



GigaVUE Cloud Suite for VMware ESXi - Deployment Guide

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 for VMware ESXi - Deployment Guide	1
Change Notes	3
Contents	4
Get Started	12
Overview of GigaVUE Cloud Suite for VMware	12
Architecture	13
GigaVUE Fabric Components and Deployment Options	14
GigaVUE-FM Orchestrated Fabric Deployment	14
Third Party Orchestrated Fabric Deployment	15
Benefits:	15
GigaVUE-FM Orchestrated Fabric Deployment	15
Third Party Orchestrated Fabric Deployment	16
Deployment Mode for VMware ESXi	17
Key Concepts	17
Secure Communication between GigaVUE Fabric Components	17
GigaVUE-FM acts as the PKI	19
Bring Your Own CA	19
Supported Platforms	19
Supported Components	20
Rules and Notes	20
Rediscover	20
Sharing the Same Host across Different Monitoring Domains	21
Selective Source Selection	21
Secure Tunnels	21
Precryption	22
How Gigamon Precryption Technology Works	23
Why Gigamon Precryption	24
Key Features	24
Key Benefits	25
How Gigamon Precryption Technology Works	25
Supported Platforms	27
Prerequisites	28
Prefiltering	29
Deployment Overview	30
Deployment Planning	30

Identify the Deployment Method	31
GigaVUE-FM Orchestration	31
Third Party Orchestration	31
Identify the vCenter Environment	31
Define GigaVUE-FM Deployment Location	32
Identify the Traffic Acquisition Method	32
Identify the VMware Hosts and Clusters	32
Identify the Network Configuration	32
Identify the Required vCenter Privileges	33
Deployment Prerequisites	33
Licensing for GigaVUE Cloud Suite for VMware ESXi	34
Default Trial Licenses	34
Volume-Based License	35
Points to Note (VMware vCenter)	41
Supported Hypervisors for VMware	41
Recommended Form Factor for VMware vCenter (Instance Types)	42
Minimum Virtual Computing Requirements	43
Network Firewall Requirements	43
Reduced Control-Plane Ports	43
GigaVUE-FM	45
UCT-V Controller	45
UCT-V	46
GigaVUE V Series Node	46
GigaVUE V Series Proxy (Optional)	46
UCT-V OVS Controller	47
UCT-V OVS Module	47
GigaVUE FM	47
UCT-V Controller	48
UCT-V	49
GigaVUE V Series Node	50
GigaVUE V Series Proxy (Optional)	51
UCT-V OVS Controller	52
UCT-V OVS Module	52
Legacy Control-Plane Ports	53
Reduced Control-Plane Ports	55
GigaVUE-FM	57
UCT-V Controller	57
UCT-V	57
GigaVUE V Series Node	58
GigaVUE V Series Proxy (Optional)	58
UCT-V OVS Controller	58
UCT-V OVS Module	58
GigaVUE FM	59
UCT-V Controller	60

UCT-V	61
GigaVUE V Series Node	62
GigaVUE V Series Proxy (Optional)	63
UCT-V OVS Controller	63
UCT-V OVS Module	64
Legacy Control-Plane Ports	64
GigaVUE FM	65
UCT-V	66
GigaVUE V Series Node	66
Giga VUE V Series Proxy(Optional)	68
GigaVUE V Series Node	69
GigaVUE V Series Proxy(Optional)	69
Configure Tokens	70
Prerequisite	70
Rules and Notes	70
Create Token	71
Revoke Tokens	71
Export Token	72
Create a Token with Custom User Group	72
Integrate Private CA	72
Rules and Notes	72
Generate CSR	73
Upload CA Certificate	73
Install GigaVUE-FM on VMware ESXi	74
Before You Install	74
Prerequisites for GigaVUE-FM	74
VMware ESXiHardware Requirements	74
Recommended Resource Requirements for Scaled Environments	75
For Standalone GigaVUE-FM	75
For GigaVUE-FM High Availability Instances	76
Rules, Notes, and Limitations	77
Reference Values for Disk Space Estimation	77
Recommended Form Factor for VMware vCenter (Instance Types)	78
Deploy OVA file on VMware vCenter	78
Troubleshooting	81
Install GigaVUE-FM using ISO Image File on VMware ESXi Host	82
Prerequisites	82
Installation Steps	82
What to do Next	84
Install GigaVUE-FM using ISO Image File on VMware vCenter	84
Initial GigaVUE-FM Configuration	85
Perform Initial Configuration	86
Perform Network Configurations	86
Add Additional vNIC, vCPUs, and Increase Memory	87

Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi	88
Important GigaVUE-FM Upgrade Guidance	94
GigaVUE-FM Upgrade Guidance for GigaVUE Cloud Suite	94
GigaVUE-FM Upgrade Guidance for Application Intelligence	95
Guidance for FMHA when Encryption Mode is Certificate based Authentication	95
Deploy	96
Deploy GigaVUE Cloud Suite for VMware ESXi using GigaVUE FM	
Orchestration	96
Acquire Traffic using Platform Tapping	96
Install GigaVUE-FM on VMware ESXi	96
Required VMware Virtual Center Privileges	97
Default Login Credentials	98
Create Monitoring Domain for VMware ESXi	98
Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM	102
Verify Your Deployment	109
Acquire Traffic using Customer Orchestrated Source	109
Install GigaVUE-FM on VMware ESXi	109
Required VMware Virtual Center Privileges	109
Default Login Credentials	111
Create Monitoring Domain for VMware ESXi	111
Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM	114
Verify Your Deployment	121
Deploy GigaVUE Cloud Suite for VMware ESXi using Third Party	
Orchestration - Generic Mode	121
Acquire Traffic using UCT-V - By Creating Monitoring Domain	121
Install GigaVUE-FM on VMware ESXi	122
Configure Role-Based Access for Third Party Orchestration	122
Configure Tokens	124
Create a Monitoring Domain	124
Install and Register UCT-V Controller	126
Install UCT-Vs	128
Prerequisites	128
Linux Network Firewall Requirements	129
Install Linux UCT-Vs	129
Install Linux UCT-Vs using Installation Script	129
Install Linux UCT-Vs using Manual Configuration	132
Install UCT-V from Ubuntu/Debian Package	132
Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS	134
Windows Network Firewall Requirements	136
Install Windows UCT-Vs	136
Install Windows UCT-Vs using Installation Script	137

Install Windows UCT-Vs using Manual Configuration	140
Configure and Deploy V Series Node using VMware ESXi - Generic Mode	142
Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain	148
Install GigaVUE-FM on VMware ESXi	148
Configure Role-Based Access for Third Party Orchestration	148
Configure Tokens	150
Create a Monitoring Domain	150
Configure and Deploy GigaVUE Fabric Components using VMware ESXi - Generic Mode	152
Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform	159
License	159
Supportability	159
Supported Deployment Models	159
Rules and notes	160
Prerequisites	160
Get Started	162
Install Terraform	162
Deployment Using Terraform	164
a name potraDeploy GigaVUE Cloud Suite for VMware ESXi using Terraform	165
Deploy Third Party Orchestration by using Terraform	173
Best Practices	183
Verification and Troubleshooting	185
Verify the deployment	185
Troubleshoot the deployment	185
Configure	186
Configure Secure Tunnel for Third Party Orchestration	186
Precrypted Traffic	186
Mirrored Traffic	187
Prerequisites	187
Notes	187
Configure Secure Tunnel from UCT-V to GigaVUE V Series Node	187
Configure Secure Tunnel between GigaVUE V Series Nodes	188
Create Prefiltering Policy Template	191
Create Precryption Template for UCT-V	193
Rules and Notes:	193
Create Precryption Template for Filtering based on Applications	193
Create Precryption Template for Filtering based on L3-L4 details	194
Configure Monitoring Session	198
Create a Monitoring Session (VMware)	199
Monitoring Session Page (VMware)	200

Configure Monitoring Session Options (VMware ESXi)	201
Create Ingress and Egress Tunnel (VMware vCenter)	203
Create Raw Endpoint (VMware vCenter)	211
Rules and Notes	211
Configure Raw Endpoint in Monitoring Session	212
Create a New Map (VMware ESXi)	213
Example- Create a New Map using Inclusion and Exclusion Maps	220
Map Library	220
Deploy Monitoring Session	221
V Series Applications	224
Masking	229
Slicing	231
De-duplication	233
Feature Overview	233
Configure De-duplication Application	234
Distributed De-duplication	235
Application Intelligence	237
Points to Note for Application Intelligence	237
Application Visualization	239
User Defined Application	243
Application Filtering Intelligence	250
Application Metadata Intelligence	256
Extend Your AMI and AMX Deployment	264
Application Metadata Exporter	265
Hardware (AMI)	266
Hardware (Control Plane Metadata)	267
Private Cloud (VMware)	268
Rules and Notes	276
Access GigaVUE-FM	280
Create a Custom User	281
Create a Custom Role and Assign the User	281
Create a Custom User Group and Assign the Custom Role and the User	282
Verify GigaVUE V Series Node Registration	286
Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka	287
Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	288
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	289
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	289
Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)	290

Prerequisites	290
Integrate ZPA with Gigamon AMX	291
Configure Ingestor for AMX	292
Set up the Third Party Tool for AMX Integration	293
Configure Exporter to Export Attributes to Third Party tools	293
Configure Cloud Tool Exporter	294
Deploy the Monitoring Session	294
Cloud Tool Exporter - Field References	294
Configure Kafka Exporter	295
Deploy the Monitoring Session	295
Kafka Exporter - Field References	296
Verify Your Deployment	297
Prerequisites	300
Installation Steps	300
Access GigaVUE-FM	301
Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka	311
Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	311
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	312
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	313
Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)	313
Prerequisites	313
Integrate ZPA with Gigamon AMX	314
Configure Ingestor for AMX	315
Set up the Third Party Tool for AMX Integration	316
Configure Exporter to Export Attributes to Third Party tools	317
Configure Cloud Tool Exporter	317
Deploy the Monitoring Session	317
Cloud Tool Exporter - Field References	317
Configure Kafka Exporter	318
Deploy the Monitoring Session	318
Kafka Exporter - Field References	319
Prerequisite:	321
Troubleshoot AMI/AMX Deployment	341
GigaSMART NetFlow Generation	344
Examples- Configuring Application Intelligence Solution with Other Applications	347
Migrate Application Intelligence Session to Monitoring Session	349
Header Stripping	352
Limitation	353

Configure Header Stripping Application	353
Create Custom Port Template	355
Load Balancing	356
Enhanced Load Balancing	358
What's Next	362
SSL Decrypt	362
Supported Protocols, Algorithms, and Ciphers for SSL Decrypt	363
Configure SSL Decrypt	365
5G Service Based Interface Application	370
Ericsson vTAP	371
How SBI Application works	373
Supported Platforms:	374
End-to End Solution Overview and Configuration:	375
Recommended VM Specifications	383
FHA Dashboards for 5G Applications	383
5G Cloud Application	383
Supported Platforms	384
5G-Cloud Casa vTAP Support	384
5G-Cloud Oracle and Nokia SCP Support	392
Adding 5G-Cloud Application in Oracle SCP/Nokia SCP	397
5G-Cloud Ericsson SCP Support	405
Add 5G-Cloud Application in Ericsson SCP	410
5G Cloud Nokia HEP3 Support	415
Reference - Configuration Table	419
Add CSV file for IP Mapping	424
Nokia CMM Solution	429
Download Logs - 5G-Cloud / GvHTTP2 Application	438
Add CSV file for IP Mapping	438
Monitor and Troubleshoot	439
Monitor Cloud Health	439
Configuration Health Monitoring	439
Traffic Health Monitoring	440
Supported Resources and Metrics	441
Create Threshold Templates	443
Apply Threshold Template	444
Clear Thresholds	445
View Health Status	446
View Health Status of an Application	446
View Operational Health Status of an Application	446
View Health Status for Individual GigaVUE V Series Nodes	448
View Application Health Status for Individual V Series Nodes	449
View Monitoring Session Statistics	449
Visualize the Network Topology (VMware ESXi)	451

Administer GigaVUE Cloud Suite for VMware ESXi	453
Manage Monitoring Domain	453
Configure VMware Settings	454
Configure Certificate Settings	455
High Availability for Cloud Infrastructure	456
About GigaVUE-FM High Availability	456
Supportability Information	459
Recommendations, Rules, Notes, and Limitations	459
Licensing Information	462
Hostname Setups	463
Configure GigaVUE-FM HA	463
Configure Secure Communication between Fabric Components in FMHA	471
GigaVUE-FM HA Landing Page	472
Remove Standby GigaVUE-FM Instance	472
Disassemble GigaVUE-FM High Availability Group	473
GigaVUE-FM High Availability States	473
Fail over Mechanism	475
GigaVUE-FM High Availability Scenarios	476
Troubleshoot GigaVUE-FM High Availability Issues	476
Analytics for Virtual Resources	477
Virtual Inventory Statistics and Cloud Applications Dashboard	477
Debuggability and Troubleshooting	482
Sysdumps	482
Sysdumps—Rules and Notes	482
Generate a Sysdump File	482
FAQs - Secure Communication between GigaVUE Fabric Components (VMware-VSeries-ESXi)	483

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

This document highlights how to deploy GigaVUE V Series Node on the VMware ESXi platforms and set up traffic monitoring sessions using these nodes.


The GigaVUE V Series Nodes provide traffic visibility on the following VMware network types:

- vSphere Standard Switch
- vSphere Distributed Switch

Based on your setup, GigaVUE-FM creates, updates, and deletes the GigaVUE V Series Nodes on the ESXi hosts. Each GigaVUE V Series must run on the same ESXi host as the virtual machines (VMs) it monitors. It receives mirrored traffic directly from those VMs.

You can deploy only one GigaVUE V Series Node on a single ESXi host. GigaVUE-FM communicates directly with the GigaVUE V Series Nodes.

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



GigaVUE Fabric Components and Deployment Options

You can deploy GigaVUE fabric components in a VMware ESXi environment to gain visibility into traffic from virtual machines, with GigaVUE-FM providing centralized orchestration and management through VMware vCenter integration.

For VMware ESXi, the primary visibility component is the GigaVUE V Series Node. It collects mirrored traffic, applies filtering and packet processing, and forwards the optimized traffic to your monitoring and security tools. VMware ESXi deployments support visibility on vSphere Standard Switch and vSphere Distributed Switch networks.

You can deploy fabric components for VMware ESXi by using either GigaVUE-FM to orchestrate the deployment through VMware vCenter, or by deploying and registering GigaVUE V Series Nodes directly in VMware ESXi as a third-party-orchestrated workflow. Refer to the following sections for more information on deployment of GigaVUE Fabric Components in VMware ESXi.

GigaVUE-FM Orchestrated Fabric Deployment

Use this approach when you want GigaVUE-FM to manage deployment through VMware vCenter. The standard workflow is to upload the GigaVUE V Series Node image into GigaVUE-FM, create a monitoring domain for VMware ESXi, configure the GigaVUE V

Series Nodes, and then rediscover the environment.

When you create the monitoring domain in GigaVUE-FM, you define the connection to VMware vCenter and specify settings such as the traffic acquisition method, resource allocation, and the maximum number of V Series Nodes per host. Refer to GigaVUE-FM Orchestrated Fabric Deployment topic for more information on fabric components supported in VMware ESXi.

Third Party Orchestrated Fabric Deployment

Use this approach when you want to deploy the fabric components directly from the VMware ESXi environment and then register them with GigaVUE-FM for centralized visibility management. VMware ESXi third-party deployment supports multiple connectivity models, including single-uplink and separate-uplink deployments, depending on your environment and traffic-handling requirements. Refer to Third Party Orchestrated Fabric Deployment topic for more information on fabric components supported in VMware ESXi.

Benefits:

- Uses your existing orchestration workflow to deploy fabric components.
- Retains centralized management in GigaVUE-FM after deployment, including monitoring sessions and related services.
- Supports VMware ESXi private-cloud deployments, where ESXi is supported via third-party orchestration only.
- Enables secure tunnel support for ESXi, which is available only in the third-party orchestration model.
- Provides more secure registration and API access through tokens, without exposing passwords.

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

Third Party Orchestrated Fabric Deployment

You can deploy GigaVUE fabric components in a VMware ESXi environment using a third-party orchestration system instead of launching them directly through GigaVUE-FM. Once deployed, these components can register with GigaVUE-FM, which then acts as the central interface for configuring monitoring sessions and associated services.

The following table provides a brief description of the components that you can deploy using the Third Party orchestration:

Component	Description
GigaVUE-FM fabric manager (GigaVUE-FM)	GigaVUE-FM is a web-based fabric management and orchestration interface that provides a single pane of glass visibility, management, and orchestration of both the physical and virtual traffic that form the GigaVUE Cloud Suite Cloud. You are responsible for launching GigaVUE-FM from your end on the supported cloud or enterprise platforms.
UCT-V Controller (earlier known as G-vTAP Controller)	- Manages multiple UCT-Vs and orchestrates the mirrored traffic flow to GigaVUE V Series nodes. - GigaVUE-FM uses one or more UCT-V Controllers to communicate with the UCT-Vs.
GigaVUE® V Series Node	GigaVUE® V Series Node is a visibility node that: - Aggregates mirrored traffic from multiple UCT-Vs. - Applies filters and manipulates the packets using GigaSMART applications.

Component	Description
	Distributes the optimized traffic to cloud-based tools or backhauls it to GigaVUE Cloud Suite Cloud using GRE or VXLAN tunnels, when the platform supports.

Deployment Mode for VMware ESXi

Using third party orchestration, you can deploy the GigaVUE fabric components in Generic Mode.

In Generic Mode, you can specify the monitoring domain and connection name directly within your orchestrator during deployment.

You can choose one of these options:

- Allow the orchestrator to automatically create a monitoring domain under the Third Party Orchestration monitoring domain page. The GigaVUE fabric components are then deployed within this domain.
- Manually create a monitoring domain under Third Party Orchestration, and provide the monitoring domain name along with the connection name in the orchestrator's user data.

In generic mode, platform credentials are not shared with GigaVUE-FM. For detailed information about configuring and deploying fabric components using different traffic acquisition methods, refer to [Deploy GigaVUE Cloud Suite for VMware ESXi using Third Party Orchestration](#).

Key Concepts

This section provides the key concepts required to understand GigaVUE Cloud Suite for VMware (ESXi):


- [Secure Communication between GigaVUE Fabric Components](#)
- [Rediscover](#)
- [Sharing the Same Host across Different Monitoring Domains](#)
- [Selective Source Selection](#)
- [Secure Tunnels](#)
- [Precryption](#)
- [Prefiltering](#)

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](#)

- 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.

Rediscover

Sometimes, changes made to a GigaVUE V Series Node in VMware vCenter can cause its settings to go out of sync with what's shown in GigaVUE-FM. When this happens, you can use the Rediscover button in GigaVUE-FM to fix the mismatch:

Rediscovery updates the following settings:

- GigaVUE V Series Node name
- Datastore
- Management IP address
- Tunnel IP address
- Network name for Management Interface
- Network name for Tunnel Interface

NOTE: GigaVUE-FM automatically performs rediscovery every 24 hours to keep settings in sync.

To manually rediscover a node:

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, and select **Monitoring Domain**. The **Monitoring Domain** page appears.
2. Select **Actions > Rediscover**.

NOTE: You can only rediscover GigaVUE V Series Nodes that are in the **OK** state.

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.

Selective Source Selection

This feature enables you to select an individual network adapter of a virtual machine as a target while creating a map.

For details, refer to [Create a New Map](#).


Secure Tunnels

Secure Tunnel transfers the cloud captured packets from one GigaVUE V Series Node to another.

When sending traffic between two V Series Nodes, the source node captures and encapsulates the packets in PCAPng format. It then sends them to the destination V Series Node that decapsulates and processes the traffic based on your configuration.

How Does it Work!

- The source V Series forwards packets to the destination V series for further analysis.
- The destination node applies processing features such as de-duplication, application intelligence, or load balancing.
- The built-in load balancer distributes traffic across multiple V Series Nodes.
- If the load balancer sends packets to another node, it can re-encapsulate them and send them over another secure tunnel.



Supported Platforms

Secure Tunnels are supported on:

- OpenStack
- Azure
- AWS
- VMware NSX-T (only for Third Party Orchestration)
- VMware ESXi (only for Third Party Orchestration)
- Nutanix (only for Third Party Orchestration)
- Google Cloud Platform (only for Third Party Orchestration)

Precryption

License: Precryption requires a **SecureVUE Plus** license.


Gigamon Precryption™ technology¹ provides you clear-text visibility into encrypted network traffic without the need for traditional decryption. It works across virtual, cloud, and container environments, helping you get the full security stack without added complexity

In this section:

- [How Gigamon Precryption Technology Works](#)
- [Why Gigamon Precryption](#)
- [Key Features](#)
- [Key Benefits](#)
- [Precryption Technology on Single Node](#)
- [Precryption Technology on Multi-Node](#)
- [Supported Platforms](#)
- [Prerequisites](#)

How Gigamon Precryption Technology Works

Precryption technology leverages built-in Linux functionality to copy communications between the application and the encryption library, such as OpenSSL.



Key Highlights

¹**Disclaimer:** The Precryption feature allows you to capture decrypted traffic from both virtual machine (VM) and container-based environments. After capturing the traffic using (via UCT-C or UCT-V), you can send to the V Series product for further processing. You can choose to secure this traffic using encrypted tunnels between the capture point and the V Series. This option helps protect sensitive data during transit. If you don't enable encrypted tunnels, the captured (decrypted) traffic remains in plain text while moving between the source and the V Series—introducing potential exposure risks. Please note that the feature behavior and security options may change over time. Stay informed about updates to ensure you use the latest protections. By using this feature, you acknowledge and accept the current limitations and potential risks associated with the transmission of decrypted traffic.

- Captures network traffic in plain text, either before the system encrypts it or after it decrypts it.
- Does not change how encryption or transmission works.
- Avoids proxies, retransmissions, and “break-and-inspect” steps. Instead, it sends the plaintext copy to the Gigamon Deep Observability Pipeline, where tools can optimize, transform, and forward the traffic as needed.
- Runs on GigaVUE® Universal Cloud Tap (UCT) and supports hybrid and multi-cloud environments, including on-prem and virtual platforms.
- Runs independently of your applications, so you don't need to change your development lifecycle.

Why Gigamon Precryption

GigaVUE Universal Cloud Tap with Precryption technology is a lightweight, friction-free solution that eliminates blind spots present in modern hybrid cloud infrastructure.

Precryption helps you:

- Improve visibility for East-West traffic into virtual, cloud, and container platforms
- Delivers unobscured visibility into all encryption types, including TLS 1.3, without managing and maintaining decryption keys.
- Manages compliance with IT organizations, keeps communications private, architects a Zero Trust foundation, and boosts security-tool effectiveness by a factor of 5x or more.

Key Features

The following are the key features of this technology:

- Plain text visibility into communications with modern encryption (TLS 1.3, mTLS, and TLS 1.2 with Perfect Forward Secrecy).
- Plain text visibility into communications with legacy encryption (TLS 1.2 and earlier).
- Non-intrusive traffic access without agents running inside container workloads.
- Elimination of expensive resource consumption associated with traditional traffic decryption.
- Elimination of key management required by traditional traffic decryption.
- Zero performance impact based on cipher type, strength, or version.
- Support across hybrid and multi-cloud environments, including on-prem, virtual, and container platforms.
- Keep private communications private across the network with plaintext threat activity delivered to security tools.
- Integration with Gigamon Deep Observability Pipeline for the full suite of optimization, transformation, and brokering capabilities.

Key Benefits

The following are the key benefits of this technology:

- Eliminates blind spots for encrypted East-West (lateral) and North-South communications, including traffic that may not cross firewalls.
- Monitors application communications with an independent approach that enhances development team velocity.
- Extends security tools' visibility to all communications, regardless of encryption type.
- Achieves maximum traffic tapping efficiency across virtual environments.
- Leverages a 5–7x performance boost for security tools by consuming unencrypted data.
- Supports a Zero Trust architecture founded on deep observability.
- Maintains privacy and compliance adherence associated with decrypted traffic management.

How Gigamon Precryption Technology Works

This section explains how Precryption technology works on single nodes and multiple nodes in the following sections:


- [Precryption Technology on Single Node](#)
- [Precryption Technology on Multi-Node](#)

Precryption Technology on Single Node



1. An application uses an encryption library, such as OpenSSL, to encrypt a message.
2. GigaVUE Universal Cloud Tap (UCT), enabled with Precryption technology, gets a copy of this message before it is encrypted on the network.
3. The encrypted message is sent to the receiving application with unmodified encryption—no proxy, no re-encryption, no retransmissions.
4. GigaVUE UCT creates packet headers as needed, encapsulates them in a tunnel, and forwards them to GigaVUE V Series in the deep observability pipeline.
5. Gigamon optimizes, transforms, and delivers data to tools without further decryption.

Precryption Technology on Multi-Node



1. An application uses an encryption library, such as OpenSSL, to encrypt a message.
2. GigaVUE Universal Cloud Tap (UCT), enabled with Precryption technology, gets a copy of this message before it is encrypted on the network
3. Optionally, GigaVUE UCT enabled with Precryption can also acquire a copy of the message from the server end after the decryption.
4. GigaVUE UCT creates packet headers as needed, encapsulates them in a tunnel, and forwards them to GigaVUE V Series in the deep observability pipeline.
5. Gigamon optimizes, transforms, and delivers data to tools without further decryption.

Supported Platforms

VM environments: Precryption™ is supported on the following VM platforms that support UCT-V:

Platform Type	Platform
Public Cloud	• AWS • Azure • GCP (via Third Party Orchestration)
Private Cloud	• OpenStack

Platform Type	Platform
	- VMware ESXi (via Third Party Orchestration only) - VMware NSX-T (via Third Party Orchestration only) - Nutanix (via Third Party Orchestration only)

Container environments: Precryption™ is supported on the following container platforms that support UCT-C:

Platform Type	Platform
Public Cloud	- EKS - AKS - GKE
Private Cloud	- OpenShift - Native Kubernetes (VMware)

Prerequisites

Points to Note

- OpenSSL version 1.0.2, version 1.1.0, version 1.1.1, and version 3.x.
- For UCT-C, worker pods should always have libssl installed to ensure that UCT-C Tap can tap the Precryption packets from the worker pods whenever libssl calls are made from the worker pods.
- For GigaVUE-FM, add port 5671 in the security group to capture the statistics.
- In security group settings on the UCT-V Controller, enable Port 9900 to receive the statistics information from UCT-V.
- For UCT-C, add port 42042 and port 5671 to the security group.
- Precryption works only on Linux systems running Kernel version 4.18 or later.

License Prerequisite

- Precryption™ requires a SecureVUE Plus license.

Supported Kernel Version

Precryption is supported on kernel versions 4.18 and above, including 5.4+ kernels, across all Linux and Ubuntu operating systems. For the Kernel versions below 5.4, refer to the following table:

Kernel-Version	Operating System
4.18.0-193.el8.x86_64	RHEL release 8.2 (Ootpa)
4.18.0-240.el8.x86_64	RHEL release 8.3 (Ootpa)
4.18.0-305.76.1.el8_4.x86_64	RHEL release 8.4 (Ootpa)

Kernel-Version	Operating System
4.18.0-348.12.2.el8_5.x86_64	RHEL release 8.5 (Ootpa)
4.18.0-372.9.1.el8.x86_64	RHEL release 8.6 (Ootpa)
4.18.0-423.el8.x86_64	RHEL release 8.7 Beta (Ootpa)
4.18.0-477.15.1.el8_8.x86_64	RHEL release 8.8 (Ootpa)
5.3.0-1024-kvm	Ubuntu 19.10
4.18.0-305.3.1	Rocky Linux 8.4
4.18.0-348	Rocky Linux 8.5
4.18.0-372.9.1	Rocky Linux 8.6
4.18.0-425.10.1	Rocky Linux 8.7
4.18.0-477.10.1	Rocky Linux 8.8
4.18.0-80.el8.x86_64	CentOS 8.2
4.18.0-240.1.1.el8_3.x86_64	CentOS 8.3
4.18.0-305.3.1.el8_4.x86_64	CentOS 8.4
4.18.0-408.el8.x86_64	CentOS 8.5

For more details, refer to [Gigamon TV](#).

Prefiltering

Prefiltering allows you to filter the traffic at UCT-Vs before sending it to the GigaVUE V Series Nodes.

Create a Prefiltering policy template

In GigaVUE-FM, you can create a prefiltering policy template that includes rules and filter values.

- You can apply the template to multiple monitoring sessions.
- A monitoring session can use only one template at a time.
- You can define a maximum of 16 rules in a monitoring session.

Editing a Monitoring Session

While editing a session, you can modify the policy template with required rules and filter values. These changes apply only to that session. GigaVUE-FM does not save these changes to the template.


Guidelines for prefiltering in Next Generation UCT-Vs are:

- Prefiltering is supported on Next Generation UCT-Vs. It is not available in the classic mirroring mechanism.
- Both Linux and Windows UCT-Vs support Prefiltering.
- A monitoring session supports only one prefiltering policy. All agents in the session use the same policy.
- If multiple monitoring sessions share the same agent, and if one session does not use prefiltering, the system disables prefiltering. It applies a PassAll policy instead and forwards all traffic without filtering.

For details, refer to [Configure Precryption in UCT-V](#)

Deployment Overview

GigaVUE Cloud Suite for VMware ESXi 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 ESXi deployment. The top-level deployment activities are:



Deployment Planning

Before deploying the GigaVUE Cloud Suite for VMware ESXi, review the following planning considerations to ensure a smooth and successful installation.

- [Identify the Deployment Method](#)
- [Identify the vCenter Environment](#)
- [Define GigaVUE-FM Deployment Location](#)
- [Identify the Traffic Acquisition Method](#)

- [Identify the VMware Hosts and Clusters](#)
- [Identify the Network Configuration](#)
- [Identify the Required vCenter Privileges](#)

Identify the Deployment Method

GigaVUE Cloud Suite for VMware ESXi supports two options for deploying the fabric components:

- GigaVUE-FM Orchestration
- Third Party Orchestration

In either scenario, you need to first install and configure GigaVUE-FM.

GigaVUE-FM Orchestration

In this method, you use the GigaVUE-FM user interface to deploy V Series Nodes. After launching GigaVUE-FM and applying the required license, you create a Monitoring Domain and connect it to VMware vCenter. GigaVUE-FM then deploys GigaVUE V Series Nodes onto the selected VMware ESXi hosts. Once deployed, GigaVUE-FM monitors the state of the deployed fabric components and automatically replaces or relaunches them if any error conditions occur. This method requires that GigaVUE-FM have network connectivity to the vCenter server and that the vCenter user account has the minimum required privileges.

