, you can perform the following functions: • Add Exporter - Use to add a new Exporter to this Application Metadata Intelligence Application. A maximum of 5 exporters can be added. • Save as New Template - Use to save the current configuration as a new custom tool template. • Delete this Exporter - Use to delete the Exporter.
Template	Use to select the tool template. Refer to <i>Tool Templates</i> section in GigaVUE Fabric Management Guide for more information on tool templates and how to create custom tool templates.
Export Params	Select any one of the following options: • Application Name - Exports Application Name Attribute in AMI records. Enable this option to export Application Name attribute with network attributes • Application Family - Exports Application Family Attribute in AMI records. Enable this option to export Application Family attribute with network attributes • Application Tags - Exports Application Tag Attribute in AMI records. Enable this option to export Application Tag attribute with network attributes
Format	Select NetFlow or CEF
NetFlow: Select this option to use NetFlow	
Record / Template type	• Segregated - The application-specific attributes and the generic attributes will be exported as individual records to the tool. • Cohesive- The application-specific attributes and the generic attributes will be combined as a single record and exported to the tool. Note: It is recommended to select Cohesive from the drop-down menu, as NetFlow exports network and transport parameters only.
Enable Maximum Packet Length	Enable this option to edit the interface MTU value.
Maximum Packet Length (This option appears only when Enable Maximum Packet Length option is enabled)	Using this field, you can configure the maximum length of the packet that can be exported. Enter this value less than or equal to egress interface MTU value to avoid fragmentation. The value can range between 1280 and 9001.
Active Timeout	Enter the active flow timeout value in seconds.
Inactive Timeout	Enter the inactive flow timeout in seconds.
Version	Select the NetFlow version. The supported versions are V5, V9, IPFIX (V10).

Parameter	Description
Template Refresh Interval	Enter the time interval at which the template must be refreshed in seconds
CEF: Select this option to use CEF	
Record / Template type	- Segregated - The application-specific attributes and the generic attributes will be exported as individual records to the tool. - Cohesive- The application-specific attributes and the generic attributes will be combined as a single record and exported to the tool.
Active Timeout	Enter the active flow timeout value in seconds.
Inactive Timeout	Enter the inactive flow timeout in seconds.
APPLICATION & ATTRIBUTES: Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each of the application. (Not applicable when using NetFlow V5, V9, NetFlow IPFIX(V10), or CEF when the flow direction is Uni-Directional in the above Template drop-down menu.)	
NOTE: When you edit application attributes during an AMI deployment with background traffic, you may experience minor data loss until the update completes.	
Add Application	Click on the Add Application button. The Add Application dialog box opens. Select a Type. The available options are: Application Family: Each application is mapped only mapped to one Application Family - Select an Application Family and the Applications that needs to be filtered from the traffic. - Attributes for the selected application is displayed in the Attribute column. You can select the required attributes. Application Tag: Each application can be mapped to one or more Application Tags. - Select an Application Tag and the Applications that needs to be filtered from the traffic. - Attributes for the selected application is displayed in the Attribute column. You can select the required attributes. NOTE: The AMI exporter exports application metadata for the selected applications but exports the packet attributes for all applications. To export packet attributes only for the selected applications, filter the required traffic by using a flow map or AFI.
NETWORK & TRANSPORT PARAMETERS: Select the Network and the transport packet attributes with the respective parameters	
Data Link	Select any one of the parameters such as Source MAC address, Destination MAC Address and VLAN.
Interface	Select any one of the parameter such as Input Physical, Output

Parameter	Description
	Physical and Input Name.
IP	Select the parameter as Version if required.
IPv4	Select the required attributes. By default, Source Address, Destination Address, and Protocol are enabled.
IPv6	Select the required attributes. By default, Source Address, Destination Address, and Next Header are enabled.
Transport	Select the required attributes. By default, Source Port, Destination Port are enabled.
Counter	Select the Bytes, and Packets.
Timestamp	Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
Flow	Select the parameter as End Reason if required.
GTP-U	Select the required parameters such as QFI and TEID.
Outer IPv4	Select any one of the parameter such as Source or Destination.
Outer IPv6	Select any one of the parameter such as Source or Destination.

NOTE: In the exporter record, the generator ID represents the hostname of the GigaVUE V Series Node. Any changes to the hostname after deploying the Monitoring Domain with the AMI application will not be reflected in the exporter record.

- Click **Save**.

After configuring the Application Metadata application, follow the below steps to create an egress tunnel:

- In the canvas page, select **New > New Tunnel**, drag and drop a new tunnel template to the workspace. The **Add Tunnel Spec** quick view appears.
- Enter the **Alias**, **Description** and select **UDP** as **Type**.

NOTE: When using Application Metadata, if you create a tunnel to send the output to the tools, select UDP as the tunnel type. You can either use a single tunnel to export all metadata from the application or configure separate tunnels for each exporter.

- Enter the **MTU** value.
- L4 Destination IP Address** - Enter the IP address of the tool port. If using Application Metadata Exporter (AMX), enter the data subnet IP address of the V Series Node where AMX application is deployed. Refer to [Application Metadata Exporter](#) for details.

5. **Source L4 Port** - The port from which the connection will be established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
6. **Destination L4 Port** - The port to which the connection will be established from the source. For example, if A is the source and B is the destination, this port value belongs to B. The destination port must match the AMX ingestor port number to ensure seamless data forwarding from AMI to AMX.

NOTE: If the export format is CEF, the default value for L4 destination port is 514. If the export format is NetFlow, the default value for L4 destination port is 2055.

To deploy the configured Application Metadata Intelligence application, follow the below steps:

1. On the Monitoring Session canvas page, click **Deploy** from the **Actions** menu to deploy the Monitoring Session. The status is displayed as **Success** in the Monitoring Sessions page. The session is successfully deployed on the V Series Node.
2. You can view the Monitoring Session Deployment Report in the **SOURCES** and **V SERIES NODES** tab. 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.
 - Success—The session is not deployed on one or more instances due to V Series Node failure.
 - Failure—The session is not deployed on any of the V Series Nodes or Instances.

After adding the Application Metadata application and deploying Monitoring Session, you cannot change the **Aggregate Round Trip time** option.

NOTE: To configure Application Metadata Intelligence in physical environment, refer to [Create an Application Intelligence Session in Physical Environment](#).

Extend Your AMI and AMX Deployment

Your AMI and AMX deployment is the foundation of deep network observability. Once it is up and running, you can:

- **Integrate with Third-Party Tools:** AMI and AMX support seamless integration with popular monitoring, analytics, security, and visualization platforms such as Splunk, Elastic, CrowdStrike, Dynatrace, Datadog, and Grafana. For step-by-step integration guides, refer to [Integrate Application Metadata Intelligence and Application Metadata Exporter with Tools](#).
- **Analyze Metadata with Gigamon Insights:** Use [Gigamon Insights](#) to query your AMI and AMX metadata using plain-language prompts and get immediate, AI-powered answers — no dashboards or manual queries needed.

View Application Statistics for Application Metadata

To view the application Statistics for the Application Metadata application, follow the steps given below:

1. Click **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. Select a monitoring session from the list view, click **Actions > Edit**. The Edit Monitoring Session page appears.
3. Click the application and select **Details**. The Application quick view appears.
4. Click on **STATISTICS** tab.
5. You can view the following in the Application Metadata application statistics page:
 - a. Exporter Name - Displays the exporters created for this application.
 - b. Format - Displays the format as NetFlow or CEF, for the individual exporters.
 - c. Packet Sent/ Sec - Displays the count of packets sent per second for each exporter.

Application Metadata Exporter

Application Metadata Exporter (AMX) is an Application Intelligence component that runs on GigaVUE V Series Nodes and works with Application Metadata Intelligence (AMI) and other telemetry sources to deliver standardized metadata to your monitoring and analytics tools.

GigaVUE-FM manages both AMX and AMI and provides a central point to deploy, configure, and monitor AMX across on-premises and cloud environments.

You can use AMX to:

- Transform and export application metadata in JSON and Parquet formats to cloud-native tools/data lakes.
- Transform and export NetFlow/IPFIX, and mobility metadata in JSON format to cloud-native tools and Kafka-based data pipeline.
- Enrich mobile-network metadata for 3G/4G/5G visibility and analytics.
- Enrich AMI records with cloud and Kubernetes workload metadata with host, environment, and container context to reduce Mean Time To Identification (MTTI) and Mean Time To Resolution (MTTR).
- Enrich AMI records with user identity information using Zscaler Private Access (ZPA), adding customer, user, and service details to help pinpoint performance issues and monitor the effectiveness and compliance of network access policies.

Refer to the following topics for more detailed information on the various deployment options and scenarios:

- [Application Metadata Exporter Deployment Options](#)
- [Application Metadata Exporter - Export Scenarios](#)

Application Metadata Exporter Deployment Options

The output from the Application Metadata Intelligence or GTP Correlation Engine is sent to the AMX application, which exports it to the tools in JSON format over HTTPS/Kafka. AMX application is always deployed in GigaVUE V Series Node. The GigaSMART application sending data to the AMX application can be on GigaVUE HC Series or the GigaVUE V Series Node. Based on where the GigaSMART applications reside, there can be four deployment methods:

NOTE: For exporting GigaVUE Enriched Metadata for Cloud Workloads, AMI should be deployed in GigaVUE V Series Nodes.


- [On-Premises](#)
 - Hardware (AMI)
 - Hardware (Control Plane Metadata)
 - Virtual (VMware, Nutanix, KVM)
- [Public Cloud](#)

On-Premises

Hardware (AMI)


In hardware deployments, the Application Metadata Intelligence (AMI) runs on a physical node/cluster, and the AMX application is deployed on a GigaVUE V Series Node running on VMware ESXi, OpenStack, or Nutanix. The output from the AMI in CEF format is sent to the AMX application in GigaVUE V Series Node. GigaVUE-FM handles the orchestration of the configuration between the Gigamon devices. The following devices support the integration of AMX application:

- GigaVUE-HC1
- GigaVUE-HC3
- GigaVUE-HC1-Plus
- GigaVUE-HCT




Hardware (Control Plane Metadata)

In hardware deployments, the GTP Correlation Engine runs on a physical node/cluster, and the AMX application is deployed on a GigaVUE V Series Node running on VMware ESXi/KVM and Nutanix. The output from the GTP Correlation Engine in Flat JSON format is sent to the AMX application in GigaVUE V Series Node. The performance of the device and the application is managed by GigaVUE-FM. The GigaVUE-HC3 Gen3 devices support the integration of AMX application.




Private Cloud (VMware)

In the Private Cloud environment, the application is supported only on VMware ESXi/KVM and Nutanix and can be deployed in the VMware as shown in the diagram.



Public Cloud

In the Public Cloud environment, the application is supported on AWS and Azure platforms, and can be deployed as shown in the diagram:




Application Metadata Exporter - Export Scenarios

Refer to the following topics for more detailed information on the various ways to configure AMX:

- [Export Application Metadata in JSON format over HTTPS/Kafka](#)
- [Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka](#)
- [Export of 3G/4G/5G Control Plane Metadata by AMX](#)
- [Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka](#)
- [Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka](#)
- [Application Metadata Exporter](#)

Export Application Metadata in JSON format over HTTPS/Kafka


Application Metadata Exporter (AMX) application converts the output from the Application Metadata Intelligence (AMI) in CEF format into JSON format and sends it to the cloud tools and Kafka Consumers. The following image illustrates this end-to-end flow.



Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka

The AMX application with NetFlow Integrator functionality supports ingesting NetFlow/IPFIX flow records from various sources and devices, such as firewalls, routers, and switches, including AMI. It normalizes these records into JSON and exports them to external tools, providing a consolidated view of network traffic and NetFlow/IPFIX insights in a standard format compatible with most Network Performance Monitoring (NPM) tools.

The following image illustrates how the NetFlow Integrator collects NetFlow/IPFIX telemetry from network devices, converts it into JSON, and exports it over HTTPS or Kafka to external cloud tools. GigaVUE-FM acts as the central control system for AMX.



Supported Platforms


- VMware ESXi
- VMware NSX-T

Important Notes

- **NetFlow v9 and v5 Timestamps:**

- The NetFlow header contains system uptime (in seconds) and an export timestamp.
- Flow start time is derived using the following formula:
 - $\text{Flow_Start} = \text{Export Timestamp} - \text{System Uptime} + \text{First_Switched}$
- Flow end time is derived using the following formula:
 - $\text{Flow_End} = \text{Export Timestamp} - \text{System Uptime} + \text{Last_Switched}$
- The calculated start and end time are converted into nanoseconds.
- Unsupported attributes such as flow start/end in milliseconds, microseconds, or nanoseconds will be ignored. Instead, the export timestamp is used as both the flow start and end times, which are then converted to nanoseconds.
- **NetFlow v10/IPFIX Timestamps:**
 - NetFlow v10/IPFIX directly incorporates timestamps without the need for system uptime.
 - Even if Flow start and end timestamps are available in seconds, milliseconds, microseconds, or nanoseconds, they are always converted to nanoseconds for output.
- **Missing Flow Start/End Attributes** - For records that do not include flow start or end attributes, AMX automatically assigns both values to the export timestamp.


The image below explains how flow start and end times are calculated using system uptime, export timestamp and flow switch times.



Export of 3G/4G/5G Control Plane Metadata by AMX

The AMX application can also export the 3G/4G control plane metadata received from the GTP Correlation engine and 5G control plane metadata received from the 5G CPN engine to the cloud tools and Kafka in Flat JSON format.

The AMX application can be deployed only on a GigaVUE V Series Node and can be connected to a GTP Correlation / 5G CPN engine running on a physical node.



Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka

The metadata enrichment enhances service provider analytics, by generating metadata on 5G/4G/3G network traffic. The AMX correlates the user plane metadata produced by AMI with the control plane metadata produced by the GTP/5G correlation mobility application to produce an enriched metadata feed for the mobile networks. This data feed helps with use cases like service personalization, planning, and many others by containing information about the:

- Subscriber Session
- Over the Top Application
- Handset Type
- Location
- Flow throughput calculation attributes - DL, UL bytes, and time stamps.
- Application Protocol
- Core Network Information
- User Tunnel Information


Export of GigaVUE Enriched Metadata for Mobile Networks is supported only for GigaVUE V Series Node deployed using Third Party Orchestration on VMware ESXi.

User Plane and Control Plane traffic from the following devices are supported for exporting GigaVUE Enriched Metadata for Mobile Networks:

- GigaVUE-HC3
- GigaVUE-HC1-Plus

NOTE: For GigaVUE-HC1-Plus, the AMI application must be configured on the built-in engine to efficiently handle higher traffic loads. The plug-in engine can be used for the Control Plane traffic.

For information on Control Plane Metadata, refer to [Control Plane Metadata](#).



Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka

GigaVUE Enriched Metadata for Cloud Workloads provides comprehensive situational awareness to address security and performance pain points in a timely manner. It enriches application metadata from N/S and lateral traffic with key host environment details that allow you to find critical information as follows:

- The location of the workloads hosted and their virtual network.
- The operational environment to which the workloads belong.
- The instance types used, images, and tags that the workload contains.
- The host name, the security associations like security group name, IAM instance profile name.

Export of GigaVUE Enriched Metadata for Cloud Workloads (Virtual Machines)

Export of GigaVUE® Enriched Metadata (GEM) for Cloud Workloads is supported on the following cloud platforms:

- AWS
- Azure
- VMware (ESXi and NSX-T)

This functionality works by using the inventory API which is queried in the following intervals.


- VMware: 300 sec
- AWS: 30 sec
- Azure: 60 sec

The default inventory query interval should suffice in most cases, however the interval can be customized in extreme situations. [Contact Gigamon Support](#) for assistance.

In addition to these fixed inventory polling intervals, you can subscribe to optional services to receive automatic workload updates in AWS (Event Bridge)/Azure (Event Grid). For VMware, no additional configuration is required because dynamic updates are enabled by default.

AMX application performs the enrichment every 10 seconds. It picks the flow records, which are 15 seconds or older, to allow any delays in fetching the inventory details, and uses the IP address of the endpoints to enrich the records based on the selected attributes. Refer to the following figure for a high-level illustration of the solution. The solution can be deployed using GigaVUE-FM or Third Party Orchestration.


The following image illustrates how AMX exports GigaVUE Enriched Metadata for cloud workloads by converting flow records into JSON/Parquet format and streaming them over HTTPS or Kafka to cloud tools.



The enrichment supported depends on the type of platform. Refer to [Attributes for GigaVUE Enriched Metadata for Cloud Workloads](#) for more details.

Export of GigaVUE Enriched Metadata for Cloud Workloads Kubernetes containers - Azure Kubernetes Service (AKS)

This feature enables the enrichment of flow records with Kubernetes inventory data, unlocking deeper visibility and advanced use cases for containerized environments. In the current Kubernetes visibility solution for AKS, UCT-TAP ships tapped traffic (both east-west and north-south) to the AMI node deployed in Azure. The AMI node then forwards metadata to the AMX node, also deployed in Azure. With the new enhancement, the AMX node enriches the AMI metadata using Kubernetes inventory details fetched directly from the AKS cluster.



GigaVUE-FM queries the Kubernetes inventory every 300 seconds, which is the default interval for monitoring events between inventory fetches. The default inventory query interval should suffice in most cases; however the interval can be customized in extreme situations. Please contact Gigamon Support for assistance.

NOTE: In addition to sending regular queries every 300 seconds, GigaVUE-FM enables a watch mechanism that updates AMX with every event—create, delete, and modify.

For Kubernetes environments, GigaVUE-FM uses a controller with access to the Kubernetes API. The controller subscribes to cluster events and pushes updates in near real-time. No built-in optional services are required.

Refer to [Prerequisites for Export of GigaVUE Enriched Metadata for Cloud Workloads](#) for more detailed information.

Rules and Notes

- With any AMI entitlement (AMI SKU, ZTA bundle, or SecureVUE PLUS bundle), you are also entitled to deploy AMX.
- The GigaVUE V Series Node deployed must be entirely dedicated to the AMX application, it cannot have other applications in it.
- The monitoring session can only have Raw End Point (REP), it cannot have other applications, maps, or tunnels when using the AMX application. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more detailed information on how to add a REP to the monitoring session and how to configure it.
- When using this application for production usage, it is recommended to use large size Virtual Machines.

- If you reload the GigaVUE V Series Node after configuring the AMX application, then the Ingestor in the AMX application fails.
- When using GigaVUE Enriched Metadata for Cloud Workloads, if duplicate IP address is assigned to workloads, the metadata received from the most recently queried workload will be enriched. Hence, it is not recommended to use the feature in environments that are prone to have duplicate IP addresses.
- When using GigaVUE Enriched Metadata for Cloud Workloads, endpoints and the GigaVUE V Series Node running AMI and AMX should be co-located in the same cloud platform.

Configure Metadata Export Filter for AMX

NOTE: Metadata export filter is now available as an Early Access feature, giving you the opportunity to explore its capabilities before the general availability (GA).

An export filter can be configured to control which metadata records are exported from the Application Metadata Exporter (AMX) to downstream tools such as Kafka. By defining a filter condition, only records that satisfy the specified criteria are exported, reducing unnecessary data volume sent to the target tool. This is an optional functionality that can be used as per the requirement.

The Application Metadata Exporter generally stores all received metadata records in an internal ClickHouse database. The export filter applies a ClickHouse query condition at the exporter component. When the filter is active, only records matching the condition are exported; all other records are dropped before reaching the tool.

Rules and Notes

- The export filter can only be configured by directly editing the `/etc/amx_settings.yaml` file on the V Series Node. This option is not available through GigaVUE-FM.
- The backend does not validate the filter string. Ensure that the value of `set_export_filters` is a valid ClickHouse sub-query. An invalid query may result in no records being exported.
- If the filter is misconfigured or contains an invalid ClickHouse sub-query, AMX does not return an explicit error. Instead, no metadata records are exported. To diagnose such issues, review the AMX logs on the V Series Node for filter-related error messages. Log files are located at `/var/log/syslog/` on the V Series Node.

To configure the export filter:

1. Create or edit the `/etc/amx_settings.yaml` file on the V Series Node where AMX is deployed.
2. Define the `set_export_filters` key with a valid ClickHouse sub-query as its value. For example, `set_export_filters: "mdata['app_name'] LIKE 'http%'"`.

3. To apply the new filter settings, undeploy the monitoring session and then deploy it again.

Deploy Application Metadata Exporter in VMware NSX-T using Third Party Orchestration

This section outlines the workflow for deploying Application Metadata Exporter (AMX) application using Third Party Orchestration on VMware ESXi.

Predeployment Checklist

1. The minimum resource requirements for deploying V Series to configure AMX are as follows. You can deploy one or more V Series for deploying AMX subject to your traffic volume (flows per second). [Contact Gigamon Support](#) for sizing guidance, if needed.
 - **8 vCPUs, 16 GB memory, and 80 GB disk space** for exporting Application Metadata in JSON format.
 - **40 vCPUs, 64 GB RAM, 128 GB disk space** for exporting NetFlow/IPFIX records in JSON format.

Exporting JSON records involves file read/write operations; using SSD storage is recommended.

2. Ensure you have a valid SecureVUE Plus VBL license to deploy AMX and to export NetFlow/IPFIX records in JSON format. With any AMI entitlement (AMI SKU, ZTA bundle, or SecureVUE PLUS bundle), you are also entitled to deploy AMX.
3. Ensure you have access to the [VUE Gigamon Community portal](#) for downloading V Series OVA image.
4. V Series requires three interfaces for deploying AMX—management, data, and tool. It is recommended to allocate dedicated interface for each. Create three separate subnets: a Management subnet for connectivity to GigaVUE-FM, a Data subnet for receiving traffic such as application metadata in CEF or NetFlow/IPFIX formats, and a Tool Port Group for exporting JSON records.

Note: Preparing a network topology with the IP layout helps in configuration, monitoring, and troubleshooting.

5. Select the appropriate deployment mode for the GigaVUE V Series node based on your requirements:
 - Single physical uplink for Management, Data and Tool connectivity.
 - Single physical uplink for Management and Tool connectivity, and another uplink for Data connectivity.
 - (Recommended) Separate uplinks (three) for Management, Data, and Tool connectivity.
6. If static IP addressing is used, obtain the following details from your server team:

- IP addresses for the management, data, and tool interfaces.
 - IP address of the DNS server and the gateway connected to the tool interface.
 - Subnet Mask and MTU.
7. Verify whether remote web console access to the V Series is available. Console access is required for troubleshooting. If remote access is not available, request the server team to enable SSH access.
 8. Review the deployment guide for the tool specific configuration needed to export records from AMX. Refer to [AMX Tool Integration](#).
 9. Review the connectivity requirements for GigaVUE-FM, AMI, AMX, and the associated tools in your environment. Ensure that the required ports are open and any external connectivity dependencies are addressed.
 10. Ensure that the standalone GigaVUE-FM deployment meets the recommended requirements shown in the following table. For medium and large scale deployments, refer to [Recommended Requirements for GigaVUE-FM Deployments](#). This step is optional. A standalone GigaVUE-FM instance is required only when an existing GigaVUE-FM is not available for managing AMX.

Requirements	Small Scale Deployment
CPU	
CPU count(minimum)	2
CPU Min speed (per CPU)	2.30GHz
Memory	
Memory Size	16GB
Disk Space	
OS Disk Size (min)	40GB
Config Disk Size - 15 Days Data	80GB
Config Disk Size - 35 Days Data	160GB
Config Disk Size - Rollup 35 Days + 120 Days Rollup	200GB
NIC	1
Number of NICs	1GB
Management/Data	
Appliance Scalability	Up to 50
Number of GigaVUE HC Series/GigaVUE TA Series Devices	Up to 50
Number of GigaVUE V Series Nodes	1 Cluster with 100 Nodes and 10000 pods with 2 rules
Number of Worker Pods monitored	

Refer to the following topics to deploy AMX using Third Party Orchestration:

- [Deploy Application Metadata Exporter in VMware NSX-T using Third Party Orchestration](#)
- [\(Optional\) Configure Metadata Export Filter for AMX](#)
- [\(Optional\) Configure Role-Based Access for Third Party Orchestration](#)
- [Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source](#)
- [Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM](#)
- [Deploy Application Metadata Exporter in VMware NSX-T using Third Party Orchestration](#)
- [Create a Monitoring Session to Configure Application Metadata Exporter](#)
- [Configure Ingestor in AMX to Process Traffic](#)
- [Integrate AMX with Third Party Tool](#)
- [Configure Exporter to Export Attributes to Third Party tools](#)
- [Verify Your Deployment](#)

(Optional) Install GigaVUE-FM on VMware NSX-T

This step is optional and applies only when an existing GigaVUE-FM instance is not available. Refer to [Install GigaVUE-FM on VMware NSX-T](#) for steps to install GigaVUE-FM on VMware NSX-T and the steps to start GigaVUE-FM instance and configure it.

(Optional) Configure Metadata Export Filter for AMX

Optionally, you can configure an export filter to control which metadata records are exported from AMX. For more information, see [Configure Metadata Export Filter](#).