Third Party Orchestration

In this method, deploy the V Series Nodes directly on the VMware ESXi host by deploying an OVF or OVA file — without using the GigaVUE-FM deploy workflow. After deployment, the V Series Node registers itself with GigaVUE-FM using the Monitoring Domain name (GroupName), connection name (SubGroupName), and a JWT token.

NOTE: When registering GigaVUE V Series Nodes in GigaVUE-FM, ensure that the connection name under each Monitoring Domain is unique.

Identify the vCenter Environment

GigaVUE Cloud Suite for VMware ESXi connects to VMware vCenter to discover ESXi hosts and deploy fabric components. Before deployment, identify the following:

- The IP address or FQDN of the vCenter server to which GigaVUE-FM will connect.
- The VMware Datacenter and Cluster where GigaVUE V Series Nodes will be deployed.
- Whether the vCenter uses certificates from a trusted Certificate Authority (CA). If so, enable the Trust Store in GigaVUE-FM to validate vCenter certificates.

NOTE: Each Monitoring Domain corresponds to a single vCenter connection. To manage multiple vCenter environments, create a separate Monitoring Domain for each.

Define GigaVUE-FM Deployment Location

You can deploy GigaVUE-FM on a VMware ESXi host or on a physical or virtual server external to the VMware environment.

Points to Note:

- Ensure that GigaVUE-FM must have network connectivity to the vCenter server (IP or FQDN) and to the management interfaces of deployed V Series Nodes.
- If GigaVUE-FM is deployed inside the VMware environment, it can use the internal network to communicate with fabric components.
- If GigaVUE-FM is deployed outside the VMware environment (for example, on-premises), ensure that the required firewall ports are open between GigaVUE-FM and the vCenter server, and between GigaVUE-FM and each ESXi host.

For firewall port requirements, refer to [Network Firewall Requirements for VMware ESXi](#).

Identify the Traffic Acquisition Method

GigaVUE Cloud Suite for VMware ESXi supports the following traffic acquisition methods:

- **Platform Tapping:** This is performed by GigaVUE-FM using VMware port mirroring.
- **Customer Orchestrated Source:** In this method, traffic is tunneled directly to the V Series Node using a tunnel or raw endpoint. This is used when you want to forward traffic from sources that are not natively tapped by GigaVUE-FM.

Identify the VMware Hosts and Clusters

Before deployment, identify the VMware ESXi hosts where V Series Nodes will be deployed: Identify the VMware datacenter and cluster that contains the target ESXi hosts and determine the number of V Series Nodes required per host, based on the number of vSwitches and workload VMs to be monitored. For GigaVUE-FM Orchestration, you can import host details from a .csv file or add host information manually in the VMware Fabric Launch Configuration page. Up to 250 hosts can be selected manually.

Identify the Network Configuration

For each ESXi host, identify the network port groups required for GigaVUE V Series Node connectivity. You can deploy the GigaVUE V Series Node using the following deployment modes based on your requirement.

- Single uplink for Management, Data and Tool connectivity.

- Single uplink for Management and Tool connectivity, and another uplink for Data connectivity.
- (Recommended) Separate uplinks (three) for Management, Data, and Tool connectivity.

For Customer Orchestrated Source deployments: You must configure two data interfaces on the V Series Node, one for the tunnel network (to receive traffic) and one for the tool network (to forward traffic).

Identify the Required vCenter Privileges

The vCenter user account used in the Monitoring Domain configuration must have the minimum required vCenter privileges. These privileges enable GigaVUE-FM to manage the virtual center user with the specified roles.

Deployment Prerequisites

- Ensure the ESXi hosts have the minimum vCPU and memory resources to host the GigaVUE V Series Nodes. For details, refer to [Recommended Form Factor for VMware vCenter \(Instance Types\)](#).
- Set the vCenter character encoding to UTF-8 to support internationalized characters in the VMware vCenter environment.
- Access to GigaVUE V Series Node device OVA image file. You can download from [Gigamon Customer Portal](#).
- Verify that all the target VMs have VMware guest tools or Open VM tools if you use IP-based filtering.
- Ensure that Port 8889 is available for GigaVUE-FM to access GigaVUE V Series Nodes.
- Ensure that TCP Port 443 is open between the GigaVUE-FM instance and the ESXi host to upload the OVA files. For details, refer to [Network Firewall Requirements](#)
- Ensure a minimum of **Intel Sandy Bridge** CPU compatibility for proper operation and optimal performance to enable **VMware EVC** in VMware vCenter. Recommended option: **Skylake** and above.

Refer to the following topics for more information:

- Configure Tokens
- Default Login Credentials
- Required VMware Virtual Center Privileges

Licensing for GigaVUE Cloud Suite for VMware ESXi

This section describes the licensing model for GigaVUE Cloud Suite for VMware ESXi. GigaVUE Cloud Suite uses volume-based licensing (VBL), available as monthly subscription licenses. Licenses are offered as bundles. The available base bundle types are CoreVUE, NetVUE, and SecureVUE Plus. A default trial license is provided after GigaVUE-FM is installed, and volume-based licenses can be activated and managed from the VBL page in GigaVUE-FM.

Refer to the following topics for more information:

- [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:

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	-	-	√

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

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.

Activate Volume-Based Licenses

To activate Volume-Based Licenses,

1. On the left navigation pane, select .
2. Go to **System > Licenses**.
3. From the top navigation bar, select the **VBL** from the **Activation** drop-down.
4. Select **Activate Licenses**. The **Activate License** page appears.
5. Select **IP Address** or **Hostname** to include this information. If you exclude the IP Address or Hostname, identify the chassis or GigaSMART card by its ID when activating.
6. Download the fabric inventory file that contains information about GigaVUE-FM.
7. Select **Next**. For details, refer to the [What is a Fabric Inventory File?](#)What is a Fabric Inventory File section in *GigaVUE Licensing Guide*
8. Select **Gigamon License Portal**.
9. On the portal, upload the Fabric Inventory file.
10. Select the required license and select **Activate**. A license key is provided.
11. Record the license key or keys.
12. Return to GigaVUE-FM and select **Choose File to** upload the file.

Manage Volume-Based Licenses

This section provides information on how to manage active and inactive Volume-Based Licenses in GigaVUE-FM.

View active Volume-Based License

To view active Volume-Based License (VBL):

1. On the left navigation pane, click .
2. Go to **System > Licenses**.
3. From the top navigation bar, select the **VBL** from the **Activation** drop-down list and click **Active**.

This page lists the following information about the active Volume-Based Licenses.

Field	Description
SKU	Unique identifier associated with the license.
Bundle	Bundle to which the license belongs to.
Volume	Total daily allowance volume.
Starts	License start date.
Ends	License end date.
Type	Type of license (Commercial, Trial, Lab, and other license types).
Activation ID	Activation ID.
Entitlement ID	Entitlement ID. Entitlement ID is the permission with which the acquired license can be activated online.
Reference ID	Reference ID.
Status	License status.

NOTE: The License Type and Activation ID are displayed by default in the Active tab in the VBL page.

To display the Entitlement ID field, select the column setting configuration option to enable the Entitlement ID field.

View Inactive Volume-Based License

To view inactive Volume-Based License (VBL):

1. On the left navigation pane, click .
2. Go to **System > Licenses**.
3. From the top navigation bar, select the **VBL** from the **Activation** drop-down and click **Inactive**.

This page lists the following information about the inactive Volume-Based Licenses.

Field	Description
SKU	Unique identifier associated with the license.
Bundle	Bundle to which the license belongs to.
Ends	License end date.
Deactivation Date	Date the license got deactivated.
Revocation Code	License revocation code.
Status	License status.

NOTE: The License Type, Activation ID and Entitlement ID fields are not displayed by default in the Inactive tab of VBL page. To display these fields, select the column setting configuration option and enable these fields.

Activate Volume-Based Licenses

To activate Volume-Based Licenses,

1. On the left navigation pane, select .
2. Go to **System > Licenses**.
3. From the top navigation bar, select the **VBL** from the **Activation** drop-down.
4. Select **Activate Licenses**. The **Activate License** page appears.
5. Select **IP Address** or **Hostname** to include this information. If you exclude the IP Address or Hostname, identify the chassis or GigaSMART card by its ID when activating.
6. Download the fabric inventory file that contains information about GigaVUE-FM.
7. Select **Next**. For details, refer to the What is a Fabric Inventory File section in *GigaVUE Licensing Guide*
8. Select **Gigamon License Portal**.
9. On the portal, upload the Fabric Inventory file.
10. Select the required license and select **Activate**. A license key is provided.
11. Record the license key or keys.
12. Return to GigaVUE-FM and select **Choose File to** upload the file.

NOTE: If a VBL is deactivated after a bundle upgrade, you cannot create or edit Monitoring Sessions that include applications from the deactivated VBL during the grace period. You should manually deactivate the upgraded license during the grace period to move the inactive lower bundle license back to active status.

For detailed information on dashboards and report generation for Volume-Based Licensing refer to the following table:

For details about:	Reference section	Guide
How to generate Volume-Based License reports	Generate VBL Usage Reports	GigaVUE Administration Guide
Volume-Based License report details	Volume Based License Usage Report	GigaVUE Administration Guide
Fabric Health Analytics dashboards for Volume-Based Licenses usage	Dashboards for Volume Based Licenses Usage	GigaVUE-FM User Guide

Points to Note (VMware vCenter)

1. **VMware vCenter Setup:** These steps assume that VMware vCenter is installed and configured. For details, refer to [VMware Documentation](#).
2. **Virtual Switch Configuration:** Ensure the source virtual machine and the monitoring tool are connected to different virtual standard switches. If both are on the same switch, traffic may loop back, causing issues.
3. **Handling Latency from NGFW with DPI:** If NextGen Firewall (NGFW) with Deep Packet Inspection (DPI) is enabled to inspect your east-west traffic, expect an increase in latency due to mirrored traffic. To avoid increased latency, perform one of the following:
 - Create an exception rule for the tunneled traffic (mirrored traffic from the GigaVUE V Series Nodes to the tool).
 - Configure a private VDS that can bypass the NGFW rules for this traffic.
4. **NSX N-VDS Support:** NSX Virtual Distributed Switch (N-VDS) based segments are not supported in **VMware vCenter** Monitoring Domain. N-VDS is supported only on NSX versions less than or equal to 3.0. For details, refer to [VMware Documentation](#). To run VMware NSX-T directly on a VDS, the VDS version must be at least 7.0 or above.

Supported Hypervisors for VMware

The following 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

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	8GB	1 Management interface, 1 Tunnel interface, and 8 vTAP interfaces
Medium	8GB	4 vCPU		
Large	16GB	8 vCPU		

NOTE: When running GigaSMART applications such as Application Visualization, Application Metadata, and Application Filtering, use the **Large instance size (8 vCPU and 16 GB RAM)**. The Large instance supports up to 200k flows on VMware. You can contact your account manager or Gigamon Technical Support to identify a suitable form factor.

Minimum Virtual Computing Requirements

When deploying GigaVUE V Series Nodes, your Virtual Distributed Switch (VDS) must have a minimum number of ports available in its port groups. The required number of ports depends on how many port groups you use in each deployment.

If using a Single Port Group:

- Each GigaVUE V Series Node needs 10 ports from that port group.
- So, the port group must support: $10 \times (\text{Number of GigaVUE V Series Nodes})$

If using Two Port Groups:

When using two port groups in one deployment (typically one for Management and one for Tunnel traffic):

- Management Port Group
Needs: 1 port per GigaVUE V Series Node
- Tunnel Port Group
Needs: 9 ports per GigaVUE V Series Node

Use these calculations to ensure your port groups are sized correctly before deploying.

Network Firewall Requirements

Network Firewall Requirements for GigaVUE V Series Node deployment

Reduced Control-Plane Ports

Starting in release 6.14.02, the number of ports required for control-plane communication is reduced from 12 ports to 3 ports (443, 5671, and 9902). The reduced control-plane ports decrease the number of ports you must open between GigaVUE-FM and other fabric components while preserving management, registration, certificate exchange, and operational communication.

This change applies to control-plane traffic only. It does not change any data-path ports.

Use the minimum versions required for the reduced three-port configuration.

Component	Version
GigaVUE-FM	6.14.02
GigaVUE V Series Node	6.14.01
GigaVUE V Series Proxy	6.14.01

**Important**

- Do not upgrade fabric components to version 6.14.01 unless GigaVUE-FM is running version 6.14.02 or later.
- Fabric components version 6.14.01 are incompatible with GigaVUE-FM 6.14.00 and 6.14.01.
- Deployments that use this unsupported version combination will not support reduced control-plane port communication and may not function correctly.

Choose the section that applies to your deployment scenario:

- [New Deployments](#) – Use this section if you are deploying a new environment with GigaVUE-FM 6.14.02 or later and all supported fabric components running version 6.14.01 or later. These deployments use only the reduced control-plane ports.
- [Existing Deployments](#) – Use this section if you are upgrading an existing deployment and want to migrate from legacy control-plane ports to the reduced control-plane ports. This section identifies the ports to add and remove for each component during migration.
- [Legacy Control-Plane Ports](#) – Use this section if GigaVUE-FM is running a version earlier than 6.14.02, if any fabric component is running a version earlier than 6.14.01, or if your deployment requires backward compatibility during a mixed-version upgrade.

New Deployments

For new deployments, configure only the reduced control-plane ports. You do not need to open any legacy control-plane ports. Use the firewall requirements in this section when deploying a new GigaVUE Cloud Suite environment.

The reduced configuration uses only the three ports listed below:

Direction	Port	Purpose
Inbound/Outbound	443	HTTPS/REST management and control-plane communication, including GigaVUE-FM UI/API access, component registration, and certificate-related requests.
Inbound	5671	RMQ/RA control-plane messaging and statistics or health updates between GigaVUE-FM and fabric components.
Outbound	9902	RPC and control-plane communication, especially between UCT-V Controllers and UCT-V agent.

Existing Deployments

Existing deployments continue to operate without any firewall changes after upgrading GigaVUE-FM and fabric components. Legacy control-plane ports remain supported for backward compatibility. You can continue to use the legacy control-plane ports or migrate to the reduced control-plane ports. For information about the legacy ports, refer to [Legacy Control-Plane Ports](#)

To migrate to the reduced control-plane ports, refer to [Migration to Reduced Control-Plane Ports](#)

Migration to Reduced Control-Plane Ports

If you are migrating your existing deployment from legacy control-plane ports to reduced control-plane ports, complete the following steps:

- Upgrade the deployment while keeping the legacy ports open.
- Verify that all upgraded components are in the supported versions.
- Open the required reduced control-plane ports for GigaVUE-FM, other fabric components, and the UCT-V Agent. Use the port requirements listed in [Firewall Rule Changes for Migrating to Reduced Control-Plane Ports](#).
- Use the **Reset** option to force the monitoring domain to reconnect by using the new ports. For more information, refer to the [Manage a Monitoring Domain](#).
- Verify that all applicable components are communicating over ports 443, 5671, and 9902.
- Remove the legacy ports from your firewall or security group rules after verification is complete.

NOTE: Do not close all legacy ports and enable only the new ports before upgrading. Doing so can disrupt existing connections and is not recommended.

Firewall Rule Changes for Migrating to Reduced Control-Plane Ports

The following tables summarize the firewall changes to make when migrating an existing deployment:

GigaVUE-FM

Action	Direction	Port
Remove	Inbound	9600
Remove	Outbound	8890
Remove	Outbound	8889
Remove	Outbound	80
Remove	Outbound	9990
Add	Inbound	443
Add	Inbound	5671
Add	Outbound	443

UCT-V Controller

Action	Direction	Port
Remove	Inbound	9900
Remove	Inbound	80

Remove	Inbound	8300
Remove	Inbound	8892
Remove	Outbound	9600
Remove	Outbound	8301
Add	Inbound	443
Add	Outbound	443
Add	Outbound	5671
Add	Outbound	9902

UCT-V

Action	Direction	Port
Remove	Inbound	8301
Remove	Outbound	8892
Remove	Outbound	9900
Remove	Outbound	8300
Add	Inbound	9902
Add	Outbound	5671
Add	Outbound	443

GigaVUE V Series Node

Action	Direction	Port
Remove	Inbound	8889
Remove	Inbound	80
Remove	Outbound	8891
Remove	Outbound	8892
Remove	Outbound	9600
Remove	Outbound	8300
Add	Inbound	443
Add	Outbound	443
Add	Outbound	5671

GigaVUE V Series Proxy (Optional)

Action	Direction	Port
Remove	Inbound	8890
Remove	Inbound	80
Remove	Inbound	8300
Remove	Inbound	8891
Remove	Inbound	8892

Remove	Outbound	8889
Remove	Outbound	9600
Remove	Outbound	80
Add	Inbound	5671
Add	Inbound	443
Add	Outbound	443

UCT-V OVS Controller

Action	Direction	Port
Remove	Inbound	9900
Add	Inbound	443
Add	Outbound	9901

UCT-V OVS Module

Action	Direction	Port
Add	Inbound	9901

The following tables identify the firewall changes required to migrate from the legacy control-plane ports to the optimized control-plane ports:

GigaVUE FM

The following table lists the reduced control-plane ports required for communication with GigaVUE-FM.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR2	Purpose
Inbound	TCP	443	9600	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.
Inbound	TCP	443	9600	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	8889	GigaVUE V Series Node IP	Allows GigaVUE-FM to communicate control and

					management plane traffic to GigaVUE V Series Node.
Outbound	TCP	443	8890	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to communicate control and management plane traffic to GigaVUE V Series Proxy.
Outbound	TCP	443	9900	UCT-V Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic with UCT-V Controller.



Removed: Outbound TCP 80 (ACME challenge) is no longer required. Certificate issuance and renewal for fabric nodes (6.14.01 or later) use JWK-based authentication instead of the ACME HTTP challenge.

UCT-V Controller

The following table lists the reduced control-plane ports required for communication with UCT-V Controller.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	9900	GigaVUE-FM IP	Allows UCT-V Controller to communicate control and management plane traffic with GigaVUE-FM.
Inbound	TCP	443	8300	UCT-VSubnet	Allows UCT-V Controller to receive the certificate requests from the UCT-V.
Inbound	TCP	443	8892	UCT-V Subnet	Allows UCT-V Controller to receive the registration requests and heartbeat from UCT-V.
Inbound	TCP	443	9900	UCT-V or Subnet IP	Allows UCT-V Controller to receive traffic health updates from UCT-V.

Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	5671	NA	UCT-V Subnet	Allows the UCT-V Controller to communicate control and management plane traffic with UCT-Vs for monitoring domains created in 6.14.01 and later.
Outbound	TCP	9902	NA	UCT-V Subnet	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs for UCT-Vs with version earlier than 6.14.01.



Removed: Inbound TCP 80 (ACME challenge from GigaVUE-FM), Outbound TCP 9600 (certificate request to GigaVUE-FM), and Outbound TCP 8301 (ACME validation to UCT-V) are no longer required for 6.14.01+ deployments.

NOTE: For UCT-V Controller-to-UCT-V traffic, the controller uses 5671 for monitoring domains created in 6.14.01 and later, and 9902 for monitoring domains created in earlier releases. The UCT-V RPC port is configurable at the monitoring-domain level.

UCT-V

The following table lists the reduced control-plane ports required for communication with UCT-V.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	5671	NA	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller for monitoring domains created in 6.14.01 and later.
Inbound	TCP	9902	NA	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller for existing monitoring domains created with version

					earlier than 6.14.01.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	9900	UCT-V Controller IP	Allows UCT-V to send traffic health updates to UCT-V Controller.
Outbound (This is the port used for Third Party Orchestration)	TCP	443	8892	UCT-V Controller IP	Allows UCT-V to receive the registration requests and heartbeat to UCT-V Controller.
Outbound	TCP	443	8300	UCT-V Controller IP	Allows UCT-V to receive the certificate requests from the UCT-V Controller.



Removed: Inbound TCP 8301 (ACME challenge) is no longer required for 6.14.01 or later.

NOTE: A 6.14.01 UCT-V listens on both 5671 and 9902 during discovery. The default listening port is 5671 for monitoring domains created in 6.14.01; monitoring domains migrated from earlier releases default to 9902. This port is configurable at the monitoring-domain level.

GigaVUE V Series Node

The following table lists the reduced control-plane ports required for communication with GigaVUE V Series Node.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8889	GigaVUE-FM IP	Allows GigaVUE V Series Node to communicate control and management plane traffic with GigaVUE-FM.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	5671	8300	GigaVUE V Series Proxy	Allows GigaVUE V Series Node to send certificate request to GigaVUE V Series Proxy IP.
Outbound	TCP	443	8891	GigaVUE V Series Proxy	Allows GigaVUE V Series Proxy to receive PKI requests from the

					GigaVUE V Series Node.
Outbound	TCP	443	8892	GigaVUE V Series Proxy	Allows GigaVUE V Series Proxy to receive registration requests and heartbeat messages from GigaVUE V Series Node.
Outbound	TCP	443	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.

 **Removed:** Inbound TCP 80 (ACME challenge) is no longer required for 6.14.01 or later.

NOTE: Cross-version deployments between GigaVUE V Series Nodes and GigaVUE V Series Proxy are not supported. Ensure that the GigaVUE V Series Node and GigaVUE V Series Proxy are running the same software version. For example, a 6.13 GigaVUE V Series Node cannot communicate with a 6.14 GigaVUE V Series Proxy.

GigaVUE V Series Proxy (Optional)

The following table lists the reduced control-plane ports required for communication with GigaVUE V Series Proxy.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8890	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic with GigaVUE V Series Proxy.
Inbound	TCP	443	8891	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive PKI requests from GigaVUE V Series Node.
Inbound	TCP	443	8892	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive registration requests and heartbeat messages from GigaVUE V Series Node.
Inbound	TCP	5671	8300	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive certificate requests from GigaVUE V Series Node for the configured params and provides

					the certificate using those parameters.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.



Removed: Inbound and outbound TCP 80 (ACME challenge) is no longer required for 6.14.01 or later.

UCT-V OVS Controller

The following table lists the reduced control-plane ports required for communication with UCT-V OVS Controller.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	9900	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic. The controller supports both ports for compatibility and falls back to 9900 when required.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	9901	NA	UCT-V OVS Module IP	Allows to communicate with UCT-V OVS Modules (RPC).

UCT-V OVS Module

The following table lists the reduced control-plane ports required for communication with UCT-V OVS Module.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	9901	NA	UCT-V OVS Controller IP	Allows to receive control communication from UCT-V OVS Controller. The module listens on 9901 for backward compatibility.

Legacy Control-Plane Ports

Legacy control-plane ports remain supported for backward compatibility with earlier versions and mixed-version deployments.

When to Use Legacy Ports

- GigaVUE-FM is running version 6.10 or later but earlier than 6.14.02.
- One or more fabric components (such as UCT-V Controller, UCT-V, GigaVUE V Series Node, GigaVUE V Series Proxy) are running a version earlier than 6.14.01.
- This includes mixed-version deployments where GigaVUE-FM 6.14.02 or later communicates with fabric components running N-1 or N-2 releases.
- The deployment has not yet been migrated to the reduced control-plane port architecture.

Source	Destination	Source Port	Destination Port	Protocol	Service	Purpose
GigaVUE-FM	ESXi hosts	Any (1024-65535)	443	TCP	https	Allows GigaVUE-FM to communicate with vCenter and all ESXi hosts to import the V Series OVA files. OVA files require access to the host IP/URL for bulk deployment
	vCenter					
GigaVUE-FM	GigaVUE V Series Nodes	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)	5671	TCP	Custom TCP	Allows GigaVUE-FM to receive the traffic health updates 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	Facilitates Management connection to GigaVUE-FM .
			22		ssh	
Administrator	GigaVUE V Series Nodes	Not Applicable	22		ssh	Facilitates troubleshooting GigaVUE V Series Nodes.
Remote Source	GigaVUE V Series Nodes	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 to the Tunnel Ingress option only)
			N/A	IP 47	GRE	
			4754	UDP	UDPGRE	
GigaVUE V Series Nodes	Tool/ GigaVUE HC Series instance	Custom Port (VXLAN),N/A for GRE	4789	UDP	VXLAN	Allows GigaVUE V Series Node to communicate and tunnel traffic to the Tool
			Not Applicable	IP 47	GRE	
GigaVUE V Series Nodes	Tool/ GigaVUE HC Series instance	Not Applicable	Not Applicable	ICMP	Echo Request	(Optional) Allows GigaVUE V Series Node to health check tunnel destination traffic
					Echo Response	
GigaVUE V Series Nodes	GigaVUE-FM	Any (1024-65535)	Any (1024-65535)	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.

Reduced Control-Plane Ports

Starting in release 6.14.02, the number of ports required for control-plane communication is reduced from 12 ports to 3 ports (443, 5671, and 9902). The reduced control-plane ports decrease the number of ports you must open between GigaVUE-FM and other fabric components while preserving management, registration, certificate exchange, and operational communication.

This change applies to control-plane traffic only. It does not change any data-path ports.

Use the minimum versions required for the reduced three-port configuration.

Component	Version
GigaVUE-FM	6.14.02
GigaVUE V Series Node	6.14.01
GigaVUE V Series Proxy	6.14.01
UCT-V	6.14.01
UCT-V Controller	6.14.01



Important

- Do not upgrade fabric components to version 6.14.01 unless GigaVUE-FM is running version 6.14.02 or later.
- Fabric components version 6.14.01 are incompatible with GigaVUE-FM 6.14.00 and 6.14.01.
- Deployments that use this unsupported version combination will not support reduced control-plane port communication and may not function correctly.

Choose the section that applies to your deployment scenario:

- [New Deployments](#) – Use this section if you are deploying a new environment with GigaVUE-FM 6.14.02 or later and all supported fabric components running version 6.14.01 or later. These deployments use only the reduced control-plane ports.
- [Existing Deployments](#) – Use this section if you are upgrading an existing deployment and want to migrate from legacy control-plane ports to the reduced control-plane ports. This section identifies the ports to add and remove for each component during migration.
- [Legacy Control-Plane Ports](#) – Use this section if GigaVUE-FM is running a version earlier than 6.14.02, if any fabric component is running a version earlier than 6.14.01, or if your deployment requires backward compatibility during a mixed-version upgrade.

New Deployments

For new deployments, configure only the reduced control-plane ports. You do not need to open any legacy control-plane ports. Use the firewall requirements in this section when deploying a new GigaVUE Cloud Suite environment.

The reduced configuration uses only the three ports listed below:

Direction	Port	Purpose
Inbound/Outbound	443	HTTPS/REST management and control-plane communication, including GigaVUE-FM UI/API access, component registration, and certificate-related requests.
Inbound/Outbound	5671	RMQ/RA control-plane messaging and statistics or health updates between GigaVUE-FM and fabric components.
Inbound/Outbound	9902	RPC and control-plane communication, especially between UCT-V Controllers and UCT-V agent.

Existing Deployments

Existing deployments continue to operate without any firewall changes after upgrading GigaVUE-FM and fabric components. Legacy control-plane ports remain supported for backward compatibility. You can continue to use the legacy control-plane ports or migrate to the reduced control-plane ports. For information about the legacy ports, refer to [Legacy Control-Plane Ports](#)

To migrate to the reduced control-plane ports, refer to [Migration to Reduced Control-Plane Ports](#)

Migration to Reduced Control-Plane Ports

If you are migrating your existing deployment from legacy control-plane ports to reduced control-plane ports, complete the following steps:

- Upgrade the deployment while keeping the legacy ports open.
- Verify that all upgraded components are in the supported versions.
- Open the required reduced control-plane ports for GigaVUE-FM, other fabric components, and the UCT-V Agent. Use the port requirements listed in [Firewall Rule Changes for Migrating to Reduced Control-Plane Ports](#).
- Use the **Reset** option to force the monitoring domain to reconnect by using the new ports. For more information, refer to the [Manage a Monitoring Domain](#).
- Verify that all applicable components are communicating over ports 443, 5671, and 9902.
- Remove the legacy ports from your firewall or security group rules after verification is complete.

NOTE: Do not close all legacy ports and enable only the new ports before upgrading. Doing so can disrupt existing connections and is not recommended.

Firewall Rule Changes for Migrating to Reduced Control-Plane Ports

The following tables summarize the firewall changes to make when migrating an existing deployment:

GigaVUE-FM

Action	Direction	Port
Remove	Inbound	9600
Remove	Outbound	8890
Remove	Outbound	8889
Remove	Outbound	80
Remove	Outbound	9990
Add	Inbound	443
Add	Inbound	5671
Add	Outbound	443

UCT-V Controller

Action	Direction	Port
Remove	Inbound	9900
Remove	Inbound	80
Remove	Inbound	8300
Remove	Inbound	8892
Remove	Outbound	9600
Remove	Outbound	8301
Add	Inbound	443
Add	Outbound	443
Add	Outbound	5671
Add	Outbound	9902

UCT-V

Action	Direction	Port
Remove	Inbound	8301
Remove	Outbound	8892
Remove	Outbound	9900
Remove	Outbound	8300
Add	Inbound	9902
Add	Outbound	5671
Add	Outbound	443

GigaVUE V Series Node

Action	Direction	Port
Remove	Inbound	8889
Remove	Inbound	80
Remove	Outbound	8891
Remove	Outbound	8892
Remove	Outbound	9600
Remove	Outbound	8300
Add	Inbound	443
Add	Outbound	443
Add	Outbound	5671

GigaVUE V Series Proxy (Optional)

Action	Direction	Port
Remove	Inbound	8890
Remove	Inbound	80
Remove	Inbound	8300
Remove	Inbound	8891
Remove	Inbound	8892
Remove	Outbound	8889
Remove	Outbound	9600
Remove	Outbound	80
Add	Inbound	5671
Add	Inbound	443
Add	Outbound	443

UCT-V OVS Controller

Action	Direction	Port
Remove	Inbound	9900
Add	Inbound	443
Add	Outbound	9901

UCT-V OVS Module

Action	Direction	Port
Add	Inbound	9901

The following tables identify the firewall changes required to migrate from the legacy control-plane ports to the optimized control-plane ports:

GigaVUE FM

The following table lists the reduced control-plane ports required for communication with GigaVUE-FM.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR2	Purpose
Inbound	TCP	443	9600	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.
Inbound	TCP	443	9600	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	8889	GigaVUE V Series Node IP	Allows GigaVUE-FM to communicate control and management plane traffic to GigaVUE V Series Node.
Outbound	TCP	443	8890	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to communicate control and management plane traffic to GigaVUE V Series Proxy.
Outbound	TCP	443	9900	UCT-V Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic with UCT-V Controller.



Removed: Outbound TCP 80 (ACME challenge) is no longer required. Certificate issuance and renewal for fabric nodes (6.14.01 or later) use JWK-based authentication instead of the ACME HTTP challenge.

UCT-V Controller

The following table lists the reduced control-plane ports required for communication with UCT-V Controller.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	9900	GigaVUE-FM IP	Allows UCT-V Controller to communicate control and management plane traffic with GigaVUE-FM.
Inbound	TCP	443	8300	UCT-VSubnet	Allows UCT-V Controller to receive the certificate requests from the UCT-V.
Inbound	TCP	443	8892	UCT-V Subnet	Allows UCT-V Controller to receive the registration requests and heartbeat from UCT-V.
Inbound	TCP	443	9900	UCT-V or Subnet IP	Allows UCT-V Controller to receive traffic health updates from UCT-V.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	5671	NA	UCT-V Subnet	Allows the UCT-V Controller to communicate control and management plane traffic with UCT-Vs for monitoring domains created in 6.14.01 and later.
Outbound	TCP	9902	NA	UCT-V Subnet	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs for UCT-Vs with version earlier than 6.14.01.



Removed: Inbound TCP 80 (ACME challenge from GigaVUE-FM), Outbound TCP 9600 (certificate request to GigaVUE-FM), and Outbound TCP 8301 (ACME validation to UCT-V) are no longer required for 6.14.01+ deployments.

NOTE: For UCT-V Controller-to-UCT-V traffic, the controller uses 5671 for monitoring domains created in 6.14.01 and later, and 9902 for monitoring domains created in earlier releases. The UCT-V RPC port is configurable at the monitoring-domain level.

UCT-V

The following table lists the reduced control-plane ports required for communication with UCT-V.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	5671	NA	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller for monitoring domains created in 6.14.01 and later.
Inbound	TCP	9902	NA	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller for existing monitoring domains created with version earlier than 6.14.01.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	9900	UCT-V Controller IP	Allows UCT-V to send traffic health updates to UCT-V Controller.
Outbound (This is the port used for Third Party Orchestration)	TCP	443	8892	UCT-V Controller IP	Allows UCT-V to receive the registration requests and heartbeat to UCT-V Controller.
Outbound	TCP	443	8300	UCT-V Controller IP	Allows UCT-V to receive the certificate requests from the UCT-V Controller.



Removed: Inbound TCP 8301 (ACME challenge) is no longer required for 6.14.01 or later.

NOTE: A 6.14.01 UCT-V listens on both 5671 and 9902 during discovery. The default listening port is 5671 for monitoring domains created in 6.14.01; monitoring domains migrated from earlier releases default to 9902. This port is configurable at the monitoring-domain level.

GigaVUE V Series Node

The following table lists the reduced control-plane ports required for communication with GigaVUE V Series Node.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8889	GigaVUE-FM IP	Allows GigaVUE V Series Node to communicate control and management plane traffic with GigaVUE-FM.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	5671	8300	GigaVUE V Series Proxy	Allows GigaVUE V Series Node to send certificate request to GigaVUE V Series Proxy IP.
Outbound	TCP	443	8891	GigaVUE V Series Proxy	Allows GigaVUE V Series Proxy to receive PKI requests from the GigaVUE V Series Node.
Outbound	TCP	443	8892	GigaVUE V Series Proxy	Allows GigaVUE V Series Proxy to receive registration requests and heartbeat messages from GigaVUE V Series Node.
Outbound	TCP	443	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.



Removed: Inbound TCP 80 (ACME challenge) is no longer required for 6.14.01 or later.

NOTE: Cross-version deployments between GigaVUE V Series Nodes and GigaVUE V Series Proxy are not supported. Ensure that the GigaVUE V Series Node and GigaVUE V Series Proxy are running the same software version. For example, a 6.13 GigaVUE V Series Node cannot communicate with a 6.14 GigaVUE V Series Proxy.

GigaVUE V Series Proxy (Optional)

The following table lists the reduced control-plane ports required for communication with GigaVUE V Series Proxy.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8890	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic with GigaVUE V Series Proxy.
Inbound	TCP	443	8891	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive PKI requests from GigaVUE V Series Node.
Inbound	TCP	443	8892	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive registration requests and heartbeat messages from GigaVUE V Series Node.
Inbound	TCP	5671	8300	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive certificate requests from GigaVUE V Series Node for the configured params and provides the certificate using those parameters.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.



Removed: Inbound and outbound TCP 80 (ACME challenge) is no longer required for 6.14.01 or later.

UCT-V OVS Controller

The following table lists the reduced control-plane ports required for communication with UCT-V OVS Controller.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	9900	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic. The controller supports both ports for compatibility and falls back to 9900 when required.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	9901	NA	UCT-V OVS Module IP	Allows to communicate with UCT-V OVS Modules (RPC).

UCT-V OVS Module

The following table lists the reduced control-plane ports required for communication with UCT-V OVS Module.

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	9901	NA	UCT-V OVS Controller IP	Allows to receive control communication from UCT-V OVS Controller. The module listens on 9901 for backward compatibility.

Legacy Control-Plane Ports

Legacy control-plane ports remain supported for backward compatibility with earlier versions and mixed-version deployments.

When to Use Legacy Ports

- GigaVUE-FM is running version 6.10 or later but earlier than 6.14.02.
- One or more fabric components (such as UCT-V Controller, UCT-V, GigaVUE V Series Node, GigaVUE V Series Proxy) are running a version earlier than 6.14.01.
- This includes mixed-version deployments where GigaVUE-FM 6.14.02 or later communicates with fabric components running N-1 or N-2 releases.
- The deployment has not yet been migrated to the reduced control-plane port architecture.

The following table lists the Network Firewall / Security Group requirements for GigaVUE Cloud Suite:

NOTE: When using dual stack network, open the below mentioned ports for both IPv4 and IPv6.

GigaVUE FM

The following table specifies the inbound and outbound communication parameters—protocols, ports, and CIDRs—required for GigaVUE-FM to support secure access, registration, certificate exchange, and control-plane communication with associated components.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	443	Administrator Subnet	Allows GigaVUE-FM to accept Management connection using REST API. Allows users to access GigaVUE-FM UI securely through an HTTPS connection.
Inbound	TCP	22	Administrator Subnet	Allows CLI access to user-initiated management and diagnostics.
Inbound (This is the port used for Third Party Orchestration)	TCP	443	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive registration requests from GigaVUE V Series Node using REST API when GigaVUE V Series Proxy is not used.
Inbound (This is the port used for Third Party Orchestration)	TCP	443	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to receive registration requests from GigaVUE V Series Proxy using REST API.
Inbound	TCP	5671	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive traffic health updates from GigaVUE V Series Nodes.
Inbound	TCP	9600	GigaVUE V Series Proxy	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.
Inbound	TCP	9600	GigaVUE V Series Node	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Node.

Inbound	UDP	2056	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive Application Intelligence and Application Visualization reports from GigaVUE V Series Node.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound (optional)	TCP	8890	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to communicate control and management plane traffic to GigaVUE V Series Proxy.
Outbound	TCP	8889	GigaVUE V Series Node IP	Allows GigaVUE-FM to communicate control and management plane traffic to GigaVUE V Series Node.
Outbound	TCP	80	GigaVUE V Series Node	Allows GigaVUE-FM to send ACME challenge requests to GigaVUE V Series Node.
Outbound	TCP	80	GigaVUE V Series Proxy	Allows GigaVUE-FM to send ACME challenge requests to GigaVUE V Series Proxy.
Outbound	TCP	443	Any IP Address	Allows GigaVUE-FM to reach the Public Cloud Platform APIs.

UCT-V

The following table outlines UCT-V Controller's network communication requirements with GigaVUE-FM, detailing essential ports, protocols, and CIDRs for registration, diagnostics, certificate exchange, and orchestration traffic.

GigaVUE V Series Node

The following table outlines GigaVUE V Series Node's network communication requirements, detailing protocols, ports, and CIDRs necessary for tunneling, management, diagnostics, and secure data transfer across connected components

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	8889	GigaVUE-FM IP	Allows GigaVUE V Series Node to communicate control and management plane traffic with GigaVUE-FM.

Inbound	TCP	8889	GigaVUE V Series Proxy IP	Allows GigaVUE V Series Node to communicate control and management plane traffic with GigaVUE V Series Proxy.
Inbound	UDPGRE	4754	Ingress Tunnel	Allows GigaVUE V Series Node to receive tunnel traffic from UDPGRE Tunnel.
Inbound	TCP	22	Administrator Subnet	Allows CLI access for user-initiated management and diagnostics, specifically when using third party orchestration.
Inbound	TCP	80	GigaVUE-FM	Allows GigaVUE V Series Node to receive the ACME challenge requests from GigaVUE-FM.
Inbound	TCP	80	GigaVUE V Series Proxy IP	Allows UCT-V to receive the ACME challenge requests from the GigaVUE V Series Proxy.
Inbound (Optional - This port is used only for configuring AWS Gateway Load Balancer)	UDP (GENEVE)	6081	Ingress Tunnel	Allows GigaVUE V Series Node to receive tunnel traffic from AWS Gateway Load Balancer.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	5671	GigaVUE-FM IP	Allows GigaVUE V Series Node to send traffic health updates to GigaVUE-FM.
Outbound	UDP (VXLAN)	VXLAN (default 4789)	Tool IP	Allows GigaVUE V Series Node to tunnel output to the tool.
Outbound	IP Protocol (L2GRE)	L2GRE (IP 47)	Tool IP	Allows GigaVUE V Series Node to tunnel output to the tool.
Outbound	UDP	2056	GigaVUE-FM IP	Allows GigaVUE V Series Node to send Application Intelligence and Application Visualization reports to GigaVUE-FM.
Outbound	UDP	2055	Tool IP	Allows GigaVUE V Series Node to send NetFlow Generation traffic to an external tool.
Outbound	UDP	8892	GigaVUE V Series Proxy	Allows GigaVUE V Series Node to send certificate request to GigaVUE V Series Proxy IP.
Outbound	TCP	514	Tool IP	Allows GigaVUE V Series Node to send Application Metadata

				Intelligence log messages to external tools.
Bidirectional (optional)	ICMP	- echo request - echo reply	Tool IP	Allows GigaVUE V Series Node to send health check tunnel destination traffic.
Outbound (This is the port used for Third Party Orchestration)	TCP	443	GigaVUE-FM IP	Allows GigaVUE V Series Node to send registration requests and heartbeat messages to GigaVUE-FM when GigaVUE V Series Proxy is not used.
Outbound (Optional - This port is used only for Secure Tunnels)	TCP	11443	Tool IP	Allows to securely transfer the traffic to an external tool.

Giga VUE V Series Proxy(Optional)

The following table defines GigaVUE V Series Proxy's network communication parameters, listing essential protocols, ports, and CIDRs for registration, certificate exchange, diagnostics, and control-plane traffic with GigaVUE-FM and V Series Nodes.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	8890	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic with GigaVUE V Series Proxy.
Inbound	TCP	22	Administrator Subnet	Allows CLI access for user-initiated management and diagnostics, specifically when using third party orchestration.
Inbound	TCP	80	GigaVUE-FM	Allows GigaVUE V Series Proxy to receive the ACME challenge requests from the GigaVUE-FM.
Inbound	TCP	8300	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive certificate requests from GigaVUE V Series Node for the configured params and provides the certificate using those parameters.
Inbound	TCP	8892	GigaVUE V Series Node IP	Allows GigaVUE V Series Proxy to receive registration requests and heartbeat messages from GigaVUE V Series Node.

Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	443	GigaVUE-FM IP	Allows GigaVUE V Series Proxy to communicate the registration requests to GigaVUE-FM.
Outbound	TCP	8889	GigaVUE V Series Node IP	Allows GigaVUE V Series Proxy to communicate control and management plane traffic with GigaVUE V Series Node.

Ports for Backward Compatibility

Ensure to open these ports for backward compatibility when GigaVUE-FM is running version 6.10 or later, and the fabric components are on (n-1) or (n-2) versions.

GigaVUE V Series Node

The following table specifies the outbound communication requirement for GigaVUE V Series Node, detailing the protocol, port, and source CIDR used to send registration and heartbeat messages to the GigaVUE V Series Proxy during third-party orchestration.

Direction	Protocol	Port	Source CIDR	Purpose
Outbound (This is the port used for Third Party Orchestration)	TCP	8891	GigaVUE V Series Proxy IP	Allows GigaVUE V Series Node to send registration requests and heartbeat messages to GigaVUE V Series Proxy when GigaVUE V Series Proxy is used.

GigaVUE V Series Proxy(Optional)

The following table specifies the optional inbound communication parameter for GigaVUE V Series Proxy, detailing the protocol, port, and source CIDR required to receive security parameter requests from GigaVUE V Series Node during third-party orchestration.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound (This is the port used for Third Party Orchestration)	TCP	8891	GigaVUE V Series Node IP	Allows GigaVUE V Series Proxy to receive security parameter requests from GigaVUE V Series Node.

Configure Tokens

You must configure tokens for registering GigaVUE Fabric Components using Third Party Orchestration and registering UCT-V with GigaVUE-FM.

This feature generates tokens to verify the identity of a user for accessing the GigaVUE-FM

If you are authenticated, GigaVUE-FM allows you to generate a token based on your access privileges. You can copy the generated tokens from the GUI to access the REST APIs. Token inherits the read or write Role-Based Access (RBAC) privileges of the groups assigned to its user.

GigaVUE-FM generates multiple tokens and associates them with the corresponding user groups. If you have write access to GigaVUE-FM Security Management, you can revoke other users' tokens but not view the created tokens.

Prerequisite

You must create user groups in GigaVUE-FM. For details, refer to [Configure Role-Based Access for Third Party Orchestration](#).

Rules and Notes

- Token authentication is an additional mechanism to access GigaVUE-FM REST APIs without replacing the existing GigaVUE-FM authentication mechanism.
- Only authenticated users can create tokens.
- A token expires or becomes invalid in the following cases:
 - The token reaches the configured expiry or default period. The default value is 30 days, and the maximum value is 105 days.
 - You delete a user group associated with the token.
 - You change the password of the local user who owns the token.
 - You change the authentication type. GigaVUE-FM deletes all tokens when this occurs.
- When you back up and restore of the GigaVUE-FM, previously generated tokens become unavailable.
- During an FMHA role changeover, active GigaVUE-FM tokens are active.
- For basic authentication, GigaVUE-FM restricts activities such as creating, revoking, and reviewing of Token APIs.
- If you use an expired or invalid token, GigaVUE-FM returns a 401 Unauthorized error when accessing the REST API.

This section explains about the following:

- [Create Token](#)
- [Revoke Tokens](#)
- [Export Token](#)
- [Create a Token with Custom User Group](#)

Create Token

GigaVUE-FM allows you to create a token or multiple tokens if required.

To create a token,

1. Go to , select **Authentication > GigaVUE-FM User Management**.

The **User Management** page appears.

2. In the **User Management** page, select **Tokens**.

NOTE: If you are a user with write access, then you can view a drop-down list under **Tokens**. Select **Current User Tokens** to create a token.

3. Select **New Token**.
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.
7. Select **OK** to generate a new token.

The generated token appears on the **Tokens** page. You can copy and use the generated token to authenticate the GigaVUE-FM REST APIs.

Select the token that you want to copy, click the **Actions** button drop-down list, and select **Copy Token**. The token is copied. You can paste in the required areas.

NOTE: You cannot view the generated token. You can only copy and paste the generated token.

Revoke Tokens

You can revoke tokens that other users create only if you have write access in GigaVUE-FM Security Management.

To revoke tokens,

1. Go to , select **Authentication > GigaVUE-FM User Management**.
2. In the **User Management** page that appears, select **Tokens**.
3. From the drop-down list, select **Token Management**.

4. You can view the token created by other users.
5. Select the token that you want to revoke.
6. Select **Action**, and then select **Revoke**.

Export Token

GigaVUE-FM allows you to export selected or all the tokens in CSV and XLSX format.

- To export a token, select the token, and from the **Export Selected** drop-down list box, select either the **CSV** or **XLSX** format.
- To export all the tokens, select the token, and from the **Export All** drop-down list box, select the **CSV** or **XLSX** format .

Create a Token with Custom User Group

GigaVUE-FM allows you to create a token with different user groups.

To create a token,

1. Create a custom user.

NOTE: Do not assign this user to any existing user group.

- **Example:** orchestration
2. Set up a custom role with the required permissions.
 - **Example:** orchestrationRole
 - Assign “Third Party Orchestration” with write permission
3. Create a user group, assign the custom role to it, and add the user to the group.
 - **Example:** orchestrationGroup
 - Assign the orchestration Role to this group.
 - Add the orchestration user to this group.
4. Log in to GigaVUE-FM using the custom user, go to the token page, and select the custom group.

The token is generated.

Integrate Private CA

You can integrate your own PKI infrastructure with GigaVUE-FM.

Rules and Notes

- Always place the root CA in a separate file.
- When using multiple intermediate CAs, consider the following:
 - Include all intermediate CAs in a single file in the correct order.

- Place the last intermediate CA in the chain at the top.
- Place the preceding CAs in descending order.

To integrate private CA

- Generate a Certificate Signing Request (CSR).
- Get a signature of the Certificate Authority (CA) on the CSR.
- Upload CA Certificate to GigaVUE-FM.

Generate CSR

To create an intermediate CA certificate:

1. Go to  > **System > Certificates**.
2. In the top navigation bar, from the **PKI** drop-down list, select **CSR**. The **Generate Intermediate CA Certificate** page appears.
3. Enter details in the following fields:
 - **Country:** Enter the name of your country.
 - **Organization:** Enter the name of your organization.
 - **Organization Unit:** Enter the name of the department or unit.
 - **Common Name:** Enter the common name associated with the certificate.
4. From the **Algorithm** drop-down list, select the desired encryption algorithm used to encrypt your private key.
5. Select **Generate CSR**.

The CSR is downloaded successfully.

Upload CA Certificate

Get the CSR signed from your Enterprise PKI or any public PKI and upload the signed intermediate CA certificate to GigaVUE-FM.

To upload the signed CA certificate to GigaVUE-FM:

1. Go to  > **System > Certificates**.
2. In the top navigation bar, from the **PKI** drop-down list, select **CA**. The **CA Certificate** page appears.
3. From the **Actions** drop-down list, select **Upload CA**. The **Upload CA** pop-up appears.
4. Next to **Intermediate CA**, select **Choose File** to upload the signed intermediate CA

certificate.

5. Next to **Root CA**, select **Choose File** to upload the corresponding root or intermediate CA.

The **CA Certificate** page displays the uploaded CA certificate.

Install GigaVUE-FM on VMware ESXi

GigaVUE-FM can be installed on VMware ESXi host and VMware vCenter using OVA file or ISO Image.

NOTE: To prevent data loss, do not allow the GigaVUE-FM instance to be vMotioned. Since vCenter does not provide a mechanism for GigaVUE-FM to prevent or control vMotion events, it is your responsibility to ensure that vMotion is never initiated for the GigaVUE-FM.

- [Before You Install](#)
- [Deploy OVA file on VMware vCenter](#)
- [Install GigaVUE-FM using ISO Image File on VMware ESXi Host](#)
- [Install GigaVUE-FM using ISO Image File on VMware vCenter](#)
- [Initial GigaVUE-FM Configuration](#)

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 as follows:

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.

Hardware Requirements	
	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

Requirements	Basic Configuration	Medium Configuration	Large Configuration	Large Configuration (GigaVUE-FM Hardware Appliance 3)
** 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	160GB	1240GB	3000GB	3000GB

Configuration	Basic Configuration	Medium Configuration	Large Configuration	Large Configuration (GigaVUE-FM Hardware Appliance 3)
Days Data				
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

	Default (Basic)	Medium	Large	GigaVUE-FM High Availability Large 3000 Nodes
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-FMOVA 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.

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-FM6.14.02, and then select **Deploy OVF Template**. The Deploy OVF Template wizard opens.

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.

NOTE: It is recommended that you do not select the Power on automatically option and instead review all configurations before powering on the virtual machine.

4. Click the **Next**. The **Select a name and folder** page of the **Deploy OVF Template** wizard appears.
5. 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.
6. Click **Next**. The **Select a compute resource** page of the **Deploy OVF Template** wizard appears.
7. 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.
8. Click **Next**. The **Review details** page of the **Deploy OVF Template** wizard appears.
9. Verify the OVF template details, and then click **Next**. The **Select storage** page of the **Deploy OVF Template** wizard appears.
10. From the **Select virtual disk format** dropdown list, select **Thin Provision** as the format for the virtual disks and provisioning.
11. Select the datastore where the virtual machine's files will be stored.
12. Select the storage policy based on the requirement else default storage policy is assigned.
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.
 - 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**. The **Ready to Complete** page of the **Deploy OVF Template** wizard appears.

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

- 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

- 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:

- 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.

Prerequisites

- 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.

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](#).
4. **(Optional):** If you want to change any of the network configurations or increase memory refer to the instructions in

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-FM6.14.02, 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 (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 **Thin Provision** as the format for the virtual disks and provisioning.
 - 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 settings are correct, and then click **Finish**. The Recent Tasks pane of the vSphere Client Home page shows the progress of the deployment operation and GigaVUE-FM deployment takes approximately 15 - 20 minutes.
15. When the operation is complete, you have successfully deployed a GigaVUE-FM instance.

Important: Clear the browser cache 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 ▾			(i)
> 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

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

Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi

This section explains the available methods to upgrade the GigaVUE V Series Nodes and provides instructions on how to upgrade GigaVUE V Series Nodes.

IMPORTANT



Before upgrading the Fabric Components to version 6.10.00 or above, ensure the following actions are performed:

- Open the required ports in the cloud platform. For details, refer to Network Firewall Requirement for GigaVUE Cloud Suite.
- When using FMHA configuration, follow the steps in the [Configure Secure Communication between Fabric Components in FMHA](#) section.

Update GigaVUE V Series Nodes

Use one of the following two options:

1. Upgrade all the GigaVUE V Series Nodes in a monitoring Domain:
 - a. Select the Monitoring Domain.
 - b. Select **Actions > Upgrade Fabric**.
2. Upgrade a single or a group of GigaVUE V Series Node. When upgrading a group of GigaVUE V Series Nodes, ensure all the GigaVUE V Series Nodes deployed on the same ESXi hosts are selected.

Upgrade Considerations

- You can select an entire monitoring domain and upgrade all the GigaVUE V Series Nodes in that particular monitoring domain, or you can select an entire host and upgrade all the GigaVUE V Series Node deployed in that particular host.
- When multiple GigaVUE V Series Nodes are deployed on the same ESXi host and only if a part of GigaVUE V Series Nodes are selected in that particular host, then the **Upgrade Fabric** button is disabled.
- When upgrading GigaVUE V Series Nodes, if a host of a particular GigaVUE V Series Node is under maintenance mode, then the **Upgrade Fabric** button is disabled.
- Unselect the GigaVUE V Series Node whose host is under maintenance mode, and upgrade that GigaVUE V Series Node when the host is out of the maintenance mode.

NOTE: GigaVUE-FM supports only (n, n-1, n-2) GigaVUE V Series Node versions.

Upgrade the GigaVUE V Series Node in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, then click **Monitoring Domain**. The **Monitoring Domain** page appears.
2. Select an entire monitoring domain or an entire host and select **Actions**.

- From the drop-down list, select **Upgrade Fabric**. The **V Series Node Upgrade Task** dialog box appears.

V Series Nodes Upgrade Task

Name*

Image*

Name Server

Change Form Factors ☒

Default Form Factor ⓘ

V Series Node	Form Factor	Version
VSeries.vp-esxi-	Small, 2vCPU, 4GB RAM, 8GB ...	6.4.00
VSeries.vp-esxi-	Small, 2vCPU, 4GB RAM, 8GB ...	6.4.00

- Enter a name for the V Series Node upgrade task.
- Select the latest GigaVUE V Series Node OVA image from the **Image** drop-down list. When upgrading the GigaVUE V Series Nodes to any version equal to or greater than 6.5.00, the **Name Server** field is displayed. This field is optional. Name Server is a server that stores the mapping between the domain names and the IP address. The maximum number of name servers that can be entered is three. Enter valid IPv4 address, separated by comma.
- (Optional) Select the **Change Form Factors** check box if you want to modify the form factor (instance) size.
- Select the form factor from the Default Form Factor drop-down menu to change the form factor of all the selected V Series Nodes. You can use the **Use Current** option to use the existing form factor of the individual GigaVUE V Series Node. You can also change the form factor of an individual GigaVUE V Series Node from the **Form Factor** drop-down menu of the respective GigaVUE V Series Node. The form factor selected here overwrites the form factor selected in the **Default Form Factor**.

NOTE: All the GigaVUE V Series Node with Static IP address retain their old IP address even after the upgrade.

- Select **Upgrade** to launch the GigaVUE V Series Node upgrade.

NOTE: Both the new and the current GigaVUE V Series Nodes appear in the same Monitoring Domain until the new nodes replace the current, and the status changes to **Ok**.

You can view the status of the upgrade in the Status column of the **Monitoring Domain** page.

New

Actions

Refresh Inventory

		Monitoring Domain	Connections	Host	Name	Management IP	Tunnel IP	Type	Version	Status	
	>	md1									Upgrade Status
	▼		con1							Connected	
	▼			10.115.81.184							
					VSeries.12710.115.81.184	10.114.82.89	10.114.84.3	V Series Node	6.6.00	upgrading	
					VSeries.new10.115.81.184-1	10.114.84.88	10.114.84.93	V Series Node	6.6.00	upgrading	
	▼			10.115.81.185							
					VSeries.new10.115.81.185-1	10.114.82.143	10.114.84.86	V Series Node	6.6.00	upgrad	
					VSeries.st10.115.81.185-1	10.115.80.190	10.115.80.191	V Series Node	6.6.00	upgrading	

To view the detailed upgrade status click **Upgrade Status**, the **V Series Node Upgrade Status** dialog box appears.

V Series Node Upgrade Status

Monitoring Domain Name: esxi-md-202-13

Upgrade Tasks

▼ Upgrade_GigaVUE_VSERIES_Node | SUCCESS

Clear

Summary

☐ Success: 2 ☐ Failed: 0 ☐ In Progress: 0 Total: 2

Node Statuses

Node	Status
VSeries.vp-redscvr-10.1.1.10.10.10	OK
VSeries.vp-redscvr-10.1.1.10.10.10 renamed2	OK

▼ Upgrade | IN_PROGRESS

Summary

☐ Success: 0 ☐ Failed: 0 ☐ In Progress: 1 Total: 1

Node Statuses

Node	Status
VSeries.vp-redscvr-10.1.1.10.10.10 upgrade	launching

- Select **Clear** to delete the monitoring domain upgrade status history of successfully upgraded nodes.
- If the GigaVUE V Series Node upgrade fails or is interrupted for any reason, select **Retry** on the **V Series Node Upgrade Status** dialog box.

NOTE: You cannot modify the node configurations when you are using **Retry** option. GigaVUE-FM uses the same values defined in the initial fabric upgrade configuration.

Important GigaVUE-FM Upgrade Guidance

To upgrade GigaVUE-FM, the recommended procedure is to back up your current configuration and perform the pre-upgrade checks (memory requirements, the maximum number of images) before installing the image. You can upgrade via the GigaVUE-FM GUI or CLI. It is recommended that you plan your upgrade before performing any upgrade steps. Refer [Supported Upgrade Paths for GigaVUE-FM](#) for details on upgrade path.

Best Practice: Before performing an upgrade, in addition to taking a backup of GigaVUE-FM, also take a Sysdump, which provides valuable configuration information that is not captured in the backup. You can use the Sysdump after the upgrade to troubleshoot any upgrade-related issues with your configurations.

Always take a snapshot of GigaVUE-FM before performing a version upgrade on GigaVUE-FM to ensure quick recovery in case of upgrade failures.

GigaVUE-FM Upgrade Guidance for GigaVUE Cloud Suite

Starting from software version 6.10.00, the Secure Communication feature in GigaVUE-FM enhances security by enabling mutual Transport Layer Security (mTLS)-based authentication across GigaVUE Fabric Components. Upon upgrading to 6.10.00, this feature will be implemented. Before upgrading GigaVUE-FM and the fabric components, ensure that the necessary ports are open as mentioned in [Network Firewall Requirement for GigaVUE Cloud Suite](#).



Note:

Monitoring session migration steps for V Series Node 1 upgrade:

When you upgrade GigaVUE-FM with a Monitoring Session in V Series Node 1 to version 6.12 from an earlier release, follow these steps to ensure a successful migration:

- Undeploy the existing V Series Node 1 Monitoring Session.
- Delete any Egress VXLAN TEPs with the Multi-Tunnel feature enabled.
- Recreate the Egress VXLAN TEP with the required configuration and enable the Multi-Tunnel flag.
- Establish a Link connection between the LB application and the Egress VXLAN TEP.
- Deploy the Monitoring Session.



AMX upgrade impact from version 6.13 or lower to 6.14:

Upgrading V Series Nodes with Application Metadata Exporter (AMX) from version 6.13 (or lower) to version 6.14 involves a change in the internal session data storage format. This change may result in the loss of existing control sessions during the upgrade process. To ensure a smooth upgrade and to avoid any disruption to control session data, contact [Gigamon Technical Support](#) before upgrading your AMX nodes to version 6.14.

GigaVUE-FM Upgrade Guidance for Application Intelligence

Starting from Software version 6.5.00, Application Intelligence solution can be configured from the Monitoring Session Page. You cannot create a new Application Intelligence Session or edit an Application Intelligence Session for virtual environment from the **Application Intelligence** page.

In the Monitoring Session page, Application Filtering can be configured either with Pass All or Advanced Rules in the fifth Rule Set. You cannot configure Application Filtering with both Pass All and Advanced Rules in the fifth Rule set. Therefore, before upgrading your GigaVUE-FM to 6.5.00, ensure that none of the existing Application Intelligence Session has Application Filtering configured with both Advanced Rules as Drop and No Rule Match Pass All in the fifth rule set. Otherwise, after upgrading to 6.5.00, you must delete that particular Application Intelligence Session to continue with your migration process.

For this reason, before upgrading to 6.5.00, you must edit the fifth rule set in the Application Filtering such that either only Advanced Rule as Drop or No Rule Match Pass All exists. Else, after upgrading to 6.5.00, you must delete that particular Application Intelligence Session to continue with your migration process. Refer to [Migrate Application Intelligence Session to Monitoring Session](#) section in the respective cloud deployment guides for more detailed information on how to migrate the existing Application Intelligence Session to the Monitoring Session page.