(Optional) Configure Role-Based Access for Third Party Orchestration

If you do not want to use an **Admin Profile** to deploy GigaVUE V Series through a Third Party orchestrator, you can create the required users, roles, and user groups in GigaVUE-FM.

Create a Custom User

You can add users only if you are a user with **fm_super_admin role** or a user with either read or write access to the GigaVUE-FM security Management category.

To add users:

1. Go to **Settings** and select **Authentication > GigaVUE-FM User Management > Users**.

2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
 - **Name:** Actual name of the user
 - **Username:** User name configured in GigaVUE-FM
 - **Email:** Email ID of the user
 - **Password/Confirm Password:** Password for the user.
 - **User Group:** Select the desired User Group to associate the user.
GigaVUE-FM prompts for your password.
4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

A user role defines permission for users to perform any task or operation in GigaVUE-FM or on the managed device. You can associate a role with user.

NOTE: A user with read-only access cannot perform configurations on the screen. The menus and action buttons in the UI pages are disabled appropriately.

To create a role:

1. Go to **Settings** and select **Authentication> GigaVUE-FM User Management >Roles**.
2. Select **New Role**.
3. In the New Role page, select or enter the following details:
 - **Role Name:** Name of the role.
 - **Description:** Description of the role.
 - **Select Permission:** From the **Select Permissions** tab, select **Third Party Orchestration**, and provide write permissions.
4. Select **Apply** to save the configuration.

Create a Custom User Group and Assign the Custom Role and the User

A user group consists of a set of roles and set of tags associated with users in that group. You can associate the user to one or more groups.

To create a new user group:

1. Go to **Settings**, and then select **Authentication> GigaVUE-FM User Management >User Groups**.
2. Select **New Group**.
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:

5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source

To create a Monitoring Domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > Third Party Orchestration**, and select **Monitoring Domain**. The **Monitoring Domain** page appears.
2. On the **Monitoring Domain** page, select **New**. The **Monitoring Domain Configuration** page appears.
3. Enter a name for the **Monitoring Domain**.
4. Enter a name for the **Connection**.
5. Select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
6. Click **Save**.

Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM

Use tokens when applications or users access GigaVUE-FM through REST APIs or when registering GigaVUE V Series using Third Party Orchestration. Tokens let GigaVUE-FM verify who is making the API calls and what they are allowed to do, without exposing passwords.

You can create tokens only if you are a user with **fm_super_admin role** or a user with write access to the GigaVUE-FM security Management category.

To create a token:

1. Go to **Settings > Authentication > GigaVUE-FM User Management**. The **User Management** page appears.
2. Select **Current User Tokens** from the **Tokens** drop-down list.
3. Click **New Token**. The **New Token** pop-up appears.
4. Enter a name for the new token in the **Name** field.
5. Enter the days until the token is valid in the **Expiry** field.
6. Select the user group for which you are privileged to access the GigaVUE-FM from the **User Group** drop-down list. The token inherits the read or write RBAC privileges of this group.
7. Select **OK** to generate a new token. You can copy and use the generated token to authenticate the GigaVUE-FM REST APIs.

Deploy GigaVUE V Series Nodes using VMware NSX-T Manager

You can deploy your GigaVUE V Series Nodes using VMware NSX-T Manager. The GigaVUE V Series Nodes register themselves with GigaVUE-FM using the information provided by the user in the NSX-T Manager. When the nodes are registered with GigaVUE-FM, you can configure Monitoring Session and related services in GigaVUE-FM.

Getting Started

To register your GigaVUE V Series Nodes using VMware NSX-T Manager,

1. In GigaVUE-FM, create a Monitoring Domain.
For details, refer to [Step 3: Connect to VMware vCenter in GigaVUE-FM](#).
2. In the **VMware Configuration** page, select **No** for the **Use FM to Launch Fabric** field.

The screenshot shows the 'VMware Configuration' page with the following fields and controls:

- Monitoring Domain***: Text input field with placeholder 'Enter a monitoring domain name'.
- Connection Alias***: Text input field with placeholder 'Alias'.
- Virtual Center***: Text input field with placeholder 'Virtual Center'.
- Username***: Text input field with placeholder 'Username'.
- Password***: Text input field with placeholder 'Password'.
- NSX-T Manager***: Text input field with placeholder 'IP address or hostname'.
- NSX-T Username***: Text input field with placeholder 'NSX-T Manager username'.
- NSX-T Password***: Text input field with placeholder 'NSX-T Manager password'.
- FM Username***: Text input field with placeholder 'FM username'.
- FM Password***: Text input field with placeholder 'FM password'.
- Use External Image**: A toggle switch (currently off) and a dropdown menu with placeholder 'Select an image'.
- Use FM to Launch Fabric®**: A toggle switch (currently on).

NOTE: When creating the Monitoring Domain for deploying GigaVUE V Series Nodes, you can use the VMware NSX-T username and password with at least the NETX Partner Admin role assigned to it.

After creating a Monitoring Domain, use a VMware NSX-T user account to deploy the GigaVUE V Series Nodes with at least the NETX Partner Admin role assigned. When creating multiple domains with the same NSX-T Manager, GigaVUE-FM assigns a unique service name to every Monitoring Domain.

You can view these names on the Monitoring Domain page.

Deploy GigaVUE V Series Nodes in VMware NSX-T Manager

To deploy the node,

1. In **VMware NSX-T Manager**, go to the **Service Deployment** page and select **Deployment**.
2. From the **Partner Service** drop-down, select the service name corresponding to your Monitoring Domain.
For detailed information, refer to [Deploy a Partner Service](#) topic in the VMware documentation.
3. Select the **Deployment template** and **Deployment Specification**, and click **Configure Attributes**.
4. In the **Configure Attributes** page, enter the
 - Service VM Host Name
 - Admin user password
 - (Optional) **SSL Private Key** and the **SSL Certificate** to use a custom certificate.

For details, refer to the Secure Communication topic.

After deployment,

- NSX-T Manager displays the runtime status of node as UP.
- GigaVUE-FM updates the node status to Launching. After successful registration, the status changes to **OK**.
- The deployed node appears on the Monitoring Domain page in GigaVUE-FM.

NOTE: VMware NSX-T automatically generates the V Series Node name. Do not rename the node in vCenter.

		New Edit Delete Fabric Refresh Inventory						
		Monitoring Domain	Connection	Name	Management IP	Type	Version	Status
		nsxt-202-13-md						
			nsxt-202-45-md					Connected
			Gigamon Inc__vp-3rd-...	10.115.206.127	V Series Node	3.0.0		Ok

Compatibility and Limitations

- All GigaVUE V Series Nodes in a Monitoring Domain must run the same version. GigaVUE-FM does not support mixed versions.



- IPv6 is not supported for the gateway of the tunnel interface when deploying nodes through NSX-T Manager.
- Avoid renaming or manually managing the V Series Node from vCenter to prevent deployment issues.

Delete GigaVUE V Series Nodes and Monitoring Domain

NOTE: When you deploy V Series nodes using VMware NSX-T manager, GigaVUE-FM disables the **Delete** button for those nodes. You must first remove the Service Deployment in NSX-T Manager.

To delete a GigaVUE V Series Node deployed using VMware NSX-T Manager,

1. In NSX-T Manager, delete the associated **Policy and Service Chain**.
2. In GigaVUE-FM, delete the related Monitoring Session.
3. In NSX-T Manager, delete the V Series Node.
GigaVUE-FM automatically unregisters it from the Monitoring Domain.
4. In GigaVUE-FM, delete the Monitoring Domain.

Create a Monitoring Session to Configure Application Metadata Exporter

In GigaVUE-FM, you must do the following to configure the Application Metadata Exporter application in the Monitoring Session.

1. **Access the Monitoring Session Page:**
 - a. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**.
 - b. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
2. **Create an ingress REP:**
 - a. In the canvas, from the New expand menu, drag and drop New Raw Endpoint to the graphical workspace.
 - b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
 - c. Enter the Alias and Description details for the Ingress Raw End Point and click **Save**.
3. **Add AMX Application:**

- a. In the canvas, from the **Applications** expand menu, drag and drop AMX to the graphical workspace.
 - b. Configure the ingestor settings to process traffic. Refer to [Configure Ingestor in AMX to Process Traffic](#) for details.
 - c. Configure the required exporter settings to export attributes to Third Party tools. Refer to [Configure Exporter to Export Attributes to Third Party tools](#) for details.
 - d. Create a link between the ingress REP and Application Metadata Exporter application.
4. **Create an egress REP:**
- a. From the **New** expand menu, drag and drop New Raw Endpoint to the graphical workspace.
 - b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
 - c. Enter the Alias and Description details for the Egress Raw End Point and click **Save**.
 - d. Create a link between the Application Metadata Exporter application and egress REP



Notes:

- Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
- When deploying AMX using the Third-Party Orchestration method, Monitoring Session must also be configured in the same work flow.

Configure Ingestor in AMX to Process Traffic

The following scenarios provide more information about the different ways to configure Ingestor AMX:

- [Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka](#)
- [Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka](#)
- [Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka](#)
- [Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka](#)

Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

To configure ingestor for AMI:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.

Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

To export GigaVUE Enriched Metadata for Mobile Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which application metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.
9. Select the Enable option to enable the enrichment.
10. Select **Mobility** from the **Type** drop-down menu.
11. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Mobile Networks](#) for more details.

Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

To export GigaVUE Enriched Metadata for Cloud Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which Application Metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.
9. Select the Enable option to enable the enrichment.
10. Select **Workload** from the **Type** drop down menu.
11. Select the platform as **VMware VCenter** from the drop down list.
12. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Cloud Workloads](#) for more details.

Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

To configure ingestor for NetFlow/IPFIX:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select the input type as **NetFlow/IPFIX**. Only the elements specified in the mapping file will be included in the JSON output. Editing or customizing the mapping file is not supported. Refer to [Supported Element Types for NetFlow Integrator Element Mapping](#).
3. Enter the port number on which NetFlow/IPFIX must be ingested. The default port value is 2055.

4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor. A maximum of 12 NetFlow/IPFIX type ingestors can be configured per deployment.

NOTE: If you select NetFlow/IPFIX as the ingestor type, Metadata Enrichment is disabled by default. Metadata Enrichment is not supported for NetFlow/IPFIX inputs.

Integrate AMX with Third Party Tool

This section explains how to set up the third-party tool and configure the exporter settings required to route application metadata attributes for analysis and visualization.

Set up the Third Party Tool for AMX Integration

You must set up the third party tools to ingest and visualize application metadata.

Refer to [AMX Tool Integration](#) for configuration guidance and the verification steps required to confirm that data export and visualization are functioning correctly.

Analyze Metadata with Gigamon Insights: Use [Gigamon Insights](#) to query your AMI and AMX metadata using plain-language prompts and get immediate, AI-powered answers — no dashboards or manual queries needed.

Configure Exporter to Export Attributes to Third Party tools

AMX supports the following types of exporters:

- [Configure Cloud Tool Exporter](#)
- [Configure Kafka Exporter](#)

Configure Cloud Tool Exporter

To configure Cloud Tool Export settings in AMX:

1. On the AMX quick view in the **Cloud Tool Exports** section, enter or select the required details as described in [Cloud Tool Exporter - Field References](#).
2. Click **Save**.

Deploy the Monitoring Session

1. Click **Deploy** from the **Actions** menu.
2. The **Select nodes to deploy the Monitoring Session** dialog box appears. Select the GigaVUE V Series Node for which you wish to map the interfaces and deploy the Monitoring Session.

3. Select the interfaces for the ingress and egress REPs configured in the Monitoring Session. The interface naming includes the type and IP address of the interface, for example, DATA (ens192) IPv4: 10.114.42.135 IPv6: fcab:0:1:2040:a0a1:874e:7a66:40a2.
4. Select **Deploy**. Once deployed, the status is displayed as **Success** in the Monitoring Sessions **Sources** tab. The session is successfully deployed on the selected V Series Node.

Cloud Tool Exporter - Field References

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable Secure Endpoint option to hide the API keys from the URL. Example: - Normal Endpoint: http://examples.com/mykey/key123 - Masked Endpoint: http://examples.com/mykey/****
Headers	Enter the custom header value of the cloud tool instance. Example: Add Content-Type: application/json in a separate header field. Ensure that the Secure Keys option is enabled.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control- Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to transform NetFlow/IPFIX records to JSON format.
Labels	Click Add. Enter the following details: - Enter the Key. - Enter the Value. Note: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Zip	Enable the box to compress the output file. Note: Enable this field when using New Relic as the cloud tool.
Interval (sec)	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.

Fields	Description
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Cloud Tool. Both IPv4 and IPv6 are supported. This field is optional.

Configure Kafka Exporter

To configure Kafka Export settings in AMX:

1. On the AMX quick view in the **Kafka Exports** section, enter or select the required details as described in [Kafka Exporter - Field References](#).
2. Click **Save**.

Deploy the Monitoring Session

1. Click **Deploy** from the **Actions** menu.
2. The **Select nodes to deploy the Monitoring Session** dialog box appears. Select the GigaVUE V Series Node for which you wish to map the interfaces and deploy the Monitoring Session.
3. Select the interfaces for the ingress and egress REPs configured in the Monitoring Session. The interface naming includes the type and IP address of the interface for example, DATA (ens192) IPv4: 10.114.42.135 IPv6: fcab:0:1:2040:a0a1:874e:7a66:40a2.
4. Select **Deploy**. Once deployed, the status is displayed as **Success** in the Monitoring Sessions **Sources** tab. The session is successfully deployed on the selected V Series Node.

Kafka Exporter - Field References

Fields	Description
Alias	Enter the alias name for the Kafka Export.
Topic	The topic name to push JSON streams to, which is generally given to users part of the Kafka administration.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to export NetFlow/IPFIX ingestor types.

Fields	Description
Brokers	The URL that contains the Kafka cluster endpoints. Click  to add another broker and click  to remove an existing broker.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds. The default time interval is 30 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Kafka Broker. Both IPv4 and IPv6 are supported. This field is optional.
Producer Configurations	Click Add to enter the authentication details if a Kafka broker needs authentication. For Example: • security.protocol=SASL_SSL • sasl.mechanism=PLAIN or sasl.mechanisms=PLAIN • sasl.username=username • sasl.password=password
Message Configurations	Use this field to specify custom Kafka message settings, such as partition key selection, that GigaVUE-FM passes to the Kafka exporter within the Application Metadata Exporter application. Valid values are: • Key - partition.key and Value - flow_id, five_tuple, three_tuple, round_robin (default) • Key - partition.keyRotateInterval and Value - <seconds> (default: 3600; range: 60–86400). This key is rotated every <value> seconds, and this rotated key is used as a fallback when the selected key option (flow_id, three_tuple, or five_tuple) cannot be determined for a record. NOTE: If the five_tuple option is selected, but port attributes are not found in the record, the partition selection will be determined using the three_tuple.

Extend Your AMI and AMX Deployment

Your AMI and AMX deployment is the foundation of deep network observability. Once it is up and running, you can:

- **Integrate with Third-Party Tools:** AMI and AMX support seamless integration with popular monitoring, analytics, security, and visualization platforms such as Splunk, Elastic, CrowdStrike, Dynatrace, Datadog, and Grafana. For step-by-step integration guides, refer to [Integrate Application Metadata Intelligence and Application Metadata Exporter with Tools](#).
- **Analyze Metadata with Gigamon Insights:** Use [Gigamon Insights](#) to query your AMI and AMX metadata using plain-language prompts and get immediate, AI-powered answers — no dashboards or manual queries needed.

Verify Your Deployment

You can view the Ingestor, Exporter and Enrichment statistics for the deployed Monitoring Session in the **TRAFFIC PROCESSING** tab. For more details, refer to [View Application Metadata Exporter Statistics](#).

To troubleshoot any issues in Application Metadata Exporter deployment, refer to [Troubleshoot Application Metadata Exporter Deployment](#).

Deploy Application Metadata Exporter (AMX) on KVM Hypervisor

For using the KVM hypervisor on GigaVUE-FM, the following sections provide information on deploying a fresh installation of GigaVUE-FM in the KVM environment. This setup enables the deployment of the Application Metadata Exporter (AMX) on a GigaVUE-FM Hardware Appliance.

Install GigaVUE-FM on KVM hypervisor

Prerequisite:

1. Download the recent GigaVUE-FM image from [Gigamon Customer Portal](#) and move qcow2 file to a desired location.
2. Install and enable Cockpit services on the KVM hypervisor. To enable Cockpit View on the KVM hypervisor, follow these steps:
 - a. Update the system package list.

```
sudo apt update
```

- b. Install Cockpit and Cockpit-Machines.

```
sudo apt install -y cockpit cockpit-machines
```

- c. Enable and start the Cockpit service.

```
sudo systemctl enable --now cockpit.socket
```

Once these commands are executed, access KVM using port 9090 in your browser: `http://<your-server-IP>:9090`

To create a VM in KVM hypervisor for deploying the GigaVUE-FM, follow the instructions.

Parameters	Instructions	Mandatory field
Connection	Select System .	Yes
Installation Type	Select Cloud base image from the drop-down list.	Yes
Installation Source	Select the location of the qcow2 image.	Yes
Operating System	Select the latest Rocky Linux version.	Yes
Storage	Select Create new qcow2 volume .	Yes
Storage Limit	Enter the required storage limit (ex: 40 GiB).	Yes
Memory	Enter the required memory allocation (ex: 16 GiB).	Yes

Click **Create and edit** to edit the Virtual Machine. Modify the **Firmware** to **UEFI** and click **Install**.

Edit the VM as mentioned in the following table.

Parameters	Instructions	Mandatory field
Firmware	Select UEFI from the drop-down list	Yes
Disks		
Add new disk		
Size	Enter the disc size as required (ex: 40 GiB).	Yes
Format	Select the format (ex: qcow2).	Yes
CPU		
vCPU Maximum	Enter the maximum number of vCPU's required (ex: 10).	Yes
vCPU Count	Enter the number of CPUs to allocate (ex: 8).	Yes

NOTE: The values mentioned for CPU, size, format, storage, and memory are sample entries and can be customized as required. Refer to [Before You Install](#) to know more about the hardware and virtual computing requirements for GigaVUE-FM.

Click **Run** to deploy **GigaVUE-FM**. Copy the IP address of the newly launched GigaVUE-FM and paste it on the browser to login to the GigaVUE-FM. You can login to GigaVUE-FM using the default user name **admin** and the default password admin123A!!.

View Application Metadata Exporter Statistics

To view the application Statistics for the Application Metadata Exporter application, follow the steps given below:

1. Go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**. The **Monitoring Sessions** Overview page appears.
2. Select the monitoring session you want to view the statistics.
3. Go to the **TRAFFIC PROCESSING** tab. In the canvas page, click the  menu button on the **AMX** application and select **Details**. The Application quick view appears.
4. Click on **STATISTICS** tab.
5. To view the statistics of a particular GigaVUE V Series Nodes, select the required GigaVUE V Series Node from the **V Series Node IP** dropdown menu to view the **Ingestor**, **Exporter**, and **Attributes Enrichment** statistics.

You can download the Ingestor and Exporter statistics using the following options:

- **Download TCP Dump** (Ingestor)- Select this option to download the PCAP file (for example, cap_ens192_20260407_073719655-run-001.pcap.gz). The TCP dump capture is limited to 100 packets and runs for up to 30 seconds, or until 100 packets are captured (whichever occurs first). You can capture a TCP dump for only one AMX ingestor at a time. The TCP dump download option is available only for the following ingestor types: AMI, NetFlow/IPFIX, Mobility, and Zscaler LSS.

- **Download Records (Exporter)** - Select this option to download the JSON file (for example, Elastic-exporter-records.json.gz) of the exporter records. You can download records for a single AMX exporter or for multiple exporters. The exporter record download waits up to 30 seconds for records to be exported. If no records are produced, the operation times out.

For instructions on how to download the ingestor and exporter records, and troubleshoot AMI/AMX issues, refer to [Troubleshoot AMI/AMX Deployment](#).

For more information on health status of exporters, refer to [View Health Status](#).

You can view the following parameters in the Application Metadata Exporter application statistics page.

Exporter - Statistics to monitor exports from AMX

Parameters	Description
Name	Displays the name of exporters created for this application.
Exported Entries	Displays the number of entries available in the files that will be uploaded from AMX to the cloud tool or the Kafka consumers.
Average File Size	Displays the average size of the file.
File Upload Errors	Displays the number of file uploaded from the AMX to the cloud tool or the Kafka consumers.
File Upload Errors	Displays the number of times the file was not uploaded to the cloud tools or Kafka consumers due to errors.
Re Exported Entries (Only for Kafka Tool Exporter)	Displays the number of records re-exported from the errors directory after the Kafka tool is restored. Errored files are re-exported automatically at a rate of 10 files per 5 seconds.
Exported Entries Rolling 15 Sec	Displays the rolling average of entries exported per second over the past 15 seconds.
Exported Entries Rolling 30 Sec	Displays the rolling average of entries exported per second over the past 30 seconds.
Exported Entries Rolling 60 Sec	Displays the rolling average of entries exported per second over the past 60 seconds.
Exported MBytes	Total number of megabytes exported.
Exported MBytes Rolling 15 Sec	Rolling average of megabytes exported over the past 15 seconds.
Exported MBytes Rolling 30 Sec	Rolling average of megabytes exported over the past 30 seconds.
Exported MBytes Rolling 60 Sec	Rolling average of megabytes exported over the past 60 seconds.

Ingestor - Statistics to monitor ingestion


Parameters	Description
Name	Displays the name of the ingestor created for this application.
Packets Received	Displays the packets received from the AMI to AMX ingestor.
Packets Dropped	Displays the packets dropped when the traffic passes from AMI to AMX ingestor.
Octets Received	Displays the octets received from the AMI to AMX ingestor.
Octets Dropped	Displays the octets dropped when the traffic passes from AMI to AMX ingestor.
Records Added	Displays the number of workload information pulled from the workload environment.
Records Dropped	Displays the number of the workload details that were dropped from the workload environment.
Request Success	Displays if the workload information was fetched from the cloud environment. The default interval at which the AMX fetches the workload information is 60 seconds.
Request Authentication Errors	Displays if the workload information was not fetched from the cloud environment due to any authentication errors like wrong password.
Request Timedout Errors	Displays if the workload information was not fetched from the cloud environment even after the default interval of 60 seconds.
Request Errors	Displays when the workload information could not be fetched from the cloud environment due to issues other than authentication failures or timeout conditions.
Decode Success	Displays number of successful decoding operations performed by a system when processing NETFLOW/IPFIX packets.
Decode Fail	Displays the number of failed attempts to decode NETFLOW/IPFIX packets.
Elements Skipped	Displays number of data elements that were not processed or were omitted during the NETFLOW/IPFIX ingestion process.
Pkts Rx Rolling 15 Sec	Rolling average of packets received per second over the past 15 seconds.
Pkts Rx Rolling 30 Sec	Rolling average of packets received per second over the past 30 seconds.
Pkts Rx Rolling 60 Sec	Rolling average of packets received per second over the past 60 seconds.
Mbytes Rx	Total number of megabytes ingested.
Mbytes Rx Rolling 15 Sec	Rolling average of total megabytes received over the past 15 seconds.

Parameters	Description
Mbytes Rx Rolling 30 Sec	Rolling average of total megabytes received over the past 30 seconds.
Mbytes Rx Rolling 60 Sec	Rolling average of total megabytes received over the past 60 seconds.

Attributes Enrichment - Statistics to monitor enrichment

Parameters	Description
Name	Displays the name of the Metadata enrichment.
One Min Percent	Displays the percentage of the traffic that is enriched from the last 1 minute.
Five Min Percent	Displays the percentage of the traffic that is being enriched from the last 5 minutes.
Ten Min Percent	Displays the percentage of the traffic that is being enriched from the last 10 minutes.

The following image shows the statistical details for general AMX application.




Attributes for GigaVUE Enriched Metadata for Mobile Networks

This section describes the various attributes available within GigaVUE Enriched Metadata for Mobile Networks.

Each flow will generate only one record per export interval; however, in this section, the attributes are separated into the following categories:

- [Control Plane Attributes](#)
- [User Plane Attributes](#)



Control Plane Attributes

The control plane attributes are generated based on network signaling message transactions with successful responses. These include signaling transactions for session establishment, modification, deletion, and others. These attributes are used for enriching the user plane metadata with the corresponding control plane subscriber, device, and location information.