Guidance for FMHA when Encryption Mode is Certificate based Authentication

You must ensure that the following key usages are added to your PKI certificate before upgrading the FMHA:

- serverAuth (TLS Web Server Authentication)
- clientAuth (TLS Web Client Authentication)
- ipsecEndSystem
- ipsecTunnel
- ipsecUser

NOTE: This is not applicable for FMHA with Pre-shared key.

To configure the certificate in FMHA, refer to [Setting up Certificate \(PKI\) based Encryption in FMHA](#).

Deploy

Deploy GigaVUE Cloud Suite for VMware ESXi using GigaVUE FM Orchestration

This section describes various options to configure GigaVUE Cloud Suite for VMware ESXi to provide visibility for both physical and virtual traffic. You can configure GigaVUE Cloud Suite for VMware ESXi in the following ways, based on the traffic acquisition method and your approach to deploying fabric components.

Refer to the following deployment options for more detailed information:

- [Acquire Traffic using Platform Tapping](#)
- [Acquire Traffic using Customer Orchestrated Source](#)

Acquire Traffic using Platform Tapping

This section outlines the workflow for acquiring traffic with platform tapping and deploying GigaVUE Fabric Components using GigaVUE-FM. Refer to the following topics for instruction on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Install GigaVUE-FM on VMware ESXi](#)
- [Required VMware Virtual Center Privileges](#)
- [Default Login Credentials](#)
- [Create Monitoring Domain for VMware ESXi](#)
- [Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM](#)
- [Verify Your Deployment](#)

Install GigaVUE-FM on VMware ESXi

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

Required VMware Virtual Center Privileges

This topic lists the minimum privileges required for the GigaVUE-FM user in Virtual Center.

You assign privileges to Virtual Center users. To assign privileges, from the left pane, select **Administration** and go to **Access Control > Roles**.

Important! You must apply the roles at the vSphere Virtual Center level, not at the data center or Host levels.

The following table lists the minimum permissions required for GigaVUE-FM to manage the virtual center user with the specified roles:

Category	Required Privilege	Purpose
Datastore	Allocate space	V Series Node Deployment
Distributed Switch	VSPAN Operation	VDS Tapping
Folder	Create Folder	V Series Node Deployment
Host	Configuration Network Configuration	VSS Tapping
	Inventory Modify Cluster	Pin V Series Node to the host in cluster configurations to prevent automatic migration.
Network	- Assign network - Configure	- V Series Node Deployment/VSS Tapping - V Series Node Deployment
Resource	Assign virtual machine to resource pool	V Series Node Deployment
vApp	Import - vApp instance configuration - vApp application configuration	V Series Node Deployment
Virtual machine	Configuration - Add new disk - Add or remove device - Modify device settings - Rename	V Series Node Deployment V Series Node Deployment/VSS Tapping
	Interaction - Connect devices - Power on	V Series Node Deployment

Category	Required Privilege	Purpose
	- Power Off - Reset	
	Inventory - Create from existing - Remove	V Series Node Deployment
	Provisioning Clone virtual machine	V Series Node Deployment

Default Login Credentials

You can use the default credentials to log in to the GigaVUE V Series Node by using the default credentials.

Field	Description
GigaVUE V Series Node	You can login to the GigaVUE V Series Node by using ssh. The default username and password is: - Username: gigamon - Password: Enter the password provided during the fabric launch configuration.

Create Monitoring Domain for VMware ESXi

This chapter describes how to create a monitoring domain for deploying GigaVUE V Series Nodes in VMware vCenter environment through GigaVUE-FM. You must establish a connection between GigaVUE-FM and VMware vCenter. Creating a monitoring domain in GigaVUE-FM allows you to establish a connection between vCenter and GigaVUE-FM.

To create a monitoring domain in GigaVUE-FM for VMware vCenter,

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, and select **Monitoring Domain**. The **Monitoring Domain** page appears. The Monitoring Domain page has three tabs, namely Monitoring Domains, Hosts and Fabric.
2. In the **Monitoring Domains** tab, select **New**. The **Create Monitoring Domain Configuration** page appears.
3. In the **Create Monitoring Domain Configuration** page, enter or select the following details:

Field	Description
Monitoring Domain	Name of the monitoring domain.
Connection Alias	Name of the connection.
Virtual Center	IP address or FQDN of the vCenter.

Field	Description
	NOTE: To ensure the validity of VMware virtual central certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. For details, refer to Cloud solution in Trust Store Trust Store section in the GigaVUE Administration Guide.
Username	Username of the vCenter user with minimum privileges as described in the Required VMware vCenter Privileges section. NOTE: Whenever you change the vCenter credentials in VMware vCenter, edit the Monitoring Domain to update that in GigaVUE-FM. Otherwise, the connection status reaches an authentication failure state.
Password	vCenter password to connect to the vCenter
Traffic Acquisition Method	Select a Traffic Acquisition Method. Platform Tapping: Platform tapping is performed in two ways. • VSS: Used when a workload Virtual Machine is connected to a Virtual Standard Switch network. GigaVUE-FM creates a promiscuous network on the VSS switch for tapping traffic. • VDS: Used when a workload Virtual Machine is connected to a Virtual Distributed Switch portgroup. Port Mirroring is created on the VDS switch by GigaVUE-FM for tapping the traffic.
Resource Allocation	When deploying multiple GigaVUE V Series Node in a single host, select any one of the following options: Target VM Based: Choose this option if your deployment workload VMs attached to less than or equal to 8 vSwitches on the same ESXi host. This resource allocation type distributes the workload VMs across multiple GigaVUE V Series Nodes on the same ESXi host. Switch Based: A single GigaVUE V Series Node can tap a maximum of 8 vSwitches. Choose this option if you have traffic monitoring VMs running on ESXi hosts connected to more than 8 vSwitches in a single host. The vSwitches are mapped to the GigaVUE V Series Node in a round-robin manner. In this model, vSwitches are evenly distributed across the available GigaVUE V Series Nodes on the same host. NOTE: This field is applicable only when using Platform Tapping as the Traffic Acquisition Method NOTE: Ensure to undeploy all the Monitoring Session associated with the connection before changing the Resource Allocation type.
Maximum Number of V Series Nodes per Host	Enter the maximum number of GigaVUE V Series Nodes possible to deploy in a single host. The default value is 10.
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.

4. Click **Save**.

5. The **VMware Fabric Launch Configuration** page appears. For details on how to deploy GigaVUE V Series Nodes in the **VMware Fabric Launch Configuration** page.
6. After you deploy the fabric node, you can view the host where the monitoring domain is deployed on the **Hosts** tab and the fabric node details on the **Fabric** tab.

**Notes:**

- Ensure that all V Series Nodes within a Monitoring Domain 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.

The Monitoring Domain created in this section is listed on the **Monitoring Domain** page.

**Points to Note:**

- Whenever you change the vCenter credentials in VMware vCenter, edit the Monitoring Domain to update that in GigaVUE-FM. Otherwise, the connection status reaches an authentication failure state.
- When the Monitoring Domain is in a "Not Connected" state, updating the vCenter credentials will only refresh the information stored in the GigaVUE-FM database.

To establish the connection:

1. Navigate to the **Monitoring Domain** page.
2. Select **Actions** and select **Connect**.

By following these steps, you can ensure that your vCenter credentials are updated and the connection is established correctly.

You can perform the following actions in the Monitoring Domains and Fabric tabs of the Monitoring Domain page.

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

Actions	Description
	Configure Certificate Settings.
Edit Monitoring Domain	Used to edit a monitoring domain.
Deploy Fabric	Used to deploy GigaVUE V Series Nodes.
Upgrade Fabric	Used to upgrade GigaVUE V Series Nodes. For details, refer to Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi.
Delete Monitoring Domain	Used to delete a monitoring domain.
Delete Fabric Nodes	Used to delete a GigaVUE V Series Node in the monitoring domain.
Connect	Used to start communication between GigaVUE-FM and the VMware vCenter.
Disconnect	Used to stop communication between GigaVUE-FM and the VMware vCenter.
Edit SSL Configuration	Used to add Certificate Authority and the SSL Keys when using the Secure Tunnels.
Fabric tab	
Delete Fabric Nodes	Used to delete a specific GigaVUE V Series Node in the monitoring domain.
Rediscover	The changes made in vCenter for the GigaVUE V Series is reflected in GigaVUE-FM. Refer to Rediscover topic for more detailed information.
Power On	You can select an individual GigaVUE V Series Node and power it on. The status of the GigaVUE V Series Node is changed to Ok .
Power Off	You can select an individual GigaVUE V Series Node and power it off. If the GigaVUE V Series Node is turned off from GigaVUE-FM, then it is not considered as part of Cloud Health Monitoring and GigaVUE-FM does not try to turn it on. The status of the GigaVUE V Series Node is changed to Down .
Reboot	You can select an individual GigaVUE V Series Node and reboot it.
Rename	Used to rename the fabric nodes.
Generate Sysdump	You can select one or multiple GigaVUE V Series Nodes (Maximum 10) to generate the system files. The generation of sysdump takes a few minutes in a GigaVUE V Series Node. You can proceed with other tasks, and upon completion, the status appears in the GUI. These system files are helpful for troubleshooting. For more information, refer to Debuggability and Troubleshooting
Download Cloud Configuration File	Use this option to download the cloud configuration file (for example, md_gigamon-cloud.conf). The file contains monitoring domain connection details such as groupName, subGroupName, and token, and can be used to register fabric components in GigaVUE-FM using third party orchestration.

More information

- To view and manage the generated sysdump files, select the GigaVUE V SeriesNode and select the **Sysdump** tab in the lower pane.
- To view the certificates associated with the fabric, select the fabric nodes and select the **Certificates** tab in the lower pane.

Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM

This section provides step-by-step instructions on how to deploy GigaVUE V Series Nodes in VMware vCenter Monitoring Domain.

To deploy GigaVUE V SeriesNodes using GigaVUE-FM, follow these steps:

1. Go to the **VMware Fabric Launch Configuration** page using one of the following options:
 - **Automatic Option:** Create a monitoring domain, and the page opens automatically.
 - **Manual Option:** Open **VMware Fabric Launch Configuration** page from the **Monitoring Domain** page.
 - To launch the **VMware Fabric Launch Configuration** from the Monitoring Domain, go to **Inventory > VIRTUAL > VMware vCenter (V Series)**.
 - Select **Actions > Deploy Fabric**.

The **VMware Fabric Launch Configuration** page appears.

VMware vCenter

ESXi Fabric Launch

Monitoring Domains

Hosts

Fabric

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-1

V Series Node Image*

gigamon-gigavue-vseries-node-6.14.00-585570_...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

Select All

Select None

10.115.81.39

Common Configuration

10.115.81.39

VSeries.10.115.81.39-2v

Datstores

Datstore Clusters

Datastore*

datastore1_81_39

Name Server

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.39-2

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

DHCP

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DataNetwork

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Interface Name

Network adapter 3

Network*

karthik

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Virtual Disk Format

Thin Provision

Deployment Folder

/

User Password* (gigamon)

Confirm User Password*

Passwords do not match

VMware vCenter

ESXi Fabric Launch

Monitoring Domains

Hosts

Fabric

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-1

V Series Node Image*

gigamon-gigavue-vseries-node-6.14.00-585570_...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

Select All

Select None

10.115.81.39

Common Configuration

10.115.81.39

VSeries.10.115.81.39-2v

Datstores

Datstore Clusters

Datastore*

datastore1_81_39

Name Server

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.39-2

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

DHCP

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DataNetwork

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Interface Name

Network adapter 3

Network*

karthik

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Virtual Disk Format

Thin Provision

Deployment Folder

/

User Password* (gigamon)

Confirm User Password*

Passwords do not match

VMware vCenter

ESXi Fabric Launch

Monitoring Domains

Hosts

Fabric

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-1

V Series Node Image*

gigamon-gigavue-vseries-node-6.14.00-585570_...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

Select All

Select None

10.115.81.39

Common Configuration

10.115.81.39

VSeries.10.115.81.39-2v

Datstores

Datstore Clusters

Datastore*

datastore1_81_39

Name Server

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.39-2

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

DHCP

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DataNetwork

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Interface Name

Network adapter 3

Network*

karthik

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Virtual Disk Format

Thin Provision

Deployment Folder

/

User Password* (gigamon)

Confirm User Password*

Passwords do not match

VMware vCenter

ESXi Fabric Launch

Monitoring Domains

Hosts

Fabric

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-1

V Series Node Image*

gigamon-gigavue-vseries-node-6.14.00-585570_...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

Select All

Select None

10.115.81.39

Common Configuration

10.115.81.39

VSeries.10.115.81.39-2v

Datstores

Datstore Clusters

Datastore*

datastore1_81_39

Name Server

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.39-2

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

DHCP

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DataNetwork

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Interface Name

Network adapter 3

Network*

karthik

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Virtual Disk Format

Thin Provision

Deployment Folder

/

User Password* (gigamon)

Confirm User Password*

Passwords do not match

VMware vCenter

ESXi Fabric Launch

Monitoring Domains

Hosts

Fabric

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-1

V Series Node Image*

gigamon-gigavue-vseries-node-6.14.00-585570_...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

Select All

Select None

10.115.81.39

Common Configuration

10.115.81.39

VSeries.10.115.81.39-2v

Datstores

Datstore Clusters

Datastore*

datastore1_81_39

Name Server

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.39-2

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

DHCP

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DataNetwork

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

Interface Name

Network adapter 3

Network*

karthik

IP Type

DHCP

MTU

Enter MTU Value

Use IPv6

Disabled

2. On the **VMware Fabric Launch Configuration** page, enter or select the following details:

Field	Description
Datacenter	vCenter Data Center with the ESXi hosts to be provisioned with GigaVUE V Series Nodes.
Cluster	Cluster where you want to deploy the GigaVUE V Series Nodes.
V Series Node Image	Select the OVA file from the drop-down menu. You can also add OVA files when launching the fabric. To add OVA files: - Click on the Add button. - The Upload Image dialog box opens. In the Upload Image dialog box, click Browse to select the <i>gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova</i> file. - Select Upload to Server to upload the selected OVA image file to GigaVUE-FM server.
Form Factor	Instance size of the GigaVUE V Series Node. For details, refer to Prerequisites for V Series nodes on ESXi Integration. NOTE: Small Form Factor is not supported when using applications like Application Visualization, Application Metadata, Application Filtering. Select 80GB Disk Space, when using AMX Application.
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 unavailable 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.
Certificate NOTE: This field appears only when Enable Custom Certificates field is enabled.	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes. For more information, refer to Secure Communication between GigaVUE Fabric Components.
Hosts	Select the ESXi hosts for GigaVUE V Series Node deployment. NOTE: You can select up to 250 hosts manually. Select Import Host Info from file or Add Host Info Manually. Import Host Info from file: To import host details from a .csv file: - Download the .csv template file. - Enter the required values in the Excel sheet and save the file. - Select Browse and select the .csv file saved in the previous step.

Field	Description																		
	NOTE: To deploy more than one GigaVUE V SeriesNode on the same host, add more rows in the Excel sheet with the same host value for each extra GigaVUE V SeriesNode you want to deploy. If your GigaVUE-FM version is above 6.5 and GigaVUE V Series Nodes are on a version below 6.5, Name Server and MTU are not supported. Therefore, these fields in the .csv file must be empty. Add Host Info Manually: Select the ESXi hosts for GigaVUE V Series Node deployment. The Common Configuration drop-down wizard appears. Expand the Common Configuration drop-down wizard and update the following details to apply the configuration to all the selected hosts. You can expand the individual hosts and add or delete GigaVUE V Series Node. You can expand the individual GigaVUE V Series Node and modify the configurations applied in the Common Configuration. Common Configuration <table> <tr> <td>Datastore/Datastore Clusters</td><td>Select the network datastore or datastore cluster shared among all ESXi hosts.</td></tr> <tr> <td>V Series Node Name Prefix</td><td>Enter a prefix for the GigaVUE V SeriesNode name.</td></tr> <tr> <td>V Series Node Name Suffix</td><td>Enter a suffix for the GigaVUE V Series Node name.</td></tr> <tr> <td>Name Server</td><td>The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.</td></tr> <tr> <td>No. of V Series Nodes per Host</td><td>Enter the number of GigaVUE V Series Nodes to be deployed in each host.</td></tr> </table> Management <table> <tr> <td>Network</td><td>Management network for GigaVUE V Series Nodes.</td></tr> <tr> <td>IP Type</td><td>Select the management network IP type as Static or DHCP.</td></tr> <tr> <td> Gateway IP NOTE: This field appears only when the Management IP type is Static. </td><td>Gateway IP address of the Management Network.</td></tr> <tr> <td> Netmask Length NOTE: This field appears only when the </td><td>Management network's subnet mask value in CIDR format. For example, 21 for /21.</td></tr> </table>	Datastore/Datastore Clusters	Select the network datastore or datastore cluster shared among all ESXi hosts.	V Series Node Name Prefix	Enter a prefix for the GigaVUE V SeriesNode name.	V Series Node Name Suffix	Enter a suffix for the GigaVUE V Series Node name.	Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.	No. of V Series Nodes per Host	Enter the number of GigaVUE V Series Nodes to be deployed in each host.	Network	Management network for GigaVUE V Series Nodes.	IP Type	Select the management network IP type as Static or DHCP.	Gateway IP NOTE: This field appears only when the Management IP type is Static.	Gateway IP address of the Management Network.	Netmask Length NOTE: This field appears only when the	Management network's subnet mask value in CIDR format. For example, 21 for /21.
Datastore/Datastore Clusters	Select the network datastore or datastore cluster shared among all ESXi hosts.																		
V Series Node Name Prefix	Enter a prefix for the GigaVUE V SeriesNode name.																		
V Series Node Name Suffix	Enter a suffix for the GigaVUE V Series Node name.																		
Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.																		
No. of V Series Nodes per Host	Enter the number of GigaVUE V Series Nodes to be deployed in each host.																		
Network	Management network for GigaVUE V Series Nodes.																		
IP Type	Select the management network IP type as Static or DHCP.																		
Gateway IP NOTE: This field appears only when the Management IP type is Static.	Gateway IP address of the Management Network.																		
Netmask Length NOTE: This field appears only when the	Management network's subnet mask value in CIDR format. For example, 21 for /21.																		

Field	Description	
	Management IP type is Static.	
MTU		The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
Data Interfaces - When using Customer Orchestrated Source as the Traffic Acquisition Method, you must configure two data interfaces.		
Use IPv6		Enable to use IPv6.
NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.		
Network		Tunnel Network for the GigaVUE 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.
Netmask Length		Tunnel network's subnet mask value in CIDR format. For example, 21 for /21.
NOTE: This field appears only when the Tunnel IP type is Static.		
MTU		The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
IPv6 Prefix Length		Enter the IPv6 prefix length as 64.
NOTE: This field appears only when the Use IPv6 toggle button is enabled.		
Virtual Disk Format		Select the Virtual Disk Format from the drop-down menu
Deployment Folder		Enter the folder name in vCenter, under which the GigaVUE V SeriesNodes must be deployed.
Password		Enter the password you wish to use for the GigaVUE V Series Node.

3. Select **Deploy**. After deployment in vCenter, it appears on the **Monitoring Domain** page under the Monitoring Domain, where the GigaVUE V Series Node is deployed.

NOTE: GigaVUE-FM can process a maximum of ten GigaVUE V Series Nodes deployment requests in parallel on VMware vCenter. Each deployment request can have multiple GigaVUE V Series Nodes for deployment.

Verify Your Deployment

Once you select Deploy, GigaVUE-FM provisions and registers the GigaVUE V Series Nodes. On the Monitoring Domain page, confirm that the node status is Ok/Healthy, which indicates a successful deployment.

What to do Next

After you successfully deploy the fabric components required for traffic acquisition, configure them within the monitoring session to direct traffic from the workload virtual machines to the appropriate tools. For instructions on setting up the session, refer to [Configure Monitoring Session](#).

Acquire Traffic using Customer Orchestrated Source

This section outlines the workflow for acquiring traffic with Port Mirroring and deploying GigaVUE Fabric Components using GigaVUE-FM. Refer to the following topics for instruction on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Install GigaVUE-FM on VMware ESXi](#)
- [Required VMware Virtual Center Privileges](#)
- [Default Login Credentials](#)
- [Create Monitoring Domain for VMware ESXi](#)
- [Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM](#)
- [Verify Your Deployment](#)

Install GigaVUE-FM on VMware ESXi

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

Required VMware Virtual Center Privileges

This topic lists the minimum privileges required for the GigaVUE-FM user in Virtual Center.

You assign privileges to Virtual Center users. To assign privileges, from the left pane, select **Administration** and go to **Access Control > Roles**.

Important! You must apply the roles at the vSphere Virtual Center level, not at the data center or Host levels.

The following table lists the minimum permissions required for GigaVUE-FM to manage the virtual center user with the specified roles:

Category	Required Privilege	Purpose
Datastore	Allocate space	V Series Node Deployment
Distributed Switch	VSPAN Operation	VDS Tapping
Folder	Create Folder	V Series Node Deployment
Host	Configuration Network Configuration	VSS Tapping
	Inventory Modify Cluster	Pin V Series Node to the host in cluster configurations to prevent automatic migration.
Network	- Assign network - Configure	- V Series Node Deployment/VSS Tapping - V Series Node Deployment
Resource	Assign virtual machine to resource pool	V Series Node Deployment
vApp	Import - vApp instance configuration - vApp application configuration	V Series Node Deployment
Virtual machine	Configuration - Add new disk - Add or remove device - Modify device settings - Rename	V Series Node Deployment V Series Node Deployment/VSS Tapping

Category	Required Privilege	Purpose
	Interaction - Connect devices - Power on - Power Off - Reset	V Series Node Deployment
	Inventory - Create from existing - Remove	V Series Node Deployment
	Provisioning Clone virtual machine	V Series Node Deployment

Default Login Credentials

You can use the default credentials to log in to the GigaVUE V Series Node by using the default credentials.

Field	Description
GigaVUE V Series Node	You can login to the GigaVUE V Series Node by using ssh. The default username and password is: - Username: gigamon - Password: Enter the password provided during the fabric launch configuration.

Create Monitoring Domain for VMware ESXi

This chapter describes how to create a monitoring domain for deploying GigaVUE V Series Nodes in VMware vCenter environment through GigaVUE-FM. You must establish a connection between GigaVUE-FM and VMware vCenter. Creating a monitoring domain in GigaVUE-FM allows you to establish a connection between vCenter and GigaVUE-FM.

To create a monitoring domain in GigaVUE-FM for VMware vCenter,

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, and select **Monitoring Domain**. The **Monitoring Domain** page appears. The Monitoring Domain page has three tabs, namely Monitoring Domains, Hosts and Fabric.
2. In the **Monitoring Domains** tab, select **New**. The **Create Monitoring Domain Configuration** page appears.
3. In the **Create Monitoring Domain Configuration** page, enter or select the following details:

Field	Description
Monitoring Domain	Name of the monitoring domain.
Connection Alias	Name of the connection.
Virtual Center	IP address or FQDN of the vCenter. NOTE: To ensure the validity of VMware virtual central certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. For details, refer to Cloud solution in Trust Store Trust Store section in the GigaVUE Administration Guide.
Username	Username of the vCenter user with minimum privileges as described in the Required VMware vCenter Privileges section. NOTE: Whenever you change the vCenter credentials in VMware vCenter, edit the Monitoring Domain to update that in GigaVUE-FM. Otherwise, the connection status reaches an authentication failure state.
Password	vCenter password to connect to the vCenter
Traffic Acquisition Method	Select a Traffic Acquisition Method. Customer Orchestrated Source: If you select Customer Orchestrated Source as the tapping method, you can use a tunnel or raw endpoint where traffic is directly tunneled to GigaVUE V Series Nodes. NOTE: Select the Traffic Acquisition Method as Customer Orchestrated Source if you want to deploy an AMX application in the Monitoring Session for this Monitoring Domain.
Maximum Number of V Series Nodes per Host	Enter the maximum number of GigaVUE V Series Nodes possible to deploy in a single host. The default value is 10.
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.

- Click **Save**.
- The **VMware Fabric Launch Configuration** page appears. For details on how to deploy GigaVUE V Series Nodes in the **VMware Fabric Launch Configuration** page.
- After you deploy the fabric node, you can view the host where the monitoring domain is deployed on the **Hosts** tab and the fabric node details on the **Fabric** tab.

**Notes:**

- Ensure that all V Series Nodes within a Monitoring Domain 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.

The Monitoring Domain created in this section is listed on the **Monitoring Domain** page.



Points to Note:

- Whenever you change the vCenter credentials in VMware vCenter, edit the Monitoring Domain to update that in GigaVUE-FM. Otherwise, the connection status reaches an authentication failure state.
- When the Monitoring Domain is in a "Not Connected" state, updating the vCenter credentials will only refresh the information stored in the GigaVUE-FM database.

To establish the connection:

1. Navigate to the **Monitoring Domain** page.
2. Select **Actions** and select **Connect**.

By following these steps, you can ensure that your vCenter credentials are updated and the connection is established correctly.

You can perform the following actions in the Monitoring Domains and Fabric tabs of the Monitoring Domain page.

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.
Edit Monitoring Domain	Used to edit a monitoring domain.
Deploy Fabric	Used to deploy GigaVUE V Series Nodes.
Upgrade Fabric	Used to upgrade GigaVUE V Series Nodes. For details, refer to Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi .
Delete Monitoring Domain	Used to delete a monitoring domain.
Delete Fabric Nodes	Used to delete a GigaVUE V Series Node in the monitoring domain.
Connect	Used to start communication between GigaVUE-FM and the VMware vCenter.
Disconnect	Used to stop communication between GigaVUE-FM and the VMware

Actions	Description
	vCenter.
Edit SSL Configuration	Used to add Certificate Authority and the SSL Keys when using the Secure Tunnels.
Fabric tab	
Delete Fabric Nodes	Used to delete a specific GigaVUE V Series Node in the monitoring domain.
Rediscover	The changes made in vCenter for the GigaVUE V Series is reflected in GigaVUE-FM. Refer to Rediscover topic for more detailed information.
Power On	You can select an individual GigaVUE V Series Node and power it on. The status of the GigaVUE V Series Node is changed to Ok .
Power Off	You can select an individual GigaVUE V Series Node and power it off. If the GigaVUE V Series Node is turned off from GigaVUE-FM, then it is not considered as part of Cloud Health Monitoring and GigaVUE-FM does not try to turn it on. The status of the GigaVUE V Series Node is changed to Down .
Reboot	You can select an individual GigaVUE V Series Node and reboot it.
Rename	Used to rename the fabric nodes.
Generate Sysdump	You can select one or multiple GigaVUE V Series Nodes (Maximum 10) to generate the system files. The generation of sysdump takes a few minutes in a GigaVUE V Series Node. You can proceed with other tasks, and upon completion, the status appears in the GUI. These system files are helpful for troubleshooting. For more information, refer to Debuggability and Troubleshooting
Download Cloud Configuration File	Use this option to download the cloud configuration file (for example, md_gigamon-cloud.conf). The file contains monitoring domain connection details such as groupName, subGroupName, and token, and can be used to register fabric components in GigaVUE-FM using third party orchestration.

More information

- To view and manage the generated sysdump files, select the GigaVUE V SeriesNode and select the **Sysdump** tab in the lower pane.
- To view the certificates associated with the fabric, select the fabric nodes and select the **Certificates** tab in the lower pane.

Configure and Deploy GigaVUE Fabric Components in GigaVUE-FM

This section provides step-by-step instructions on how to deploy GigaVUE V Series Nodes in VMware vCenter Monitoring Domain.

To deploy GigaVUE V SeriesNodes using GigaVUE-FM, follow these steps:

1. Go to the **VMware Fabric Launch Configuration** page using one of the following options:
 - **Automatic Option:** Create a monitoring domain, and the page opens automatically.
 - **Manual Option:** Open **VMware Fabric Launch Configuration** page from the **Monitoring Domain** page.
 - To launch the **VMware Fabric Launch Configuration** from the Monitoring Domain, go to **Inventory > VIRTUAL > VMware vCenter (V Series)**.
 - Select **Actions > Deploy Fabric**.

The **VMware Fabric Launch Configuration** page appears.

tance:GigaVUE-FM - 6.14.01

2. On the **VMware Fabric Launch Configuration** page, enter or select the following details:

Field	Description
Datacenter	vCenter Data Center with the ESXi hosts to be provisioned with GigaVUE V Series Nodes.
Cluster	Cluster where you want to deploy the GigaVUE V Series Nodes.
V Series Node Image	Select the OVA file from the drop-down menu. You can also add OVA files when launching the fabric. To add OVA files: - Click on the Add button. - The Upload Image dialog box opens. In the Upload Image dialog box, click Browse to select the <i>gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova</i> file. - Select Upload to Server to upload the selected OVA image file to GigaVUE-FM server.
Form Factor	Instance size of the GigaVUE V Series Node. For details, refer to Prerequisites for V Series nodes on ESXi Integration. NOTE: Small Form Factor is not supported when using applications like Application Visualization, Application Metadata, Application Filtering. Select 80GB Disk Space, when using AMX Application.
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 unavailable 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.
Certificate NOTE: This field appears only when Enable Custom Certificates field is enabled.	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes. For more information, refer to Secure Communication between GigaVUE Fabric Components.
Hosts	Select the ESXi hosts for GigaVUE V Series Node deployment. NOTE: You can select up to 250 hosts manually. Select Import Host Info from file or Add Host Info Manually. Import Host Info from file: To import host details from a .csv file: - Download the .csv template file. - Enter the required values in the Excel sheet and save the file. - Select Browse and select the .csv file saved in the previous step.

Field	Description																		
	NOTE: To deploy more than one GigaVUE V SeriesNode on the same host, add more rows in the Excel sheet with the same host value for each extra GigaVUE V SeriesNode you want to deploy. If your GigaVUE-FM version is above 6.5 and GigaVUE V Series Nodes are on a version below 6.5, Name Server and MTU are not supported. Therefore, these fields in the .csv file must be empty. Add Host Info Manually: Select the ESXi hosts for GigaVUE V Series Node deployment. The Common Configuration drop-down wizard appears. Expand the Common Configuration drop-down wizard and update the following details to apply the configuration to all the selected hosts. You can expand the individual hosts and add or delete GigaVUE V Series Node. You can expand the individual GigaVUE V Series Node and modify the configurations applied in the Common Configuration. Common Configuration <table> <tr> <td>Datastore/Datastore Clusters</td><td>Select the network datastore or datastore cluster shared among all ESXi hosts.</td></tr> <tr> <td>V Series Node Name Prefix</td><td>Enter a prefix for the GigaVUE V SeriesNode name.</td></tr> <tr> <td>V Series Node Name Suffix</td><td>Enter a suffix for the GigaVUE V Series Node name.</td></tr> <tr> <td>Name Server</td><td>The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.</td></tr> <tr> <td>No. of V Series Nodes per Host</td><td>Enter the number of GigaVUE V Series Nodes to be deployed in each host.</td></tr> </table> Management <table> <tr> <td>Network</td><td>Management network for GigaVUE V Series Nodes.</td></tr> <tr> <td>IP Type</td><td>Select the management network IP type as Static or DHCP.</td></tr> <tr> <td> Gateway IP NOTE: This field appears only when the Management IP type is Static. </td><td>Gateway IP address of the Management Network.</td></tr> <tr> <td> Netmask Length NOTE: This field appears only when the </td><td>Management network's subnet mask value in CIDR format. For example, 21 for /21.</td></tr> </table>	Datastore/Datastore Clusters	Select the network datastore or datastore cluster shared among all ESXi hosts.	V Series Node Name Prefix	Enter a prefix for the GigaVUE V SeriesNode name.	V Series Node Name Suffix	Enter a suffix for the GigaVUE V Series Node name.	Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.	No. of V Series Nodes per Host	Enter the number of GigaVUE V Series Nodes to be deployed in each host.	Network	Management network for GigaVUE V Series Nodes.	IP Type	Select the management network IP type as Static or DHCP.	Gateway IP NOTE: This field appears only when the Management IP type is Static.	Gateway IP address of the Management Network.	Netmask Length NOTE: This field appears only when the	Management network's subnet mask value in CIDR format. For example, 21 for /21.
Datastore/Datastore Clusters	Select the network datastore or datastore cluster shared among all ESXi hosts.																		
V Series Node Name Prefix	Enter a prefix for the GigaVUE V SeriesNode name.																		
V Series Node Name Suffix	Enter a suffix for the GigaVUE V Series Node name.																		
Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.																		
No. of V Series Nodes per Host	Enter the number of GigaVUE V Series Nodes to be deployed in each host.																		
Network	Management network for GigaVUE V Series Nodes.																		
IP Type	Select the management network IP type as Static or DHCP.																		
Gateway IP NOTE: This field appears only when the Management IP type is Static.	Gateway IP address of the Management Network.																		
Netmask Length NOTE: This field appears only when the	Management network's subnet mask value in CIDR format. For example, 21 for /21.																		

Field	Description	
	Management IP type is Static.	
MTU		The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
Data Interfaces - When using Customer Orchestrated Source as the Traffic Acquisition Method, you must configure two data interfaces.		
Use IPv6		Enable to use IPv6.
	NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.	
Network		Tunnel Network for the GigaVUE 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.
Netmask Length		Tunnel network's subnet mask value in CIDR format. For example, 21 for /21.
	NOTE: This field appears only when the Tunnel IP type is Static.	
MTU		The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
IPv6 Prefix Length		Enter the IPv6 prefix length as 64.
	NOTE: This field appears only when the Use IPv6 toggle button is enabled.	
Virtual Disk Format		Select the Virtual Disk Format from the drop-down menu
Deployment Folder		Enter the folder name in vCenter, under which the GigaVUE V SeriesNodes must be deployed.
Password		Enter the password you wish to use for the GigaVUE V Series Node.

3. Select **Deploy**. After deployment in vCenter, it appears on the **Monitoring Domain** page under the Monitoring Domain, where the GigaVUE V Series Node is deployed.

NOTE: GigaVUE-FM can process a maximum of ten GigaVUE V Series Nodes deployment requests in parallel on VMware vCenter. Each deployment request can have multiple GigaVUE V Series Nodes for deployment.

Verify Your Deployment

Once you select Deploy, GigaVUE-FM provisions and registers the GigaVUE V Series Nodes. On the Monitoring Domain page, confirm that the node status is Ok//Healthy, which indicates a successful deployment.

What to do Next

After you successfully deploy the fabric components required for traffic acquisition, configure them within the monitoring session to direct traffic from the workload virtual machines to the appropriate tools. For instructions on setting up the session, refer to [Configure Monitoring Session](#).

Deploy GigaVUE Cloud Suite for VMware ESXi using Third Party Orchestration - Generic Mode

The GigaVUE fabric components can be deployed in generic mode for VMware ESXi platform. In Generic Mode, you can deploy GigaVUE V Series Nodes using your own orchestration system. Provide the monitoring domain name and connection name in the user data during deployment. GigaVUE-FM automatically creates the Monitoring Domain under the Third Party Orchestration monitoring domain page. It assigns the deployed fabric components to that domain after the deployment completes.

For VMware ESXi in Generic Mode, you can configure traffic acquisition using the following methods:

- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)
- [Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain](#)

Acquire Traffic using UCT-V - By Creating Monitoring Domain

In generic mode, you can also create a Monitoring Domain under Third Party Orchestration and provide the Monitoring Domain name and the connection name in the user data that used by your orchestrator.

This section outlines the workflow to acquire traffic directly from virtual machines and send it to GigaVUE V Series Nodes using UCT-V. Refer to the following topics for instruction on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Install GigaVUE-FM on VMware ESXi](#)
- [Configure Role-Based Access for Third Party Orchestration](#)
- [Configure Tokens](#)
- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)
- [Install and Register UCT-V Controller](#)
- [Install UCT-Vs](#)
- [Configure and Deploy V Series Node using VMware ESXi - Generic Mode](#)

Install GigaVUE-FM on VMware ESXi

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

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. For more details, refer to [Change Your Password](#).
 - **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.

Configure Tokens

You must configure tokens for registering GigaVUE Fabric Components like UCT-V with GigaVUE-FM. Refer to [Configure Tokens](#) for more information on creating tokens.

Create a Monitoring Domain

To create a monitoring domain in Third Party Orchestration:

1. Go to **Inventory > VIRTUAL > Third Party Orchestration**, and then click **Monitoring Domain**. The **Monitoring Domain** page appears.
2. In the Monitoring Domain page, select **New**. The **Monitoring Domain Configuration** page appears.

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

Field	Description
Monitoring Domain	An alias used to identify the monitoring domain. A monitoring domain consists of set of connections.
Connection Alias	An alias used to identify the connection.
Traffic Acquisition Method	Select a tapping method. UCT-V : UCT-Vs are deployed on your VMs to acquire the traffic and forward the acquired traffic to the GigaVUE V Series nodes. If you select UCT-V as the tapping method, you must configure the UCT-V Controller to communicate to the UCT-Vs from GigaVUE-FM. The default MTU value is 1450.
Uniform Traffic Policy (When Traffic Acquisition Method is Customer Orchestrated Source)	Enable this option if you wish to use the same monitoring session configuration for the GigaVUE V Series Node deployed under this monitoring domain. Enable this check box when using packet mirroring configuration for GCP. NOTE: Once the monitoring session is deployed for the monitoring domain you cannot enable or disable this option.
Traffic Acquisition Tunnel MTU (When Traffic Acquisition Method is UCT-V)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the UCT-V to the GigaVUE V Series Node. The MTU values for the respective platforms when using IPv4 tunnels: AWS - 8950 Azure - 1450 OpenStack - 1450 Nutanix - 1250 GCP - 1410 When using IPv4 tunnels, the MTU must be 50 bytes less than the native MTU of the respective platform. The MTU values for the respective platforms when using IPv6 tunnels: AWS - 8930 Azure - 1430 OpenStack - 1430 Nutanix - 1230 GCP - 1390 When using IPv6 tunnels, the MTU must be 70 bytes less than the native MTU of the respective platform.
Enable IPv6 Preference	Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes.

4. Click **Save**.**Notes:**

- Ensure that all V Series Nodes within a single Monitoring Domain are running 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 domain page:

Actions	Description
Edit Monitoring Domain	Use to edit a monitoring domain.
Delete Monitoring Domain	Use to delete a Monitoring Domain.
Edit SSL Configuration	Use to add Certificate Authority and the SSL Keys when using the Secure Tunnels.
Generate Sysdump	You can select one or multiple GigaVUE V Series Nodes (Maximum 10) to generate the system files. The generation of sysdump takes a few minutes in a GigaVUE V Series Node. You can proceed with other tasks, and upon completion, the status appears in the GUI. These system files are helpful for troubleshooting. For more information, refer to Debuggability and Troubleshooting
Manage Certificates	You can use this button to perform the following actions: • Re-issue- Certificates can be reissued to address security compromises, key changes, or configuration updates, like validity period adjustments. • Renew- Renewing a certificate just extends its 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. Refer to Configure Certificate Settings for more details.
Download Cloud Configuration File	Use this option to download the cloud configuration file (for example, md_gigamon-cloud.conf). The file contains monitoring domain connection details such as groupName, subGroupName, and token, and can be used to register fabric components in GigaVUE-FM using third party orchestration.

To view and manage the generated sysdump files, select the GigaVUE V Series Node and select the **Sysdump** tab in the lower pane.

To view the certificates associated with the fabric, select the fabric nodes and click the **Certificates** tab in the lower pane.

Install and Register UCT-V Controller

Recommended Instance Type

The following table lists the recommended instance type for deploying UCT-V Controller.

Compute Instances	vCPU	Memory
UCT-V Controller	2vCPU	4GB

To install UCT-V Controller in ESXi:

1. Go to the [Gigamon Community Portal](#) and download the UCT-V Controller OVA package for your target GigaVUE-FM version.
2. Extract the archive to get the .ovf and .vmdk files, these are what you upload during VM creation.
3. Log into the VMware ESXi web interface.
4. Right-click the ESXi Host, Cluster, or Datacenter where you want to deploy the UCT-V Controller and select Create/Register VM.
5. Choose Deploy a virtual machine from an OVF and OVA file, click **Next**.
6. Give the VM a unique name and upload the UCT-V Controller's .ovf and .vmdk files.
7. Select a datastore, then configure deployment options (DHCP vs. static IP for the management interface, disk provisioning)
8. It's recommended to leave "Power on automatically" unchecked until you've reviewed the configuration.
9. In Additional Settings, set hostname and an admin password, configure networking (nameserver, MTU if needed), and enter registration details: the token, GigaVUE-FM RemoteIP/RemotePort, and the Monitoring Domain and Connection name.
10. Review and click **Finish** to complete the deployment.

To register UCT-V Controller using a configuration file:

1. Download the cloud configuration file **gigamon-cloud.conf** from the Monitoring Domain Actions tab of GigaVUE-FM.
2. Edit or enter the required details in the file to register UCT-V Controller with GigaVUE-FM. A sample of the file is given as follows:

```
Registration:
  groupName: <Monitoring Domain Name>
  subGroupName: <VPC Name>
  token: <Token>
  remoteIP: <IP address of the GigaVUE-FM>
  sourceIP: <IP address of UCT-V Controller> (Optional Field)
  remotePort: 443
```

3. Log in to the UCT-V Controller.
4. Upload the configuration file in UCT-V Controller.
5. Restart the UCT-V Controller service using the following command.

```
$ sudo service uctv-cntlr restart
```

Verify UCT-V Controller Registration

1. In GigaVUE-FM, go to **Inventory > VIRTUAL > Third Party Orchestration**. The **Monitoring Domains** page appears.
2. Select your monitoring domain, then open the **Connections** tab to see the connection-level detail view.
3. Confirm the UCT-V Controller shows up as a row with its Management IP populated.
4. Verify whether the Status column displays as Ok/Healthy, confirming the controller successfully registered and is sending heartbeats.

Install UCT-Vs

You can install UCT-V on both Linux and Windows environments. For details, refer to the following topics:

- [Linux UCT-V Installation](#)
- [Windows UCT-V Installation](#)

Linux UCT-V Installation

You can install UCT-V on various Linux distributions using Debian or RPM packages.

You must have sudo/root access to edit the UCT-V configuration file. Establish an SSH connection to the virtual machine and ensure you have permission to execute the sudo command.

You may need to modify the network configuration files for dual or multiple network interface configurations to ensure that the extra NIC/Network interface initializes at boot time.

Refer to the following sections:

- [Linux Network Firewall Requirements](#)
- [Install and Register UCT-V Controller](#)
- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)

Prerequisites

- UCT-V is a standalone service. By default, most modern Linux operating systems come pre-installed with all the necessary packages for the UCT-V to function without additional configuration.
- Before registering Linux UCT-V, you should generate a token and place it in the **/etc/gigamon-cloud.conf** configuration file. Refer to .

Linux Network Firewall Requirements

If Network Firewall requirements or security groups are configured in your environment, you must open the following ports for the virtual machine.

Direction	Port	Protocol	CIDR	Purpose
Inbound	9902	TCP	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller

You can use the following commands to add the Network Firewall rule.

```
sudo firewall-cmd --add-port=9902/tcp
sudo firewall-cmd --runtime-to-permanent
```

Install Linux UCT-Vs

You can install the UCT-Vs either from Debian or RPM packages using one of the following options:

- [Install Linux UCT-Vs using Installation Script](#)
- [Install Linux UCT-Vs using Manual Configuration](#)

Refer to the following sections for more detailed information and step-by-step instructions.

Install Linux UCT-Vs using Installation Script

Using installation script, you can complete installation.

Perform the following steps:

1. To install UCT-V from Ubuntu/Debian:

- Download the UCT-V **6.14.00** Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance, contact [Gigamon Technical Support](#).
- Copy this package to your instance and Install the package with root privileges. For example,

```
$ ls gigamon-gigavue-uctv-6.14.00-amd64.deb
$ sudo dpkg -i gigamon-gigavue-uctv-6.14.00-amd64.deb
```

2. To install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS:

- a. Download the UCT-V6.14.00 RPM (.rpm) package from the [Gigamon Customer Portal](#). For assistance, contact [Gigamon Technical Support](#).
- b. Copy this package to your instance and install the package with root privileges. For example,

```
$ ls gigamon-gigavue-uctv-6.14.00-x86_64.rpm
```

```
$ sudo rpm -i gigamon-gigavue-uctv-6.14.00-x86_64.rpm
```

- Use the command given below to perform pre-check, installation, and configuration functionalities.

sudo uctv-wizard

NOTE: The installation script is not provided with the Debian or RPM packages. You can use the script (installation_wizard.sh/uctv-wizard) only after the UCT-V is installed.

Refer to the table below to know more about **uctv-wizard** command usage options and functionalities:

Options	Use Command	Description
pre-check	sudo uctv-wizard pre-check	Checks the status of the required packages and firewall requirements. - If any package is missing, it displays an appropriate message with the missing package details. - If installation includes all the packages, it displays a success message indicating that UCT-V is ready for configuration.
pkg-install	sudo uctv-wizard pkg-install NOTE: The uctv-wizard install command requires access to a repository, either public (internet-based) or local, that hosts prerequisite packages for installation. If no repository is accessible, you must manually install the required packages.	Displays the missing package and version details. To proceed with the installation, you can choose between the following: If you wish to skip the prompts and proceed with the system update, enter your option as y . The console interface installs the missing packages and restarts the UCT-V service. Enter N if you wish to install it manually. For details, refer to Install Linux UCT-Vs using Manual Configuration .
configure	sudo uctv-wizard configure	First, it checks for any existing configured file in the tmp directory (file named gigamon-cloud.conf in the C:\Users\<username>\AppData\Local (Windows) or /etc/uctv/uctv.conf (Linux) location). If available, UCT-V uses that configuration. If unavailable, UCT-V automatically adds the interface configuration in uctv.conf file, excluding the loopback (lo) interface, with all permissions enabled (source ingress, source egress, and destination). You can add the required policy for the available port if a firewall is installed. If you wish to skip the prompts to add the

Options	Use Command	Description
		required firewall policy, enter your option as y . The console interface adds the firewall rules automatically. Enter N if you wish to configure manually. For details, refer to Install Linux UCT-Vs using Manual Configuration section.
uninstall	<code>sudo uctv-wizard uninstall</code>	Automatically stops the UCT-V service, removes the firewall rules, and uninstalls the UCT-V.

**Notes:**

- Use the command below to view all the log messages generated from uctv-wizard. These log messages are stored at **/var/log/uctv-installation.log**
`sudo vi /var/log/uctv-installation.log`
- Use the command below to know the usage descriptions for the individual operations.
`sudo uctv-wizard help`

Linux UCT-V Installation Scenarios

- Zero Touch Installation** - When using a cloud-integrated script to deploy UCT-V in a virtual machine, no interference is required as the script installs and configures everything automatically.
- One Touch Installation** - When using .deb or .rpm packages with all prerequisite packages in place, UCT-V determines that all dependencies are met. It performs auto-configuration and restarts the service.
- Two Touch Installation** - When using .deb or .rpm packages with missing prerequisite packages, the platform displays a warning message about the missing packages. You need to install the missing packages using the **sudo uctv-wizard pkg-install** command.

Install Linux UCT-Vs using Manual Configuration

- [Install UCT-V from Ubuntu/Debian Package](#)
- [Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS](#)

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

- Download the UCT-V **6.14.00** Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance contact [Gigamon Technical Support](#).
- Copy this package to your instance.

3. Install the package with root privileges. For example,

```
$ ls gigamon-gigavue-uctv-6.14.00-amd64.deb
$ sudo dpkg -i gigamon-gigavue-uctv-6.14.00-amd64.deb
```

4. Modify the file **/etc/uctv/uctv.conf** to configure and register the source and destination interfaces.

The following examples registers eth0 as the mirror source for both ingress and egress traffic and eth1 as the destination for this traffic:

NOTE: When you have an active, successful monitoring session deployed, any modification to the UCT-V config file made after the initial setup requires a UCT-V restart and an inventory refresh or sync from GigaVUE-FM to pick up the new changes and re-initiate the traffic mirroring. GigaVUE-FM does a periodic sync on its own every 15 minutes.

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

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

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

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

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

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

Example 4—Configuration example to monitor ingress traffic at iface 'eth0' and egress traffic at iface 'eth1' and use iface 'eth2' to transmit the mirrored packets.

```
# eth0    mirror-src-ingress
# eth1    mirror-src-egress
# eth2    mirror-dst
```

Example 5—Configuration example to monitor traffic at iface 'lo' that is always registered as bidirectional traffic regardless of the config and use iface 'eth0' to transmit the mirrored packets.

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

NOTE: Ensure that the configuration for a single interface is provided on a single line.

5. Save the file.

6. Restart the UCT-V service.

```
$ systemctl restart uctv.service
```

The UCT-V status is displayed as running. Verify the status using the following command:

```
$ systemctl status uctv.service
```

Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS

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

1. Download the UCT-V**6.14.00** RPM (.rpm) package from the [Gigamon Customer Portal](#).
For assistance contact [Gigamon Technical Support](#).
2. Copy this package to your instance.
3. Install the package with root privileges. For example,

```
$ ls gigamon-gigavue-uctv-6.14.00-x86_64.rpm
```

```
$ sudo rpm -i gigamon-gigavue-uctv-6.14.00-x86_64.rpm
```

4. Modify the **/etc/uctv/uctv.conf** file to configure and register the source and destination interfaces.

The following example registers the eth0 as the mirror source for both ingress and egress traffic and registers eth1 as the destination for this traffic as follows:

NOTE: When you have an active, successful monitoring session deployed, any changes to the UCT-V config file made after the initial setup require an UCT-V restart and an inventory refresh or sync from GigaVUE-FM to pick up the new changes and re-initiate the traffic mirroring. GigaVUE-FM does a periodic sync on its own every 15 minutes.

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

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

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

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

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

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

Example 4—Configuration example to monitor ingress traffic at iface 'eth0' and egress traffic at iface 'eth1' and use iface 'eth2' to transmit the mirrored packets.

```
# eth0    mirror-src-ingress
# eth1    mirror-src-egress
# eth2    mirror-dst
```

Example 5—Configuration example to monitor traffic at iface 'lo' that is always registered as bidirectional traffic regardless of the config and use iface 'eth0' to transmit the mirrored packets.

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

NOTE: Ensure that the configuration for a single interface is provided on a single line.

5. Save the file.
6. Restart the UCT-V service.

```
$ sudo service uctv restart
```

The UCT-V status is displayed as running. Verify the status with the following command:

```
$ sudo service uctv status
```



Notes:

- When UCT-V fails to start due to a “**start-limit-hit**” (caused by repeated restarts within 10 minutes), you should correct the underlying issue first. To clear the failure and allow UCT-V to restart, run the following command:

```
sudo systemctl reset-failed uctv.service
```

Post Deployment Check:

After installing UCT-V, you can perform the following to verify the version of UCT-V:

1. Enter the command:

```
sudo uctvl uctv-show
```

2. Manually execute the following command:

```
export LD_LIBRARY_PATH=/usr/lib/uctv/ssl-lib64/
```

Windows UCT-V Installation

Windows UCT-V allows you to select the network interfaces by subnet/CIDR and modify the corresponding monitoring permissions in the configuration file. This options helps you get granular control over the monitored and mirrored traffic.

Refer to the following sections for the Windows UCT-V installation:

- [Windows Network Firewall Requirements](#)
- [Install and Register UCT-V Controller](#)
- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)

Windows Network Firewall Requirements

If your environment uses network firewall rules or security groups, you must open specific ports for the virtual machine.

Ensure to use UCT-V Wizard with the configure option to configure Windows local firewall rules. Refer to [Install Windows UCT-Vs using Installation Script](#) for more details.

Install Windows UCT-Vs

Rules and Notes:

- VXLAN is the only tunnel type supported for Windows UCT-V.
- Loopback Interface is not supported for Windows UCT-V.

- Before registering Windows UCT-V, generate a token and place it in the **C:\ProgramData\uctv\gigamon-cloud.conf** configuration file. Refer to .
- Do not manually edit the UCT-V Windows service registry settings. Use the supported installation and upgrade workflows instead.

You can install the UCT-Vs with MSI package using one of the following options:

- [Install Windows UCT-Vs using Installation Script](#)
- [Install Linux UCT-Vs using Manual Configuration](#)



The Windows UCT-V MSI is a self-contained package that includes all necessary dependencies. However, during set up, it automatically installs the following components:

- **Visual C++ Redistributable 2019 (x86)**
- **Npcap (v1.81 OEM)**

Before installing the Windows Agent, ensure that Npcap is not already present on the system. If an existing version of Npcap is present, uninstall it manually to avoid conflicts and ensure compatibility with the version bundled in the UCT-V.

Refer to the following sections for more detailed information and instructions.

Install Windows UCT-Vs using Installation Script

1. Download the Windows UCT-V **6.14.00** MSI package from the [Gigamon Customer Portal](#). For assistance, contact [Gigamon Technical Support](#).
2. Install the downloaded MSI package as **Administrator**.
The UCT-V service starts automatically.

- Use the command below to perform pre-check, adapter setup, adapter restore, and configuration functionalities.

uctv-wizard

Refer to the table below to know more about **uctv-wizard** command usage options and functionalities:

Options	Use Command	Description
pre-check	uctv-wizard pre-check	Checks the network adapter properties and firewall requirements. It notifies the user if the network adapter's send buffer size is smaller than the required size for the Windows UCT-V and if any firewall rules need to be added. NOTE: We recommend to Increase the send buffer size of network adapters to 128 MB during the UCT-V installation to optimize performance and minimize traffic disruption.
adapter-setup	uctv-wizard adapter-setup	Checks the compatible network adapters, increases the send buffer size and restarts the service. Before changing the buffer size, the existing configuration is saved as a backup. You can choose between the following: - If you wish to skip the prompts for changing the buffer size of compatible network adapters, enter the option as y. - Enter N if you wish to set it up manually. For details, refer to Install Linux UCT-Vs using Manual Configuration.
adapter-restore	uctv-wizard adapter-restore	Using this command, you can restore the backup copy of the network adapter buffer size configuration saved in the in the uctv-wizard adapter-setup step. NOTE: You need to manually restart the network adapters for changes to take effect immediately. You can choose between the following: If you wish to skip the prompts for restoring the buffer size of the compatible network adapters, enter the option as y.

Options	Use Command	Description
		Enter N if you wish to restore it manually. For details, refer to
configure	uctv-wizard configure	First, it checks for any existing configured file in the tmp directory (file named gigamon-cloud.conf in the C:\Users\<username>\AppData\Local location). If available, UCT-V will use that configuration. If unavailable, UCT-V automatically adds the interface configuration in uctv.conf file, excluding the loopback (lo) interface, with all permissions enabled (source ingress, source egress, and destination). You can add the required policy for the available port if a firewall is installed. - If you wish to skip the prompts to add the required firewall policy, enter your option as y. The console interface adds the firewall rules automatically. - Enter N if you wish to configure manually.
uninstall	uctv-wizard uninstall	Automatically stops the UCT-V service, removes the firewall rules, and uninstalls the UCT-V.

**Notes:**

- The log messages generated from uctv-wizard are stored at **/C:\ProgramData\uctv\uctv-installation.txt**
- Use the command below to know the usage descriptions for the individual operations.
uctv-wizard help

Windows UCT-V Installation Scenarios

- Zero Touch Installation:** When using a cloud integrated script to deploy UCT-V in a virtual machine, zero interference is required as the script installs and configures everything automatically.
- One Touch Installation:** When using a .msi package with all prerequisite packages in place, UCT-V determines that all dependencies are met. It performs auto-configuration and restarts the service.

Install Windows UCT-Vs using Manual Configuration

1. Download the Windows UCT-V **6.14.00** MSI package from the [Gigamon Customer Portal](#). For assistance, contact [Gigamon Technical Support](#).
2. Install the downloaded MSI package as **Administrator**. The UCT-V service starts automatically.

3. Modify the file **C:\ProgramData\Uct-v\uctv.conf** to configure and register the source and destination interfaces.

NOTE: When you have an active, successful monitoring session deployed, any changes to the UCT-V config file made after the initial setup require a UCT-V restart and an inventory refresh or sync from GigaVUE-FM to pick up the new changes and re-initiate the traffic mirroring. GigaVUE-FM does a periodic sync on its own every 15 minutes.



Following are the rules to modify the UCT-V configuration file:

- Interface is selected by matching its CIDR address with config entries.
- For the VMs with single interface (*.conf file modification is optional*):
 - if neither mirror-src permissions is granted to the interface, both mirror-src-ingress and mirror-src-egress are granted to it.
 - mirror-dst is always granted implicitly to the interface.
- For the VMs with multiple interfaces:
 - mirror-dst is granted explicitly in the config file. Only the first matched interface is selected for mirror-dst, all other matched interfaces are ignored.
 - if none interfaces is granted any mirror-src permission, all interfaces are granted mirror-src-ingress and mirror-src-egress.

Example 1—Configuration example to monitor ingress and egress traffic at interface 192.168.1.0/24 and use the same interface to send out the mirrored packets.

For IPv4:

```
192.168.1.0/24 mirror-src-ingress mirror-src-egress mirror-dst
```

For IPv6:

```
2001:db8:abcd:ef01::/64 mirror-src-ingress mirror-src-egress mirror-dst
```

Example 2—Configuration example to monitor ingress and egress traffic at interface 192.168.1.0/24 and use the interface 192.168.2.0/24 to send out the mirrored packets.

For IPv4:

```
192.168.1.0/24 mirror-src-ingress mirror-src-egress
```

```
192.168.2.0/24 mirror-dst
```

For IPv6:

```
2001:db8:abcd:ef01::/64 mirror-src-ingress mirror-src-egress
```

```
2001:db8:abcd:ef01::2/64 mirror-dst
```

4. Save the file.

5. Restart the Windows UCT-V using one of the following actions:
 - From the command prompt, run **sc stop uctv** and **sc start uctv**.
 - From the Windows Task Manager, restart the UCT-V.

You can verify the status of the UCT-V in the Service tab of the Windows Task Manager.

Verify UCT-V Deployment

1. In GigaVUE-FM, go to **Inventory > VIRTUAL > Third Party Orchestration**. The Monitoring Domains page appears.
2. Select your monitoring domain, then open the **Connections** tab to see the connection-level detail view.
3. Confirm the deployed UCT-V agent shows up with its Monitoring Domain, IP address, Registration time, Last heartbeat time, Agent mode, and Status.
4. Confirm the Status column reads Ok/Healthy, indicating that the agent is registered successfully with GigaVUE-FM through the UCT-V Controller and is sending heartbeats.

Configure and Deploy V Series Node using VMware ESXi - Generic Mode

This topic provides instructions on how to deploy the fabric components for VMware ESXi.

NOTE: When registering GigaVUE V Series Nodes in GigaVUE-FM, make sure that the connection name under each Monitoring Domain is unique. When GigaVUE-FM version is 6.10.00 or above and the Fabric Components are on (n-1) or (n-2) versions, you must create a Username and Password instead of using tokens in the registration data. For more details, refer to the Configure Role-Based Access for Third-Party Orchestration section in the 6.9 Documentation.

Recommended Instance Type

The following table lists the recommended instance type for deploying the fabric components:

Compute Instances	vCPU	Memory
GigaVUE V Series Node	4vCPU	8GB

You can deploy the GigaVUE V Series Node using the following deployment modes based on your requirement.

- Single uplink for Management, Data and Tool connectivity.
- Single uplink for Management and Tool connectivity, and another uplink for Data connectivity.
- (Recommended) Separate uplinks (three) for Management, Data, and Tool connectivity.

Prerequisites

The following prerequisites are required for deployment of the GigaVUE V Series Node in VMware ESXi.

- **Configuring port groups:** Create a Management Port Group for connectivity with GigaVUE-FM, a Data Port Group to receive data from the H Series node, and a Tool Port Group for connectivity with the tools.
- **Configuring virtual switch:** Create unique virtual switches for each port group. Refer to [VMware Documentation](#) for more information.
- Download the OVA files from the [Gigamon Community portal](#) and extract it to get the OVF and VMDK files.

Deploy and Register GigaVUE V SeriesNode

To deploy and register:

1. Log into the VMware ESXi web interface.
2. Right-click the ESXi Host, Cluster, or data center on which you want to deploy the GigaVUE V Series Node and then select **Create/Register VM**. The **New Virtual Machine** wizard appears.
3. In the **Select Creation Type** page, select Deploy a virtual machine from an OVF and OVA file option.
4. Click **Next**. The **Select OVF and VMDK files** page appears.
5. In the **Select OVF and VMDK files** page, enter a unique name for the virtual machine and upload the .ovf and .vmdk files from your local machine.
6. Click **Next** and the **Select storage** page appears.
7. Select a datastore where the virtual machine's files will be stored.
8. Click **Next**. The **Deployment Options** page appears.
9. In the **Deployment Options** page, select the management port, data port, and tool port as referenced in the prerequisites.
 - a. Select the **Deployment Type** from the list:
 - Do Not Use DHCP – Select this option if you want to use static IP addresses for the management, data, and tool ports.
 - Management, Data and Tool Port DHCP – Select this option if you want to use dynamic IP addresses for the management, data, and tool ports.
 - Management Port DHCP – Select this option if you want to use dynamic IP address only for the management port.
 - Tool Port DHCP - Select this option if you want to use dynamic IP address only for the tool port.
 - Data Port DHCP – Select this option if you want to use dynamic IP address only for the data port.