Attribute	Description
imsi	International Mobile Subscriber Identity. A 3GPP-defined unique private identifier for a mobile user on a 4G network.
imei	International Mobile Equipment Identity. A 3GPP-defined unique identifier for a mobile device on a 4G network.
msisdn	Mobile Station Integrated Services Digital Network. A 3GPP-defined identifier for a mobile device on a 4G network.
supi	Subscription Permanent Identifier. A 3GPP-defined unique identifier for a mobile user on a 5G network.
gpsi	Generic Public Subscription Identifier. A 3GPP-defined public identifier for a mobile user on a 5G network.
pei	Permanent Equipment identifier. A 3GPP-defined unique identifier for a mobile device on a 5G network.
ctrl_tun_access_ipv4_addr	Control Plane Access Tunnel IPv4 Address. MME identifier for S11, SGW identifier for S5/S8-C, AMF identifier for N11 and SGSN for Gn, GP
ctrl_tun_access_ipv6_addr	Control Plane Access Tunnel IPv6 Address. MME identifier for S11, SGW identifier for S5/S8-C, AMF identifier for N11 and SGSN for Gn, GP

ctrl_tun_access_teid	Control Plane Access Tunnel Endpoint Identifier
ctrl_tun_core_ipv4_addr	Control Plane Core Tunnel IPv4 Address. SGW identifier for S11, PGW identifier for S5/S8-C, SMF identifier for N11 and GGSN for Gn,Gp
ctrl_tun_core_ipv6_addr	Control Plane Core Tunnel IPv6 Address. SGW identifier for S11, PGW identifier for S5/S8-C, SMF identifier for N11 and GGSN for Gn,Gp
ctrl_tun_core_teid	Control Plane Core Tunnel Endpoint Identifier
user_tun_up_link_ipv4_addr	User Plane Uplink Tunnel IPv4 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_ipv6_addr	User Plane Uplink Tunnel IPv6 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_teid	User Plane Uplink Tunnel Endpoint Identifier
user_tun_down_link_ipv4_addr	User Plane Downlink Tunnel IPv4 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_ipv6_addr	User Plane Downlink Tunnel IPv6 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_teid	User Plane Downlink Tunnel Endpoint Identifier
if_name	3GPP Control Plane Interface Name
name	Name of the event on the 3GPP control plane interface
cause_code	Outcome of the event on the 3GPP control plane interface. Mostly set to 16 to indicate a successful event
sm_context_ref	Session Management Context
ebi	EPS(Evolved Packet System) Bearer ID
lbi	Linked Bearer Identity
pdu_session_id	PDU(Packet Data Unit) Session ID
apn	Access Point Name
dnn	Data Network Name
ue_ipv4_addr	User Equipment IPv4 Address
ue_ipv6_addr	User Equipment IPv6 Address
ue_ipv4v6_addr	User Equipment IPv4v6 Address

addr	
ue_non_ip_addr	User Equipment non-IP Address
ue_addr_type	User Equipment IP address type (IPv4 or IPv6)
qci	QoS Class Identifier
qfi	QoS Flow Identifier
five_qi	5G QoS Identifier
cgi_mcc	Mobile Country Code from Cell Global Identity
cgi_mnc	Mobile Network Code from Cell Global Identity
cgi_lac	Local Area Code from Cell Global Identity
cgi_cell_id	Cell Identification from Cell Global Identity
sai_mcc	Mobile Country Code from Service Area Identifier
sai_mnc	Mobile Network Code from Service Area Identifier
sai_lac	Local Area Code from Service Area Identifier
sai_sac	Service Area Code from Service Area Identifier
tai_mcc	Mobile Country Code from Tracking Area Identity
tai_mnc	Mobile Network Code from Tracking Area Identity
tai_tac	Tracking Area Code from Tracking Area Identity
ecgi_mcc	Mobile Country Code from E-UTRAN Cell Global Identifier
ecgi_mnc	Mobile Network Code from E-UTRAN Cell Global Identifier
ecgi_cell_id	Cell Identification from E-UTRAN Cell Global Identifier
lai_mcc	Mobile Country Code from Location Area Identity
lai_mnc	Mobile Network Code from Location Area Identity
lai_lac	Local Area Code from Location Area Identity
enode_id_mcc	Mobile Country Code from Evolved Node
enode_id_mnc	Mobile Network Code from Evolved Node
macro_enode_id	Evolved Node ID
ncgi_mcc	Mobile Country Code from NR Cell Global Identity
ncgi_mnc	Mobile Network Code from NR Cell Global Identity
ncgi_nr_cell_id	Cell Identification from NR Cell Global Identity
rat_type	Type of Radio Access Technology
snssai_sst	Slice/Service Type from Single Network Slice Selection Assistance Information
snssai_sd	Slice Differentiator id from Single Network Slice Selection Assistance Information

User Plane Attributes

The user plane attributes are generated by Application Metadata Intelligence (AMI), and include information such as IP addresses, protocols, timestamps, and byte/packet counters. These attributes are generated from the network flows and do not require DPI capabilities. The attributes for the mobile network traffic are generated for the inner IP flow by bypassing the outer headers

Attribute	Description
ts	Timestamp of the metadata generated in UTC
vendor	Identifying Gigamon as the vendor providing the metadata
version	Version number of the GigaSMART software release from the AMI application
generator	Identifies the Gigamon device that generates the user plane metadata. For example - HC3
src_ip	Source IPv4 address of the inner flow
dst_ip	Destination IPv4 address of the inner flow
src_mac	Source MAC address of the inner flow
dst_mac	Destination MAC address of the inner flow
mpls	MPLS Label if the flow has the MPLS header
protocol	Layer 4 protocol of the inner flow
src_port	Layer 4 protocol source port of the inner flow
dst_port	Layer 4 protocol destination port of the inner flow
device_inbound_interface	Traffic receiving port on the Gigamon device example HC3
tcp-sport	Source TCP port
tcp_dport	Destination TCP port
udp_sport	Source UDP port
udp_dport	Destination UDP port
src_bytes	Source bytes from client to server in the flow including the outer and the inner IP header
dst_bytes	Destination bytes from client to server in the flow including the outer and the inner IP header
src_inner_bytes	Source bytes from the client to the server in the flow with the inner IP header (without the outer header)
dst_inner_bytes	Destination bytes from the server to the client in the flow with the inner IP header (without the outer header)
src_packets	Source packets from client to server in the flow including the header
dst_packets	Destination packets from client to server in the flow including header
total_bytes	Total bytes is sum of source and destination bytes
total_packets	Total packets is sum of source and destination packets
id	Unique flow id assigned to a flow


ingress_vlan_id	Vlan id in the packet received for the flow
src_ipv6	Source IPv6 address of the inner flow
dst_ipv6	Destination IPv6 address of the inner flow
ip_version	IP version v4, v6
ip_cos_id	QOS value in the IPv4 header
ip_dscp	DSCP value in the IPv4 header
ip_ttl	Time to live in the IPv4 header
ip6hop_limit	Hop limit in the IPv6 header
flow_labelv6	flow label in the IPv6 header
ip6_ds	DSCP value in the IPv6 header
tcp_flags	TCP flag in the TCP header
gre_key	Generic routing encapsulation header key in the GRE header
ip_hdr_len	IPv4 header length
ip_frag_flags	IPv4 fragmentation flags
ipv6frag_flags	IPv6 fragmentation flags
ip_frag_id	IPv4 fragment id
ipv6frag_id	IPv6 fragment id
ip_frag_offset	IPv4 fragment offset
ipv6frag_offset	IPv6 fragment offset
ipv4opt	IPv4 options
ip_precedence_id	IPv4 precedence id
ip6precedence_id	IPv6 precedence id
ip_tot_len	IPv4 total length
ip6tot_len	IPv6 total length
ipv6hdr_len	IPv6 header length
payload_len_id	Payload length for the flow excluding the L3 IPv4 header
next_hdr_v6	Layer 4 protocol
icmp_type_v4	ICMP message type IPv4
icmp_code_v4	ICMP response code IPv4
icmp_code_v6	ICMP response code IPv6
icmp_type_v6	ICMP message type IPv6
tcp_ack_id	TCP ack id
tcp_hdr_len	TCP header length
tcp_seq_no	TCP sequence number
tcp_urgent_ptr	TCP urgent pointer
tcp_win_size	TCP window size
udp_msg_len	UDP message length
ip6traffic_class	IPv6 traffic class same as QoS in IPv4

flow_start_usec	Inner flow start time in microseconds in UTC
flow_end_usec	Inner flow end time in microseconds in UTC
flow_start_sec	Inner flow start time in seconds in UTC
flow_end_sec	Inner flow end time in seconds in UTC
start_time	Inner flow start time in milliseconds in UTC
end_time	Inner flow end time in milliseconds in UTC
egress_intf_id	egress IP interface-id for the Gigamon device sending the metadata
sys_up_time_first	Difference between the flow start time and the gigaSMART uptime in milliseconds
sys_up_time_last	Difference between the flow end time and the gigaSMART uptime in milliseconds
end_reason	Inner flow end reason – TCP ack, reset, inactive, etc
tcpflagsyn	TCP flag SYN from TCP header
tcpflagsynack	TCP flag SYNACK from TCP header
tcpflagfin	TCP flag FIN (finish)from TCP header
tcpflagrst	TCP flag RST (reset) from TCP header
Labels	Labels are added to the metadata by configuration in AMX. There can be multiple labels configured and added. Each label is configured as a key and a string value.
Labels configured in the JSON example in the section below	
label_event_type	Label identifying the vendor generating the event in the AMX exporter
label_deployment	Label identifying the deployment address of the AMX exporter
label_traffic	Label identifying the traffic region where the AMX is deployed
MSS	Maximum payload size a TCP segment can carry without the headers. The MSS value is exported at the beginning of a flow. Its exported only when bidirectional flows are configured.
aggr-window-size	The actual TCP window size minimum, mean, and maximum values, calculated separately for both sender and receiver packets.
zero-window	This captures TCP zero window event statistics for both sender and receiver, indicating when the TCP receive buffer is full and data transmission is temporarily paused due to flow control.


Attributes for GigaVUE Enriched Metadata for Cloud Workloads

This section describes the various attributes available within GigaVUE Enriched Metadata for Cloud Workloads for the following platforms:

- [AWS](#)
- [Azure](#)
- [VMware vCenter](#)
- [Azure Kubernetes Service \(AKS\)](#)
- [Zscaler Private Access \(ZPA\) Attributes](#)



- [Azure Kubernetes Service \(AKS\)](#)



The following table defines the purpose and scope of each referenced dataset, detailing configuration, resource, and operational parameters across AWS, Azure, and Kubernetes environments to support multi-cloud infrastructure documentation.

AWS

Attribute	Description	Source	Destination
aws_instance_id	ID of the EC2 instance	✓	✓
aws_instance_type	Type of EC2 instance	✓	✓
aws_availability_zone	Availability Zone of the EC2 instance	✓	✓
aws_image_id	ID of the base image use for EC2 instance	✓	✓
aws_private_dns_name	Private DNS hostname name assigned to the EC2 instance	✓	✓
aws_private_ip	Private IPv4 address assigned to the EC2 instance	✓	✓
aws_public_	Public DNS name assigned to the EC2 instance	✓	✓

dns_name			
aws_public_ip	Public IPv4 address assigned to the EC2 instance	✓	✓
aws_instance_state	Current state of the EC2 instance	✓	✓
aws_subnet_id	ID of the subnet in which the EC2 instance is running	✓	✓
aws_vpc_id	ID of the VPC in which the EC2 instance is running	✓	✓
aws_iam_instance_profile_arn	Amazon Resource Name (ARN) of the IAM instance profile	✓	✓
aws_iam_instance_profile_id	ID of the IAM instance profile	✓	✓
aws_security_group_name	Security group name assigned for the EC2 instance	✓	✓
aws_security_group_id	Security group ID assigned for the EC2 instance	✓	✓
aws_network_if_id	ID of the network interface	✓	✓
aws_network_if_attach_id	ID of the network interface attachment	✓	✓
aws_tags	Tags assigned to the EC2 instance in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
aws_flat_tags	Tags assigned to the EC2 instance extracted to top level keys and values	✓	✓
aws_owner_id	ID of the AWS account that created the network interface	✓	✓
workload_platform	Platform Name - AWS	✓	✓

Azure

Attribute	Description	Source	Destination
azure_subscription_id	Subscription ID of the Azure account that created the VM	✓	✓
azure_resource_group	Resource Group of the VM	✓	✓
azure_vm_id	Complete path of the Azure VM. This includes subscription	✓	✓
azure_	Instance ID of the Azure VM	✓	✓

instance_id			
azure_instance_type	Instance type of the Azure VM	✓	✓
azure_availability_zone	Availability zone of the Azure VM	✓	✓
azure_image_id	Image ID of the Azure VM	✓	✓
azure_private_dns_name	Private DNS name assigned to the VM instance	✓	✓
azure_private_ip	Private IP address assigned to the VM instance	✓	✓
azure_public_dns_name	Public DNS name assigned to the VM instance	✓	✓
azure_public_ip	Public IP address assigned to the VM instance	✓	✓
azure_instance_state	Current state of the VM instance	✓	✓
azure_subnet_id	Subnet ID of the VM instance	✓	✓
azure_vnet_id	ID of the VNET in which the VM instance is running	✓	✓
azure_network_if_id	ID of the network interface	✓	✓
azure_network_if_attach_id	ID of the network interface attachment	✓	✓
azure_tags	Tags assigned to the VM instance in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
azure_flat_tags	Tags assigned to the VM instance extracted to top level keys and values	✓	✓
azure_owner_id	Owner of the VM instance	✓	✓
azure_location	Resource Location	✓	✓
azure_scaleset_name	Scale set name	✓	✓
azure_scaleset_mode	Scale set orchestration mode	✓	✓
workload_platform	Platform Name - Azure	✓	✓

VMware vCenter

Attribute	Description	Source	Destination
vmware_vm_name	Name of the guest VM	✓	✓

vmware_vm_network	Network name of the guest VM	✓	✓
vmware_vm_status	Current status of the guest VM	✓	✓
vmware_vm_tags	Tags assigned to the guest VM in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
vmware_vm_flat_tags	Tags assigned to the guest VM extracted to top level keys and values	✓	✓
vmware_vm_os_family_name	OS type of the guest VM (eg: Linux)	✓	✓
vmware_vm_os_distro_name	OS distribution name of the guest VM	✓	✓
vmware_vm_os_distro_version	OS distribution version of the guest VM	✓	✓
vmware_vm_os_distro_pretty_name	OS distribution pretty name of the guest VM	✓	✓
vmware_vm_dns_host_name	Host DNS name of the guest VM	✓	✓
vmware_vm_dns_host_domain_name	Host DNS domain name of the guest VM	✓	✓
vmware_host_name	Host name of the guest VM	✓	✓
vmware_host_ip	Host IP address of the guest VM	✓	✓
vmware_datacenter_name	Datacenter name of the guest VM	✓	✓
vmware_cluster_name	Cluster name of the guest VM	✓	✓
vmware_vcenter_name	vCenter name of the guest VM	✓	✓
workload_platform	Platform Name - VMware	✓	✓

Azure Kubernetes Service (AKS)

Attribute	Description	Source	Destination
k8s_cluster_name	Name identifying the Kubernetes cluster	✓	✗
k8s_cluster_version	Version of Kubernetes running on the cluster	✓	✗
k8s_cluster_status	Operational state of the Kubernetes cluster	✓	✗
k8s_node_name	Identifier for the node within the cluster	✓	✗
k8s_node_internal_ip	Internal IP address assigned to the node	✓	✗
k8s_node_external_ip	External IP address of the node	✓	✗
k8s_node_os_version	Operating system version installed on the node	✓	✗
k8s_node_os_image	Name of the OS image used by the	✓	✗

	node		
k8s_node_kernel_version	Kernel version of the node	✓	✗
k8s_node_cri_version	Version of the container runtime interface used on the node	✓	✗
k8s_node_status	Health and availability status of the node	✓	✗
k8s_node_labels	Node labels parsed into key-value pairs for metadata enrichment	✓	✗
k8s_pod_name	Name assigned to the Kubernetes pod	✓	✗
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✓	✗
k8s_pod_namespace	Namespace in which the pod is deployed	✓	✗
k8s_pod_status	Operational state of the pod	✓	✗
k8s_pod_ip	IP address assigned to the pod	✓	✗
k8s_service_name	Name of the Kubernetes service	✓	✗
k8s_service_type	Type of service	✓	✗
k8s_container_name	Name of the container running inside the pod	✓	✗
k8s_container_image	Image used to create the container	✓	✗
k8s_container_port	Container port details	✓	✗
k8s_cluster_name	Name identifying the Kubernetes cluster	✗	✓
k8s_cluster_version	Version of Kubernetes running on the cluster	✗	✓
k8s_cluster_status	Operational state of the Kubernetes cluster	✗	✓
k8s_node_name	Identifier for the node within the cluster	✗	✓
k8s_node_internal_ip	Internal IP address assigned to the node	✗	✓
k8s_node_external_ip	External IP address of the node	✗	✓
k8s_node_os_version	Operating system version installed on the node	✗	✓
k8s_node_os_image	Name of the OS image used by the node	✗	✓
k8s_node_kernel_version	Kernel version of the node's operating system	✗	✓
k8s_node_cri_version	Version of the container runtime interface used on the node	✗	✓
k8s_node_status	Health and availability status of the node	✗	✓

k8s_node_labels	Node labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_name	Name assigned to the Kubernetes pod	✗	✓
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_namespace	Namespace in which the pod is deployed	✗	✓
k8s_pod_status	Operational state of the pod	✗	✓
k8s_pod_ip	IP address assigned to the pod	✗	✓
k8s_service_name	Name of the Kubernetes service	✗	✓
k8s_service_type	Type of service	✗	✓
k8s_container_name	Name of the container running inside the pod	✗	✓
k8s_container_image	Docker image used to create the container	✗	✓
k8s_container_port	Container port details	✗	✓

Zscaler Private Access (ZPA) Attributes

Attribute	Description
z_iss_username	Username of Zscaler Private Access user.
z_iss_client_private_ip	Private IP address of the Zscaler Private Access client
z_iss_client_public_ip	Public IP address of the Zscaler Private Access. client.
z_iss_customer	Name of Zscaler Private Access customer.
z_iss_connection_id	ID of the Zscaler Private Access connection.
z_iss_connection_status	Status of the Zscaler Private Access connection.
z_iss_connection_status_code	Status code of the Zscaler Private Access connection.
z_iss_connector	Name of the Zscaler Private Access application connector.
z_iss_hostname	Name of the host device as given by the Zscaler client connector.
z_iss_client_country_code	Country code of the Zscaler Private Access client location.
z_iss_policy	Name of the Zscaler Private Access policy.
z_iss_app_group	Group name of the Zscaler Private Access application.
z_iss_policy_processing_time	Processing time of the Zscaler Private Access policy.
z_iss_idp	Identity provider used by Zscaler Private Access.
z_iss_client_city	City of the Zscaler Private Access client location.
z_iss_platform	Client platform of Zscaler Private Access.
z_iss_connection_setup_time	Connection setup time for Zscaler Private Access.

Supported Element Types for NetFlow Integrator Element Mapping

This section outlines the element types (v5, v9, and v10) supported by the AMX application for element mapping from the ingest NetFlow/IPFIX flow records to JSON on the following platforms. These elements are predefined and cannot be configured:

- VMware ESXi
- VMware NSX-T

Elements	Description	Supported for NetFlow v5	Supported for NetFlow v9	Supported for NetFlow v10/IPFIX
type	Type of flow message	✓	✓	✓
sampler_address	Address of the device that generated the packet	✓	✓	✓
sampling_rate	Sampling rate of the flow	✓	✓	✓
export_time_ns	Export time in Nano seconds	✓	✓	✓
time_received_ns	Timestamp in nanoseconds of when the message was received	✓	✓	✓
sequence_num	Sequence number of the flow packet	✓	✓	✓
source_id	An identifier that uniquely identifies the flow source. For IPFIX it refers to the observation domain ID.	✗	✓	✓
etype	Ethernet type (0x86dd for IPv6...)	✓	✓	✓
src_addr	Source address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_SRC_ADDR(8), IPV6_SRC_ADDR(27) • NetFlow v10/IPFIX - sourceIPv4Address(8), sourceIPv6Address(27)	✓	✓	✓
dst_addr	Destination address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_DST_ADDR(12), IPV6_DST_ADDR(28) • NetFlow v10/IPFIX - destinationIPv4Address(12), destinationIPv6Address(28)	✓	✓	✓
ip_tos	IP Type of Service • NetFlow v5 - tos • NetFlow v9 - SRC_TOS(5) • NetFlow v10/IPFIX - ipClassOfService(5)	✓	✓	✓

¹The numbers within the enclosed brackets represent the element ID in IPFIX/NetFlow v10 and the field type ID in NetFlow v9.

ip_dscp	Differentiated Services Code Point ipDiffServCodePoint (195)	x	✓	✓
ip_precedence	The value of the IP Precedence ipPrecedence (196)	x	✓	✓
proto	Protocol (UDP, TCP, ICMP...) - NetFlow v5 – proto - NetFlow v9 - PROTOCOL (4) - NetFlow v10/IPFIX - protocolIdentifier (4)	✓	✓	✓
ip_ttl_max	Maximum TTL value observed for packets of the flow - NetFlow v9 - MAX_TTL (53) - NetFlow v10/IPFIX - maximumTTL (53)	x	✓	✓
src_port	Source port (when UDP/TCP/SCTP) - NetFlow v5 – srcport - NetFlow v9 - L4_SRC_PORT (7) - NetFlow v10/IPFIX - sourceTransportPort (7)	✓	✓	✓
dst_port	Destination port (when UDP/TCP/SCTP) - NetFlow v5 – dstport - NetFlow v9 - L4_DST_PORT (11) - NetFlow v10/IPFIX - destinationTransportPort (11)	✓	✓	✓
bytes	Number of bytes in a flow - NetFlow v5 – dOctets - NetFlow v9 - IN_BYTES (1), OUT_BYTES (23) - NetFlow v10/IPFIX - octetDeltaCount (1), postOctetDeltaCount (23)	✓	✓	✓
packets	Number of packets in a flow - NetFlow v5 – dPkts - NetFlow v9 - IN_PKTS (2), OUT_PKTS (24) - NetFlow v10/IPFIX - packetDeltaCount (2), postPacketDeltaCount (24)	✓	✓	✓
bytes_total	Running byte counter for a permanent flow - NetFlow v9 - IN_PERMANENT_BYTES (85) - NetFlow v10/IPFIX - octetTotalCount(85)	x	✓	✓
packets_total	Running packet counter for a permanent flow - NetFlow v9 - IN_PERMANENT_PKTS (86) - NetFlow v10/IPFIX - packetTotalCount	x	✓	✓

	(86)			
flow_direction	Flow direction - NetFlow v9 - DIRECTION (61) - NetFlow v10/IPFIX - flowDirection(61)	x	✓	✓
in_if	Input interface - NetFlow v5 – input - NetFlow v9 - INPUT_SNMP (10) - NetFlow v10/IPFIX - ingressInterface (10)	✓	✓	✓
out_if	Output interface - NetFlow v5 – output - NetFlow v9 - OUTPUT_SNMP (14) - NetFlow v10/IPFIX - egressInterface (14)	✓	✓	✓
time_flow_start_ns	Time the flow started in nanoseconds. Refer to Important Notes . - NetFlow v5 – System uptime and first - NetFlow v9 - System uptime and FIRST_SWITCHED (22) - NetFlow v10 - flowStartXXX (150, 152, 154, 156)	✓	✓	✓
time_flow_end_ns	Time the flow ended in nanoseconds. Refer to Important Notes . - NetFlow v5 – System uptime and last - NetFlow v9 - System uptime and LAST_SWITCHED (21) - NetFlow v10/IPFIX - flowEndXXX (151, 153, 155, 157)	✓	✓	✓
tcp_flags	TCP flags - NetFlow v5 – tcp_flags - NetFlow v9 - TCP_FLAGS (6) - NetFlow v10/IPFIX - tcpControlBits (6)	✓	✓	✓
ip_length_min	Minimum IP packet length - NetFlow v9 - MIN_PKT_LENGTH (25) - NetFlow v10/IPFIX - minimumIpTotalLength (25)	x	✓	✓
ip_length_max	Maximum IP packet length - NetFlow v9 - MAX_PKT_LENGTH (26) - NetFlow v10/IPFIX - maximumIpTotalLength (26)	x	✓	✓
next_hop	Nexthop IP address - NetFlow v5 – nexthop - NetFlow v9 - IPV4_NEXT_HOP (15), IPV6_NEXT_HOP (62) - NetFlow v10/IPFIX - ipNextHopIPv4Address (15), ipNextHopIPv6Address (62)	✓	✓	✓

ipv4_next_hop	Nexthop IPv4 address - NetFlow v9 - IPV4_NEXT_HOP (15) - NetFlow v10/IPFIX - ipNextHopIPv4Address (15)	x	✓	✓
ipv6_next_hop	Nexthop IPv6 address - NetFlow v9 - IPV6_NEXT_HOP (62) - NetFlow v10/IPFIX - ipNextHopIPv6Address (62)	x	✓	✓
flow_id	An identifier of a flow that is unique within an Observation Domain - NetFlow v9 - flowId (148) - NetFlow v10/IPFIX - flowId (148)	x	✓	✓
firewall_event	Indicates a firewall event. Allowed values are listed in the firewall Event registry. firewallEvent(233)	x	✓	✓
icmp_type	Type of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpTypeXXX (176, 178), icmpTypeCodeXXX (32, 139)	x	✓	✓
icmp_code	Code of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpCodeXXX (177, 179) icmpTypeCodeXXX (32, 139)	x	✓	✓
flow_end_reason	The reason for flow termination. Values are listed in the flowEndReason registry.	x	✓	✓
application_id	Specifies an Application ID - NetFlow v9 - APPLICATION TAG (95) - NetFlow v10/IPFIX - applicationId (95)	x	✓	✓
application_description	Specifies the description of an application - NetFlow v9 - APPLICATION DESCRIPTION (94) - NetFlow v10/IPFIX - applicationDescription (94)	x	✓	✓
application_name	Specifies the name of an application - NetFlow v9 - APPLICATION NAME (96) - NetFlow v10/IPFIX - applicationName (96)	x	✓	✓
tcp_window_size	The window field in the TCP header tcpWindowsSize (186)	x	✓	✓
viptela_vpn_id	Private enterprise number for vIPtela Inc. (41916)	x	x	✓

Troubleshoot AMI/AMX Deployment

This topic helps troubleshoot AMI/AMX deployments on GigaVUE V Series Nodes and focuses specifically on scenarios where AMI is not exporting IPFIX/CEF records and AMX is not receiving traffic or is not exporting JSON records to downstream tools.