- b. Select Thin/Thick in the **Disk Provisioning** field.
 - c. Unselect the **Power on automatically** checkbox. This is selected by default. It is recommended to disable the Power on automatically option and review all configuration before switching on the virtual machine.
10. Click **Next**. The **Additional Settings** page appears.
11. Do the following configuration in the **Additional Settings** page.

- a. In the **System** section, enter the hostname of the V Series node instance in the **Hostname** field and create a new admin password for the V Series node instance in the **Administrative Login Password** field.

This credential will be used for V series SSH access. The default username is gigamon. If the deployment fails, you can login to V Series SSH/Console and check the logs for troubleshooting.

- b. In the **Network Connectivity** section, enter the required fields based on the selected network configuration.

NOTE: Ensure to unselect the **Management Port DHCP** checkbox if you want to use static IP address for the management port. If you select the Management Port DHCP checkbox, dynamic IP address will be configured for the management port even if you have selected Do Not Use DHCP option in the Configuration page.

- c. Enter the required value in the **Management Port MTU size in bytes** field. The default value is 1500B.
 - d. Enter the DNS server address in the **Nameserver** field to resolve the domain name of the tool destination URL.
 - e. In the **Optional Parameters** section, enter the Monitoring Domain name in the **GroupName** field and connection name in the **SubGroupName** field. The Monitoring Domain and connection name corresponds to the domain name and connection that you created during the Create Monitoring Domain configuration.
 - f. Enter the token created in GigaVUE-FM in the **JWT Token used for registration** field. Refer to [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#).
 - g. Enter the GigaVUE-FM IP address and remote port in the **RemoteIP** and **RemotePort** fields respectively. In case of any error, you should re-deploy the V Series.
 - h. In the **Custom node properties** field enter app_mode=linux_apps (mandatory).
12. Click **Next** and the **Ready to complete** page appears.

13. Review all the entered information and modify the configuration if required.
14. Click **Finish**. When the operation completes, you have successfully deployed a GigaVUE V Series Node.

Verify GigaVUE V Series Node Deployment

During the initial bring up, the V Series reboots multiple times for initialization. After few minutes, you can check the status of the deployment in GigaVUE-FM. If the status is failed, you can check the logs to perform troubleshooting.

To check the logs:

1. Log in to V Series Node via console or SSH to management IP address.
2. Enter the below path in the terminal to check the registration details.

```
tail -f /var/log/vseries-node-reg.log
```

OVF Package Files

Form Factor	Supported Ports	File Name	Comments
Small (2vCPU, 4GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file1.ovf	Use these files when deploying GigaVUE V Series Node via VMware vCenter.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file2.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file3.ovf	
Small (2vCPU, 4GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file4.ovf	Use these when deploying GigaVUE V Series Node via VMware NSX-T Manager.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file5.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 2 Network Ports	vseries-node-file6.ovf	
Small (2vCPU, 4GB Memory, and 8GB Disk space)		vseries-node-file7.ovf	Use these files when deploying GigaVUE V Series Node via VMware ESXi without vCenter.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file8.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file9.ovf	
Larger (8vCPU, 16GB Memory, and 80GB)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file12.ovf	Use these files when deploying GigaVUE V Series Node via

Form Factor	Supported Ports	File Name	Comments
Disk space)			VMware vCenter and if you wish to configure AMX application.
Larger (8vCPU, 16GB Memory, and 80GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file15.ovf	Use these files when deploying GigaVUE V Series Node via VMware ESXi without vCenter and if you wish to configure AMX application. NOTE: This file supports form factor with higher range of CPU, memory and disk space.
	Mgmt Port , Data Port, and Tool Port	vseries-node-file16.ovf	minipc - Virtual Small Form Factor

Assign Static IP address for GigaVUE V Series Node

By default, the GigaVUE V Series gets assigned an IP address using DHCP.

The static IP addresses are assigned to the GigaVUE V Series node in the following scenarios:

- When you have assigned or selected DHCP for the port groups during deployment and want to change it to static IP address after deployment.
- When you have assigned static IP address for the port groups during deployment and want to update the assigned static IP address. Assign Static IP address for UCT-V Controller and V Series (Same snippet file for both)

To assign a static IP address, perform the following steps:

1. Navigate to **/etc/netplan/** directory.
2. Create a new **.yaml** file.

NOTE: Do not use the default 50-cloud-init.yaml file.

3. Update the file as shown in the following sample:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    <interface>:                # Replace with your actual interface name (e.g., eth0)
      dhcp4: no
      dhcp6: no
      addresses:
        - <IPV4/24>            # e.g., 192.168.1.10/24
```

```

- <IPV6/64>                # e.g., 2001:db8:abcd:0012::1/64
nameservers:
  addresses:
    - <DNS_IPV4>            # e.g., 8.8.8.8
    - <DNS_IPV6>            # e.g., 2001:4860:4860::8888
  routes:
    - to: 0.0.0.0/0
      via: <IPV4_GW>        # e.g., 192.168.1.1
    - to: ::/0
      via: <IPV6_GW>        # e.g., 2001:db8:abcd:0012::fffe

```

Example netplan config:

```

network:
  version: 2
  renderer: NetworkManager
  ethernets:
    ens3:
      addresses:
        - 10.114.53.24/21
      dhcp4: no
      dhcp6: no
      accept-ra: false
      routes:
        - to: 10.114.48.1/32
          scope: link
        - to: default
          via: 10.114.48.1

```

4. Save the file.

5. Apply the configuration.

```
$ sudo netplan apply
```

6. Restart the GigaVUE V Series service.

```
$ sudo service vseries-node restart
```

The deployed GigaVUE V Series registers with the GigaVUE-FM. After successful registration the GigaVUE V Series sends heartbeat messages to GigaVUE-FM every 30 seconds. If one heartbeat is missing, the fabric component status appears as **Unhealthy**. If more than five heartbeats fail to reach GigaVUE-FM, GigaVUE-FM tries to reach the GigaVUE V Series. If that fails as well then GigaVUE-FM unregisters the GigaVUE V Series and removes from GigaVUE-FM.

What to do Next

After you successfully deploy the fabric components required for traffic acquisition, configure them within the monitoring session to direct traffic from the workload virtual machines to the appropriate tools. For instructions on setting up the session, refer to [Configure Monitoring Session](#).

Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain

In generic mode, you can also create a Monitoring Domain under Third Party Orchestration and provide the Monitoring Domain name and the connection name in the user data that used by your orchestrator.

This section outlines the workflow to tunnel traffic directly to GigaVUE V Series Nodes without deploying UCT-V or its controllers. using Customer Orchestrated Source. Refer to the following topics for instruction on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Install GigaVUE-FM on VMware ESXi](#)
- [Configure Role-Based Access for Third Party Orchestration](#)
- [Create a Monitoring Domain](#)
- [Configure and Deploy GigaVUE Fabric Components using VMware ESXi - Generic Mode](#)
- [Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain](#)

Install GigaVUE-FM on VMware ESXi

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

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. For more details, refer to [Change Your Password](#).
- **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.

Configure Tokens

You must configure tokens for registering GigaVUE Fabric Components like UCT-V with GigaVUE-FM. Refer to [Configure Tokens](#) for more information on creating tokens.

Create a Monitoring Domain

To create a monitoring domain in Third Party Orchestration:

1. Go to **Inventory > VIRTUAL > Third Party Orchestration**, and then click **Monitoring Domain**. The **Monitoring Domain** page appears.
2. In the Monitoring Domain page, select **New**. The **Monitoring Domain Configuration** page appears.

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

Field	Description
Monitoring Domain	An alias used to identify the monitoring domain. A monitoring domain consists of set of connections.
Connection Alias	An alias used to identify the connection.
Traffic Acquisition Method	Select a tapping method. The available options are: Customer Orchestrated Source : If you select the Customer Orchestrated Source option, the mirrored, tunneled or the raw traffic from your workloads is directed directly to the GigaVUE V Series Nodes, and you need not configure the UCT-Vs and UCT-V Controllers.
Uniform Traffic Policy (When Traffic Acquisition Method is Customer Orchestrated Source)	Enable this option if you wish to use the same monitoring session configuration for the GigaVUE V Series Node deployed under this monitoring domain. Enable this check box when using packet mirroring configuration for GCP. NOTE: Once the monitoring session is deployed for the monitoring domain you cannot enable or disable this option.
Traffic Acquisition Tunnel MTU (When Traffic Acquisition Method is UCT-V)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the UCT-V to the GigaVUE V Series Node. The MTU values for the respective platforms when using IPv4 tunnels: AWS - 8950 Azure - 1450 OpenStack - 1450 Nutanix - 1250 GCP - 1410 When using IPv4 tunnels, the MTU must be 50 bytes less than the native MTU of the respective platform. The MTU values for the respective platforms when using IPv6 tunnels: AWS - 8930 Azure - 1430 OpenStack - 1430 Nutanix - 1230 GCP - 1390 When using IPv6 tunnels, the MTU must be 70 bytes less than the native MTU of the respective platform.
Enable IPv6 Preference	Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes.

4. Click **Save**.**Notes:**

- Ensure that all V Series Nodes within a single Monitoring Domain are running 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 domain page:

Actions	Description
Edit Monitoring Domain	Use to edit a monitoring domain.
Delete Monitoring Domain	Use to delete a Monitoring Domain.
Edit SSL Configuration	Use to add Certificate Authority and the SSL Keys when using the Secure Tunnels.
Generate Sysdump	You can select one or multiple GigaVUE V Series Nodes (Maximum 10) to generate the system files. The generation of sysdump takes a few minutes in a GigaVUE V Series Node. You can proceed with other tasks, and upon completion, the status appears in the GUI. These system files are helpful for troubleshooting. For more information, refer to Debuggability and Troubleshooting
Manage Certificates	You can use this button to perform the following actions: • Re-issue- Certificates can be reissued to address security compromises, key changes, or configuration updates, like validity period adjustments. • Renew- Renewing a certificate just extends its 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. Refer to Configure Certificate Settings for more details.
Download Cloud Configuration File	Use this option to download the cloud configuration file (for example, md_gigamon-cloud.conf). The file contains monitoring domain connection details such as groupName, subGroupName, and token, and can be used to register fabric components in GigaVUE-FM using third party orchestration.

To view and manage the generated sysdump files, select the GigaVUE V Series Node and select the **Sysdump** tab in the lower pane.

To view the certificates associated with the fabric, select the fabric nodes and click the **Certificates** tab in the lower pane.

Configure and Deploy GigaVUE Fabric Components using VMware ESXi - Generic Mode

This topic provides instructions on how to deploy the fabric components for VMware ESXi.

NOTE: When registering GigaVUE V Series Nodes in GigaVUE-FM, make sure that the connection name under each Monitoring Domain is unique. When GigaVUE-FM version is 6.10.00 or above and the Fabric Components are on (n-1) or (n-2) versions, you must create a Username and Password instead of using tokens in the registration data. For more details, refer to the Configure Role-Based Access for Third-Party Orchestration section in the 6.9 Documentation.

Recommended Instance Type

The following table lists the recommended instance type for deploying the fabric components:

Compute Instances	vCPU	Memory
GigaVUE V Series Node	4vCPU	8GB

You can deploy the GigaVUE V Series Node using the following deployment modes based on your requirement.

- Single uplink for Management, Data and Tool connectivity.
- Single uplink for Management and Tool connectivity, and another uplink for Data connectivity.
- (Recommended) Separate uplinks (three) for Management, Data, and Tool connectivity.

Prerequisites

The following prerequisites are required for deployment of the GigaVUE V Series Node in VMware ESXi.

- **Configuring port groups:** Create a Management Port Group for connectivity with GigaVUE-FM, a Data Port Group to receive data from the H Series node, and a Tool Port Group for connectivity with the tools.
- **Configuring virtual switch:** Create unique virtual switches for each port group. Refer to [VMware Documentation](#) for more information.
- Download the OVA files from the [Gigamon Community portal](#) and extract it to get the OVF and VMDK files.

Deploy and Register GigaVUE V SeriesNode

To deploy and register:

1. Log into the VMware ESXi web interface.
2. Right-click the ESXi Host, Cluster, or data center on which you want to deploy the GigaVUE V Series Node and then select **Create/Register VM**. The **New Virtual Machine** wizard appears.
3. In the **Select Creation Type** page, select Deploy a virtual machine from an OVF and OVA file option.

4. Click **Next**. The **Select OVF and VMDK files** page appears.
5. In the **Select OVF and VMDK files** page, enter a unique name for the virtual machine and upload the .ovf and .vmdk files from your local machine.
6. Click **Next** and the **Select storage** page appears.
7. Select a datastore where the virtual machine's files will be stored.
8. Click **Next**. The **Deployment Options** page appears.
9. In the **Deployment Options** page, select the management port, data port, and tool port as referenced in the prerequisites.
 - a. Select the **Deployment Type** from the list:
 - Do Not Use DHCP – Select this option if you want to use static IP addresses for the management, data, and tool ports.
 - Management, Data and Tool Port DHCP – Select this option if you want to use dynamic IP addresses for the management, data, and tool ports.
 - Management Port DHCP – Select this option if you want to use dynamic IP address only for the management port.
 - Tool Port DHCP - Select this option if you want to use dynamic IP address only for the tool port.
 - Data Port DHCP – Select this option if you want to use dynamic IP address only for the data port.
 - b. Select Thin/Thick in the **Disk Provisioning** field.
 - c. Unselect the **Power on automatically** checkbox. This is selected by default. It is recommended to disable the Power on automatically option and review all configuration before switching on the virtual machine.
10. Click **Next**. The **Additional Settings** page appears.
11. Do the following configuration in the **Additional Settings** page.
 - a. In the **System** section, enter the hostname of the V Series node instance in the **Hostname** field and create a new admin password for the V Series node instance in the **Administrative Login Password** field.

This credential will be used for V series SSH access. The default username is gigamon. If the deployment fails, you can login to V Series SSH/Console and check the logs for troubleshooting.
 - b. In the **Network Connectivity** section, enter the required fields based on the selected network configuration.

NOTE: Ensure to unselect the **Management Port DHCP** checkbox if you want to use static IP address for the management port. If you select the Management Port DHCP checkbox, dynamic IP

address will be configured for the management port even if you have selected Do Not Use DHCP option in the Configuration page.

- c. Enter the required value in the **Management Port MTU size in bytes** field. The default value is 1500B.
 - d. Enter the DNS server address in the **Nameserver** field to resolve the domain name of the tool destination URL.
 - e. In the **Optional Parameters** section, enter the Monitoring Domain name in the **GroupName** field and connection name in the **SubGroupName** field. The Monitoring Domain and connection name corresponds to the domain name and connection that you created during the Create Monitoring Domain configuration.
 - f. Enter the token created in GigaVUE-FM in the **JWT Token used for registration** field. Refer to [Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain](#).
 - g. Enter the GigaVUE-FM IP address and remote port in the **RemoteIP** and **RemotePort** fields respectively. In case of any error, you should re-deploy the V Series.
 - h. In the **Custom node properties** field enter app_mode=linux_apps (mandatory).
12. Click **Next** and the **Ready to complete** page appears.
 13. Review all the entered information and modify the configuration if required.
 14. Click **Finish**. When the operation completes, you have successfully deployed a GigaVUE V Series Node.

Verify GigaVUE V Series Node Deployment

During the initial bring up, the V Series reboots multiple times for initialization. After few minutes, you can check the status of the deployment in GigaVUE-FM. If the status is failed, you can check the logs to perform troubleshooting.

To check the logs:

1. Log in to V Series Node via console or SSH to management IP address.
2. Enter the below path in the terminal to check the registration details.

```
tail -f /var/log/vseries-node-reg.log
```

OVF Package Files

Form Factor	Supported Ports	File Name	Comments
Small (2vCPU, 4GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file1.ovf	Use these files when deploying GigaVUE V Series Node via VMware vCenter.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file2.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file3.ovf	
Small (2vCPU, 4GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file4.ovf	Use these when deploying GigaVUE V Series Node via VMware NSX-T Manager.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file5.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 2 Network Ports	vseries-node-file6.ovf	
Small (2vCPU, 4GB Memory, and 8GB Disk space)		vseries-node-file7.ovf	Use these files when deploying GigaVUE V Series Node via VMware ESXi without vCenter.
Medium (4vCPU, 8GB Memory, and 8GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file8.ovf	
Large (8vCPU, 16GB Memory, and 8GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file9.ovf	
Larger (8vCPU, 16GB Memory, and 80GB Disk space)	Mgmt Port, Tool Port, and 8 Network Ports	vseries-node-file12.ovf	Use these files when deploying GigaVUE V Series Node via VMware vCenter and if you wish to configure AMX application.
Larger (8vCPU, 16GB Memory, and 80GB Disk space)	Mgmt Port , Data Port, and Tool Port	vseries-node-file15.ovf	Use these files when deploying GigaVUE V Series Node via VMware ESXi without vCenter and if you wish to configure AMX application. NOTE: This file supports form factor with higher range of CPU, memory and disk space.
	Mgmt Port , Data Port, and Tool Port	vseries-node-file16.ovf	minipc - Virtual Small Form Factor

Assign Static IP address for GigaVUE V Series Node

By default, the GigaVUE V Series gets assigned an IP address using DHCP.

The static IP addresses are assigned to the GigaVUE V Series node in the following scenarios:

- When you have assigned or selected DHCP for the port groups during deployment and want to change it to static IP address after deployment.
- When you have assigned static IP address for the port groups during deployment and want to update the assigned static IP address. Assign Static IP address for UCT-V Controller and V Series (Same snippet file for both)

To assign a static IP address, perform the following steps:

1. Navigate to **/etc/netplan/** directory.
2. Create a new **.yaml** file.

NOTE: Do not use the default 50-cloud-init.yaml file.

3. Update the file as shown in the following sample:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    <interface>:                # Replace with your actual interface name (e.g., eth0)
      dhcp4: no
      dhcp6: no
      addresses:
        - <IPV4/24>              # e.g., 192.168.1.10/24
        - <IPV6/64>              # e.g., 2001:db8:abcd:0012::1/64
      nameservers:
        addresses:
          - <DNS_IPV4>            # e.g., 8.8.8.8
          - <DNS_IPV6>            # e.g., 2001:4860:4860::8888
      routes:
        - to: 0.0.0.0/0
          via: <IPV4_GW>          # e.g., 192.168.1.1
        - to: ::/0
          via: <IPV6_GW>          # e.g., 2001:db8:abcd:0012::fffe
```

Example netplan config:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    ens3:
      addresses:
        - 10.114.53.24/21
      dhcp4: no
```

```
dhcp6: no
accept-ra: false
routes:
  - to: 10.114.48.1/32
    scope: link
  - to: default
    via: 10.114.48.1
```

4. Save the file.

5. Apply the configuration.

```
$ sudo netplan apply
```

6. Restart the GigaVUE V Series service.

```
$ sudo service vseries-node restart
```

The deployed GigaVUE V Series registers with the GigaVUE-FM. After successful registration the GigaVUE V Series sends heartbeat messages to GigaVUE-FM every 30 seconds. If one heartbeat is missing, the fabric component status appears as **Unhealthy**. If more than five heartbeats fail to reach GigaVUE-FM, GigaVUE-FM tries to reach the GigaVUE V Series. If that fails as well then GigaVUE-FM unregisters the GigaVUE V Series and removes from GigaVUE-FM.

What to do Next

After you successfully deploy the fabric components required for traffic acquisition, configure them within the monitoring session to direct traffic from the workload virtual machines to the appropriate tools. For instructions on setting up the session, refer to [Configure Monitoring Session](#).

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.

Select the topic that matches the orchestration method and traffic acquisition model used in the deployment.

GigaVUE-FM Orchestration

- [a name potraDeploy GigaVUE Cloud Suite for VMware ESXi using Terraform](#)

Third-Party Orchestration

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

a name potraDeploy GigaVUE Cloud Suite for VMware ESXi using Terraform

Platform Mirroring on VMware ESXi uses the GigaVUE-FMorchestration deployment model. Use the Terraform deployment files to create or update the monitoring domain, vCenter connection, GigaVUE V Series deployment, and monitoring session in GigaVUE-FM. You can refer to theVMware ESXIsample configuration as a reference when building your deployment files.

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 in the VMware ESXi documentation](#).
- 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](#).

NOTE: If GigaVUE-FM is not installed or is not reachable from the deployment system, Terraform cannot create or update the VMware ESXi deployment objects.

2. Ensure that the VMware ESXi environment and vCenter are available.

Terraform uses the VMware ESXi user-created deployment files to apply deployment values such as vCenter, datacenter, cluster, host, datastore, and network. For more details refer to [Prerequisites for VSN VMware-ESXi](#).

3. Download the GigaVUE V Series image from the community portal.

Ensure that the image is accessible from the deployment system. Terraform uses this path during deployment.

4. Install and verify Terraform 1.14 or later

Run the `terraform version` command and confirm that respective 1.14 or later version 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 VMware ESXi Deployment Files

Before you run the deployment, create the Terraform configuration files for your VMware ESXi environment. Sample configuration files for the VMware ESXi deployment path are available in the [Terraform Registry](#).

Use the Gigamon Terraform provider documentation on the Terraform Registry as your reference when writing the configuration. The provider documentation lists all supported resources, required and optional arguments, and usage examples.

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

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

1. [Single-File Configuration](#)
2. [Module-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 deployment file can include the following blocks in any .tf file:

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 VMware ESXi object that Terraform must create or manage. Define the resources in the following order:

- Monitoring domain
- vCenter connection
- GigaVUE V Series node configuration
- Monitoring session

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.

Example:

```
module "monitoring_domain" {
```

```

    source      = "./modules/monitoring-domain"
    fm_ip       = var.fm_ip
    domain_alias = var.domain_alias
  }

  module "vseries_nodes" {
    source              = "./modules/vseries-nodes"
    monitoring_domain_id = module.monitoring_domain.id
    vcenter_address    = var.vcenter_address
    # Add required variables for your environment.
  }

  module "monitoring_session" {
    source              = "./modules/monitoring-session"
    monitoring_domain_id = module.monitoring_domain.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 Registry](#) documentation for the full list of resources and arguments for each module.

Run the Terraform Deployment File

After you finish the deployment-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. Validate the configuration.

- Optional: 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 VMware ESXi objects, including:
 - [Monitoring domain](#)
 - [vCenter connection](#)
 - [GigaVUE V Series deployment](#)
 - [Monitoring session](#)
- 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 VMware monitoring domain and the connection, applies the GigaVUE V Series deployment values, and configures the monitoring session objects 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 VMware ESXi 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.

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 VMware ESXi monitoring domain, the vCenter connection, the GigaVUE V Series deployment, 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.

©

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 configuration](#)[Module-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 Secure Tunnel for Third Party Orchestration

You can configure the Secure tunnels on:

- [Precrypted Traffic](#)
- [Mirrored Traffic](#)

Precrypted Traffic

Secure tunnels help protect sensitive precrypted traffic by framing packets in PCAPng format and transmitting them through a TLS socket. When you enable secure tunnel for both precrypted and mirrored traffic, the system creates two separate TLS sessions.

Recommendation: Always enable secure tunnels for precrypted traffic to ensure secure data transmission.

For more information about PCAPng, refer to [PCAPng Application](#).

Mirrored Traffic

You can also enable the Secure Tunnel for mirrored traffic. By default, Secure Tunnel is disabled.

Refer to the following sections for Secure Tunnel Configuration:

- [Configure Secure Tunnel for Third Party Orchestration](#) in UCT-V
- [Configure Secure Tunnel for Third Party Orchestration](#)

Prerequisites

- While creating Secure Tunnel, you must provide the following details:
 - SSH key pair
 - CA Certificate chain
- Port 11443 is enabled in security group settings. For details, refer to [Network Firewall Requirement](#).

Notes

- Protocol versions IPv4 and IPv6 are supported.
- If you wish to use IPv6 tunnels, your GigaVUE-FM and the fabric components version must run version 6.6.00 or above.
- For UCT-V with a version lower than 6.6.00, if the secure tunnel is enabled in the Monitoring Session, secure mirror traffic is transmitted over IPv4, regardless of IPv6 preference.
- After configuring secure tunnels, if a Monitoring Domain contains only one GigaVUE V Series Node and that GigaVUE V Series Node reboots or restarts, then you must manually add the SSL Key pairs to the Monitoring Domain again. For details, refer to [Edit SSL Configuration](#).

Configure Secure Tunnel from UCT-V to GigaVUE V Series Node

To configure a secure tunnel in UCT-V, you must configure one end of the tunnel to the UCT-V and the other end to GigaVUE V Series Node. You must configure the CA Certificate chain in UCT-V and the private keys and SSL certificates in GigaVUE V Series Node. Refer to the following steps for configuration:

S. No	Task	Refer to
1.	Upload a CA Certificate chain in a single file	You must upload a CA Certificate chain to UCT-V Controller to establish a connection with the GigaVUE V Series Node. To upload the CA Certificate chain using GigaVUE-FM, follow the steps given below: 1. Go to Inventory > Resources > Security > CA List. 2. Select New to add a new Authority. The Add Custom Authority page appears. 3. Enter or select the following information. - o Alias - Alias name of the CA. - o File Upload - Choose the certificate from the desired location. 4. Select Save. NOTE: Ensure that the Intermediate CA Certificate(s) are included in the CA certificate chain file in the correct signing order, followed by the Root CA Certificate at the end. For more information, refer to the section Adding Certificate Authority.
2.	Upload a SSL Key Pair	You must add a SSL key pair to GigaVUE V Series Node. To add SSL Key Pair, follow the steps in the section SSL Decrypt
3.	Select the SSL Key Pair and CA Certificate chain, after deploying the fabric components.	You must select the added SSL Key Pair and CA Certificate chain in GigaVUE V Series Node after creating a Monitoring Domain configuring the fabric components in GigaVUE-FM. Refer to Edit SSL Configuration for more detailed information on how to select the added SSL Key Pair and CA Certificate chain in GigaVUE V Series Node.
4.	Enable the secure tunnel	You should enable the secure tunnel feature to establish a connection between the UCT-V and GigaVUE V Series Node. To enable the secure tunnel feature follow these steps: 1. In the Edit Monitoring Session page, click Options. The Monitoring Session Options page appears. 2. Enable the Secure Tunnel button. You can enable secure tunnel for both mirrored and precrypted traffic.

Configure Secure Tunnel between GigaVUE V Series Nodes

You can create secure tunnel:

- Between two GigaVUE V Series Nodes.
- From one GigaVUE V Series Node to multiple GigaVUE V Series Nodes.

You must have the following details before you start configuring secure tunnels between two GigaVUE V Series Nodes:

- IP address of the tunnel destination endpoint (Second GigaVUE V Series Node).
- SSH key pair (pem file).

To configure secure tunnel between two GigaVUE V Series Nodes, refer to the following steps:

S. No	Task	Refer to
1.	Upload a CA Certificate chain in a single file	You must upload a CA Certificate chain to UCT-V Controller to establish a connection with the GigaVUE V Series Node. To upload the CA Certificate chain using GigaVUE-FM, follow the steps given below: 1. Go to Inventory > Resources > Security > CA List. 2. Select New to add a new Authority. The Add Custom Authority page appears. 3. Enter or select the following information. - o Alias - Alias name of the CA. - o File Upload - Choose the certificate from the desired location. 4. Select Save. NOTE: Ensure that the Intermediate CA Certificate(s) are included in the CA certificate chain file in the correct signing order, followed by the Root CA Certificate at the end. For more information, refer to the section Adding Certificate Authority.
2.	Upload an SSL Key Pair	You must add an SSL key pair to GigaVUE V Series node. To add an SSL Key Pair, follow the steps in the section Upload SSL Keys.
3	Select the SSL Key Pair and CA Certificate chain, after deploying the fabric components.	You must select the added SSL Key Pair and CA Certificate chain in GigaVUE V Series Node after creating a Monitoring Domain configuring the fabric components in GigaVUE-FM. Refer to Edit SSL Configuration for more detailed information on how to select the added SSL Key Pair and CA Certificate chain in GigaVUE V Series Node.
4	Create a secure tunnel between UCT-V and the first GigaVUE V Series Node	You should enable the secure tunnel feature to establish a connection between the UCT-V and the first GigaVUE V Series Node. To enable the secure tunnel feature follow these steps: 1. In the Edit Monitoring Session page, click Options. The Monitoring Session Options page appears. 2. Enable the Secure Tunnel button. You can enable secure tunnel for both mirrored and preencrypted traffic.
5	Create an Egress tunnel from the first GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session.	You must create a tunnel for traffic to flow out from the first GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session. Refer to Create Ingress and Egress Tunnels for more detailed information on how to create tunnels. To create the egress tunnel, follow these steps: 1. After creating a new Monitoring Session, or click Actions > Edit on an existing Monitoring Session, the GigaVUE-FM canvas appears.

S. No	Task	Refer to
		2. In the canvas, select New > New Tunnel, drag and drop a new tunnel template to the workspace. The Add Tunnel Spec quick view appears. 3. On the New Tunnel quick view, enter or select the required information as described in the following table: • Alias - The name of the tunnel endpoint. • Description - The description of the tunnel endpoint. • Type - Select TLS-PCAPNG for creating egress secure tunnel • Traffic Direction - Choose Out (Encapsulation) for creating an egress tunnel from the V Series node to the destination. Select or enter the following values: - o MTU- The default value is 1500. - o Time to Live - Enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255. The default value is 64. - o DSCP - Enter the Differentiated Services Code Point (DSCP) value. - o Flow Label - Enter the Flow Label value. - o Source L4 Port- Enter the Source L4 Port value - o Destination L4 Port - Enter the Destination L4 Port value. - o Flow Label - o Cipher- Only SHA 256 is supported. - o TLS Version - Select TLS Version 1.3. - o Selective Acknowledgments - Choose Enable to turn on the TCP selective acknowledgments. - o SYN Retries - Enter the value for number of times the SYN has to be tried. The value ranges from 1 to 6. - o Delay Acknowledgments - Choose Enable to turn on delayed acknowledgments. • IP Version - The version of the Internet Protocol. IPv4 and IPv6 are supported. • Remote Tunnel IP - Enter the interface IP address of the second GigaVUE V Series Node (Destination IP). 4. Select Save.
6	Select the added SSL Key Pair after deploying the fabric components in the second GigaVUE V Series Node	You must select the added SSL Key Pair in the second GigaVUE V Series Node. Select the second GigaVUE V Series Node and follow the steps given in Edit SSL Configuration .
7	Create an ingress tunnel in the	You must create a ingress tunnel for traffic to flow in from the first

S. No	Task	Refer to
	second GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session for the second GigaVUE V Series Node.	GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session. Refer to Create a Monitoring Session to know about Monitoring Session. To create the ingress tunnel, follow these steps: 1. After creating a new Monitoring Session, or click Actions > Edit on an existing Monitoring Session, the GigaVUE-FM canvas appears. 2. In the canvas, select New > New Tunnel, drag and drop a new tunnel template to the workspace. The Add Tunnel Spec quick view appears. 3. On the New Tunnel quick view, enter or select the required information as described in the following table: - o Alias - The name of the tunnel endpoint. - o Description - The description of the tunnel endpoint. - o Type - Select TLS-PCAPNG for creating egress secure tunnel. - o Traffic Direction - Choose In (Decapsulation) for creating an ingress tunnel that receives traffic from V Series node 1. Select or enter the values as described in Step 6. - o IP Version - The version of the Internet Protocol. IPv4 and IPv6 are supported. - o Remote Tunnel IP - Enter the interface IP address of the first GigaVUE V Series Node (Destination IP). 4. Select Save.