Scenario 1: AMI not exporting IPFIX/CEF records (exporter side)

Download exporter records

To download the exporter records:

1. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > VMware** (You can select the cloud platform which you have configured. This scenario explains the VMware platform). The Monitoring Session page appears.
2. Select the monitoring session that you want to troubleshoot.
3. Go to the **Traffic Processing** tab. The canvas appears with the monitoring session components
4. Select the AMI application and click Three-dots menu to view the application details.
5. Select the **Statistics** tab and the Exporter section appears.
6. Select the required records and click **Download Records**. The Exporter Mapping dialog appears.
7. Select the V Series Node and the tunnels for exporters from the dropdown list.
8. Click **Download Exporter**. The PCAP file (for example, Pcap__e1_link3_2026-03-12_174220.pcapng) is downloaded to your local machine. The progress of the PCAP file download appears in the browser.

NOTE: If there are no packets or records during the capture window, the system displays the message **No file captured**.

9. Verify whether the PCAP file contains valid IPFIX/CEF records for the specified exporter instance.

Scenario 2: AMX not receiving AMI/NetFlow records (ingestor side)

Verify topology and mapping

To verify topology and mapping:

1. Confirm whether the AMI exporter is configured to send records to the data interface of the correct AMX instance.
2. Use the AMX interface mapping in FM (management/data/tool) to ensure the correct interface is used for receiving records.

Download PCAP file using Download TCP Dump

To download the PCAP file:

1. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > VMware**. The Monitoring Session page appears. You can select the cloud platform which you have configured. This scenario explains the VMware platform.
2. Select the monitoring session that you want to troubleshoot.
3. Go to the **Traffic Processing** tab. The canvas appears with the monitoring session components
4. Select the AMX application and click Three-dots menu to view the application details.
5. Select the **Statistics** tab and select the V Series node for the AMX instance from the V Series Node IP dropdown list.
6. In the **Ingestor** section, select the required record and click **Download TCP Dump**. The **Ingestor Interface Mapping** dialog appears.
7. Select the interface mapping for the ingestor. The AMX interface naming includes the type and IP address of the interface:
 - For example, Zscaler ingestor - MGMT (ens160) IPv4: 10.114.26.181 IPv6: fcab:0:1:2040:a0a1:874e:7a66:40a2
 - Other ingestors (AMI, NetFlow/IPFIX, Mobility) - DATA (ens192) IPv4: 10.114.42.135 IPv6: fcab:0:1:2040:a0a1:874e:7a66:40a2
8. click **Download TCP Dump**. The PCAP file is downloaded to your local machine. The progress of the PCAP file download appears in the browser.

NOTE: The PCAP file is stored temporarily on the Application Metadata Exporter virtual machine and is automatically deleted after download.

9. Verify whether (IPFIX/CEF/NetFlow/JSON-over-UDP/GTPC) records are visible on the ingestor port. If the PCAP is empty or only shows partial traffic, this suggests a datapath issue upstream of AMX (AMI export or fabric).

Scenario 3: AMX not exporting JSON records (exporter side)

Check exporter status and counters

To check exporter status:

1. In the Exporter section, verify the number of files processed versus files skipped.
2. Check the error counters (if exposed via UI).

Download exporter records

To download the exporter records:

1. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > VMware**. The Monitoring Session page appears. You can select the cloud platform which you have configured. This scenario explains the VMware platform.
2. Select the monitoring session that you want to troubleshoot.
3. Go to the **Traffic Processing** tab. The canvas appears with the monitoring session components
4. Select the AMX application and click Three-dots menu to view the application details.
5. Select the **Statistics** tab and select the V Series node from the V Series Node IP dropdown list.
6. In the Exporter section, verify the number of Files processed vs files skipped.
7. Select the required records and click **Download Records**. The exporter JSON file is downloaded to your local machine. The progress of the PCAP file download appears in the browser.

NOTE: If there are no packets or records during the capture window, the system displays the message **No file captured**.

8. Validate the JSON structure and key fields required by the tool (e.g., Splunk, Elastic).
9. Check for malformed or truncated records.

GigaSMART NetFlow Generation

NetFlow Generation is a simple and effective way to increase visibility into traffic flows and usage patterns across systems. The flow-generated data can be used to build relationships and usage patterns between nodes on the network.

Refer to the following topic for step-by-step instructions on how to configure NetFlow for Virtual Environment:

- [Configure Application Metadata Intelligence for Virtual Environment](#)- For SecureVUE Plus Base Bundle
- [Create NetFlow Session for Virtual Environment](#) - ForNetVUE Base Bundle

Create NetFlow Session for Virtual Environment

NOTE: This configuration is applicable only when using NetVUE Base Bundle.

To create an NetFlow session, follow these steps:

1. Drag and drop **Application Metadata** from **APPLICATIONS** to the graphical workspace.
2. Click the Application Metadata application and select **Details**. The Application quick view appears.

3. In the Application quick view, enter or select the following details in the **General** tab:

Parameter	Description																							
Name	Enter a name for the application.																							
Description	Enter the description.																							
Application Metadata Settings																								
Flow Direction	Enable or Disable Bi-Directional Flow behavior. Bi-Directional is enabled by default. Disable this option for Uni-Directional Flow behavior.																							
Timeout	Specify the traffic flow inactivity timeout, in seconds. The session will be removed due to inactivity when no packets match.																							
Data Link	If you want to include the VLAN ID along with the 5-tuple to identify the traffic flow, select the Data Link and enable the VLAN option.																							
Observation ID	Enter a value to identify the source from where the metadata is collected. The range is from 0 to 255. The calculated value of Observation Domain Id in Hexadecimal is 00 01 02 05 , and in Decimal is 66053 .																							
Advanced Settings																								
Number of Flows	The number of flows supported by the application. Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms. <table border="1"> <thead> <tr> <th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr> </thead> <tbody> <tr> <td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> <tr> <td rowspan="4">AWS</td><td>AMD - Large (c5n.2xlarge)</td><td>300k</td></tr> <tr> <td>AMD - Medium (t3a.xlarge)</td><td>100k</td></tr> <tr> <td>ARM - Large (c7gn.2xlarge)</td><td>100k</td></tr> <tr> <td>ARM - Medium (m7g.xlarge)</td><td>200k</td></tr> <tr> <td rowspan="2">Azure</td><td>Large (Standard_D8s_V4)</td><td>500k</td></tr> <tr> <td>Medium (Standard_D4s_v4)</td><td>100k</td></tr> <tr> <td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> </tbody> </table> - Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k.. - If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.	Cloud Platform	Instance Size	Maximum Number of Flows	VMware	Large (8 vCPU and 16GB RAM)	200k	AWS	AMD - Large (c5n.2xlarge)	300k	AMD - Medium (t3a.xlarge)	100k	ARM - Large (c7gn.2xlarge)	100k	ARM - Medium (m7g.xlarge)	200k	Azure	Large (Standard_D8s_V4)	500k	Medium (Standard_D4s_v4)	100k	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows																						
VMware	Large (8 vCPU and 16GB RAM)	200k																						
AWS	AMD - Large (c5n.2xlarge)	300k																						
	AMD - Medium (t3a.xlarge)	100k																						
	ARM - Large (c7gn.2xlarge)	100k																						
	ARM - Medium (m7g.xlarge)	200k																						
Azure	Large (Standard_D8s_V4)	500k																						
	Medium (Standard_D4s_v4)	100k																						
Nutanix	Large (8 vCPU and 16GB RAM)	200k																						

NOTE: When using NetVUE Base Bundle, Multi-Collect, Fast Mode, and Aggregate round-trip time fields are disabled.

4. In the Application quick view, enter or select the following details in the **Exporters** tab:

Parameter	Description
Exporter Name	Enter a name for the Exporter.
Actions	Using this option, you can perform the following functions: • Add Exporter - Use to add a new Exporter to this Application Metadata Intelligence Application • Apply Template - Use to select the tool template. Refer to Create Custom Tool Templates for more details on what are tool templates and to create custom tool templates. • Save as New Template - Use to save the current configuration as a new custom tool template. • Delete this Exporter - Use to delete the Exporter.
APPLICATION ID	Enable to export the data with Application Id.
Format	Select NetFlow
NetFlow: Select this option to use NetFlow	
Record / Template type	• Segregated - The application-specific attributes and the generic attributes will be exported as individual records to the tool. • Cohesive- The application-specific attributes and the generic attributes will be combined as a single record and exported to the tool.
Active Timeout	Enter the active flow timeout value in seconds.
Inactive Timeout	Enter the inactive flow timeout in seconds.
Version	Select the NetFlow version.
Template Refresh Interval	Enter the time interval at which the template must be refreshed in seconds
APPLICATION & ATTRIBUTES: Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each application. (Not applicable when using Netflow V5, V9, Netflow IPFIX(V10), or CEF when the flow direction is Uni-Directional in the above Template drop-down menu.)	
Add Application	Click on the Add Application button. The Add Application dialog box opens. Select a Type. The available options are: • Application Family: Each application is mapped only mapped to one Application Family ■ Select an Application Family and the Applications that needs to be filtered from the traffic. ■ Attributes for the selected application is displayed in the Attribute column. You can select the required attributes. • Application Tag: Each application can be mapped to one or more Application Tags. ■ Select an Application Tag and the Applications that needs

Parameter	Description
	to be filtered from the traffic. Attributes for the selected application is displayed in the Attribute column. You can select the required attributes.
NETWORK & TRANSPORT PARAMETERS: Select the Network and the transport packet attributes with the respective parameters	
Data Link	Select any one of the parameters such as Source MAC Address, Destination MAC Address and VLAN.
Interface	Select any one of the parameter such as Input Physical, Output Physical and Input Name.
IP	Select the parameter as Version if required.
IPv4	Select the required attributes. By default, Source Address, Destination Address, and Protocol are enabled.
IPv6	Select the required attributes. By default, Source Address, Destination Address, and Next Header are enabled.
Transport	Select the required attributes. By default, Source Port, Destination Port are enabled.
Counter	Select the Bytes, and Packets.
Timestamp	Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
Flow	Select the parameter as End Reason if required.
GTP-U	Select the required parameters such as QFI and TEID.
Outer IPv4	Select any one of the parameter such as Source or Destination.
Outer IPv6	Select any one of the parameter such as Source or Destination.

5. Click **Save**.

Examples- Configuring Application Intelligence Solution with Other Applications

This sections provides information on how applications like Application Filtering Intelligence, Application Metadata Intelligence and Application Metadata Exporter can be used with other applications in the monitoring session.

Refer to the following topics for more detailed information:

- [Slicing and Masking with Application Filtering Intelligence](#)
- [De-duplication with Application Metadata Intelligence](#)

Slicing and Masking with Application Filtering Intelligence

When the traffic passes through the Application Filtering Intelligence, application metadata is created. You can use the Slicing and Masking application along with Application Filtering application slice, mask, or slice and mask the filtered packets before sending them to the destination tunnel endpoint.

NOTE: When combining Slicing and Masking operations, the offset range of the Masking must be lesser than the offset value entered for the Slicing operation, as the Slicing operation is performed first.

Follow the steps below to configure Application Filtering Intelligence with Masking and Slicing:

1. Drag and drop **New Map / Application Filtering** from **New** to the graphical workspace.
2. Click the map and select **Details**. The Application quick view appears.
3. Configure Application Filtering Intelligence using the steps given in [Configure Application Filtering Intelligence for Virtual Environment](#).
4. Drag and drop **Slicing** from **Applications** to the graphical workspace.
5. Click the application and select **Details**. The Application quick view appears.
6. Configure Slicing application using the steps given in [Slicing](#).
7. Drag and drop **Masking** from **Applications** to the graphical workspace.
8. Click the application and select **Details**. The Application quick view appears.
9. Configure Masking application using the steps given in [Masking](#).
10. Drag and drop **New Tunnel** from **New** to the graphical workspace.
11. Click the tunnel and select **Details**. The Application quick view appears.
12. Select the **Type** as L2GRE/VXLAN. Select the Traffic Direction as **Out**. Refer to Create Ingress and Egress Tunnel section in the respective Cloud Deployment guides for step-by-step instructions on how to configure Tunnels.
13. Enter Source L4 Port and Destination L4 Port.
14. After placing the required items in the canvas, hover your mouse on the applications, click the red dot, and drag the arrow over to another item (map, application, or tunnel).

The filtered traffic will be sent to the Slicing application, the sliced traffic will be sent to Masking application and then to the destination tunnel Endpoint.

De-duplication with Application Metadata Intelligence

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output. Using de-duplication with Application Metadata Intelligence lets you eliminate the duplicate packets in the

Application Metadata output, only forwarding a packet once and thus reducing the processing load on your tools.

Follow the steps below to configure Application Metadata Intelligence with De-duplication:

1. Drag and drop **Application Metadata** from **Applications** to the graphical workspace.
2. Click the application and select **Details**. The Application quick view appears.
3. Configure Application Metadata Intelligence using the steps given in [Configure Application Metadata Intelligence for Virtual Environment](#).
4. Drag and drop **dedup** from **Applications** to the graphical workspace.
5. Click the application and select **Details**. The Application quick view appears.
6. Configure de-duplication application using the steps given in [De-duplication](#).
7. Drag and drop **New Tunnel** from **New** to the graphical workspace.
8. Click the tunnel and select **Details**. The Application quick view appears.
9. Select the **Type** as UDP.
10. Enter Source L4 Port, Destination L4 Port, and Destination IP. Refer to Create Ingress and Egress Tunnel section in the respective Cloud Deployment guides for step-by-step instructions on how to configure Tunnels.
11. After placing the required items in the canvas, hover your mouse on the map, click the red dot, and drag the arrow over to another item (map, application, or tunnel).

The duplicate packets are removed before sending the traffic to AMI. This will reduce the load on Application Metadata application which in turn can avoid exporting the duplicated Metadata to the tool.

Migrate Application Intelligence Session to Monitoring Session

Starting from Software version 6.5.00, you must configure the Application Intelligence solution from Monitoring Session Page. After upgrading to 6.5.00, you cannot create a new Application Intelligence Session or edit an existing Application Intelligence Session for a virtual environment from the **Application Intelligence** page.

The following actions are available only when using the existing Application Intelligence Session:

- View Details
- Delete
- Forced Delete

We recommend to migrate the existing sessions to Monitoring Session for full functionality. GigaVUE-FM seamlessly migrates all your virtual Application Intelligence sessions and their connections. If migration fails, all sessions return to their original states.



Points to Note:

- You must have write access for the **Traffic Control Management** Resource in GigaVUE-FM to perform this migration. For details, refer to Create Roles section In GigaVUE Administration Guide
- The migration does not proceed:
 - If any of the existing Application Intelligence Session is in PENDING or SUSPENDED. Resolve the issue and start the migration process.
 - If any of the existing Application Intelligence Session is in FAILED state due to incorrect configuration. Resolve the issue and start the migration process.
 - If an existing Monitoring Session has the same name as the Application Intelligence Session. Change the existing Monitoring Session name to continue with the migration process.
- You cannot continue the session if any of the existing Application Intelligence Session has Application Filtering configured with Advanced Rules as Drop Rule and No Rule Match Pass All in the 5th rule set. In the Monitoring Session, the fifth Rule Set supports either Pass All or Advanced Rules as Drop. Delete this session and start the migration.
- When migrating the Application Intelligence Session, in rare scenarios, the migration process might fail after the pre-validation. In such cases, all the Application Intelligence Session roll back to the Application Intelligence page. Contact Technical Support for assistance.

Migrate your existing Application Intelligence Session to Monitoring Session Page

Follow these steps:

1. In the left navigation pane, select **Traffic > Solutions > Application Intelligence**. You cannot create a new Application Intelligence Session from this page. When you have an existing virtual Application Intelligence Session in the above page, the **Migrate Virtual Application Intelligence** dialog box appears.
2. Review the message and select **Migrate**. The **Confirm Migration** dialog box appears with the list of Application Intelligence Session that you need to migrate.
3. Review the list and select **Migrate**. GigaVUE-FM verifies the requirements and then migrates the Application Intelligence Sessions to the Monitoring Session Page.
4. Select **Go to Monitoring Session Page**.

You can view that all the virtual Application Intelligence Sessions in the Application Intelligence page are migrated to the Monitoring Session Page.

Post Migration Notes for Application Intelligence

After migrating Application Intelligence session to Monitoring Session page, you must consider the following:

1. Secure Tunnels in the Options page

If you wish to enable Secure tunnels after migrating the Application Intelligence Session, follow these steps:

- a. Enable Secure Tunnels in the **Options** page.

For more information, refer to [Configure Monitoring Session Options](#).

- b. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.

The **Monitoring Sessions** page appears.

- c. Select the Monitoring Session for which you enabled Secure Tunnels.

- d. Select **Actions > Undeploy**.

The monitoring session is undeployed.

- e. Select the Monitoring Session for which you enabled Secure Tunnels.

- f. Select **Actions > Edit**.

The **Edit Monitoring Session** Canvas page appears.

- g. Add the Application Intelligence applications.

- h. Modify the Number of Flows as per the below table.

Cloud Platform	Instance Size	Maximum Number of Flows
VMware	Large (8 vCPU and 16GB RAM)	200k

- Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
- If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.

- i. Select **Deploy**.

For details, refer to Application Intelligence section in the GigaVUE V Series Applications Guide

2. Temporary Loss of Statistics with Version Mismatch

When GigaVUE-FM version is 6.5.00, and the GigaVUE V Series Node version is below 6.5.00, after migrating the Application Intelligence Session to the Monitoring Session and redeploying the monitoring session, a momentary loss in the statistical data of the Application Visualization application is possible while redeploying the monitoring session.

3. Configuration Changes Post-Migration

To make configuration changes post-migration, make sure that the GigaVUE V Series Node version is greater than or equal to 6.3.00.

De-duplication

De-duplication application targets, identifies, and eliminates duplicate packets, blocking unnecessary duplication and sending optimized flows to your security and network monitoring tools. De-duplication lets you detect and choose the duplicate packets to count or drop in a network analysis environment.

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output. They can also appear when packets are gathered from multiple collection points along a path. The de-duplication application lets you eliminate these packets, only forwarding a packet once and thus reducing the processing load on your tools.

Refer to the following topics:

- [Feature Overview](#)
- [Configure De-duplication Application](#)
- [Distributed De-duplication](#)

Feature Overview

There are two actions that can be specified for handling the duplicate packets detected:

- drop, which drops the duplicate packets
- count, which counts the duplicate packets, but does not drop them

A time interval can be configured within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination.

For example, if two of the same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates.


For IPv4 and IPv6 packets, to determine if a packet is considered to be a duplicate, parts of the IP headers (Layer 3 and Layer 4), as well as part of the payload are compared.

For non-IP packets, a packet is considered to be a duplicate if it is identical.

Configure De-duplication Application

To add a de-duplication application:

1. Drag and drop **Dedup** from **APPLICATIONS** to the graphical workspace.
2. Click the Dedup application and select **Details**. The Application quick view appears.



The screenshot shows the 'Application' configuration window for the 'Dedup' application. The window has a title bar with 'X', 'Application', and a 'Save' button. On the left is a sidebar with a grid icon and a 'dedup' icon. The main area contains the following configuration fields:

Application	Dedup ⓘ
Alias	dedup
Action	☐ Count ☒ Drop
IP Tclass	☒ Include ☐ Ignore
IP TOS	☒ Include ☐ Ignore
TCP Sequence	☒ Include ☐ Ignore
VLAN	☐ Include ☒ Ignore
Timer	50000

3. In the Application quick view, enter the information as follows:

Parameter	Description
Alias	Enter a name for the application.
Action	Specifies whether duplicate packets are to be counted or dropped as follows: - o Count– The de-duplication application counts the duplicate packets, but does not drop them. - o Drop– The de-duplication application drops the duplicate packets. The default is drop.
IP Tclass IP TOS TCP Sequence VLAN	These options are useful when applying de-duplication operations to packets in a NAT environment. Different NAT implementations can change certain packet header fields (for example, the TCP sequence number). If you want to be able to detect duplicates without requiring that these fields match (ToS field, TCP sequence number, VLAN ID), you can disable the corresponding option. - o IP Tclass – Ignore or include IPv6 traffic class. Use for IPv6. The default is include. - o IP TOS – Ignore or include the IP ToS bits when detecting duplicates. Use for IPv4. The default is include. - o TCP Sequence – Ignore or include the TCP Sequence number when detecting duplicates. The default is include. - o VLAN – Ignore or include the VLAN ID when detecting duplicates. The default is ignore. Include means the field will be included when the application compares packets. Ignore means the field will be ignored when the application compares packets.
Timer <Value: 10-500000 µs>	Configures the time interval within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination. The default is 50,000µs. For example, if two same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates. NOTE: Retransmissions are not counted as duplicates.

4. Click **Save**.

Distributed De-duplication

In distributed de-duplication, when you set up a Monitoring Session with a de-duplication app, the traffic is first sent to a component which distributes the traffic based on a consistent mechanism to ensure that packets from a particular traffic go to the same de-duplication instance. The packets are shared across the GigaVUE V Series Nodes. The distributed de-duplication more efficient that the existing de-duplication.

NOTE: From version 6.9, Traffic Distribution option is renamed to Distributed De-duplication.

From version 6.5, distributed de-duplication is supported on Azure, AWS, GCP, Nutanix, VMware NSX-T, VMware ESXi, and OpenStack. An enhanced configuration profile for the load balancer will be set by default with no option for modification. The default profile will use source and destination IP addresses, source and destination ports as the configuration for calculating the hash value for traffic distribution.

Important:

1. False traffic health alarms could be raised due to distribution of traffic across GigaVUE V Series Nodes.
2. Statistics are displayed for all the applications. Distributed De-duplication requires additional entities which will be listed in statistics page.

Limitations

- When using distributed de-duplication application, if **Prefer IPv6** option is enabled when configuring GigaVUE V Series Node using GigaVUE-FM, IPv6 tunnel will be created only between GigaVUE V Series Node and UCT-V and the rest of the tunnels created are IPv4 tunnels. Refer to Configure GigaVUE Fabric Components section in the respective cloud deployment guide for more detailed information on how to deploy GigaVUE V Series Node with IPv6 tunnels.
- When using distributed de-duplication application, if **Enable IPv6 Preference** option is enabled when configuring GigaVUE V Series Node using Third Party Orchestration, IPv6 tunnel will be created only between GigaVUE V Series Node and UCT-V and the rest of the tunnels created are IPv4 tunnels. Refer to *Create Monitoring Domain* section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration Guide for more detailed information on how to create a monitoring domain to register GigaVUE fabric components with IPv6 tunnels.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on Traffic health monitoring and how to create threshold template, refer to Traffic Health Monitoring section in the respective cloud deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to Monitor Cloud Health section in the respective cloud deployment guides.


You can view the statistics of the application in the **STATISTICS** tab.

GENEVE Decapsulation

The GENEVE Decapsulation application is used to acquire and strip GENEVE headers. To route the traffic through the third-party network appliances seamlessly, the AWS gateway load balancer with a VPC adds GENEVE header to packets as they are forwarded to a third-party network appliance. Each appliance is expected to terminate the GENEVE tunnel and process the GENEVE encapsulated traffic. When the GigaVUE-FM directs the acquisition of the customer traffic, the packets are encapsulated and forwarded as GENEVE tunnels that are terminated in GigaVUE V Series nodes.

To add a GENEVE application:

1. Drag and drop **GENEVE** from **APPLICATIONS** to the graphical workspace.
2. Click the GENEVE application and select **Details**. The Application quick view appears.



3. Enter an alias for the GENEVE application.
4. Click **Save**.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on Traffic health monitoring and how to create threshold template, refer to *Traffic Health Monitoring* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to *Monitor Cloud Health* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the statistics of the application in the **STATISTICS** tab.