Create Prefiltering Policy Template

GigaVUE-FM allows you to create a prefiltering policy template with one or more rules. You can configure a rule with one or more filters. A single monitoring session supports a maximum of 16 rules.

To create a prefiltering policy template,

1. Go to **Traffic > Resources > Prefiltering** and select **UCT-V**.
2. Select **New**.
3. In the **Template Name** field, enter the name of the template,
4. In the **Rule Name** field, enter the name of a rule.
5. Select one of the following options:
 - Pass: Allows the traffic.

- Drop: Blocks the traffic..

NOTE: If no prefilter rules are defined, traffic is implicitly allowed. When rules are defined, an implicit drop rule applies. Traffic that does not match any specified rule is dropped.

6. Select one of the following options:

- **Bi-Directional:** Allows the traffic in both directions of the flow. A single Bi-direction rule requires 1 Ingress and 1 Egress rule.
- **Ingress:** Filters incoming traffic.
- **Egress:** Filters outgoing traffic.

NOTE: When using loopback interface in Linux UCT-V, you can use only Bi-directional.

7. Select a priority value from 1-8.

- 1: Select the value as 1 to pass or drop a rule in top priority.
- 2-8 Select the value as 2, 3, 4 to 8, where 8 indicates a rule with the least priority.

Drop rules are added first based on the priority and then pass rules are added.

8. Select one of the following options as **Filter Type**:

- L3
- L4

9. Select one of the following options **Filter Name**:

- ip4Src
- ip4Dst
- ip6Src
- ip6Dst
- Proto: Applies to both ipv4 and ipv6.

10. Select one of the following options for **Filter Relation**:

- Not Equal to
- Equal to

11. In the **Value** field, enter the source or destination port.

12. Select **Save**.

NOTE: Select + to add more rules or filters or select - to remove a rule or a filter.

To enable prefiltering, refer to [Monitoring Session Options](#).

Create Precryption Template for UCT-V

GigaVUE-FM allows you to filter packets during Precryption in the Data Acquisition at the UCT-V level. This filtering is based on L3/L4 5 tuple information (5-tuple filtering) and the applications running on the workload virtual machines.

Rules and Notes:

- Selective Precryption works with GigaVUE-FM and the fabric components version 6.8.00 or above.
- When a single UCT-V is associated with two different Monitoring Sessions with contrasting pass and drop rules, then instead of prioritizing a single rule, GigaVUE-FM passes all the traffic.
- Once the templates are associated with a Monitoring Session, the changes made in the template are not reflected in the Monitoring Session.

Refer to the section the following sections for more detailed information:

- [Create Precryption Template for Filtering based on Applications](#)
- [Create Precryption Template for Filtering based on L3-L4 details](#)

Create Precryption Template for Filtering based on Applications

The application filter allows you to select the applications for which you apply Precryption in the Monitoring Session Options page.

To create,

1. Step Go to **Traffic > Resources > Precryption**.

The **Precryption Policies** page appear

2. Step Select the **APPLICATION** tab.
3. Select **Add**.

The New Precryption Template page appears.

4. Select **csv** as the **Type**, if you wish to add applications using a .csv file.
 - a. Download the sample .csv file and edit it.
 - b. Save your .csv file.
 - c. Select **Choose File** and upload the file.

5. Select **Manual** as the **Type** if you wish to add the applications manually.
6. Enter the **Application Name** select + icon to add more applications.
7. Select **Save**.

You can view the added applications in the **APPLICATION** tab.

You can delete a selected application or you can delete all the application using the **Actions** button.

Create Precryption Template for Filtering based on L3-L4 details

To create,

1. Go to **Traffic > Resources > Precryption**. The **Precryption Policies** page appears.
2. Select the **L3-L4** tab.

3. Perform the following steps:

- a. In the **Template** field, enter a name for the template.
- b. In the **Rule Name** field, enter a name for the rule.
- c. For **Action**, select one of the following options:

- **Pass:** Passes the traffic.
- **Drop:** Drops the traffic.

NOTE: In the absence of a Precryption rule, traffic is implicitly allowed. However, the defined rules include an implicit pass all rule. Should the traffic not conform to any of the specified rules, it is passed.

- d. For **Direction**, select one of the following options:

- **Bi-Directional:** Allows the traffic in both directions of the flow. A single Bi-direction rule should consist of 1 Ingress and 1 Egress rule.
- **Ingress:** Filters the traffic that flows in.
- **Egress:** Filters the traffic that flows out.

- e. In the **Priority** field, select one of the following values:

- 1: Select to pass or drop a rule in top priority.
- 2 to 8: Select to decide priority where 8 is used for setting a rule with the least priority.

NOTE: Drop rules are added based on the priority, and then pass rules are added.

f. Select **Filter Type** from the following options:

- L3
- L4

NOTE: You can use L4 Filter Type only with L3.

For L3, perform the following:

i. Select **Filter Name** from the following options:

- IPv4 Source
- IPv4 Destination
- IPv6 Source
- IPv6 Destination
- Protocol: It is common for both IPv4 and IPv6.

ii. Select **Filter Relation** from any one of the following options:

- Not Equal to
- Equal to

iii. Enter or Select the Value based on the selected **Filter Name**.

NOTE: When using **Protocol** as **Filter Name**, select **TCP** from the drop-down menu.

For L4, perform the following:

i. Select the **Filter Name** from the following options:

- Source Port
- Destination Port

ii. Select the **Filter Relation** from any one of the following options:

- Not Equal to
- Equal to

iii. Enter the source or destination port value.

4. Select **Save**.

NOTE: Select + to add more rules or filters. Select - to remove a rule or a filter.

The template is successfully created. To enable Precryption, refer to [Monitoring Session Options](#) section.

You can delete a selected template or you can delete all the templates using the **Actions** button.

You can also edit a selected template using **Actions > Edit**.

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 you want to monitor. You can also choose to monitor egress, ingress, or all traffic.

When a new target VM is added to your environment, GigaVUE-FM automatically detects it. If the detected VM matches the defined selection criteria, your monitoring session displays it.

Similarly, when a traffic monitoring target VM is removed, the monitoring sessions are updated to show the removed instance.

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

Limitations:

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

Refer to the following topics for details:

- [Create a Monitoring Session \(VMware\)](#)
- [Monitoring Session Page \(VMware\)](#)
- [Create Raw Endpoint \(VMware vCenter\)](#)
- [Create a New Map \(VMware ESXi\)](#)
- [V Series Applications](#)

- [Interface Mapping](#)
- [Deploy Monitoring Session](#)
- [View Monitoring Session Statistics](#)
- [Visualize the Network Topology \(VMware ESXi\)](#)
- [Configure VMware Settings](#)

Create a Monitoring Session (VMware)

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.

Monitoring Session Page (VMware)

The Monitoring Session page provides a comprehensive view of each monitoring session in your VMware environment. This topic provides detailed insights into session configuration, traffic analysis, and network topology.

Tabs Menu

Tab	Description
Overview	Displays key information about the selected monitoring session, such as connections, tunnel and health status, deployment state, and information related to Application Intelligence statistics. You can view traffic trends (ingress/egress) hourly, daily, weekly, or monthly. You can filter statistics by associated elements. For more information, refer to View Monitoring Session Statistics
Sources	Lists all monitored sources and targets. You can view/edit connection details, check deployment status, number of targets, and source health. For OVS Mirroring, the source tab also displays hypervisor and instance details. 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.
Traffic Processing	Allows you to add and configure applications, tunnel endpoints, raw endpoints, and maps. View stats for each application, apply threshold templates, enable user-defined apps, and toggle distributed de-duplication. For more information, refer to Configure Monitoring Session Options (VMware ESXi)
V Series Nodes	Shows the V Series Nodes tied to the session. In the split view, you can view the node name, health, deployment status, host VPC, version, and management IP. You can also change the interfaces mapped to an individual GigaVUE V Series Node. For more information, refer to the Interface Mapping .
Topology	Visualizes the monitored network fabric, showing connections, subnets, and instances. You can explore topology by selecting specific connections, offering a clear visualization of the monitored network elements. For more information, refer to Visualize the Network Topology (VMware ESXi) .

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

Actions Button

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.

Sidebar

You can use the  icon on the left side of the Monitoring Session page to view the Monitoring Sessions list.

Select  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 (VMware ESXi)

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 vCenter)

GigaVUE-FM 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 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 1: 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	

Field	Description	
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.	
	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	Unique value which is used to identify the VXLAN. The value

Field	Description	
	Identifier	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 Ericsson SCP Support. Applicable Platforms: OpenStack, Third Party Orchestration, 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	

Field	Description	
	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 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.

Field	Description	
	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.
	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		

Field	Description	
Tunnel page. This allows you to conveniently configure secure tunnels on your physical device . For details, refer to Secure Tunnels section.		
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.

Create Raw Endpoint (VMware vCenter)

This topic provides instructions on configuring a RAW Endpoint (REP) in the Monitoring Session.

Rules and Notes

Ingress REP works only when you select the Traffic Acquisition Method as **Customer Orchestrated Source** while configuring the Monitoring Domain.

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

Consideration

When deploying GigaVUE V Series Nodes in the Monitoring Domain, the number of interfaces varies based on the Traffic Acquisition Method. See the table below for details.

Traffic Acquisition Methods - Customer Orchestrated Source

Display Name	Interface Name	Role	Comments
Network Adapter 1	ens160	Management	
Network Adapter 2	ens192	Data	Supports Tunnel and RAW endpoint. Available for Ingress and Egress REP
Network Adapter 3	ens224	Data	Supports Tunnel and RAW endpoint. Available for Ingress and Egress REP

Traffic Acquisition Methods - Platform Tapping

Display Name	Interface Name	Role	Comments
Network Adapter 1	ens160	Management	
Network Adapter 2	ens192	Data	Supports Tunnel and Egress RAW endpoint.
Network Adapter 3 - 10	-	Data	Reserved and used for platform tapping (Port Mirroring)

Configure Raw Endpoint in Monitoring Session

Raw End Point (REP) is used to pass traffic from an interface. REP is used to ingress data from a physical interface attached to GigaVUE V Series Nodes. You can optionally use this end point to send traffic to the applications deployed in the monitoring session.

To configure,

1. From the New expand menu, drag and drop **New Raw Endpoint** to the graphical workspace.
2. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
3. Enter the Alias and Description details for the Raw End Point and select **Save**.
4. Perform the following steps to deploy the Monitoring Session after adding the Raw Endpoint:
 - a. From the **Actions** drop-down list on the **TRAFFIC PROCESSING** page, select **Deploy**. The **Deploy Monitoring Session** dialog box appears.
 - b. Select the V Series Nodes to deploy the Monitoring Session.
 - c. From the drop-down menu of the selected individual V Series Nodes, select the interfaces for each REPs and the TEPs deployed in the Monitoring Session.
 - d. Select **Deploy**.
5. Select **Export** to download all or selected V Series Nodes in the CSV and XLSX formats.

Create a New Map (VMware ESXi)

A map filters traffic that flows through GigaVUE V Series Nodes. It includes one or more rules, each defining what traffic to match. Traffic can match one or more rules in the map.

Refer to this table to configure traffic filter maps and enable automatic target selection in GigaVUE V Series Nodes.

Parameters to create a map

Parameter	Description
Rules	A rule (R) includes filtering conditions that traffic must match. It also defines traffic direction (ingress or egress) and target.
Priority	Sets the rule execution order. Use values from 1 (highest) to 5 (lowest).
Pass	Sends traffic from the VM to the destination.
Drop	Drops the traffic from the virtual machine 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 be included for monitoring. This map is used only for target selection.

Exclusion Map	An exclusion map determines the instances to be excluded from monitoring. This map is used only for target selection.
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.
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 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.

Create a Map

To create a 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 required information as described below.

- **Name:** Name of the new map
- **Description:** Details of the map
- **Application Filtering:** Enable this option if you wish to use Application Filtering Intelligence. Enabling this option allows you to filter traffic based on Application name or family. For details, refer to [Application Filtering Intelligence](#).
- **Selected Virtual Machines:** Using this option, you can select an individual Network adapter of a virtual machine as a target. Both ingress and egress traffic of the network adapter chosen are redirected for monitoring. You can also view and filter the list of virtual machines available. When using this option, you cannot use Automatic Target Selection (ATS)

- **Virtual Machine List:** Save the list or filter.

To save,

- Select **Virtual Machine List**.

The Virtual Machine List quick view opens.

- Select the virtual machines you wish to use as the target.
- Select **Apply** in the Virtual Machine List quick view to save your changes.

To filter,

- Select **Filter** and select one of the following criteria:
 - Data Center
 - Cluster
 - Host Name
 - VM Name
 - VM Tag Category
 - VM Tag Name
- After selecting the details, select **Apply** in the filter dialog box to apply the filters.

The list of virtual machines appears based on the filter criteria.

- Select the virtual machines you wish to use as the target, and select **Apply** in the Virtual Machine List quick view to save your changes.
- VMware tools are not required to discover targets, as GigaVUE-FM can discover targets with ATS using the tags attached to the VMs.
- Targets are selected when you provide the VM's node name or the hostname as selection criteria. A host is selected when the hostname matches all the active targets.

- Pass and Drop rule selection with Automatic Target Selection (ATS) differs 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 and perform the following:

a. Create a new rule set:

- Click **Actions > New Ruleset**.
- Enter a **Priority** value from 1 to 5 for the rule with 1 being the highest and 5 is the lowest priority.
- Enter the Application Endpoint in the Application EndPoint ID field.
- Select a required condition from the drop-down list.
- Select the rule to **Pass** or **Drop** through the map.

b. Create a new rule:

- Click **Actions > New Rule**.
- Select a required condition from the drop-down list.

iii. Select  and select **Add Condition**.

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 [Example- Create a New Map using Inclusion and Exclusion Maps](#).

You can also perform the following actions on the Monitoring session canvas.

- To edit a map, select the  menu button of the required map on the canvas and select **Details**, or select **Delete** to delete the map.
- To apply the 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 rules 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.

To create a map,

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. In the **RULESETS** tab, enter the priority and Application Endpoint ID.
3. 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.
4. Select the Expand icon at the bottom of the Monitoring session canvas. The Inclusion Maps and Exclusion Maps section appears.
5. Drag and drop a new map template to the Inclusion Maps region. The New Map quick view appears. Enter the Name and Description of the map.
 - a. In the **GENERAL** tab, enter the name as Inclusionmap1 and enter the description. In the **RULESETS**, enter the priority and Application Endpoint ID.
 - b. Select the condition as VM Name and enter the VM Name as **target-1**. Then the instance with VM name **target-1-1**, **target-1-2**, and **target-1-3** will be included.
6. Drag and drop a new map template to the Exclusion Maps region. The New Map quick view appears. 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	✓	✓

GigaSMA RT 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
					party Orchestration)		
Application Metadata Intelligence	✓	✓	✓	✓	✓	✓	✓
Application Visualization	✓	✓	✓	✓	✓	✓	✓
De- duplication	✓	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Encapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	x	✓	✓	✓	✓	✓
VXLAN Tunnel Encapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on	✓	✓	✓	✓	✓	✓	✓

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
the retrospective 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.	✓	x	✓	✓	✓	✓	✓
VXLAN Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment	✓	✓	✓	✓	✓	✓	✓

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
Guide.							
ERSPAN Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✓
UDPGRE Tunnel Decapsulation Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✗
GigaSMART NetFlow Generation	✓	✓	✓	✓	✓	✓	✓

GigaSMA RT 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
GRE-In-UDP Tunnel Decapsulation	x	x	✓	✓	x	✓	x
GENEVE Decapsulation	✓	x	x	x	✓	x	x
Header Stripping	✓	✓	✓	✓	✓	✓	✓
Load Balancing	✓	✓	✓	✓	✓	✓	✓
Masking	✓	✓	✓	✓	✓	✓	✓
SSL Decrypt	✓	✓	✓	✓	x	✓	✓
Secure Tunnels Encapsulation	x	✓	x	✓	✓	✓	✓
Secure Tunnels Decapsulation	x	✓	x	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓	✓	✓
5G-Cloud Application	x	x	✓	✓	x	✓	x
5G-Service Based Interface Application	x	x	✓	✓	✓	✓	x

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.

Application	Masking
Alias	masking
Protocol	none
Offset	
Pattern	
Length	

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 and a main configuration area. The configuration fields are as follows:

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.

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.



Field	Value
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.

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
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>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	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																	
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	REGEXP	\fupload\(create_file new_slice upload_slice)\?.*upload_token=.*
	cts-server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEXP	(.*\.)?gigamon\.com
	mime_type	MIME Type	Content type of Request or the Web page	Both, Client to Server or Server to Client	REGEXP	http
	cts-user_agent	User Agent	Software / Browser used for request	Client to Server Only	REGEXP	mozilla
	cts-referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEXP	http://gigamon.com/
	stc-server_agent	Server Agent	Software used for the server	Server to Client Only	REGEXP	NWS_TCloud_PX
	stc-location	Redirect Location	Destination address where the client	Server to Client Only	REGEXP	.*\football\.*

			is redirecte d to			
	cts- cookie	Cookie (Raw)	Raw value of the HTTP Cookie header line	Client to Server Only	REGEXP	.*tEstCoOkie.*
	content	Conten t	Message body content	Both, Client to Server or Server to Client	REGEXP	.*GIGAMON.* mindata = 206 Refer Mindata
http2	cts-uri	Reque st URI	Partially Normaliz ed URL (path + request)	Client to Server Only	REGEXP	\/fupload\/(create_file new_ slice upload_slice)\?.*upload_ token=.*
	cts- server	Server Name	Web Server Name from URI or Host	Client to Server Only	REGEXP	(.*\.)?gigamon\.com
	cts- user_ agent	User Agent	Software / Browser used for request	Client to Server Only	REGEXP	mozilla
	cts- referer	Referer URI	Source address where client got the URI	Client to Server Only	REGEXP	http:\/gigamon.com\/
ssl	commo n_name	Domai n Name	Domain name from Client		REGEXP	(.*\.)?gigamon\.com

			Hello message or the certificate			
	stc-subject_alt_name	Subject Alt Name(s)	List of host names which belong to the same certificate	Server to Client Only	REGEXP	(.*\.)?gigamon\.com
rtmp	cts-page_url	Page URL	URL of the webpage where the audio/video content is streamed	Client to Server Only	REGEXP	http://www.music.tv/recorded/1234567
tcp	stream	Payload Data	Data payload for a packet, excluding the header.		REGEXP	.*GIGAMON.* mindata = 70 Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEXP String	80-4350
udp	stream	Payload Data	Data payload for a		REGEXP	.*GIGAMON.* mindata = 100

			packet, excluding the header			Refer Mindata
	port	Server Port	Server (listen) port number		UINT16 RANGE as REGEXP String	80-4350
sip	user_agent	User Agent	Software used	Both, Client to Server or Server to Client	REGEXP	GVUE-release 6.2.0
icmp	code	Message Code	Code of the ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEXP String	200
	typeval	Message Type	Type of ICMP message	Both, Client to Server or Server to Client	UINT8 as REGEXP String	10
ip	address	Server IP Addresses	IP address of the server		IPV4 as REGEXP String	62.132.12.30\24
	dscp	DSCP Value	DSCP from Different		UINT8 as REGEXP	33

			ia ted Service (DS) Field in IP header		String	
	resolv_ name	DNS Name	Server's DNS name		REGEXP	gigamon.com
ipv6	address	Server IP Addres s	IP address of the server		IPV6 as REGEXP String	2001:0:9d38:6ab8:307b:16a 4:9c66:5f4 2001:0:9d38::9c66:5f4/64
	dscp	DSCP Value	DSCP from Different ia ted Service (DS) Field in IP header		UINT8 as REGEXP 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.																		
	<table><tr><th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr><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>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr></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	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																
	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. You can view the list of attributes/applications available in the fast mode by navigating to the app editor under AMI feature in the GigaVUE-FM. If all the attributes of the application are disabled then the application itself is disabled.																		

- 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 Even 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																				
	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. <table><tr><th>Protocol</th><th>Attribute</th></tr><tr><td>http</td><td>rtt</td></tr><tr><td>icmp</td><td>rtt</td></tr><tr><td>icmp6</td><td>rtt</td></tr><tr><td>ssh</td><td>rtt</td></tr><tr><td>tcp</td><td>rtt</td></tr><tr><td>tcp</td><td>rtt_app</td></tr><tr><td>telnet</td><td>rtt</td></tr><tr><td>wsp</td><td>connect_rtt</td></tr><tr><td>wsp</td><td>query_rtt</td></tr></table>			Protocol	Attribute	http	rtt	icmp	rtt	icmp6	rtt	ssh	rtt	tcp	rtt	tcp	rtt_app	telnet	rtt	wsp	connect_rtt	wsp	query_rtt
Protocol	Attribute																						
http	rtt																						
icmp	rtt																						
icmp6	rtt																						
ssh	rtt																						
tcp	rtt																						
tcp	rtt_app																						
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).
Template Refresh Interval	Enter the time interval at which the template must be refreshed in seconds

Parameter	Description
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 Physical and Input Name.
IP	Select the parameter as Version if required.

Parameter	Description
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.
- 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-HCI-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 ESXi 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 port groups (ESXi/vCenter) or subnets (other platforms): a Management Port Group for connectivity to GigaVUE-FM, a Data Port Group 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:

- (Optional) Install GigaVUE-FM using ISO Image File on VMware ESXi Host
- (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 V Series Node and Register with GigaVUE-FM
- 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 using ISO Image File on VMware ESXi Host

This step is optional and applies only when an existing GigaVUE-FM instance is not available.

To install:

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:
 - Enter a name for the virtual machine.
 - Select the **Guest OS family** as Linux.
 - Select the **Guest OS version** as Other Linux (64-bit).
 - 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.
- b. In the **CD/DVD Drive**, select **Datastore ISO file**. The Datastore browser appears.
- c. Select the ISO image uploaded in Step 3. 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.

Access GigaVUE-FM

1. Log in to GigaVUE-FM using the IP address available on the ESXi web Interface.
2. Use the following default login Credentials:
 - Username: admin
 - Password: admin123A!!
3. On first login, you will be asked to change the password. Enter the current password, then set and confirm a new one.

(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 V Series Node and Register with GigaVUE-FM

The following image illustrates a scenario for deploying and registering a GigaVUE V Series node. In this example, a GigaVUE HC Series device monitors north-south traffic and forwards Application Metadata to a SIEM/observability tool that accepts only JSON records. The entire solution is managed through GigaVUE-FM.

The GigaVUE HC Series uses Application Metadata Intelligence to generate application metadata in CEF format. Application Metadata Exporter, running on the GigaVUE V Series node, converts these CEF records to JSON before exporting them to the tool over HTTPS or Kafka.

NOTE: AMI is supported on both GigaVUE HC Series and GigaVUE V Series platforms. You can also deploy a separate GigaVUE V Series node to generate application metadata for monitoring east-west traffic.



To deploy and register:

1. Log into the vSphere or VMware ESXi web interface.
2. Right-click the ESXi Host, Cluster or data center (if using vSphere) on which you want to deploy the GigaVUE V Series Node and then select **Create/Register VM**. The **New Virtual Machine** wizard appears.
3. In the **Select Creation Type** page, select Deploy a virtual machine from an OVF and OVA file option. Downloaded the AMX supported OVA from the [VUE Gigamon Community portal](#). Extract the OVA to obtain the OVF and VMDK files, then select **vseries-node-file-ovf15.ovf** (AMX-supported OVF that brings up three interfaces) for the AMX deployment.
4. Click **Next**. The **Select OVF and VMDK files** page appears.
5. In the **Select and VMDK files** page, enter a unique name for the virtual machine and upload the .ovf and .vmdk files from your local machine.
6. Click **Next** and the **Select storage** page appears.
7. Select a datastore where the virtual machine's files will be stored.
8. Click **Next**. The **Deployment Options** page appears.
9. In the **Deployment Options** page, select the management port, data port, and tool port as referenced in the prerequisites.
 - a. Select the **Deployment Type** from the list:
 - Do Not Use DHCP – Select this option if you want to use static IP addresses for the management, data, and tool ports.
 - Management, Data and Tool Port DHCP – Select this option if you want to use dynamic IP addresses for the management, data, and tool ports.
 - Management Port DHCP – Select this option if you want to use dynamic IP address only for the management port.
 - Tool Port DHCP - Select this option if you want to use dynamic IP address only for the tool port.
 - Data Port DHCP – Select this option if you want to use dynamic IP address only for the data port.
 - b. Select Thin/Thick in the **Disk Provisioning** field.
 - c. Unselect the **Power on automatically** check box. This is selected by default. It is recommended to disable the Power on automatically option and review all configuration before switching on the virtual machine.
10. Click **Next**. The **Additional Settings** page appears.
11. Do the following configuration in the **Additional Settings** page.

- a. In the **System** section, enter the hostname of the V Series node instance in the **Hostname** field and create a new admin password for the V Series node instance in the **Administrative Login Password** field.

This credential will be used for V series SSH access. The default username is gigamon. If the deployment fails, you can login to V Series SSH/Console and check the logs for troubleshooting.

- b. In the **Network Connectivity** section, enter the required fields based on the selected network configuration.

NOTE: Ensure to unselect the **Management Port DHCP** checkbox if you want to use static IP address for the management port. If you select the Management Port DHCP checkbox, dynamic IP address will be configured for the management port even if you have selected Do Not Use DHCP option in the Configuration page.

- c. Enter the required value in the **Management Port MTU size in bytes** field. The default value is 1500B.
 - d. Enter the DNS server address in the **Nameserver** field to resolve the domain name of the tool destination URL.
 - e. In the **Optional Parameters** section, enter the Monitoring Domain name in the **GroupName** field and connection name in the **SubGroupName** field. The Monitoring Domain and connection name corresponds to the domain name and connection that you created during the [Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source](#) configuration.
 - f. Enter the token created in GigaVUE-FM in the **JWT Token used for registration** field. Refer to [Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM](#).
 - g. Enter the GigaVUE-FM IP address and remote port in the **RemoteIP** and **RemotePort** fields respectively. In case of any error, you should re-deploy the V Series.
 - h. In the **Custom node properties** field enter app_mode=linux_apps (mandatory).
12. Click **Next** and the **Ready to complete** page appears.
 13. Review all the entered information and modify the configuration if required.
 14. Click **Finish**. When the operation completes, you have successfully deployed a GigaVUE V Series Node.

Verify GigaVUE V Series Node Registration

During the initial bring up, the V Series reboots multiple times for initialization. After few minutes, you can check the status of the deployment in GigaVUE-FM. If the status is failed, you can check the logs to perform troubleshooting.

To check the logs:

1. Log in to V Series Node via console or SSH to management IP address.
2. Enter the below path in the terminal to check the registration details.

```
tail -f /var/log/vseries-node-reg.log
```

A single AMX instance processes traffic within the published performance KPIs. Deploy additional instances when traffic volume exceeds these thresholds or when packet drops occur, or CPU or memory utilization remains consistently high, to maintain stable and efficient performance. For details regarding performance KPIs, [Contact Technical Support](#).

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 5 - Configure Identity Enrichment Metadata using Zscaler Private Access \(ZPA\)](#)

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.

Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

Prerequisites

Before configuring identity enrichment using Zscaler Private Access (ZPA), ensure that the following prerequisites are met.

Zscaler Infrastructure

The following Zscaler components must already be set up in your environment. These are standard ZPA deployment requirements. For detailed instructions, refer to the Zscaler documentation.

- **Acquire Zscaler Tenant Access:** Ensure that you have admin access to the Zscaler Private Access admin portal (ZPA tenant). Refer to [Zscaler Tenant Access](#) for more information.
- **Integrate Identity Provider (IdP):** An IdP (such as Okta, Microsoft Azure AD, or any supported IdP) must be configured and integrated with ZPA. ZPA uses the IdP to authenticate users before granting access to private applications. Refer to [About External Identity Providers](#) for more information.
- **Configure App Connector Group :** At least one App Connector group must be created and have active App Connectors registered to it. You can use an existing App

Connector group or create a new one before proceeding. For instructions, refer to the Zscaler documentation. Refer to [Configuring App Connector Groups](#) for more information.

- **Install App Connector:** The Zscaler App Connector package must be downloaded and installed on a Linux-based VM (Ubuntu, plain Linux, or similar) in the domain where your private applications reside. Refer to the Zscaler App Connector deployment guide for platform-specific instructions (Linux, Docker, GCP, Kubernetes, and so on). Refer to [App Connectors](#) for more information.
- **Define Private Applications:** Application segments must be created for all private applications you want to monitor (for example, internal web apps, Jira, databases). Each application segment must include the application IP address and port range, and must be linked to a server group and an App Connector group. Refer to [Configuring Defined Application Segments](#) for more information.

Gigamon Infrastructure

- **Deploy VSN Node with AMX on ESXi:** The Application Metadata Exporter (AMX) must be deployed on a V Series Node (VSN). In this release, only the ESXi platform is supported. The IP address of the VSN node on which AMX is running, as it is required when configuring the ZPA Log Receiver. Refer to [Deploy V Series Node and Register with GigaVUE-FM](#) for more information.
- **Configure AMI:** Application Metadata Intelligence (AMI) must be active and generating flow records. The AMI metadata (source IP, destination IP, source port, destination port) is used to match and enrich flows with ZPA log data. For more information on configuring AMI, refer to [Configure Application Metadata Intelligence for Virtual Environment](#).