Header Stripping

Header Stripping application efficiently examines the packets for specified headers like GTP, ISL, ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-tag, VXLAN, FM6000Ts, GENEVE, and generic and remove them before sending the packet to the appropriate security and analysis tools. Each packet is examined for the packet forwarding addition, and it also ensures that the headers are removed before sending the packet to the tools. This application is useful when working with tools that either cannot recognize these headers or must engage in additional processing to adjust for them.

Furthermore, the presence of protocols like GTP, ISL, ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-tag, VXLAN, GENEVE, and FM6000Ts in the packet can restrict or limit the ability to apply filtering and flow-based load balancing to the traffic as it is forwarded to specific tools. To address each of these challenges, Header Stripping of these protocols is required.

List of Protocols that are supported for stripping:

- GTP
- ISL
- ESPRAN
- MPLS
- MPLS+VLAN
- VLAN
- VN-tag
- VXLAN
- FM6000Ts
- GENEVE
- Generic



Note:

Header Stripping for the GENEVE protocol is currently supported for the following platforms:

- VMware vCenter
- VMware NSX-T
- Nutanix

Limitation

The GENEVE protocols currently supported for header stripping are the Transparent Ethernet Bridge (0x6558), IPv4 (0x0800), and IPv6 (0x86DD).

Configure Header Stripping Application

To configure the Header Stripping application in GigaVUE-FM, follow the steps given below:

1. Drag and drop **Header Stripping** from **APPLICATIONS** to the graphical workspace.
2. Click the Header Stripping application and select **Details**. The Application quick view appears.

3. In the application quick view, enter or select the required information as described in the following table:

Field	Description
Alias	Enter the alias name for the application.
Protocol	Select the type of protocol.
VLAN: Use this option to strip the VLAN header from the packets. You can either strip the outer VLAN header or the entire VLAN header. When choosing VLAN as your protocol for stripping, enter the following details:	
VLAN Header	The VLAN Header that should be stripped. The supported minimum value is 0, and the maximum value is 16777215. The default value is 0.
VXLAN: Use this option to strip VXLAN (Virtual eXtensible Local Area Network) headers. You can strip either matching VXLAN headers or all VXLAN headers. When choosing VXLAN as your protocol for stripping, enter the following details:	
VXLAN ID	The VXLAN ID that should be stripped. The default value is outer.
FM6000Ts: Use this option to strip FM6000Ts time stamp headers. Packets entering the application from other devices may contain FM6000 timestamps. FM6000 is an Intel chip used for time stamping. FM6000 has a hardware timestamp in the packet. When choosing FM6000Ts as your protocol for stripping, enter the following details:	
Time Stamp Format	The format of the time stamp you wish to strip. The only supported format for the time stamp is None .
ESPRAN: Use this option to strip ERSPAN Type II and Type III headers. When choosing ESPRAN as your protocol for stripping, enter the following details:	
ESPRAN FlowID	Specify an ERSPAN flow ID between 0 to 1023. A flow ID of zero is a wildcard value that matches all flow IDs.
GENEVE: Use this option to strip the GENEVE header from the packets. The default L4 Destination Port value is 6081. To change the default L4 destination port value, follow the steps given below:	
1. Expand the Custom Port section. 2. You can either select an existing template with the port details or provide them directly in the pop-up window. - o Using Template - If you wish to use an existing template, select the template from the Template drop-down menu. To create a new template, refer to Create Custom Port Template.	
Note: You cannot modify an existing template by editing the values directly in the Monitoring Session. However, you can modify the values in the Monitoring Session and save it as a new template.	
o Without Template - You can directly enter your port details in the L4 Destination Port field. Click the Save as New Template button to save this existing configuration as a template. You can view the newly created template in the Custom Port Template page.	

Generic: Using this option to strip any header without worrying about the header level. When choosing generic as your protocol for stripping, enter the following details:	
Ah1	The anchor header (AH1), after which the header to be stripped occurs.
Offset	Based on the selected Offset, enter the following details: • Offset Range: If you wish to use offset range as your offset, then enter the following details: a. Offset Range Value: Specify the offset of the header occurrence from the above anchor header. The minimum supported value is one, and the maximum supported value is 1500. b. Header Count: Specify the number of headers the application should remove from the offset. The minimum supported value is one, and the maximum is 32. c. Custom Len: The length (in bytes) of the header that should be stripped. d. Ah2: The next possible standard header that occurs immediately after the header. • Start / End: If you wish to use start or end as your offset, then enter the following details: a. Header Count: Specify the number of headers that the application should remove from the offset.. The minimum supported value is one, and the maximum is 32. b. Custom Len: The length (in bytes) of the header that should be stripped. The minimum supported value is one, and the maximum supported value is 1500. c. Ah2: The next possible standard header that occurs immediately after the header.

4. Click **Save**.

Create Custom Port Template

GigaVUE-FM enables you to specify custom L4 destination port values when using the Header Stripping application. You can create a template and utilize it to define custom port values when configuring the Header Stripping application.

To create a custom port template for the Header Stripping application:

1. Go to **Inventory > Resources > Custom Port Template**. The **Custom Port Template** page appears.
2. Click **New** to create a new template. The **New Custom Port Template** dialog box appears.
3. Enter a name for the template.
4. Enter valid port numbers between the range 1025 to 65535 in the **L4 Destination Port** field. Enter only comma separated integer values. You can provide a maximum of 10 ports. For example: 1034,1098,1039
5. Click **Save**.

You can view the created template on the **Custom Port Template** page. If a template is associated with a Monitoring Session, you can view the name of the Monitoring Session to which the template is associated in the Affected Entity column.

You can use the **Actions** buttons to perform the following:

- **Edit:** You can use this option to edit a template. If a template is associated with a Monitoring Session, editing it will be reflected in all the Monitoring Sessions associated with that template.
- **Delete:** You can use this option to delete a template. You cannot delete a template associated to a Monitoring Session.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on traffic health monitoring and how to create a threshold template, refer to *Traffic Health Monitoring* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to *Monitor Cloud Health* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the application statistics in the **STATISTICS** tab.

Load Balancing

Load balancing application performs stateless distribution or Enhanced Load Balancing of the packets between different endpoints. Stateless load balancing distributes the processed traffic to multiple tool ports or tunnel endpoints based on hash values generated from predefined protocol fields in the packet.

When Enhanced Load Balancing is enabled and an endpoint fails, the traffic is redistributed for the failed endpoint. When the failed endpoint recovers, the redistributed traffic is restored to the recovered endpoint. The traffic across other endpoints remain undisturbed during this process.

To add a load balancing application:

1. Drag and drop **Load Balancing** from **APPLICATIONS** to the graphical workspace.
2. Click the load balancing application and select **Details**. The Application quick view appears.



X Application Save

Application

Load Balancing

Alias

lb

Hash Fields

ipOnly

Field Location

outer

Load balancing groups

+

-

Application Endpoint ID 2

Weight 1-100

3. In the Application quick view, enter the information as follows:

Metric	Description
Alias	Enter a name for the load balancing application
Stateless	Select this option to enable Stateless Load Balancing
Enhanced Load Balancing	Select this option to enable Enhanced Load Balancing and select the ELB profile.
Hash Field	• ipOnly : The source IP and destination IP addresses. • ipAndPort: The source IP and destination IP addresses, and Layer 4 source port and destination port numbers. • fiveTuple: The source IP and destination IP addresses, source port and destination port numbers, and protocol field in the IP header. • gtpuTeid The GTP-u tunnel identifier (ID). NOTE: There is no inner or outer field location for GTPU-TEID. • greFlowid: The flow identifier location. NOTE: The default field location value is outer. When configuring 5G-Cloud Ericsson SCP Transparent, greFlowid hash field should be selected. Refer to 5G-Cloud Ericsson SCP Support.
Field Location	• Outer: The first occurrence of header or field. For example, IP Only outer is the first IP header in the packet, which could be IPv4 or IPv6. • Inner: The second occurrence of header or field. The supported IP encapsulation types are: IP-in-IP, VXLAN, GTP, GRE, and ERSPAN.
Load balancing groups	Add or remove an application with the Endpoint ID and Weight value (1-100). A load balancing group can have minimum of two endpoints. Endpoint with higher weight receives more traffic.

4. Click **Save**.

NOTE: When you configure the Load Balancing Application in enhanced mode, you can associate it with only a single enhanced load balancing profile. However, you have the flexibility of changing the association to different profile as needed.

Enhanced Load Balancing

Enhanced Load Balancing redistributes the traffic from a failed endpoint to another endpoint which is a part of the ELB group. When the failed endpoint recovers, the redistributed traffic is restored to the recovered endpoint. The traffic across other endpoints remain undisturbed during this process. Hashing is performed on the packets based on outer and/or inner headers.

To enable to enhanced load balancing application:

1. Create an Enhanced Load Balancing Profile.
 - a. In the Monitoring domain page, click Settings, then click **Enhanced Load Balancing**
 - b. Click **New** and then create a profile.

Fields	Values
Hash fields	The various hash options are - ip - ip-src - ip-dst - l4-port - l4-portsrc - l4-portdst - gtpoteid
Position	Select either inner/outer location of the hash field to be matched with the incoming packet..
Hash-mask	When the hash-fields ip-src/ip-dst are defined along with the IP, then ip-src/ip-dst hash-mask will overwrite the IP hash mask. For load balancing, you must apply this mask before hashing the IP. The default options and their default values are: - IPv4 : 255.255.255.255 - IPv6 :FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF

2. Go to monitoring session page and
 - a. Drag and drop **Load Balancing** from **APPLICATIONS** to the graphical workspace.
 - b. Click the load balancing application and select **Details**. The Application quick view appears.

× Load Balancing: lb
 Save

DETAILS
 THRESHOLDS
 HEALTH STATUS
 STATISTICS

Application Load Balancing

Alias* lb

☐ Stateless
 ☒ Enhanced Load balancing

ELB profiles* Select... ▼

Load balancing groups

Application Endpoint ID	2	Weight	1	⊕ ⊖
Application Endpoint ID	3	Weight	1	⊕ ⊖

⏪

c. In the Application quick view, enter the information as follows:

Metric	Description
Alias	Enter a name for the load balancing application
Enhanced Load Balancing	Select this option to enable Enhanced Load Balancing and select the ELB profile.
Load balancing groups	Add or remove an application with the Endpoint ID and Weight value (1-100). A load balancing group can have minimum of two endpoints.

3. Click **Save**.

NOTE: When you configure the Load Balancing Application in enhanced mode, you can associate it with only a single enhanced load balancing profile. However, you have the flexibility of changing the association to different profile as needed.

What's Next

You can configure the traffic health monitoring for this application in the **THRESHOLDS** tab. You can select an existing template from the Threshold Templates drop-down menu or provide the threshold values. For more details on Traffic health monitoring and how to create threshold template, refer to *Traffic Health Monitoring* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the configuration health status and the traffic health status of the application in the **HEALTH STATUS** tab. For more details on configuration health and traffic health, refer to *Monitor Cloud Health* section in the respective GigaVUE Cloud Suite deployment guides.

You can view the statistics of the application in the **STATISTICS** tab.

5G Service Based Interface Application

5G-Core is a service-based architecture in which many control plane network functions are available and communication across these network functions happens through the HTTP2 protocol. These HTTP2 transactions are mirrored by specific network functions and are in JSON encoded format.

The 5G-Service Based Interface (SBI) Application synthesizes the HTTP2 transactions with proper L2, L3, and L4 headers from the JSON encoded data that it receives from UDP-GRE. Once the headers are synthesized and a complete HTTP2 transaction is formed, the packets are sent to the egress TEP and then sent to virtual probes.

Ericsson vTAP

Ericsson 5G Core network functions (**PCC** (Packet Core Controller), **SC** (Signaling Controller), **CCxx** (Cloud Core), etc.) mirror SBI and non-SBI traffic using **Ericsson vTAP**. The vTAP captures this traffic, encapsulates it using GRE-in-UDP with destination UDP port 4754 and GRE Protocol ID 0x8037, and forwards it to Gigamon V Series Nodes.

The Ericsson vTAP payload consists of PCAPng Enhanced Packet Blocks, which the V Series PCAPng and 5G SBI applications process to reconstruct complete HTTP/2 transactions.

These vTAP packets are not raw packets that monitoring tools can understand. Although Ericsson includes rich session information, it is encoded in JSON and cannot be forwarded directly to tools. Therefore, V Series must synthesize tool facing packets by adding the necessary L2/L3/L4 headers and reconstructing valid TCP flows before forwarding it to the external tools over VXLAN or L2GRE egress tunnels.

How Packet Ordering Works

Each Ericsson vTAP packet carries a timestamp in the PCAPng Enhanced Packet Block that records when the packet was observed on the network (in EPOCH format). The 5G-EVP application uses this timestamp to reorder packets before forwarding them.

The application sorts packets using a bucket-based mechanism:

1. **Receive** - The application receives packets from Ericsson vTAPs over UDP-GRE or DTLS tunnels. Multiple receive threads handle incoming traffic in parallel.
2. **Classify into flows** - Each packet is assigned to an egress flow based on its flow identifiers. The application supports up to 512 concurrent egress flows (configurable).
3. **Place into time buckets** - The application places each packet into a time bucket based on its observed timestamp. Each bucket spans a configurable interval (1–5 ms). The application maintains a rolling set of buckets that together define the sorting window. For example, 100 buckets × 1 ms = a 100 ms sorting window.
4. **Sort and transmit** - When a bucket expires (falls outside the sorting window), the application sorts all packets in that bucket by their observed timestamp and transmits them in order to the downstream application.
5. **Handle late packets** - Packets that arrive with timestamps outside the current sorting window are handled based on the configured policy:
6. **Forward (recommended)** - Sends the packet to the downstream application without reordering.
7. **Discard** - Drops the packet at the 5G-EVP application.

NOTE: The maximum sorting window is 100 ms (numBuckets × bucketInterval must not exceed 100). The application validates this constraint and returns an error if the values exceed the limit.

DTLS Support

Ericsson can send vTAP traffic over DTLS (Datagram Transport Layer Security) for encrypted delivery. When you enable DTLS on the 5G-EVP application:

- Ericsson sends encrypted traffic on UDP port 4755 (instead of 4754 for non-secure traffic).
- The 5G-EVP application decrypts the traffic using the configured private key and certificate.
- The application forwards cleartext packets to the downstream 5G-SBI application over UDP-GRE on port 4754. The downstream applications see no difference.

Upload the DTLS private key and certificates using the same process as other SSL key pairs. In GigaVUE-FM, go to Inventory > Security > SSL Key.

Supported Modes

The 5G-EVP application supports the following combinations of packet ordering and DTLS:

Packet Ordering	DTLS	Supported
Disabled	Disabled	No
Disabled	Enabled	Yes
Enabled	Disabled	Yes
Enabled	Enabled	Yes

NOTE: You cannot disable both packet ordering and DTLS simultaneously. At least one must be enabled for the 5G-EVP application to function.

Ericsson vTAP Domain Tagging

Ericsson vTAP components mirror control plane and selected user plane traffic, for example, from Ericsson PCC nodes, and send it to the GigaVUE V Series using UDP-GRE with PCAPng/JSON metadata. Previously, the 5G-SBI application treated all mirrored flows uniformly, requiring monitoring tools to independently infer the Ericsson domain (for example, pc-mm, pc-sm, sc-scp) and the associated 3GPP interface.

With Domain Tagging, the PCAPng application on V Series reads the Ericsson vTAP metadata, identifies the Ericsson domain for each flow, and assigns a VLAN ID based on a configurable domain to VLAN map. When you enable domain classification in the PCAPNG app and VLAN tagging on the VXLAN or L2GRE egress tunnel, V Series tags each packet with the appropriate VLAN so probes can separate domains simply by filtering on VLAN.


You can also reserve VLANs for:

- Traffic mapped to the **others** domain (for example, unsupported or unexpected Ericsson domains)

- Traffic mapped to the **unclassified** domain (for example, traffic for which domain information is not yet available)

Encoded messages from the Ericsson vTAP contain Kubernetes Pod ephemeral IP addresses for NF instances, which are internal to the Ericsson deployment and not accessible outside the Kubernetes cluster. Because these IPs cannot be used externally, the 5G SBI application always maps the NF identifiers (NF instance ID and FQDN) in the encoded message to the configured ingress IP address using the string-to-IP mapping table.

The NF instance ID and FQDN ID must be provided in the form of CSV file. You can upload the CSV file through GigaVUE-FM.




How SBI Application works

In the GigaVUE V Series Node, the SBI application receives the HTTP2 transaction messages as JSON encoded data from the PCAPng application, which in turn receives the data from the UDP-GRE TEP.

The SBI application:

1. Parses the JSON data to extract source and destination information, ports, and message type (request/response).
2. Synthesizes a complete HTTP/2 transaction with L2, L3, and L4 headers, along with the HTTP/2 headers and body from the original transactions.
3. Sends the synthesized packets to the egress TEP, which forwards them to the appropriate virtual probes.

The following image shows the block diagram of the data flow in the V Series Node containing the SBI application.



Supported Platforms:

The application is supported on the following platforms:

- VMware ESXi
- OpenStack
- Third Party Orchestration

Rules and Notes

- The maximum number of HTTP2 headers (in the synthesized HTTP2 transactions) that is supported is 64.
- The PCAPng application that is linked to 5G-SBI application (on the right side) should only be linked to UDP-GRE TEP with key value 1 on the left side. If it is linked to other UDP-GRE TEPs (key values other than 1), then the behavior cannot be defined and leads to unexpected results.
- The maximum number of NF and FQDN entries supported is 4K.
- In 5G-SBI application, the V Series node can log the following details to CSV files:

- **Transaction details** - Represents the transaction or flow of request and response packets into the application. The details of the flow or transaction are recorded in the CSV file for 5 minutes or 60 minutes based on the configuration.
- **Flow statistics details** - Represents the packet and flow statistics in 60 seconds time interval.


These files help you to understand the records and traffic efficiently. The files are named as per the date and time in which the files were created. When the number and size of files grow, the application automatically detects the old files and deletes them.

- Ericsson vTAP domain tagging can be configured only in primary mode when configuring Pcapng application. Refer to [PCAPNG Processing and Flow Tracking](#).

End-to End Solution Overview and Configuration:

This section explains how Ericsson vTAP traffic flows through the Gigamon V Series, from ingress UDP-GRE packets to domain tagged traffic on the egress tunnels. In GigaVUE-FM, you must do the following steps in the Monitoring Session of a Monitoring Domain in the V Series:

- [Ingress: Ericsson vTAP UDP-GRE Feed](#)
- [PCAPNG Processing and Flow Tracking](#)
- [Configure 5G-SBI Application](#)
- [Egress Tunnels and VLAN Tag Insertion](#)



Ingress: Ericsson vTAP UDP-GRE Feed

At ingress, Ericsson vTAP mirrors traffic from Ericsson 5G Core network functions and sends it to the Gigamon V Series as UDP-GRE packets. A specific GRE protocol ID (0x8037) marks the traffic as Ericsson vTAP, and the GRE key indicates what kind of payload is inside (JSON metadata or PCAPNG) and which protocol it carries. Based on this GRE key, the V Series recognizes the packets as Ericsson vTAP traffic and selects the appropriate processing path.

For more details on how to configure UDP-GRE Ingress TEP, refer to the below sections:

- [Create Ingress and Egress Tunnel \(VMware vCenter\)](#)
- [Create Ingress and Egress Tunnels \(OpenStack\)](#)
- [Create Ingress and Egress Tunnel \(Third Party Orchestration\)](#)

PCAPNG Processing and Flow Tracking

In this Ericsson vTAP design, PCAPng classifies flows by Ericsson domain and inserts the corresponding VLAN tag on the egress traffic.

The PCAPNG application on the V Series serves as the configuration point for Ericsson domain classification and VLAN mapping (app mode, domainClassification, flowTimeout, and domainTable). The 5G SBI application performs the actual parsing of Ericsson vTAP metadata, flow and domain determination, and SBI failure detection. The dataplane and egress tunnel configuration apply the VLAN tags and forward the traffic.

To configure a PCAPng application:

1. In the canvas, drag and drop the PCAPng application and select **Details**. The PCAPng application quick view appears.
2. On the application quick view, enter or select the required information as described in the following table:

Field	Description
Application	The name Pcapng appears by default.
Alias	Enter an Alias as required.
App Mode	Set App Mode to Primary to view and configure domain related options.  Notes: • Primary mode is required to configure domain classification. When you set App Mode to Secondary, Domain Classification, Domain VLAN Mapping, and Flow Timeout fields are not visible. • Only one Primary PCAPng instance can be configured. If a Primary PCAPng instance already exists, any new PCAPng instance is automatically set to Secondary, and the mode is disabled for editing.

Field	Description
Domain Classification	Enable the Domain Classification option to classify Ericsson vTAP traffic by domain.
Domain VLAN Mapping	- From the Domain VLAN Mapping list, select the VLAN mapping profile that maps Ericsson domains to VLAN IDs for per-domain tagging on the egress tunnel. - Click Add New to create a new 5G-Apps configuration. Refer to Adding CSV file for IP Mapping section in 5G-Service Based Interface Application for details. When Domain Classification is enabled, the selected Domain VLAN Mapping is used to tag packets per Ericsson domain with the configured VLAN ID (and, if applicable, a separate VLAN ID for failed SBI transactions) before sending traffic to the probes.
Flow Timeout	Set the Flow Timeout value in seconds to define how long an inactive Ericsson vTAP flow is kept before it is closed and its resources are released. Valid range: 360-1860 Default value: 660

NOTE: In the PCAPng configuration, the Domain VLAN Mapping and Flow Timeout fields are displayed only when App Mode is set to Primary and the Domain Classification check box is enabled.

Configure 5G-SBI Application

In V Series, 5G-SBI application receives all the mirrored traffic from 5G-Ericsson.

In GigaVUE-FM, the **Vendor Integrations** field determines how incoming data is processed. When Ericsson vTap is selected as the vendor integration, the system processes packets accordingly.

You can add a 5G-SBI application to:

- **New Monitoring Session** - Add the 5G-SBI application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
- **Existing session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

To configure a 5G-SBI application:

1. In the canvas, drag and drop the 5G-SBI application and select **Details**. The 5G-SBI application quick view appears.

2. On the application quick view, enter or select the required information as described in the following table:

Field	Description
Application	The name 5g-sbi appears by default.
Alias	The name sbi5gAppTemplate appears by default.
Vendor Integrations	Select the option ericssonVTap from the drop-down list.
SBI Operational Mode	Select the SBI Monitoring Mode for 5G-SBI Network traffic. L7json is selected by default.
Destination IP	The Destination IP toggle determines how synthesized SCP mediated SBI transactions are tagged. You can choose from the following options: • SCP - Sets the synthesized transaction's destination IP to the SCP IP address. This option is selected by default. • Destination Network Function - Sets the synthesized transaction's destination IP to the IP address of the destination network function.
Ericsson vTAP Version	Select 1 or 2 from the drop-down list box.
MTU Packet Size	Select how packets should be sent to the tool: • Full Packet - Sends the entire synthesized packet as a single packet without segmentation. The packet is fragmented at the egress TEP. • MTU Size - Segments packets based on the configured tool port MTU value before sending them to the tool. You can specify a value between 1200 and 9200 bytes.
Incomplete SBI Transaction Domain Tag	Enable this option to tag unclassified SBI transactions with the dedicated 'Failed SBI Transaction' VLAN ID defined in the Domain table. The probe uses this tag to identify and troubleshoot these transactions more effectively. NOTE: This option appears only when at least one Primary PCAPng file is configured on a V Series node.
Indexed Headers	Enable the check box to index the HTTP2 headers in the 5G-SBI application.
Compressed Headers	Enable the check box to compress the HTTP2 headers in the 5G-SBI application.
TCP Flows	Specify the maximum number of TCP flows based on the VM form factor: • Small - Supports 128 to 2048 TCP flows • Medium - Supports 128 to 8192 TCP flows • Large - Supports 128 to 16384 TCP flows

TCP Flow Timeout	Specify the flow range value (in seconds) that determines how long a TCP flow remains valid in the application. • Minimum value: 0 • Maximum Value: 7200 • Default value: 1800
Number of Transaction Flows	Specify the maximum number of streams or flows based on the VM form factor: • Small - Supports 1 to 20480 Transaction flows • Medium - Supports 1 to 51200 Transaction flows • Large - Supports 1 to 122880 Transaction flows
Request Timeout	Specify the time (in seconds) for which the request packet waits for the response packet in a stream. • Minimum value: 1 • Maximum Value: 300 • Default value: 10
Response Timeout	Specify the time (in seconds) for which the response packet waits for the request packet in a stream. • Minimum value: 1 • Maximum Value: 300 • Default value: 2
IP Mapping	
Network Function Instance Mapping	Select the required CSV file from the drop-down list with required Network Function Instance ID (NFID) instance mapping. • Select an alias: Choose the alias name associated with the uploaded NFID table (CSV file). • Add a new alias: If no alias exists, click Add New to create one. Refer to Adding CSV file for IP Mapping for details.
FQDN Mapping	Select the required CSV file from the drop-down list with required FQDN mapping. • Select an alias: Choose the alias name associated with the uploaded FQDN table (CSV file). • Add a new alias: If no alias exists, click Add to create one. Refer to Adding CSV file for IP Mapping section for details.
Logging	
Transaction Log	Enable or disable 5G-SBI transaction logging. When enabled, transaction logs are written and rotated based on the configured Transaction Log File Interval.
Transaction Log Interval(Minutes)	Set the transaction log file interval in minutes. This defines how often a new Transaction Log CSV file is

	created. - Valid options: 5 or 60 minutes - Default: 60 This setting is used only when Transaction Log is enabled.
Stats Log	Enable or disable statistics logging for the 5G-SBI application. When enabled, statistics are written to the configured log folder and managed based on the Log Folder Size setting.
Log Folder Size(MB)	Specify the memory size (in MB) allocated to the log folder. This sets the maximum allowable size for the 5G-SBI log folder and applies to both transaction and statistics logs stored in the configured Log Folder Location.
Log Folder Location	Enter the file system path where 5G-SBI log files should be stored. This is the location for both transaction and statistics logs. Example: /home/user/log

Egress Tunnels and VLAN Tag Insertion

When the V Series classifies Ericsson vTAP traffic by domain in the PCAPNG application, it forwards packets to the probes through egress tunnels (VXLAN or L2GRE).

When creating an Egress tunnel in the Monitoring Session Canvas, enable the **Domain Tagging** option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping.

NOTE: This setting is available only when **Domain Classification** is enabled in the associated PCAPng application. Refer to [PCAPNG Processing and Flow Tracking](#) for details.

When **Domain Classification** is enabled in the PCAPNG app and **Domain Tagging** is enabled on the egress tunnel, the V Series does the following:

- Looks up the Ericsson domain for each packet.
- Maps it to a VLAN ID from the Domain Table (including any failed SBI or special case VLANs)
- Inserts that VLAN tag on the egress packet

This ensures each probe interface only receives the Ericsson domains it is configured to monitor.

For more details on how to configure VXLAN/L2GRE Egress , refer to the below sections:

- [Create Ingress and Egress Tunnel \(VMware vCenter\)](#)
- [Create Ingress and Egress Tunnels \(OpenStack\)](#)

- [Create Ingress and Egress Tunnel \(Third Party Orchestration\)](#)

Deploy the Monitoring Session

After adding all the required elements to the canvas, to deploy the session do the following:

1. From the **Actions** menu, select **Deploy**.

After successful deployment on all the V Series Nodes, the status appears as **Success** on the **Monitoring Sessions** page.

2. **View the Deployment Report**

- a. You can view the Monitoring Session Deployment Report in the **SOURCES** and **V SERIES NODES** tab.
- b. When you select the **Status** link, the Deployment Report is displayed.
- c. When the deployment is incorrect, the Status column displays one of the following errors:
 - **Success:** Not deployed on one or more instances due to V Series Node failure.
 - **Failure:** Not deployed on all V Series Nodes or Instances.



Notes:

- After you deploy a Monitoring Session, you cannot edit its associated 5G SBI or PCAPNG application configurations. To modify them, undeploy the Monitoring Session, update the configuration, and redeploy.
- If a PCAPNG application instance is set as the Primary for a Monitoring Session, you cannot delete it while it holds that role. To delete or replace it, undeploy the Monitoring Session or reassign the Primary.

Adding CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field.
3. From the **Type** drop-down list, select one of the following. You can also choose to download the template using the **Download Template** option and add the required entries.

- **5G-SBI Network Function Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping in 5G-Ericsson.

Header details with an example:

- **instanceID** - e4d909c2-5d2f-4f7f-9bcb-12a3b6c4d789
- **ipAddress** - 1.1.1.1

- **5G-SBI FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

Header details with an example:

- **fqdnID** - amf.local
- **ipAddress** - 1.1.1.1
- **nfType** - SCP

- **Domain VLAN Mapping** - Add the CSV file containing a valid domain and a valid VLAN ID for IP mapping.

Refer the table below for example header details of the CSV file.

domainName	vlanID	failedSBITransactionVlan ID
pc-mm	100	200
pc-sm	300	400
ccsm-5gc	500	
ccsm-cnhs	600	
ccrc-nrf	700	
nrfagent-nrf	800	
ccrc-nssf	900	
ccpc-pcf	1000	
ccpc-provisioning	1100	
ccdm-udr-5g	1200	
ccdm-provisioning	1300	
ccdm-udr-legacy	1400	
cces-nef	1500	
cces-ees	1600	
cces-apigm	1700	
sc-bsf	1800	
sc-scp	1900	

domainName	vlanID	failedSBITransactionVlan ID
sc-sepp	2000	
others	2100	
unclassified	2200	

- Click **Choose File** in **File Name** field to upload the CSV file into GigaVUE-FM.
- Click **Validate** to validate the CSV file.
- Click **Save** to add the CSV file.

**Notes:**

- Failed SBI Transactions are applicable only for **pc-mm** and **pc-sm** domains.
- Each domain in the Domain table must use a unique VLAN ID and Failed SBI Transaction ID to allow the probes to distinguish traffic for each Ericsson domain.

Recommended VM Specifications

The configuration of the 5G-SBI application should be performed on large VMs.

- Configuration specifications for large VM:
 - vCPUs: 8
 - Memory: 16 GB
 - Disk Space: 80 GB
- Supported Ports: 1 Management interface, 1 ingress, and 1 egress interfaces.
- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to [Configure GigaVUE Fabric Components using VMware ESXi](#) topic for more details.

FHA Dashboards for 5G Applications

After configuring the 5G-SBI application, you can monitor the statistics for Ericsson vTAP by the reports displayed in the Dashboard. To access the details, refer to [FHA Dashboards for 5G Applications](#).

Interface Mapping

You can remap interfaces for individual GigaVUE V Series Nodes within a Monitoring Session.

Note: When using Raw and Tunnel In, Interface Mapping is mandatory before you deploy the Monitoring Session.

To perform interface mapping,

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.

The **Monitoring Sessions** landing page appears.

2. Navigate to the **V SERIES NODES** tab and select **Interface Mapping**.

The **Deploy Monitoring Session** dialog box appears.

3. Select the GigaVUE V Series Nodes to which you wish to map the interface.
4. From the drop-down menu of the GigaVUE V Series Nodes, select the interfaces for the following deployed in the Monitoring Session:
 - REPs (Raw Endpoints)
 - TEPs (Tunnel Endpoints)
5. Select **Deploy**.

NOTE: The updated mappings take effect when deployed.

Monitor and Troubleshoot

Monitor Cloud Health

GigaVUE-FM allows you to monitor the traffic and configuration health status of the monitoring session and its individual components. This section provides detailed information on how to view the traffic and configuration health status of the monitoring session and its individual components. Refer to the following topics for more detailed information on configuration health, traffic health and how to view the health status:

- [Configuration Health Monitoring](#)
- [Traffic Health Monitoring](#)
- [View Health Status](#)
- [View Monitoring Session Statistics \(Azure\)](#)
- [Visualize the Network Topology \(Azure\)](#)

Configuration Health Monitoring

The configuration health status provides detailed information about the configuration and deployment status of the deployed monitoring session.

It supports specific fabric components and features on the respective cloud platforms.

Configuration Health Monitoring	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware	GigaVUE Cloud Suite for Nutanix
GigaVUE V Series Nodes	✓	✓	✓	✓	✓
UCT-V	✓	✓	✓	✗	✗
VPC Mirroring	✓	✗	✗	✗	✗
OVS Mirroring and VLAN Trunk Port	✗	✗	✓	✗	✗

Refer to the [View Health Status](#) section to view the configuration health status.

Traffic Health Monitoring

GigaVUE-FM monitors the traffic health of the entire Monitoring Session and each individual GigaVUE V Series Node in that session. It checks for issues like packet drops or traffic overflows.

When it detects a problem, GigaVUE-FM updates the health status of the related Monitoring Session. It monitors traffic health in near real-time.

The GigaVUE V Series Node tracks traffic levels. If traffic goes above or below the configured threshold, it alerts GigaVUE-FM. Then, GigaVUE-FM then uses this data to calculate traffic health.

If you deploy GigaVUE-FM and GigaVUE V Series Nodes in different cloud platforms, you must add the GigaVUE-FM public IP address as the Target Address in the Data Notification Interface on the Event Notifications page.

For details, refer to the [Configuration Settings](#) section in the *GigaVUE Administration Guide*.

This feature supports GigaVUE V Series Nodes on the respective cloud platforms:

For V Series Nodes:

- AWS
- Azure
- OpenStack
- VMware
- Third Party Orchestration

For instructions on creating and applying threshold templates across a Monitoring Session or an application, and viewing the traffic health status, refer to the following topics:

- [Supported Resources and Metrics](#)
- [Create Threshold Templates](#)
- [Apply Threshold Template](#)
- [Clear Thresholds](#)

Consideration to configure a threshold template

- By default, Threshold Template is not configured to any Monitoring Session. If you wish to monitor the traffic health status, then create and apply threshold template to the Monitoring Session.
- Editing or redeploying the Monitoring Session reapplies all the threshold policies associated with that Monitoring Session.
- Deleting the Monitoring Session clears all the threshold policies associated with that Monitoring Session.
- Threshold configuration is applied before deploying a Monitoring Session and remains even if the session is undeployed.
- After applying threshold template to a particular application, you need not deploy the Monitoring Session again.

Supported Resources and Metrics

The following table lists the resources and the respective metrics supported for traffic health monitoring. Refer to [Create Threshold Templates](#) and [Apply Threshold Template](#) sections for details on Threshold types and Threshold events.

Resource	Metrics	Threshold types	Trigger Condition
Tunnel End Point	1. Tx Packets	1. Difference	1. Over

	2. Rx Packets 3. Tx Bytes 4. Rx Bytes 5. Tx Dropped 6. Rx Dropped 7. Tx Errors 8. Rx Errors	2. Derivative	2. Under
RawEnd Point	1. Tx Packets 2. Rx Packets 3. Tx Bytes 4. Rx Bytes 5. Tx Dropped 6. Rx Dropped 7. Tx Errors 8. Rx Errors	1. Difference 2. Derivative	1. Over 2. Under
Map	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Slicing	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Masking	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Dedup	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
HeaderStripping	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
TunnelEncapsulation	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
LoadBalancing	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
SSLDecryption	1. Tx Packets 2. Rx Packets	1. Difference 2. Derivative	1. Over 2. Under

	3. Packets Dropped		
Application Metadata	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
AMX	1. Tx Packets 2. Rx Packets 3. Packets Dropped 4. Ingestor - Rx packets 5. Ingestor - Packets Dropped 6. Ingestor - Rx Octets 7. Ingestor - Octets Dropped 8. Ingestor - Records Dropped 9. Workload - Records Dropped 10. Workload - Req Auth Errors 11. Workload - Req Timedout Errors 12. Workload - Req Errors 13. Exporter - Avg File Size 14. Exporter - File Uploads 15. Exporter - File Uploads Errors 16. Enrichment - One Minute Percent	1. Difference 2. Derivative	1. Over 2. Under
Geneve	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under

5G-SBI	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
SBIPOE	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
PCAPNG	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under

Create Threshold Templates

To create threshold templates:

1. Go to **Inventory > Resources > Threshold Templates**.

The **Threshold Templates** page appears.

2. Select **Create** to open the New Threshold Template page.
3. Enter the appropriate information for the threshold template as described in the following table:

Field	Description
Threshold Template Name	The name of the threshold template.
Thresholds	
Traffic Element	Select the resource for which you wish to apply the threshold template. Ex: TEP, REP, Maps, Applications like Slicing, De-dup etc
Time Interval	Frequency at which the traffic flow needs to be monitored.
Metric	Metrics that need to be monitored. For example: Tx Packets, Rx Packets.
Type	Difference: The difference between the stats counter at the start and end time of an interval, for a given metric. Derivative: Average value of the statistics counter in a time interval, for a given metric.
Condition	Over: Checks if the statistics counter value is greater than the 'Set Trigger Value'. Under: Checks if the statistics counter value is lower than the 'Set Trigger Value'.
Set Trigger Value	Value at which a traffic health event is raised, if statistics counter goes below or above this value, based on the condition configured.
Clear Trigger Value	Value at which a traffic health event is cleared, if statistics counter goes below or above this value, based on the condition configured.

4. Select **Save**.
The newly created threshold template is saved, and it appears on the **Threshold** templates page.

Apply Threshold Template

You can apply your threshold template across the entire Monitoring Session and also to a particular application.

Apply Threshold Template to Monitoring Session

To apply the threshold template across a Monitoring Session, follow these steps:

1. On the left pane in GigaVUE-FM, select **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Session** page appears.
2. In the **TRAFFIC PROCESSING** tab, select **Options>Thresholds** menu.
3. From the **Select Template** drop-down list, select the template you wish to apply across the Monitoring Session.
4. Select **Apply**.

NOTE: You can apply the Threshold configuration to a Monitoring Session before it is deployed. Furthermore, undeploying the Monitoring Session does not remove the applied Thresholds.

Apply Threshold Template to Applications

Applying threshold template across Monitoring Session does not overwrite the threshold value applied specifically for an application. When a threshold value is applied to a particular application, it overwrites the existing threshold value for that particular application.

To apply the threshold template to a particular application in the Monitoring Session, follow these steps:

1. On the **Monitoring Session** page, select **TRAFFIC PROCESSING** tab. The Monitoring Session canvas page appears.
2. Select on the application for which you wish to apply or change a threshold template and select **Details**. The **Application** quick view opens.
3. Select the **Thresholds** tab.
4. Select the template you wish to apply from the Threshold Template drop-down menu or enter the threshold values manually.
5. Select **Save**.

Clear Thresholds

You can clear the thresholds across the entire Monitoring Session and also to a particular application.

Clear Thresholds for Applications

To clear the thresholds of a particular application in the Monitoring Session, follow these steps:

1. On the **Monitoring Session** page, select the **TRAFFIC PROCESSING** tab. The Monitoring Session canvas page appears.
2. Select the application for which you wish to clear the thresholds and click **Details**. The **Application** quick view opens.
3. Select the **Thresholds** tab.
4. Select **Clear All** and then select **Save**.

Clear Thresholds across the Monitoring Session

To clear the applied thresholds across a Monitoring Session, follow these steps:

1. On the left navigation pane in GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Sessions** landing page appears.
2. Select the Monitoring Session and navigate to **TRAFFIC PROCESSING > Options > Thresholds**,
3. Select **Clear Thresholds**.
4. On the **Clear Threshold** pop-up appears, select **Ok**.

NOTE: Clearing thresholds at Monitoring Session level does not clear the thresholds that were applied specifically at the application level. To clear thresholds for a particular application, refer to [Clear Thresholds for Applications](#)

View Health Status

You can view the health status of the Monitoring Session on the Monitoring Session details page. The health status of the Monitoring Session is healthy only if both the configuration health and traffic health are healthy.

View Health Status of an Application

To view the health status of an application across an entire Monitoring Session,

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
2. Select a Monitoring Session and navigate to the **TRAFFIC PROCESSING** tab.
3. Select the application for which you wish to see the health status and select **Details**. The quick view page appears.
4. Select the **HEALTH STATUS** tab.

This displays the application's **Configuration Health**, **Traffic Health**, and the **Operational Health**, along with the thresholds applied to each.

NOTE: The secure tunnel status is refreshed every 5 minutes, and the GigaVUE-FM does not display UCT-V secure tunnel status that is older than 7 minutes. If the secure tunnel in the UCT-V is removed, it takes up to 7 minutes to reset the status on the GigaVUE-FM.

View Operational Health Status of an Application

When you configure the Application Metadata Exporter to use **Kubernetes** as the workload platform, the V Series Node transmits failure and error events to GigaVUE-FM, which processes them and updates the node's health status on the Monitoring Session page. When interacting with Kubernetes workloads, the system may encounter errors while retrieving resources such as pods, services, nodes, or endpoints. Refer to [Errors](#) for additional error details.

Operational events for Exporter:

Refer below for message format and messages that indicate common issues that can occur during the operations:

Format: <Server Type>_<Message>

Server Types: CLOUD EXPORT, KAFKA

Message	Description
UPLOAD_MAX_TRIES_EXCEED	Upload retries exceeded the maximum limit Example: CLOUDEXPORTER_UPLOAD_MAX_TRIES_EXCEED
REACHABILITY_FROM_AMX_TO_TOOLS	AMX failed to reach the tool (Cloud Exporter server or Kafka server) Example: CLOUDEXPORTER_REACHABILITY_FROM_AMX_TO_TOOLS
NO_IP_ADDRESS	No IP address was configured on the interface Example: CLOUDEXPORTER_NO_IP_ADDRESS
EXPORTER_UPLOAD_ERROR	Upload to the exporter failed Example: CLOUDEXPORTER_EXPORTER_UPLOAD_ERROR
EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED	The exporter health-monitor probe or upload to the tool timed out. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED
EXPORTER_ERR_EVENT_CONTEXT_CANCELED	The exporter health-monitor probe or upload was canceled before completion, typically due to shutdown or reconfiguration. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_

Message	Description
	CANCELED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_CANCELED
EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL	The hostname of the configured Kafka broker or Cloud Exporter endpoint could not be resolved due to a DNS failure. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL KAFKA_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL
EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE	The exporter probe failed at the network layer, for example, due to no route to the host or an unreachable network. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE KAFKA_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE
EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR	The exporter probe failed at the transport layer, such as connection refusal, TCP reset, or TLS handshake failure. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR KAFKA_EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR

Operational events for Enrichment:

Refer below for message format and messages that indicate common issues that can occur during the operations:

Format: <Operation Type>_<Message>

Operation Types: GETSERVICES, GETPODS, GETNODES, GETENDPOINTS, WATCHALL

Message	Description
K8S_AUTHORIZATION_FAILURE	The request was denied due to insufficient permissions Example: GETPODS_K8S_AUTHORIZATION_FAILURE
K8S_AUTHENTICATION_FAILURE	Authentication failed. Verify your credentials Example: GETPODS_K8S_AUTHENTICATION_FAILURE
K8S_UNHANDLED_ERROR	An unspecified error occurred. Check the error description Example: GETPODS_K8S_UNHANDLED_ERROR

View Health Status for Individual GigaVUE V Series Nodes

You can also view the health status of an individual GigaVUE V Series Node. To view the configuration health status and traffic health status of the V Series Nodes:

1. On the Monitoring Session page, select the required Monitoring Session from the list view.
2. In the **Overview** tab, view the health status of the required GigaVUE V Series Node from the chart options.

View Application Health Status for Individual V Series Nodes

To view the application configuration and traffic health status of the GigaVUE V Series Nodes:

1. On the Monitoring Session page, select the required Monitoring Session from the list view.
2. On the Overview tab, select the GigaVUE V Series Node from the All V Series Nodes drop-down menu.

The list view displays the list of applications for the selected GigaVUE V Series Node and the health status of each application.

View Monitoring Session Statistics

The Monitoring Session **OVERVIEW** page lets you analyze the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis.

You can view the high level information of the selected Monitoring Session such as, connections, tunnel details, health status, deployment status, and information related to Application Intelligence statistics. You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab.

You can view the statistics by applying different filters as per the requirements of analyzing the data. GigaVUE-FM allows you to perform the following actions on the Monitoring Session Statistics page:

- You can view the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis.
- You can filter the traffic and view the statistics based on factors such as **Incoming, Outgoing, Ratio (Out/In), Incoming Packets, Outgoing Packets, Ratio (Out/In) Packets**. You can select the options from the drop-down list box in the **TOTAL TRAFFIC** section of the **OVERVIEW** page.
- You can also view the statistics of the Monitoring Session deployed in the individual V Series Nodes. To view the statistics of the individual GigaVUE V Series Node, select the name of the **V Series Node** for which you want to view the statistics from the GigaVUE V Series Node drop-down list on the bottom left corner of the **OVERVIEW** page.



Raw EndPoint (REP) is a part of the monitoring session but can also receive the bypassed traffic that is not filtered by the map, so it is recording more packets



than expected. For example, if the map has a rule as IPv4, but the REP can receive the bypassed (non-ipv4) traffic. The recorded number of packets from the V Series node can be more than expected.

Visualize the Network Topology (VMware NSX-T)

You can have multiple connections in GigaVUE-FM. Each connection can have multiple Monitoring Sessions configured within it. The Topology tab provides a visual representation of the monitored elements within a selected connection and Monitoring Session.


To view the topology in GigaVUE-FM:

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Sessions** landing page appears.
2. Create a Monitoring Session or select an existing Monitoring Session,
3. Open the **TOPOLOGY** tab.
4. From the **Connection** list on the Topology page, select a connection.

The topology view of the monitored subnets and instances in the selected session is displayed.

5. From **View**, select one of the following instance types:

- Fabric
- Monitored



6. (Optional) Hover over the subnet or VM group icons to view details such as the

subnet ID, subnet range, and the total number of subnets and instances.

7. Select the subnet or VM group icons to explore the subnets or instances within the group.

In the Topology page, you can also perform the following:

- Use the **Filter** button to filter the instances based on the VM name, VM IP, OS Type, 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 monitored instances.
- Apply Navigation controls, such as:
 - Use the arrows at the bottom-right corner to move the topology page up, down, left, or right.
 - Use + or - icons to zoom in and zoom out of the topology view.
 - Select the **Fit View** icon to fit the topology diagram according to the width of the page.

Administer GigaVUE Cloud Suite for VMware NSXT

You can perform the following administrative tasks in GigaVUE-FM for GigaVUE Cloud Suite for VMware NSX-T.

Create Service Chain and NSX-T Group

You need to create a VMware NSX-T group and a service chain to redirect network traffic to GigaVUE Cloud Suite.

- The VMware NSX-T group defines which virtual machines (VMs) are monitored.
- The service chain links the group to GigaVUE Cloud Suite using a service profile.

Create Service Chain

Using the instructions given below, you can create a service chain that defines traffic redirection for selected VMs. For details on creating security policies, refer to the *Service Composer* chapter of the *NSX Administration Guide* in the VMware documentation.

To create the service chain in VMware NSX-T,

1. In the VMware NSX Manager, go to **Security > Settings > Network Introspection**.
2. Go to the **SERVICE CHAINS** tab and select **ADD CHAIN**.
3. On the New Service Chain, perform the following:
 - a. In the **Name** and **Description** fields, enter a name and description for the service chain, respectively.
 - b. For **Service Segments**, select a service segment.
 - c. Select **Forward Path**.
 - d. On the **Set Forward Path** dialog box, select a Service Profile for Forward Path.
 - e. For **Reverse Path**, select or deselect the **Inverse Forward Path** to define traffic direction.
 - f. For **Failure Policy**, specify whether to allow or block the service chain.
4. Select **Save**.

A Service Chain is created.

Create Group

You need to create a group to forward NSX-T network traffic to the GigaVUE Cloud Suite.

To create the group, do the following in the NSX-T:

1. In the VMware NSX manager, select **Inventory > Groups**.
2. On the **Groups** page, select **ADD GROUP**.
3. On the **New Group** page, enter or select the values as follows.
 - a. Enter a name for the new group.
 - b. Select **Set Members**.
The **Select Members** dialog box appears.
4. Select **Save**.
A group is created and appears on the **Groups** page. Add the description and define group membership using **Membership Criteria, Members, IP/MAC Addresses**, or **AD Groups**.

Create and Publish a Policy

A Policy is a set of rules defined to filter the traffic. You need to create a policy to define how traffic flows from the selected VMs to the configured tunnel endpoint.

To create and publish a policy in NSX-T:

1. In the VMware NSX manager, select **Security > Service Chain Management > Network Introspection (E-W)**.
2. Select **ADD POLICY**.

3. On the New Policy, enter or select the values as follows:
 - a. Enter a name for the policy.
 - b. Select the **Sources** of the traffic.
 - c. Select the **Destinations** of the traffic.
 - d. Select the **Services** for the traffic.
 - e. For the **Applied To** field, select the appropriate groups.



Points to Note:

- When using the same NSX-T manager to create multiple Monitoring Domains, if you prefer to associate a single vCenter with each Monitoring Domain, ensure that you select only the members of the vCenter related to that specific Monitoring Domain.
- All the workload VMs from the vCenters that are selected in the **Applied To** group sends traffic to the GigaVUE V Series Node. Ensure to select only the workload VMs from the vCenters that are associated with the Monitoring Domain, else the GigaVUE V Series Node receives traffic from the vCenters that are not associated with the Monitoring Domain.

- f. On the **Action** field, specify whether to redirect the traffic or not.
4. Select **Publish**.

On publishing the rule or policy, you can view the traffic flow from the GigaVUE V Series Nodes to the tunnel endpoint.

Configure VMware Settings

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

To configure the VMware Settings,

1. Go to **Inventory > VIRTUAL > VMware NSX-T (V Series)**,
2. Select **Settings > Advanced Settings** to edit the VMware V Series NSX-T settings.

Advanced Settings

Maximum number of vCenter connections allowed	20
Refresh interval for VM target selection inventory (secs)	300
Refresh interval for fabric deployment inventory (secs)	86400
Traffic distribution tunnel range start	8000
Traffic distribution tunnel range end	8512
Traffic distribution tunnel MTU	9001
Maximum V Series node up wait time in minutes	5

Refer to the following table for details:

Settings	Description
Maximum number of vCenter connections allowed	Specifies the maximum number of vCenter connections you can establish in GigaVUE-FM
Refresh interval for VM target selection inventory (secs)	Specifies the frequency for updating the state of target VMs in VMware vCenter
Refresh interval for fabric deployment inventory (secs)	Specifies the frequency for updating the state of GigaVUE-FM fabrics deployed in VMware vCenter
Traffic distribution tunnel range start	Specifies the start range value of the tunnel ID.
Traffic distribution tunnel range end	Specifies the closing range value of the tunnel ID.
Traffic distribution tunnel MTU	Specifies the Tunnel MTU value.
Maximum V Series Node up wait time	Specifies the maximum amount of time taken for the GigaVUE V Series Node state to go to OK.

Configure Certificate Settings

To configure certificate settings:

1. Go to **Inventory > VIRTUAL**.
2. Select your cloud platform.
3. Select **Settings > Certificate Settings**.
The **Certificate Settings** page appears.

4. From the **Algorithm** drop-down list, select the algorithm that determines the cryptographic security of the certificate.

NOTE: If selecting RSA 8192, the certificate generation may take longer due to the increased key size.

5. In the **Validity** field, enter the total validity period of the certificate.
6. In the **Auto Renewal** field, enter the number of days before expiration of the auto-renewal process should begin.
7. Select **Save**.

High Availability for Cloud Infrastructure

This section provides details about the GigaVUE-FM High Availability (HA) feature in Cloud environment and describes how to configure, upgrade, and troubleshoot the feature.


Refer to the following topics for details:

- [About GigaVUE-FM High Availability](#)
- [Supportability Information](#)
- [Recommendations, Rules, Notes, and Limitations](#)
- [Licensing Information](#)
- [Hostname Setups](#)
- [Configure GigaVUE-FM HA](#)
- [Configure Secure Communication between Fabric Components in FMHA](#)
- [GigaVUE-FM HA Landing Page](#)
- [Remove Standby GigaVUE-FM Instance](#)
- [Disassemble GigaVUE-FM High Availability Group](#)
- [GigaVUE-FM High Availability States](#)
- [Fail over Mechanism](#)
- [GigaVUE-FM High Availability Scenarios](#)
- [Troubleshoot GigaVUE-FM High Availability Issues](#)

About GigaVUE-FM High Availability

The GigaVUE-FM High Availability (HA) feature supports a highly available fabric management environment with minimal interruption. The GigaVUE-FM HA architecture consists of a minimum of three GigaVUE-FM instances that run together as a highly available group. The highly available group provides protection from failure of any one of the members in the group. The communication between the GigaVUE-FM instances in the highly available group can be encrypted using Pre-shared Secret Key or Certificate based Authentication.

The following image shows the high-level architecture of the GigaVUE-FM HA feature.



GigaVUE-FM allows you to add more than three GigaVUE-FM instances to the HA group (up to 7). The hardware requirements and resource allocation process for the additional instances are the same as that of the active and standby instances. The additional instances can be configured either as supporting instances or as active-eligible instances:

Supporting instances

Supporting instances are GigaVUE-FMs that run minimal services, thereby providing more resources to OpenSearch services. These instances:


- Do not distribute application requests, as all application services are disabled by default.
- Do not participate in distributed services
- Provide enhanced scalability.
- Are excluded from active instance election

Active-Eligible instances

Active-Eligible instances are GigaVUE-FMs that distribute application requests and run distributed services. These instances:

- Enhance fault-tolerant and high available system.
- Are eligible for active instance election.

Refer to the following diagram for the roles in the HA group.



NOTE: When you add a GigaVUE-FM instance as either a supporting instance or a standby instance, the current configuration of that GigaVUE-FM instance will be removed once it is added to the existing GigaVUE-FM HA cluster.

Gigamon cloud solution deployments in GigaVUE-FM HA are supported on both public and private platforms. These deployments can operate in one of the following modes:

- GigaVUE-FM Orchestrated
- Third-party Orchestrated (Generic / Integrated)

In either scenario, you need to first launch GigaVUE-FM and run.

GigaVUE-FM Orchestration

Using GigaVUE-FM orchestration, you can deploy GigaVUE Fabric Components through the GigaVUE-FM user interface. After launching GigaVUE-FM and applying the required license, you can deploy components such as the GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller.

Once the deployment is complete, GigaVUE-FM continuously monitors the health and status of the deployed components. If any error condition occurs, GigaVUE-FM automatically relaunches or replaces the affected component to ensure system reliability.

Third-party Orchestration

With this method, you deploy GigaVUE Fabric Components using external orchestration solutions. GigaVUE-FM does not manage the deployment process directly, but you can still use it to monitor the health of the deployed components.

To enable monitoring, you need to provide an additional configuration file during deployment. This file allows components like the GigaVUE V Series Node, GigaVUE V Series Proxy, UCT-V Controller, UCT-V, UCT-C, and GCB to register with the Monitoring Domain in GigaVUE-FM. Refer to the following sections for more details:

- [Deploy UCT-C Solution in Kubernetes](#)
- [Deploy GCB Controller Service and Pods](#)

Supportability Information

The GigaVUE-FM HA feature is supported on the following platforms:

Orchestration Method	Platforms/components Supported
GigaVUE-FM Orchestration	• AWS • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix
Third-party Orchestration	Generic Mode • AWS • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix • GCP • UCT-C • GCB Integrated Mode • AWS • Azure • OpenStack

Recommendations, Rules, Notes, and Limitations

Keep in mind the following recommendations, rules and notes when you configure the GigaVUE-FM HA feature:

Recommendations

- Do not configure GigaVUE-FM High Availability (HA) across multiple regions or multi-cloud environments. This setup is not recommended and may lead to instability.
- You can deploy HA across multiple availability zones within the same cloud region. Ensure that you deploy a GigaVUE-FM instance in each availability zone.

- Use the orchestrated upgrade procedure to upgrade the GigaVUE-FM instances.
- It is recommended to configure reverse DNS configuration while configuring FMHA with FQDN.
- When creating a GigaVUE-FM HA group, it is highly recommended to create the High Availability group using DNS name. GigaVUE-FM HA cannot manage the individual GigaVUE-FM instances if the IP address changes (if GigaVUE-FM IP(s) are dynamic).

Rules and Notes

- GigaVUE-FM cannot support more than 7 active-eligible nodes (this is due to the underlying constraint in MongoDB).
- The GigaVUE-FM instances in the HA group must be identical in terms of system configuration such as hard disk, memory, and network interfaces, which include domain server, NTP server, and name server. Domain server, NTP server, and name server should be configured in all GigaVUE-FM instances.
- The GigaVUE-FMs in the high availability group can be accessed using the IPv4/IPv6 address (DNS name) that is used to form the High Availability group.

NOTE: Do not access all the GigaVUE-FM instances at the same time, as this will impact the performance of the HA group.

- You can deploy the GigaVUE-FM HA virtual machines on a WAN link with a maximum latency of 200 ms.
- You cannot add a GigaVUE-FM Hardware Appliance and a GigaVUE-FM virtual machine in the same HA group.
- The **Reload All** Or **Reload any FM** commands will work only if there is an active GigaVUE-FM instance in the HA group.
- You must ensure that NTP is configured across all FM HA nodes.
- The GigaVUE-FM HA Cluster functions like standalone GigaVUE-FM's when a fail-over occurs during long-running operations, such as Fabric Launch, Monitoring Session Deployment, or Fabric Upgrade.
- When FMHA is configured with a reverse proxy, access to GigaVUE-FM requires non-SAML authentication, such as local authentication. Use the following sample URL to access GigaVUE-FM:
`https://<proxy_ip>/admin`
- If GigaVUE-FM has multiple IPv6 addresses, ensure FMHA uses the permanent IPv6 as the source IP. GigaVUE-FM excludes temporary IPv6 addresses from the IP set, which causes connection failures and FMHA formation issues.

Shrinking the size of the HA group:

- When reducing the size of a HA group that has both active-eligible and supporting instances, it is always recommended to remove the supporting instances first so that the active-eligible instances remain intact in the HA group. If active-eligible instances are removed first from the HA group, then the HA group will become unstable. This, however, depends on the number of active-eligible and supporting instances in the HA group. Consider a GigaVUE-FM HA group with 7 instances. The number of active and supporting instances in a HA group may be as follows:

HA group 1	With 3 active instances and 4 supporting instances	The recommendation is specifically applicable for HA group 1 in which there are only 3 active eligible instances in the group. Removing the active-eligible instances may cause the HA group to become unstable.
HA group 2	With 5 active instances and 2 supporting instances.	The recommendation is not applicable for HA group 2 as removing the active-eligible instances will not impact the stability of the HA group.

For a HA group to be stable, it must have a minimum of two active-eligible instances in the group.

- When reducing the size of HA group from 7 instance to 5 instance or 3 instance, the quorum is re-computed and adjusted accordingly. This is to ensure that the HA group remains intact with the available set of instances.

- When reducing the size of the HA cluster, during removal of an instance make sure to give time for the OpenSearch to rebalance and return to a complete (100%) healthy state before removing the next instance. If you remove the instances subsequently one after the other, the health state of the HA group will turn red as there will be no time for shard rebalancing amongst the available instances. Login to GigaVUE-FM CLI as a root user and execute the following command to check the ES cluster health:

```
curl -XGET "http://localhost:9200/_cluster/health?pretty"
```

Supporting GigaVUE-FM Instances:

- When a supporting instance is added to the HA group, its system information will not be available immediately and will be available only after a minute.
- When the HA group is upgraded, application services of the supporting instances will be turned on automatically for the upgrade and once GigaVUE-FM HA upgrade is completed it will be turned off automatically. Application services will be turned off automatically irrespective of the upgrade completion status, that is, either success or failure.



Notes:

- For supporting instances, system information will be populated into MongoDB as the Application Service is down by default (scheduler populates the system information every minute into MongoDB). Hence any changes in the system hostname or GigaVUE-FM version will be reflected in the GUI only after a minute.
- In case of suspended mode, MongoDB will go into write protected mode. Therefore, any system related updates performed in the supporting nodes will not be reflected in the GUI.
- If the supporting instance is not reachable, all system related information that is populated in the GUI reflects the previous known values and not the latest values for Application Status, Up Time, and other system related information.
- Whenever the application services are up and running, events will be logged in the Events page (with the severity level - Info)

Licensing Information

You must install a Prime license on the active GigaVUE-FM instance to configure a High Availability group.

If the Prime license expires or if you accidentally delete the license, the existing configurations will still be present in the GigaVUE-FMs that are part of the HA group, but

you will not be able to perform any new configurations. Moreover, if you disassemble the HA group, you cannot reconfigure the HA group without installing a valid Prime license. You can re-use the same Prime license to form the HA group again.

Hostname Setups

The GigaVUE-FM instances are not required to be in the same subnet, but still must be able to communicate with each other.

In addition, ensure that the instances have unique host names. You must be able to ping a GigaVUE-FM instance from the other instances using the hostname or the IP address.

To add a GigaVUE-FM instance to a GigaVUE-FM HA group:

- The GigaVUE-FM instances in the HA group must be reachable to each other.
- If host names are used to configure the HA group, the host names of the GigaVUE-FM instances must be resolvable through a DNS server.

Configure GigaVUE-FM HA

To configure the GigaVUE-FM HA feature, you must have access to at least three authenticated GigaVUE-FM instances that reside on a trusted network. All GigaVUE-FM instances must run the same software version. The interfaces in the GigaVUE-FM instances must be up and must be assigned IPv4/IPv6 addresses. You can also choose to use DNS host names.

NOTE: You can configure the GigaVUE-FM HA feature only if you have administrative privileges. You must have Prime license installed for forming the GigaVUE-FM HA group.


Refer to the following sections:

- [Create and Add GigaVUE-FM Instances to HA Group](#)
- [Connect GigaVUE-FM High Availability to a Monitoring Domain using GigaVUE-FM Orchestration](#)
- [Connect GigaVUE-FM High Availability to a Monitoring Domain using Third Party Orchestration](#)
- [Integrate GigaVUE-FM High Availability with UCT-C](#)
- [Integrate GigaVUE-FM High Availability with Gigamon Containerized Broker](#)

Create and Add GigaVUE-FM Instances to HA Group

To add instances and to configure the HA group:

1. Go to  High Availability. The launch page displays the required prerequisites to create a HA group.



2. Click **Create**. In the **Group Name** field, enter a unique name for the HA group.

NOTE: Avoid using functionality names as the group alias when creating an FMHA group. Functionality names typically refer to actions or operations within the system and can cause confusion and conflicts. Choose a unique and descriptive name instead.

3. Choose any one of the following options to encrypt the communication between nodes:
 - **Pre-shared Secret Key** – Encrypts communication between GigaVUE-FM nodes using an application-managed shared secret key. Refer to [Encryption of Communication between FMHA Nodes](#).
 - **Certificate Based** – Encrypts communication between GigaVUE-FM nodes using manually uploaded certificates. Refer to [Encryption of Communication between FMHA Nodes](#).
4. In the **Add GigaVUE-FM Instance** section, the details of the GigaVUE-FM instance from which the HA group is created will be fetched automatically.

Edit the Primary Node

1. On the first instance, click the ellipses icon and select **Edit Node**.
2. Enter the following details: (These fields are optional)
 - Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
 - **Tunnel IP Address/DNS Name:** IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.

- **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name:** If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
 - Enter the Entity ID and the IdP Metadata URL.
3. Click the **save** icon to save the configuration. Refer to [Additional Fields](#) for details about other fields displayed on the HA page.

High Availability

Create New

There must be atleast three active eligible nodes that are reachable to each other.

Group Name * FMHA1

Secure Communication

☒ Pre-Shared Key ☐ Certificate Based

Add GigaVUE-FM Instances

Add New Node

IP ADDRESS / DNS NAME	TUNNEL IP ADDRESS / D...	EXTERNAL IP ADDRESS / ...	ROLE	USERNAME	PASSWORD	ENTITY ID	IDP METADATA URL
fcab:0:1:2200:96ae:3c8...	—	—	Active Eligible	admin	*****	—	—

Add New Node

1. Click **Add New Node** to add the GigaVUE-FM instances. Enter the following details:
 - **IP Address/DNS Name** - IP address/DNS name of the GigaVUE-FM instances (for communication between the GigaVUE-FMs and to configure the HA group). It is always recommended to use the DNS name while creating the GigaVUE-FM High Availability group.
 - (Optional) Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
 - **Tunnel IP Address/DNS Name:** IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.
 - **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name:** If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
 - When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.

**Notes:**

- When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.
- When using unique DNS for HA cluster configuration, ensure it always resolves to only one IP address—not multiple IPs.

- **Role** - Role of the GigaVUE-FM instance. Can be configured as below. Refer to [Active-Eligible instances](#) for details.
 - **Active Eligible**- provides resiliency and load distribution.
 - **Supporting Node**- provides scalability and performance.
 - **Username** - GigaVUE-FM GUI Username.
 - **Password** - GigaVUE-FM GUI password.
 - (Optional) **Entity ID** - Unique alphanumeric ID assigned to the GigaVUE-FM instance
 - (Optional) **IdP Metadata URL** - URL pointing to the IdP's XML configuration
2. Click **Save** to validate the details entered for the instances.

Use the **Add New Node** button to continue to add the GigaVUE-FM HA instances. You can click the **Edit Node** icon to change the external IP address. After adding the instances to the HA group, click **Submit**. The HA group is configured and will be displayed as follows. Refer to [Additional Fields](#) for details about other fields displayed on the HA page.

High Availability

Healthy
This HA group is Healthy. The members are protected from failure.

No supporting nodes in cluster ①

Group Name: fmha

No of Nodes: 3

Encryption Authentication Method: Pre-Shared Key Change ①

Tunnel Status: ✓ All tunnels are healthy

+ Add New Node Delete HA Group

STATUS	IP ADDRESS / D...	TUNNEL IP ADD...	ROLE	EXTERNAL IP A...	ENTITY ID	HOST NAME	APPLICATION STA...	SOFTWARE VERSI...	SYSTEM UPTIME	IDP METADATA ...	
Standby		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...
Standby		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...
Active		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...

Go to page: 1 of 1

Note: Please click [Refresh](#) to obtain the latest data.

**Notes:**

- When switching between the Pre-shared Secret Key and Certificate-Based options for encrypting communication between nodes, it's important to refresh the GigaVUE-FM to see the updated tunnel status. By clicking the link in the **Tunnel Status**, GigaVUE-FM displays the health status of the FMHA nodes.
- When creating a HA group, all the existing data on the first configured GigaVUE-FM instance will be preserved, allowing the system to continue using the data after it is fully set up. However, the data on the subsequent nodes will be lost during the setup.

The HA page also displays the following details:

Table 3: Additional Fields

Field	Description
Status and Reachability	Status of the GigaVUE-FM instances. Can be: • Active • Standby • Supporting Refer to Active-Eligible instances for details. Reachability: • - indicates that the instances are reachable. • - indicates that the instances are unreachable.
Host Name	Host name of the GigaVUE-FM instance.
Application Status	Status of application services that are responsible for handling the user requests: • Up: Indicates service is ready to handle the user requests • Down: Indicates service is not ready to handle the user requests with one or more underlying services being down. • In Progress: A state that appears between up and down indicating that the services are initializing.
Software Version	Software version of the GigaVUE-FM instance.
System Uptime	Time since the instance is up and active.
Third Party Authentication URL	Third party authentication URL required for SSO configuration.
Upgrade Status	Upgrade status of GigaVUE-FM. This is displayed only when the upgrade is triggered. Otherwise, this field is empty.

You can also remove the instances from the HA group and reduce the size of the HA group as required. However, while removing the instances ensure to retain the quorum¹ number of instances.

¹Quorum number of instances is the minimum number of instances that must be active-eligible in a HA group. In a HA group with five instances, the quorum number is 3 (based on the formula $n/2+1$).

NOTE: To scale components such as Fabric Health Analytics, topology visualization and to improve the performance of GigaVUE-FM by fully utilizing the resource available, it is recommended to configure three active eligible instances and additional supporting instances.

The following are the behavioral changes observed in the HA group with active, standby and supporting instances:

Task	Scenario	Behavior
Upgrade Operation	Few instances not being up (either active-eligible or supporting instances not up)	Upgrade will not be initiated.
	Few supporting instances are removed from the HA group after formation of the group	Upgrade will be initiated, and the supporting instances present in HA group will be upgraded first followed by the active-eligible instances.
	All instances are up and running	Upgrade will be initiated in an active instance. Upgrade will be done in a rolling fashion starting with supporting instances and ending with the active-eligible instances.
System Configurations	All instances are active-eligible	All GigaVUE-FM's GUI will be accessible, and user can configure the system through the same
	HA with active-eligible and few Supporting instances	For supporting instances, GUI will not be available (as application service is shut down). An option is provided in the GUI to turn on the application service and its dependent service(s) and the status can be seen in HA page. You can then configure the system related configuration from the logged in GigaVUE-FM instance once the application status of the supporting instance is up.
Background processes		
-FM Service	All GigaVUE-FMs are active-eligible	Services will be equally

Distribution	instances.	distributed/redistributed among the instances based on the availability of the instances.
	HA group with active-eligible instances and supporting instances.	Services will be distributed/re-distributed only to active eligible instances
-Load Balancer	All GigaVUE-FMs are active-eligible instances.	All GigaVUE-FM API requests will be distributed to all instances in the HA group.
	With supporting instances.	All GigaVUE-FM API requests will be distributed only to active-eligible instances.

Connect GigaVUE-FM High Availability to a Monitoring Domain using GigaVUE-FM Orchestration

When fabric nodes need to communicate with GigaVUE-FM using public IP addresses, you should enable the **Use Public IP** option during Monitoring Domain creation. When **Use Public IP** option is enabled, GigaVUE-FM pushes its public IP address to Fabric Nodes. Fabric Nodes use this IP to communicate back to FM. FM does not initiate communication using its private IP. This setting is essential in cloud or hybrid environments where private IPs are not reachable, such as across cloud platforms. It ensures that GigaVUE-FM can send notifications and accept registrations using an externally accessible IP. Refer to [Supportability Information](#) for the list supported platforms.

Connect GigaVUE-FM High Availability to a Monitoring Domain using Third Party Orchestration

Fabric Nodes can be configured with either the IP address (remoteIP) or the Fully Qualified Domain Name (FQDN) (remoteAddress) of GigaVUE-FM to establish communication. The node sends registration requests to GigaVUE-FM until it successfully registers. In High Availability mode, if a node fails to register with one GigaVUE-FM instance, it automatically attempts registration with the next available GigaVUE-FM in its list until successful. Refer to [Supportability Information](#) for the list supported platforms.

Use the following data to manually configure **/etc/gigamon-cloud.conf** with registration configurations to register the fabric components with GigaVUE-FM.

```
Registration:
  groupName: <Monitoring Domain Name>
```

```
subGroupName: <Connection Name/Alias>
token: <Token>
remoteAddress: <remoteIPs/FQDN of the GigaVUE-FM>
remotePort: 443
sourceIP: <IP address of the fabric node>
nameServer: <DNS server's IP address>
usePublicIP: <true/false>
preferV6: <true/false><If set to true IPv6 will be preferred over IPv4, used only for DNS
resolution>
```



Notes:

- Configure either remoteIP or remoteAddress—not both. Using both prevents the node from registering with GigaVUE-FM.
- Remote IP will continue to work as expected. However, if you plan to use FQDN, update the configuration to use the Remote Address field instead.
- When configuring a remote address, you can specify multiple IP addresses. However, if you use a FQDN, ensure that you configure only one address
- FQDN configuration requires a valid nameserver.
- Ensure that the FQDN configured for Third Party Orchestration resolves to all IP addresses in the High Availability cluster.

Integrate GigaVUE-FM High Availability with UCT-C

UCT-C supports integration with GigaVUE-FM in high availability mode. You can configure the controller with either a Fully Qualified Domain Name (FQDN) or a comma-separated list of GigaVUE-FM IP addresses in the Helm configuration. The controller iterates through the provided FQDN or IP list and attempts to register with the master GigaVUE-FM node. If the GigaVUE-FM IP changes, you must update the configuration again. This enhancement enables UCT-C to maintain connectivity and recover from GigaVUE-FM failover events in HA deployments.

Below are the communication parameters:

- **fm_fqdn** – Specifies the GigaVUE-FM HA FQDN value used by the UCT-C Controller for communication with the GigaVUE-FM HA group. When this parameter is set, the controller uses the IP addresses resolved from the FQDN.
- **fmha_ip_list** – Specifies a comma-separated list of GigaVUE-FM HA IP addresses used by the UCT-C Controller for communication with the GigaVUE-FM HA group. If fm_fqdn is not set, the controller uses the IP addresses provided in fmha_ip_list by default.

Integrate GigaVUE-FM High Availability with Gigamon Containerized Broker

GigaVUE-FM High Availability group can be configured in an Gigamon Containerized Broker through a DNS Server in which the GigaVUE-FM instances are launched and deployed in the Gigamon Containerized Broker environment.

For GCB to make use of high availability feature of GigaVUE-FM, you must configure the FQDN (Fully Qualified Domain Name) of the GigaVUE-FM.

In a standalone GigaVUE-FM, the GCB uses the FQDN name of the GigaVUE-FM (if configured). It uses the legacy method of using the configured IP address when:

- The FQDN is not configured or
- The GCB fails to resolve to an IP address.

For example:

GCB Controller YAML file:

```
env:
  - name: GCB_CNTLRL_EXT_IP_DNS
    value: "10.xxx.xx.xx" (IP address of the DNS server - external to the
    Kubernetes cluster)
  - name: FM_FQDN
    value: "fm.myorg.com" (FMs FQDN that is used for DNS lookups)
```

NOTE: The FQDN identifies either a standalone GigaVUE-FM or an FMHA cluster that includes multiple GigaVUE-FM instances. It returns the IP addresses (IPv4, IPv6, or both) for all GigaVUE-FMs associated with that FQDN.

For more details on the GigaVUE-FM High Availability configuration, refer to the *GigaVUE Administration Guide*.

Configure Secure Communication between Fabric Components in FMHA

IMPORTANT: Before upgrading the Fabric Components to version 6.10 or later, complete the following steps after upgrading GigaVUE-FM to version 6.10 or later.

Follow these steps:

1. Access the active GigaVUE-FM via CLI.
2. Archive the stepCA directory using the following commands:

```
sudo su
cd /var/lib
tar -cvf /home/admin/stepca.tar stepca
```
3. Set the permissions of the tar file using the following commands:

```
chmod 666 /home/admin/stepca.tar
```
4. Copy the tar file to all standby instances in the **/home/admin/ directory** using scp:

```
scp /home/admin/stepca.tar <standby-node>:/home/admin/
```
5. Download the **runstepca_fmha** script from the Community Portal.
6. Log in to the standby instance using CLI.

- Copy the script in the standby instance in the **/home/admin directory** and execute it using the following command:

```
sh /home/admin/runstepca_fmha
```


GigaVUE-FM HA Landing Page

When you log in to any instance in the HA group, the Dashboard page appears. On specific pages, you can use the instance dropdown list in the page header to view information for a particular instance.

Remove Standby GigaVUE-FM Instance

To remove or replace a standby GigaVUE-FM instance from the GigaVUE-FM HA group follow the below listed steps:

- Go to  High Availability.
- Click the ellipsis on the standby GigaVUE-FM instance widget in the High Availability page as shown in [Figure 2 Remove Standby GigaVUE-FM Instance](#).



The screenshot shows the 'High Availability' dashboard. At the top, a green checkmark indicates the group is 'Healthy'. Below this, a table lists the nodes in the HA group. The first node is in 'Standby' status. A dropdown menu is open for this node, showing options: 'Edit', 'Reboot', and 'Remove from group'. A red arrow points to the 'Remove from group' option.

STATUS	IP ADDRESS / D...	TUNNEL IP ADD...	ROLE	EXTERNAL IP A...	ENTITY ID	HOST NAME	APPLICATION STA...	SOFTWARE VERSI...	SYSTEM UPTIME	IDP METADATA ...	
Standby		—	Active Eligible		FM3	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...
Active		—	Active Eligible		FM1	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...
Standby		—	Active Eligible		FM2	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...

Figure 2 Remove Standby GigaVUE-FM Instance

- Select the **Remove from group** option and confirm your changes. The selected standby GigaVUE-FM instance is removed from the GigaVUE-FM HA group.

NOTE: The status of the GigaVUE-FM HA group changes from **Healthy** to **At Risk**. You will not be allowed to remove the other standby GigaVUE-FM instance after the HA status changes to **At Risk**. To remove the standby node, if there is any network latency in GigaVUE-FM and the node is not removed correctly, use the command **fmcs**

remove force in the GigaVUE-FM CLI of the standby node.

Disassemble GigaVUE-FM High Availability Group

To completely disassemble the GigaVUE-FM HA group:

1. Go to  High Availability.

NOTE: You cannot remove an active GigaVUE-FM instance or disassemble the GigaVUE-FM HA group by logging in from a standby GigaVUE-FM instance.

2. Click the **Delete HA group** option on the High Availability page and confirm your changes. The GigaVUE-FM HA group will be disassembled and each of the GigaVUE-FM instances become standalone GigaVUE-FM instances.

NOTE: Executing the above steps disassembles the GigaVUE-FM HA group completely. The database of the individual GigaVUE-FM HA group members (active, standby, and supporting instances) will be erased and reinitiated with the default database. If you had performed orchestrated configuration flows (such as the Flexible Inline Arrangement or Intent Based Orchestration) you must manually remove and recreate these if they do not backup the GigaVUE-FM database before choosing Delete HA Group.

If you cannot access the active GigaVUE-FM instance's GUI, use the following CLI command in all the GigaVUE-FM instances to disassemble the HA group:

/opt/fmcs/bin/fmcs leave force


Before disassembling the instances, take a backup of the configuration. After the GigaVUE-FM instances are disassembled from the HA group, the configurations present in the instances will be completely removed.

GigaVUE-FM High Availability States

The GigaVUE-FM HA state depends on the status of the three GigaVUE-FM instances. The following table lists the various states of the GigaVUE-FM HA group. You can view the HA group state from **Administration > High Availability** in the GigaVUE-FM GUI.

Table 4: High Availability States

State	Number of GigaVUE-FM Instances	Description
Healthy	Three GigaVUE-FM instances are up and running	One GigaVUE-FM instance is in active state. Other two instances are in standby state.
At Risk	Two GigaVUE-FM instances are up and running	One GigaVUE-FM instance is in active state. Another GigaVUE-FM instance is in standby state. The third GigaVUE-FM instance has either not joined the HA group or has left the HA group. NOTE: You must recover the standby GigaVUE-FM instance within 3 days (72 hours). Failure to do so will cause the instance to move out of the High Availability group.
Incomplete	One GigaVUE-FM instance is up and running	Only one GigaVUE-FM instance is in active state. The other two GigaVUE-FM instances have either not joined the HA group or have left the HA group.
Standalone	One GigaVUE-FM instance is up and running	HA is not configured on the GigaVUE-FM instance.
Suspended	Two or more GigaVUE-FM instances are up and running	HA is configured, but an active GigaVUE-FM instance is yet to be elected. NOTE: The GigaVUE-FM high availability group can move to a Suspended state from the At Risk state when the HA group has only one active-eligible and reachable GigaVUE-FM instance. Events will not be captured for the status of the GigaVUE-FM instance that went down (which causes the HA group to move into a suspended state) as there is no active/leader instance in the High Availability group.



In case of FMHA going to suspended mode from At-Risk mode, the reachability status of node which went down (due to which HA went into suspended mode) won't be audited as writes are not accepted if there is no active/leader instance in the DB cluster

Figure 3 High Availability States

Fail over Mechanism

The active GigaVUE-FM instance in the high availability group may fail at times resulting in one of the standby instances to take over and become the active instance. This process is called failover.

The following table provides the reasons for failover:

Reason for Failover	Description
Reloading the active GigaVUE-FM instance	An active GigaVUE-FM instance is reloaded (using the Reboot option) to bring back the HA group to healthy state again.
Planned downtime of the active GigaVUE-FM instance	An active GigaVUE-FM instance is brought down due to various reasons, for example to upgrade to a newer software version.
Monitoring session deployment during HA leader switchover	The new leader automatically redeploys Monitoring Sessions that failed during the switchover.
Node registration during HA leader switchover	Registration requests may temporarily fail or be delayed. Fabric nodes automatically retry registration. The new leader's IP address is configured as the RMQ notification target on all fabric nodes.
Certificate handling or fabric deployment/upgrade during HA leader switchover	The newly elected leader resumes certificate handling and continues any in-progress fabric deployment or upgrade from the last known state. No manual intervention is required.

GigaVUE-FM High Availability Scenarios

The High Availability page displays the current state of the GigaVUE-FM HA group. When a failover occurs, the HA group state changes in the GUI.

The following table lists the GUI changes for the various scenarios:

Scenario	Changes in GUI
What happens to the High Availability page immediately after a failover?	The High Availability page may not update immediately or may not show all the GigaVUE-FM instances. NOTE: Refresh the page after a minute to view the new active GigaVUE-FM instance. However, the page will get updated automatically after 5 minutes.
What happens to an active GigaVUE-FM instance when it fails (either by itself or if failover is triggered manually)?	- One of the standby GigaVUE-FM instances changes to the active state. The GigaVUE-FM instance that was initially in the active state changes to the standby state. It takes a few seconds for this transition. - The GigaVUE-FM GUI of the new active instance will have all the menus and dashboards.
What happens to the GigaVUE-FM instances that were previously in standby state?	- One of the standby GigaVUE-FM instances changes to the active state - The other standby GigaVUE-FM instance remains in the standby state.
What happens to the embedded devices when a new Active instance takes over?	There are no changes to the devices except that they are being managed by the new active GigaVUE-FM instance.
How do you trigger a failover?	Click on the 'Reboot' option on the current active GigaVUE-FM instance to trigger a failover.

Troubleshoot GigaVUE-FM High Availability Issues

Use the following table to troubleshoot issues that you might encounter while working with the HA feature.

Problem	Solution
Unable to add a license to GigaVUE-FM HA group after a failover Reason: Using the MAC Address in the About page to generate the license	Always use the Challenge MAC Address in the Licenses page of the active GigaVUE-FM instance to generate licenses. Add the licenses to the HA group.
Unable to join the GigaVUE-FM HA group after changing the password Reason: Not logging out of GigaVUE-FM after changing the password and before joining the HA group	Always logout of GigaVUE-FM after changing the password and login again with the new password before joining the HA group.

(missing or bad snippet)(missing or bad snippet)

Analytics for Virtual Resources

Analytics in GigaVUE-FM is a standalone service that provides data visualization capabilities. Using Analytics ¹, you can create visual elements such as charts that are embedded as visualizations. The visualizations are grouped together in dashboards.

You can also create search objects using Analytics. Dashboards, Visualizations and Search Objects are called Analytics objects. For details, refer to [Analytics](#).

Rules and Notes:

- You cannot edit or delete these default dashboards. However, you can clone the dashboards and visualizations. Refer to the Clone Dashboard section in GigaVUE-FM Installation and Upgrade Guide for more details.
- Use the **Time Filter** option to select the required time interval for which you need to view the visualization.

Virtual Inventory Statistics and Cloud Applications Dashboard

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly.

¹Analytics uses the OpenSearch front-end application to visualize and analyze the data in the OpenSearch database of GigaVUE-FM.

For details, refer to the [Analytics](#) section in *GigaVUE Fabric Management Guide* for details on how to create a new dashboard, clone a dashboard, create a new visualization, and other information about the Discover page and Reports page.

How to access the dashboards

1. Go to  -> **Analytics -> Dashboards.**
2. Select the required dashboard to view the visualizations.

The following table lists the various virtual dashboards:

Dashboard	Displays	Visualizations	Displays
Inventory Status (Virtual)	Statistical details of the virtual inventory based on the platform and the health status. You can view the following metric details at the top of the dashboard: • Number of Monitoring Sessions • Number of V Series Nodes • Number of Connections • Number of GCB Nodes You can filter the visualizations based on the following control filters: • Platform • Health Status	<i>V Series Node Status by Platform</i>	Number of healthy and unhealthy V Series Nodes for each of the supported cloud platforms.
		<i>Monitoring Session Status by Platform</i>	Number of healthy and unhealthy monitoring sessions for each of the supported cloud platforms
		<i>Connection Status by Platform</i>	Number of healthy and unhealthy connections for each of the supported cloud platforms
		<i>GCB Node Status by Platform</i>	Number of healthy and unhealthy GCB nodes for each of the supported cloud platforms
V Series Node Statistics	Displays the Statistics of the V Series node such as the CPU usage, trend of the receiving and transmitting packets of the V Series node. You can filter the visualizations based on the following control filters: • Platform • Connection • V Series Node	<i>V Series Node Maximum CPU Usage Trend</i>	Line chart that displays maximum CPU usage trend of the V Series node in 5 minutes interval, for the past one hour. NOTE: The maximum CPU Usage trend refers to the CPU usage for service cores only. Small form factor V Series nodes do not have service cores, therefore the CPU usage is reported as

Dashboard	Displays	Visualizations	Displays
			0.
		<i>V Series Node with Most CPU Usage For Past 5 minutes</i>	Line chart that displays Maximum CPU usage of the V Series node for the past 5 minutes. NOTE: You cannot use the time based filter options to filter and visualize the data.
		<i>V Series Node Rx Trend</i>	Receiving trend of the V Series node in 5 minutes interval, for the past one hour.
		<i>V Series Network Interfaces with Most Rx for Past 5 mins</i>	Total packets received by each of the V Series network interface for the past 5 minutes. NOTE: You cannot use the time based filter options to filter and visualize the data.
		<i>V Series Node Tunnel Rx Packets/Errors</i>	Displays the reception of packet at the Tunnel RX. This is the input to V Series Node, Grouping by tunnel identifier comprising {monDomain, conn, VSN, tunnelName}, before aggregation.
		<i>V Series Node Tunnel Tx Packets/Errors</i>	TX is for output tunnels from VSN. V Series Node Tunnel Tx Packets/Errors
Dedup	Displays visualizations related to Dedup application. You can filter the visualizations based on the following control filters: - Platform - Connection - V Series Node	<i>Dedup Packets Detected/Dedup Packets Overload</i>	Statistics of the total de-duplicated packets received (ipV4Dup, ipV6Dup and nonIPDup) against the de-duplication application overload.

Dashboard	Displays	Visualizations	Displays
		<i>Dedup Packets Detected/Dedup Packets Overload Percentage</i>	Percentage of the de-duplicated packets received against the de-duplication application overload.
		<i>Total Traffic In/Out Dedup</i>	Total incoming traffic against total outgoing traffic
Tunnel (Virtual)	Displays visualizations related to the tunneled traffic in both bytes as well as the number of packets. You can select the following control filters, based on which the visualizations will get updated: • Monitoring session: Select the required monitoring session. The cloud platform, monitoring domain and connection within the monitoring domain that is used by the V Series node are shown in square brackets, comma-separated, after the name, to distinguish the whole path to it. • V Series node: Management IP of the V Series node. Choose the required V Series node from the drop-down. • Tunnel: Select any of the tunnels shown in the Tunnel drop-down. The direction for each tunnel is shown with the prefix in or out. The following statistics are displayed for the tunnel: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets	<i>Tunnel Bytes</i>	Displays received tunnel traffic vs transmitted tunnel traffic, in bytes. • For input tunnel, transmitted traffic is displayed as zero. • For output tunnel, received traffic is displayed as zero.
		<i>Tunnel Packets</i>	Displays packet-level statistics for input and output tunnels that are part of a monitoring session.
App (Virtual)	Displays Byte and packet level statistics for the applications for the chosen monitoring session on the selected V Series node. You can select the following control	<i>App Bytes</i>	Displays received traffic vs transmitted traffic, in Bytes.

Dashboard	Displays	Visualizations	Displays
	filters, based on which the visualizations will get updated: • Monitoring session • V Series node • Application: Select the required application. By default, the visualizations displayed includes all the applications. By default, the following statistics are displayed: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Errored Packets • Dropped Packets		
		<i>App Packets</i>	Displays received traffic vs transmitted traffic, as the number of packets.
End Point (Virtual)	Displays Byte and packet level statistics for the un-tunneled traffic deployed on the V Series nodes. The following statistics that are shown for Endpoint (Virtual): • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets The endpoint drop-down shows <V Series Node Management IP address : Network Interface> for each endpoint. You can select the following control filters, based on which the visualizations will get updated: • Monitoring session • V Series node • Endpoint: Management IP of the V Series node followed by the Network Interface (NIC)	<i>Endpoint Bytes</i>	Displays received traffic vs transmitted traffic, in Bytes.
		<i>Endpoint Packets</i>	Displays received traffic vs transmitted traffic, as the number of packets.

NOTE: The Tunnel (Virtual), App (Virtual) and Endpoint (Virtual) dashboards do not show data from the previous releases if the *Monitoring Session [Platform :*

Domain : Connection] dashboard filter is applied. This is because, this filter relies on the new attributes in the OpenSearch database, which are available only from software version 5.14.00 and beyond.

Remove Gigamon Service from NSX-T and GigaVUE-FM

To clean up the Gigamon Deep Observability Pipeline from VMware NSX-T and GigaVUE-FM, perform the following steps in order:

Step 1: Remove the Rule and Policy

To delete the network monitoring services:

1. Select **Security > Service Chain Management > E-W Network Introspection**. The E-W Network Introspection page appears.
2. On the E-W Network Introspection page, locate the policy created for Gigamon tapping.
3. Expand the policy and perform the following:
 - a. To delete the rule associated with the policy, select the  icon on the Rule Column and then select **Delete**
 - b. To delete the policy, select the  icon on the policy and then select **Delete**.

Step 2: Remove the Service Chain

To delete the network monitoring services:

1. In NSX-T Manager, select **Security > Settings > Network Introspection**.
2. Select the **SERVICE CHAINS** tab.
3. Located the service chain associated with the Gigamon service profile,
4. Select  available next to the service chain and select **Delete**.

Step 3: Delete the Monitoring Session

To delete the Monitoring session from GigaVUE-FM:

1. From the left pane, select **Traffic > VIRTUAL > Orchestrated Flows > VMware**. The Monitoring Sessions pertaining to all VMware deployment appears.
2. Select the NSX-T related Monitoring Session and select **Actions > Undeploy**. The Monitoring Session is Undeployed.
3. Select the Monitoring Session again and select **Actions > Delete**.

This action deletes the corresponding service profile and the profile from the NSX-T manager console.

Step 4: Delete the Monitoring Domain

To delete the Monitoring Domain and the GigaVUE V Series Node deployed in GigaVUE-FM:

1. From the left navigation pane, select **Inventory > VIRTUAL > VMware NSX-T > Monitoring Domain**.
The Monitoring Domain page appears with the deployed GigaVUE V Series Nodes.
2. Select the appropriate **Monitoring Domain**,
3. Select **Actions > Delete Monitoring Domain**.

This action deletes the corresponding **Service Deployment** from the NSX-T Manager.


GigaVUE V Series Deployment Clean up

If a GigaVUE V Series node installation fails or the service is not fully removed, you must clean up the deployment before attempting a reinstall. Follow the steps below to remove all related components from NSX-T Manager and GigaVUE-FM.


Step 1: Delete Rule and Policy in NSX-T Manager

To remove Service Profiles:

1. From NSX-T Manager, navigate to **Security > Network Introspection (E-W)**.
2. In the **SERVICE PROFILES** tab, select the service name of the Monitoring Domain from the **Partner Service** drop-down.




3. Delete all the existing Service Profiles.




Remove Service Deployments

To remove Service Deployments:

1. From NSX-T Manager, navigate to **System > Service Deployments**.
2. From the **DEPLOYMENT** tab, select the service name of the Monitoring Domain from the **Partner Service** drop-down.




3. Delete all the existing Service Deployments.



To remove the Service Deployments through NSX-T API:

1. Log in to Postman.




2. Get the Service ID. **GET** `https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/`
3. Get the ID of the Service Deployments. **GET** `https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/<Service_ID>/service-deployments/`
4. Delete all Service Deployments. **DELETE** `https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/<Service_ID>/service-deployments/<Service_Deployment_ID>`

Remove Service Reference

To remove Service References through NSX-T API:

1. Log in to Postman.




2. Get the Service Reference ID. **GET** `https://<NSX_Manager_IP>/policy/api/v1/infra/service-references/`
3. Delete the Service Reference. **DELETE** `https://<NSX_Manager_IP>/policy/api/v1/infra/service-references/<Service_Reference_ID>`

Remove Service Manager

To remove Service Manager through NSX-T API:

1. Log in to Postman.




2. Get the Service Manager ID.**GET** https://<NSX_Manager_IP>/api/v1/serviceinsertion/service-managers/
3. Delete the Service Manager.**DELETE** https://<NSX_Manager_IP>/api/v1/serviceinsertion/service-managers/<Service_Manager_ID>

Remove Vendor Template and Service Definition

To remove Vendor Template and Service Definition through NSX-T API:

1. Log in to Postman.



2. Get the Service ID.**GET** https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/
3. Get the Vendor Templates' ID.**GET** https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/<Service_ID>/vendor-templates/
4. Delete the Vendor Templates.**DELETE** https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/<Service_ID>/vendor-templates/<Vendor_Template_ID>
5. Delete the Service.**DELETE** https://<NSX_Manager_IP>/api/v1/serviceinsertion/services/<Service_ID>

Debuggability and Troubleshooting

Use the following information to help diagnose and resolve GigaVUE V Series Nodes issues.

Sysdumps

A sysdump is a log and system data package generated when a GigaVUE V Series Node experiences a crash (such as kernel, application, or hardware failure). These files are essential for debugging.

You cannot download sysdump files if the associated fabric component is deleted or unreachable.

Sysdumps—Rules and Notes


Consider the following points before you generate sysdumps:

- You can generate only one sysdump file at a time for a GigaVUE V Series Node.
- You cannot generate a sysdump file when generation of another sysdump file is in progress.
- The limit of sysdump files available per GigaVUE V Series Node is six. When you generate a seventh sysdump file, the file overwrites the first sysdump file.
- You can download only one sysdump file per GigaVUE V Series Node at a time.
- You can delete sysdump files in bulk for a GigaVUE V Series Node.
- To ensure efficient usage, the system limits the number of simultaneous sysdump generation requests to 10 GigaVUE V Series Nodes.
- GigaVUE V Series Node sysdumps are not stored in Fabric Manager but generated and stored on the GigaVUE V Series Node itself.

Generate a Sysdump File

To generate a sysdumps file:

1. Perform one of the following:
 - Go to **Inventory > VIRTUAL > VMware NSX-T > Monitoring Domain**.
 - Go to **Inventory > VIRTUAL > VMware ESXi > Monitoring Domain**.
 - Go to **Inventory > VIRTUAL > Third Party Orchestration > Monitoring Domain**.
 - Go to **Inventory > VIRTUAL > Nutanix > Monitoring Domain**.The **Monitoring Domain** page appears.
2. Select the required node, and use one of the following options to generate a sysdump file:
 - Select **Actions > Generate Sysdump**.
 - In the lower pane, go to **Sysdump**, and select **Actions > Generate Sysdump**.
3. View the latest status, click **Refresh**.



Other Actions

- To download a sysdump file, select the file in the lower pane, and then click **Actions > Download**.
- To delete a sysdump file,
 1. Select the file in the lower pane.
 2. Select the desired sysdump file.
 3. Select **Actions > Delete**.
- To bulk delete, select all the sysdump files, and then select **Actions > Delete All**.

FAQs - Secure Communication between GigaVUE Fabric Components (VMware-VSeries-NSX-T)

This section addresses frequently asked questions about Secure Communication between GigaVUE Fabric Components and GigaVUE-FM. Refer to Secure Communication between GigaVUE Fabric Components section for more details.

1. Is there a change in the upgrade process for GigaVUE-FM and GigaVUE V Series Node?

No. The upgrade process remains unchanged across all supported upgrade paths. You can upgrade your nodes without any additional steps. The upgrade results in the automatic deployment of the appropriate certificates based on the node versions

GigaVUE-FM	GigaVUE V Series Nodes	Custom Certificates Selected (Y/N)	Actual Node Certificate
6.10	6.10	Y	GigaVUE-FM PKI Signed Certificate
6.10	6.9 or earlier	Y	Custom Certificate
6.10	6.9 or earlier	N	Self-Signed Certificate

2. What is the new authentication type used between GigaVUE-FM and the GigaVUE Fabric Components? Is backward compatibility supported?

Backward compatibility is supported, ensuring that fabric components running on version 6.9 or earlier remain compatible with GigaVUE-FM 6.10. The following authentication types are supported across different versions:

GigaVUE-FM	GigaVUE Fabric Components	Authentication
6.10	6.10	Tokens + mTLS Authentication (Secure Communication)
6.10	6.9 or earlier	User Name and Password

3. What are the new ports that must be added to the security groups?

The following table lists the port numbers that must be opened for the respective fabric components:

Component	Port
GigaVUE-FM	9600
GigaVUE V Series Node	80, 8892
GigaVUE V Series Proxy	8300, 80, 8892
UCT-V Controller	8300, 80
UCT-V	8301, 8892, 9902 For more details, refer to Network Firewall Requirements .

4. Is the registration process different for deploying the fabric components using Third-Party Orchestration?

Yes. Beginning with version 6.10, you must use tokens in the gigamon-cloud.conf file instead of the username and password. To generate the token in GigaVUE-FM, go to **Settings > Authentication > User Management > Token**. For more details, refer to [Configure Tokens](#).

Example Registration Data for UCT-V:

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      token: <Token>
      remoteIP: <IP address of the UCT-V Controller 1, <IP address of the UCT-V
Controller 2>
      sourceIP: <IP address of UCT-V> (Optional Field)
```

5. Are there any changes to the UCT-V manual installation and upgrade process?

Starting from version 6.10, you must add tokens during manual installation and upgrades.

- Create a configuration file named gigamon-cloud.conf with the token and place it in the /tmp directory during UCT-V installation
- After installing UCT-V, you can add the configuration file in the /etc directory.

Important! Without this token, UCT-V cannot register with GigaVUE-FM.

6. Can I use my PKI infrastructure to issue certificates for the Fabric Components?

Direct integration of your PKI with GigaVUE-FM is not supported. However, you can provide your Intermediate Certificate Authority (CA) to sign the node certificate.

7. What happens to the existing custom certificates introduced in the 6.3 release?

The custom certificate feature is not supported for the fabric components with version 6.10 or higher, even if a custom certificate is selected in the Monitoring Domain. However, this feature remains available for older versions.

- When upgrading from version 6.9 or earlier with custom certificates upgrades to version 6.10, the system automatically generates and deploys certificates signed by GigaVUE-FM.
- If deploying version 6.9 or earlier components from a 6.10 GigaVUE-FM, custom certificates are still applied.

8. How to issue certificates after upgrading the fabric components to 6.10?

When the upgrade process begins, GigaVUE-FM transmits the certificate specifications to the new fabric components using the launch script. The fabric components utilize these specifications to generate their own certificates.

9. Is secure communication supported in FMHA deployment?

Yes, it is supported. However, you must follow a few manual steps before upgrading the fabric components to 6.10. For details, refer to [Configure Secure Communication between Fabric Components in FMHA](#).

NOTE: This step is essential if you are using cloud deployments in FMHA mode and would like to deploy or upgrade the fabric components to version 6.10 or later.