Integrate ZPA with Gigamon AMX

After the Zscaler infrastructure and Gigamon infrastructure are set up, configure a Log Receiver in the ZPA admin portal to stream ZPA logs to AMX. This is the key step that establishes the integration between ZPA and Gigamon.

The following table lists the specific options you must select when configuring a log receiver to integrate with AMX. For navigation steps and detailed instructions, refer to [Configuring Log Receiver](#) topic in Zscaler documentation.

Parameters	Description
Log Receiver	
Name	Enter a name for the log receiver.
Domain or IP Address	Enter the domain or IP address to the IP address of the VSN node on which AMX is deployed (not the ESXi host IP address).
TCP Port	Enter the same port that is configured as the ZPA LSS ingestor port on AMX.
App Connector Groups	Select an app connector group that was created.

Parameters	Description
Log Stream	
Log Type	Select User Activity as the log type.
Log Template	Select the Custom template.
Log Stream Content	Add the new user activity log ("ConnectionSetupTime": %j {ConnectionSetupTime} }\n) in the content.

Configure Ingestor for AMX

You can configure metadata enrichment for ZPA only. Optionally, you can combine it with GEM for Cloud Workloads for VMware vCenter.

To export identity enrichment metadata using ZPA :

1. Configure an ingestor for AMX
 - a. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
 - b. In the **Type** field, select **AMI** as the input type.
 - c. Enter the port number on which application metadata must be ingested. The default port value is 514.
 - d. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
2. Configure identity metadata enrichment using ZPA
 - a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
 - b. Select the **Enable** option to enable the enrichment.
 - c. In the **Type** field, select **Identity** as the input type.
 - d. Select **Zscaler Private Access** from the **Platform** dropdown list.
 - e. Select the attributes for identity enrichment from the **Attributes** tab.
 - f. In the **Source Information** tab, enter the port number on which AMX receives the ZPA LSS logs.
 - g. Click  to add another ZPA source information to add multiple inputs to the AMX application and click  to remove an existing source information.
3. (Optional) Configure GigaVUE Enriched Metadata workload for ESXi

- a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
- b. Select the Enable option to enable the enrichment.
- c. In the **Type** field, select **Workload** as the input type.
- d. Select **VMware vCenter** from the **Platform** dropdown list.

NOTE: Only VMware vcenter platform is applicable as additional workload along with ZPA.

- e. Select the required workload attributes for ESXi.
- f. In the Source Information tab, enter the Key-Value pair details as given:
 - Key - url and Value - Primary URL for the vCenter Server SDK. For example, https://<vCenter-IP-or-FQDN>/sdk.
 - Key - username and Value - VMware vCenter username.
 - Key - password and Value - vCenter password.
- g. Click  to add another ESXi source information to add multiple inputs to the AMX application and click  to remove an existing source information.

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 AML 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: - o Enter the Key. - o Enter the Value.

Fields	Description
	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.
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.
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.

Fields	Description
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.

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 in VMware ESXi using FM Orchestration

This section outlines the work flow for deploying Application Metadata Exporter application on VMware ESXi platform using GigaVUE-FM Orchestration.

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 port groups (ESXi/vCenter) or subnets (other platforms): a Management Port Group for connectivity to GigaVUE-FM, a Data Port Group 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 GigaVUE-FM Orchestration:

- [Install GigaVUE-FM using ISO Image File on VMware ESXi](#)
- [Upload GigaVUE V Series Node Image into GigaVUE-FM](#)
- [\(Optional\) Configure Metadata Export Filter for AMX](#)
- [\(Optional\) Integrate Private CA](#)
- [Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source](#)
- [Prerequisites](#)
- [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](#)

Install GigaVUE-FM using ISO Image File on VMware ESXi

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

Prerequisites

- 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.

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.

Access GigaVUE-FM

1. Log into GigaVUE-FM using the IP address available on the ESXi web interface.
2. Use ith default credentials: In the VM console, enter the default credentials:
 - Username: admin
 - Password: admin123A!!
3. On first login, you will be asked to change the password. Enter the current password (admin123A!!), then set and confirm a new one.
4. 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](#).

Upload GigaVUE V Series Node Image into GigaVUE-FM

This step is optional. You can upload the GigaVUE V Series Node image to GigaVUE-FM even while deploying the GigaVUE V Series Node.

To upload the V Series image into GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, and select **Settings > OVA Files**.
2. In the **OVA Files** page, select **Browse**, and then select the image file. For example, *gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova*. You can upload three OVA files to GigaVUE-FM for VMware vCenter.
3. Select **Upload to Server**. You have uploaded the OVA image file to the GigaVUE-FM server.

(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) Integrate Private CA

You can integrate your own PKI infrastructure with GigaVUE-FM.

Rules and Notes

- Always place the root CA in a separate file.
- When using multiple intermediate CAs, consider the following:
 - Include all intermediate CAs in a single file in the correct order.
 - Place the last intermediate CA in the chain at the top.
 - Place the preceding CAs in descending order.

To integrate private CA

- Generate a Certificate Signing Request (CSR).
- Get a signature of the Certificate Authority (CA) on the CSR.
- Upload CA Certificate to GigaVUE-FM.

Generate CSR

To create an intermediate CA certificate:

1. Go to  > **System > Certificates**.
2. In the top navigation bar, from the **PKI** drop-down list, select **CSR**. The **Generate Intermediate CA Certificate** page appears.
3. Enter details in the following fields:
 - **Country:** Enter the name of your country.
 - **Organization:** Enter the name of your organization.
 - **Organization Unit:** Enter the name of the department or unit.
 - **Common Name:** Enter the common name associated with the certificate.

4. From the **Algorithm** drop-down list, select the desired encryption algorithm used to encrypt your private key.
5. Select **Generate CSR**.

The CSR is downloaded successfully.

Upload CA Certificate

Get the CSR signed from your Enterprise PKI or any public PKI and upload the signed intermediate CA certificate to GigaVUE-FM.

To upload the signed CA certificate to GigaVUE-FM:

1. Go to  > **System > Certificates**.
2. In the top navigation bar, from the **PKI** drop-down list, select **CA**. The **CA Certificate** page appears.
3. From the **Actions** drop-down list, select **Upload CA**. The **Upload CA** pop-up appears.
4. Next to **Intermediate CA**, select **Choose File** to upload the signed intermediate CA certificate.
5. Next to **Root CA**, select **Choose File** to upload the corresponding root or intermediate CA.

The **CA Certificate** page displays the uploaded CA certificate.

Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source

To create a Monitoring Domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > VMware vCenter (V Series)**, and select **Monitoring Domain**. The **Create Monitoring Domain Configuration** page appears.
2. On the **Monitoring Domain** page, select **New**. The **VMware Configuration page** appears.
3. Enter the following fields:
 - a. **Monitoring Domain** - Enter a name for the monitoring domain.
 - b. **Connection** - Enter a name for the connection.
 - c. **Virtual Center** - Enter the IP address or FQDN of the vCenter
 - d. **Username** - Enter the username of the vCenter user.
 - e. **Password** - Enter the password to connect to the vCenter.
4. Select the Traffic Acquisition Method as Customer Orchestrated Source.
5. Click **Save**. Once you save the changes, the **Virtual Fabric Launch Configuration** page appears. The configuration details are explained in the following section.

Deploy V Series Node

This section provides step-by-step instructions on how to deploy GigaVUE V Series Nodes in VMware vCenter Monitoring Domain.

To deploy GigaVUE V SeriesNodes using GigaVUE-FM, follow these steps:

1. Go to the **VMware Fabric Launch Configuration** page using one of the following options:
 - **Automatic Option:** Create a monitoring domain, and the page opens automatically.
 - **Manual Option:** Open **VMware Fabric Launch Configuration** page from the **Monitoring Domain** page.
 - a. To launch the **VMware Fabric Launch Configuration** from the Monitoring Domain, go to **Inventory > VIRTUAL > VMware vCenter (V Series)**.
 - b. Select **Actions > Deploy Fabric**.

The **VMware Fabric Launch Configuration** page appears.

VMware vCenter

ESXi Fabric Launch

Monitoring Domain

Settings

Deploy

Cancel

VMware Fabric Launch Configuration

Datacenter*

DataCenter

Cluster*

Cluster-2_Scaled

V Series Node Image*

gigamon-gigavue-vseries-node-6.13.0...

Add

Form Factor*

Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates

☐

Select All

Select None

10.115.81.126

Common Configuration

10.115.81.126

VSeries.10.115.81.126-1

☒ Datastores

☐ Datastore Clusters

Datastore*

datastore1 (2)

Name Server ⓘ

Use Comma to Separate

V Series Node Name*

VSeries.10.115.81.126-1

Management

Interface Name

Network adapter 1

Network*

VM Network

IP Type

Static

IP Address*

10.115.85.38

Gateway IP*

10.115.80.1

CIDR Prefix* ⓘ

21

MTU

Enter MTU Value

Data interfaces

Interface Name

Network adapter 2

Network*

DPortGroup_tunnel

IP Type

Static

IP Address*

172.21.1.220

Configure
V Series Applications

305

2. On the **VMware Fabric Launch Configuration** page, enter or select the following details:

Field	Description																				
Datacenter	Select the vCenter datacenter with the ESXi hosts to be provisioned with GigaVUE V Series Nodes.																				
Cluster	Select the cluster where you want to deploy the GigaVUE V Series Nodes.																				
V Series Node Image	Select the OVA file uploaded in the Upload GigaVUE V Series Node Image into GigaVUE-FM , from the dropdown list.																				
Form Factor	Select Large, 8vCPU, 16GB RAM, 80GB Disk option.																				
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 unavailable 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.																				
Certificate NOTE: This field appears only when Enable Custom Certificates option is enabled.	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes. For more information, refer to Secure Communication between GigaVUE Fabric Components .																				
Hosts	Select the ESXi hosts for GigaVUE V Series Node deployment. The Common Configuration dropdown wizard and the specific host dropdown wizard appears. Expand the individual host and configure the GigaVUE V Series Node. <table border="1"> <tr> <th colspan="2">ESXi Host</th></tr> <tr> <td>Datastore</td><td>Network datastore shared among all ESXi hosts.</td></tr> <tr> <td>Name Server</td><td>The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.</td></tr> <tr> <td>V Series Node Name</td><td>Enter a name for the GigaVUE V SeriesNode.</td></tr> <tr> <th colspan="2">Management</th></tr> <tr> <td>Network</td><td>Management network for GigaVUE V Series Nodes.</td></tr> <tr> <td>IP Type</td><td>Select the management network IP type as Static or DHCP.</td></tr> <tr> <td>IP Address</td><td>Enter the IP address of the management port.</td></tr> <tr> <td>Gateway IP</td><td>Gateway IP address of the management network.</td></tr> <tr> <td colspan="2"> NOTE: This field appears only when the Management IP type is </td></tr> </table>	ESXi Host		Datastore	Network datastore shared among all ESXi hosts.	Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.	V Series Node Name	Enter a name for the GigaVUE V SeriesNode.	Management		Network	Management network for GigaVUE V Series Nodes.	IP Type	Select the management network IP type as Static or DHCP.	IP Address	Enter the IP address of the management port.	Gateway IP	Gateway IP address of the management network.	NOTE: This field appears only when the Management IP type is	
ESXi Host																					
Datastore	Network datastore shared among all ESXi hosts.																				
Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.																				
V Series Node Name	Enter a name for the GigaVUE V SeriesNode.																				
Management																					
Network	Management network for GigaVUE V Series Nodes.																				
IP Type	Select the management network IP type as Static or DHCP.																				
IP Address	Enter the IP address of the management port.																				
Gateway IP	Gateway IP address of the management network.																				
NOTE: This field appears only when the Management IP type is																					

Field	Description
	Static.

Field	Description	
	CIDR Prefix NOTE: This field appears only when you select the IP Type as Static.	Enter the CIDR prefix of the IP address.
	MTU (Optional)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
	Data Interfaces - When using Customer Orchestrated Source as the Traffic Acquisition Method, you must configure two data interfaces, one for data port and other for tool port.	
	Interface 1	
	Network	Tunnel Network for the GigaVUE V Series Nodes.
	IP Type	Select the tunnel network IP address type as Static or DHCP.
	IP Address	Enter the IP address of the data port.
	Gateway IP	Gateway IP address of the tunnel network.
	CIDR Prefix NOTE: This field appears only when you select the IP Type as Static.	Enter the CIDR prefix of the IP address.
	MTU (Optional)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
	Use IPv6 NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.	Enable this option to use IPv6 address.
	Interface 2	
	Network	Tunnel Network for the GigaVUE V Series Nodes.
	IP Type IP Address Gateway IP CIDR Prefix	Select the tunnel network IP address type as Static or DHCP. Enter the IP address of the tool port. Gateway IP address of the tool network. Enter the CIDR prefix of the IP address.

Field	Description
	NOTE: This field appears only when you select the IP Type as Static.
MTU (Optional)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
Use IPv6	Enable this option to use IPv6 address.
	NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.
Virtual Disk Format	Select the Virtual Disk Format from the dropdown list.
Deployment Folder	Enter the folder name in vCenter, under which the GigaVUE V SeriesNodes must be deployed.
Password	Enter the password you wish to use for the GigaVUE V Series Node.

3. Select **Deploy**. After deployment in vCenter, it appears on the **Monitoring Domain** page under the Monitoring Domain, where the GigaVUE V Series Node is deployed.

NOTE: GigaVUE-FM can process a maximum of ten GigaVUE V Series Nodes deployment requests in parallel on VMware vCenter. Each deployment request can have multiple GigaVUE V Series Nodes for deployment.

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 5 - Configure Identity Enrichment Metadata using Zscaler Private Access \(ZPA\)](#)

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.

Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)

Refer to [Application Metadata Exporter - Export Scenarios](#) topic for more details and conceptual information.

Prerequisites

Before configuring identity enrichment using Zscaler Private Access (ZPA), ensure that the following prerequisites are met.

Zscaler Infrastructure

The following Zscaler components must already be set up in your environment. These are standard ZPA deployment requirements. For detailed instructions, refer to the Zscaler documentation.

- **Acquire Zscaler Tenant Access:** Ensure that you have admin access to the Zscaler Private Access admin portal (ZPA tenant). Refer to [Zscaler Tenant Access](#) for more information.
- **Integrate Identity Provider (IdP):** An IdP (such as Okta, Microsoft Azure AD, or any supported IdP) must be configured and integrated with ZPA. ZPA uses the IdP to authenticate users before granting access to private applications. Refer to [About External Identity Providers](#) for more information.
- **Configure App Connector Group :** At least one App Connector group must be created and have active App Connectors registered to it. You can use an existing App Connector group or create a new one before proceeding. For instructions, refer to the Zscaler documentation. Refer to [Configuring App Connector Groups](#) for more information.
- **Install App Connector:** The Zscaler App Connector package must be downloaded and installed on a Linux-based VM (Ubuntu, plain Linux, or similar) in the domain where your private applications reside. Refer to the Zscaler App Connector deployment guide for platform-specific instructions (Linux, Docker, GCP, Kubernetes, and so on). Refer to [App Connectors](#) for more information.
- **Define Private Applications:** Application segments must be created for all private applications you want to monitor (for example, internal web apps, Jira, databases). Each application segment must include the application IP address and port range, and must be linked to a server group and an App Connector group. Refer to [Configuring Defined Application Segments](#) for more information.

Gigamon Infrastructure

- **Deploy VSN Node with AMX on ESXi:** The Application Metadata Exporter (AMX) must be deployed on a V Series Node (VSN). In this release, only the ESXi platform is supported. The IP address of the VSN node on which AMX is running, as it is required when configuring the ZPA Log Receiver. Refer to [Prerequisites](#) for more information.
- **Configure AMI:** Application Metadata Intelligence (AMI) must be active and generating flow records. The AMI metadata (source IP, destination IP, source port, destination port) is used to match and enrich flows with ZPA log data. For more information on configuring AMI, refer to [Configure Application Metadata Intelligence for Virtual Environment](#).

Integrate ZPA with Gigamon AMX

After the Zscaler infrastructure and Gigamon infrastructure are set up, configure a Log Receiver in the ZPA admin portal to stream ZPA logs to AMX. This is the key step that establishes the integration between ZPA and Gigamon.

The following table lists the specific options you must select when configuring a log receiver to integrate with AMX. For navigation steps and detailed instructions, refer to [Configuring Log Receiver](#) topic in Zscaler documentation.

Parameters	Description
Log Receiver	
Name	Enter a name for the log receiver.
Domain or IP Address	Enter the domain or IP address to the IP address of the VSN node on which AMX is deployed (not the ESXi host IP address).
TCP Port	Enter the same port that is configured as the ZPA LSS ingestor port on AMX.
App Connector Groups	Select an app connector group that was created.
Log Stream	
Log Type	Select User Activity as the log type.
Log Template	Select the Custom template.
Log Stream Content	Add the new user activity log ("ConnectionSetupTime": %j {ConnectionSetupTime} }\n) in the content.

Configure Ingestor for AMX

You can configure metadata enrichment for ZPA only. Optionally, you can combine it with GEM for Cloud Workloads for VMware vCenter.

To export identity enrichment metadata using ZPA :

1. Configure an ingestor for AMX
 - a. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
 - b. In the **Type** field, select **AMI** as the input type.
 - c. Enter the port number on which application metadata must be ingested. The default port value is 514.
 - d. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
2. Configure identity metadata enrichment using ZPA
 - a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
 - b. Select the **Enable** option to enable the enrichment.
 - c. In the **Type** field, select **Identity** as the input type.
 - d. Select **Zscaler Private Access** from the **Platform** dropdown list.
 - e. Select the attributes for identity enrichment from the **Attributes** tab.
 - f. In the **Source Information** tab, enter the port number on which AMX receives the ZPA LSS logs.

- g. Click  to add another ZPA source information to add multiple inputs to the AMX application and click  to remove an existing source information.
3. (Optional) Configure GigaVUE Enriched Metadata workload for ESXi
 - a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
 - b. Select the Enable option to enable the enrichment.
 - c. In the **Type** field, select **Workload** as the input type.
 - d. Select **VMware vCenter** from the **Platform** dropdown list.

NOTE: Only VMware vcenter platform is applicable as additional workload along with ZPA.

 - e. Select the required workload attributes for ESXi.
 - f. In the Source Information tab, enter the Key-Value pair details as given:
 - Key - url and Value - Primary URL for the vCenter Server SDK. For example, https://<vCenter-IP-or-FQDN>/sdk.
 - Key - username and Value - VMware vCenter username.
 - Key - password and Value - vCenter password.
 - g. Click  to add another ESXi source information to add multiple inputs to the AMX application and click  to remove an existing source information.

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 AML 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

Fields	Description
	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: - o Enter the Key. - o 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.
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.
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.

Fields	Description
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.

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 Timeout 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.
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

addr	
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
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_dns_name	Public DNS name assigned to the EC2 instance	✓	✓
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	Instance ID of the Azure VM	✓	✓
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)	✓	✓	✓
ip_dscp	Differentiated Services Code Point ipDiffServCodePoint (195)	✗	✓	✓
ip_precedence	The value of the IP Precedence ipPrecedence (196)	✗	✓	✓
proto	Protocol (UDP, TCP, ICMP...) - NetFlow v5 – proto - NetFlow v9 - PROTOCOL (4)	✓	✓	✓

¹The numbers within the enclosed brackets represent the element ID in IPFIX/NetFlow v10 and the field type ID in NetFlow v9.

	• 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(86)	x	✓	✓
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)	✗	✓	✓
ip_length_max	Maximum IP packet length - NetFlow v9 - MAX_PKT_LENGTH (26) - NetFlow v10/IPFIX - maximumIpTotalLength (26)	✗	✓	✓
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)	✗	✓	✓
ipv6_next_hop	Nexthop IPv6 address - NetFlow v9 - IPV6_NEXT_HOP (62) - NetFlow v10/IPFIX - ipNextHopIPv6Address (62)	✗	✓	✓
flow_id	An identifier of a flow that is unique within an Observation Domain - NetFlow v9 - flowId (148) - NetFlow v10/IPFIX - flowId (148)	✗	✓	✓

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>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	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																	
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 to be filtered from the traffic. ■ Attributes for the selected application is displayed in the

Parameter	Description
	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 \(VMware ESXi\)](#)

- 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.

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.

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 - gtpuportid
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.

SSL Decrypt

License: For information on licensing, refer to the) section of respective GigaVUE Cloud Suite Deployment Guide.

SSL Decrypt application delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. Secure Socket Layer (SSL) is a cryptographic protocol that adds security to TCP/IP communications such as Web browsing and email. The protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them.

SSL encryption secures traffic between a client and a server, such as a Web server. SSL decryption uses keys to decode the traffic between the client and server.

SSL and Transport Layer Security (TLS) protocols consist of a set of messages exchanged between a client and server to set up and tear down the SSL connection between them. To set up the connection, the client and server use the Public Key Infrastructure (PKI) to exchange the bulk encryption keys needed for data transfer.

IMPORTANT: To use SSL Decrypt application in GigaVUE-FM 6.3.00, install new GigaVUE-FM 6.3.00 image. Refer to *GigaVUE-FM Installation and Upgrade Guide* for step-by-step instructions on how to install GigaVUE-FM. SSL Decrypt application does not work if you upgrade from any previous GigaVUE-FM version to GigaVUE-FM 6.3.00.

Keep in mind the following when using SSL Decrypt application:

1. On updating the keys, service, or key maps which are already used in a monitoring session, the monitoring session is dynamically updated, and you need not re-deploy the monitoring session. You can also see if the updated keys, services, or key maps were successfully updated to the monitoring session and the respective GigaVUE V Series Nodes on the **All Events** page. Refer to [Overview of Events](#) section in the *GigaVUE Administration Guide* for detailed information on Events.
2. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the key will be removed from the Key Map. If that key is the only available entry in the Key Map, then it will not be removed.
3. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session that is not deployed, then the key will be removed from the Key Map and if that key is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
4. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the service will be removed from the Key Map. If that service is the only available entry in the Key Map, then it will not be removed.
5. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is not deployed, then the service will be removed from the Key Map and if that service is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
6. In VMware NSX-T platform, the throughput of SSL Decrypt application is improved to 480 Mbps.

Refer to the following topics for more detailed information:

- [Supported Protocols, Algorithms, and Ciphers for SSL Decrypt](#)
- [Configure SSL Decrypt](#)

Supported Protocols, Algorithms, and Ciphers for SSL Decrypt

The supported protocols are as follows:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2

The supported authentication (Au) is as follows:

- RSA

The supported key exchange (Kx) is as follows:

- RSA

The supported encryption algorithms (Enc) are as follows:

- NULL
- RC4
- DES
- 3DES
- AES (including GCM mode)
- CAMELLIA
- SEED
- IDEA

The supported compression algorithm is as follows:

- NULL

The supported digest algorithms are as follows:

- MD5
- SHA1
- SHA2

The supported key sizes are 128, 256, 512, 1024, 2048, and 4096.

The supported TLS extensions are as follows:

- Extended Master Secret, RFC 7627
- Encrypt-then-MAC, RFC 7366

The following table lists the supported ciphers:

Table 2: Supported Ciphers for SSL Decrypt

Cipher Name	Kx	Au	Enc	Bits	Mac
TLS_RSA_WITH_NULL_MD5	RSA	RSA	NULL	0	MD5
TLS_RSA_WITH_NULL_SHA	RSA	RSA	NULL	0	SHA
TLS_RSA_EXPORT_WITH_RC4_40_MD5	RSA_EXPORT	RSA_EXPORT	RC4_40	40	MD5
TLS_RSA_WITH_RC4_128_MD5	RSA	RSA	RC4_128	128	MD5
TLS_RSA_WITH_RC4_128_SHA	RSA	RSA	RC4_128	128	SHA
TLS_RSA_EXPORT_WITH_RC2_CBC_40_MD5	RSA_EXPORT	RSA_EXPORT	RC2_CBC_40	40	MD5
TLS_RSA_WITH_IDEA_CBC_SHA	RSA	RSA	IDEA_CBC	128	SHA
TLS_RSA_EXPORT_WITH_DES40_CBC_SHA	RSA_EXPORT	RSA_EXPORT	DES40_CBC	40	SHA

Cipher Name	Kx	Au	Enc	Bits	Mac
TLS_RSA_WITH_DES_CBC_SHA	RSA	RSA	DES_CBC	56	SHA
TLS_RSA_WITH_3DES_EDE_CBC_SHA	RSA	RSA	3DES_EDE_CBC	168	SHA
TLS_RSA_WITH_AES_128_CBC_SHA	RSA	RSA	AES_128_CBC	128	SHA
TLS_RSA_WITH_AES_256_CBC_SHA	RSA	RSA	AES_256_CBC	256	SHA
TLS_RSA_WITH_CAMELLIA_128_CBC_SHA	RSA	RSA	CAMELLIA_128_CBC	128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	256	SHA
TLS_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	128	SHA
TLS_RSA_WITH_NULL_SHA256	RSA	RSA	NULL	0	SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	128	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	256	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	128	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	256	SHA384

All algorithms used for SSL Decrypt application are FIPS 140-3 compliant.

All key URLs must point to an RSA private key stored in the PEM or PKCS12 format, as follows:

- <http://keyserver.domain.com/path/keyfile.pem>
- <https://keyserver.domain.com/path/keyfile.pem>
- <ftp://keyserver.domain.com/path/keyfile.pem>
- <tftp://keyserver.domain.com/path/keyfile.pem>
- [scp://username\[:password\]@keyserver.domain.com/path/keyfile.pem](scp://username[:password]@keyserver.domain.com/path/keyfile.pem)

The supported applications are as follows:

- HTTPS
- FTPS
- SMTP, IMAP, and POP3 with StartTLS

Configure SSL Decrypt

To configure SSL Decrypt Application using GigaVUE-FM follow the steps given below:

- [Upload SSL Key Pairs](#)
- [Create SSL Service](#)

- [Key Mapping](#)
- [SSL Key Pair Store](#)
- [Add SSL Decrypt to Monitoring Session](#)

Upload SSL Key Pairs

To upload an SSL private key, do the following:

1. Go to **Inventory > Resources > Security > SSL Keys**.
2. Click **Add**. The **Create SSL Key Pair** page appears.
3. Enter the following details:

Field	Description
Key Alias	Enter a name for the key.
Comment	Enter a description
Key Type	Select the either of the key type:
PEM	
PassPhrase (optional)	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.
Private key	Enter the Private Key using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
Certificate	Enter the Certificate using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
PKCS12	
PassPhrase	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.

4. Click **Save**.

NOTE: SSL Decrypt application does not support HSM.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Key pair, select the key from the list in the SSL Key Pairs page and click the Edit button.
Delete	To delete a SSL Key pair, select the key from the list in the SSL Key pairs page and click the Delete button.
Delete all	Use this button to delete all the SSL Key pairs in the SSL Key pairs page.
View certificate	To view the certificate associated with the particular SSL Key Pair, select the key from the list in the SSL Key Pairs page and click the View Certificate button.

Create SSL Service

After uploading a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number.

Prerequisite

Before creating a service, upload a private key as described in [Upload SSL Key Pairs](#)

To create a SSL service, do the following:

1. Go to **Inventory > Resources > Security > SSL Service**. The SSL Services page appears.
2. Click **Add**. The **Create SSL Service** page appears.
3. On the Create SSL Service page, enter the following details:

Field	Description
Alias	Enter a name for the SSL Service.
Default Service	Enable this to use default service.
Server IP Address	Enter the IP address of the server in which the service runs.
Server Port	Enter the port number of the server.

4. Click **Save**.

The following actions can also be performed from the SSL Service Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the Edit button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the Delete button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

Key Mapping

After adding the SSL Service, now you map the private key with the service using Key Mapping.

To map a key with the service, follow the steps given below,

1. Go to **Inventory > Resources > Security > SSL Key Mapping**. The SSL Key Pair Mapping page appears.
2. Click **Add**.
3. Enter the Key Mapping Alias.
4. Select the SSL Service and Key Alias from the drop-down.
5. Click **Save**.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the Edit button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the Delete button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

SSL Key Pair Store

SSL Key Pair Store is a repository, that allows you to save all the key under a single location. You can create multiple key stores and in each key store you can store multiple keys.

1. Go to **Inventory > Resources > Security > SSL Key Store**. The SSL Key Pair Store page appears.
2. Click **Add**.
3. Enter the **Key Store Alias** and **Comment**.
4. Click **Save**.


The following actions can also be performed from the SSL Key Pair Store Page:

Field	Description
Edit	To edit a SSL Key pair Store, select the Key Store from the list in the SSL Key Pair Store page and click the Edit button.
Delete	To delete a SSL SSL Key Pair Store, select the SSL Key Pair Store from the list in the SSL Key Pair Store page and click the Delete button.
Delete all	Use this button to delete all the SSL Key Pair Store in the SSL Key Pair Store page.

Add SSL Decrypt to Monitoring Session

After mapping your keys with service, to add GigaSMART applications to GigaVUE V Series Node, follow the steps given below,

1. Drag and drop **SSL Decrypt** from APPLICATIONS to the graphical workspace.
2. Click the SSL Decrypt application and select **Details**.



Application	SSL Decrypt
Alias	ssl-decrypt
Enable	☒
Key Map	v
In Port ⓘ	0
Out Port ⓘ	0
Session Timeout (sec)	300
Pending Session Timeout (sec)	60
Tcp Syn Timeout (sec)	20
Decrypt Fail Action	☐ Pass ☒ Drop
Key Cache Timeout (sec)	10800
Ticket Cache Timeout (sec)	10800
Non-ssl Traffic	☐ Pass ☒ Drop

3. Enter the following details in the Application quick view:

Fields	Description
Alias	Enter the alias name for the application.
Enable	Enable the box to enable SSL Decryption.
Key Map	Select the Key Map from the list of available Key Maps. Refer to Key Mapping for more details on how to map the key to SSL Service.
In Port	Enter the source port number from which the traffic should be fetched.
Out Port	Enter the destination port number to which the decrypted traffic should be delivered.
Session Timeout	Enter the value in seconds after which the session should be timeout. The default value is 300 seconds.
Pending Session Timeout	Enter the value in seconds after which the session must timeout if the session is in pending state
Tcp Syn Timeout	Enter the value in seconds after which the session must timeout when the session does not synchronize TCP.
Decrypt Fail Action	Select Pass to allow the traffic to pass through the application when the decryption fails and select Drop to drop the traffic before passing through the application when the decryption fails.
Key Cache Timeout (sec)	Enter the value in seconds until which the key cache information can be reused for resumption.
Ticket Cache Timeout (sec)	Enter the value in seconds until which the ticket cache information can be reused for resumption.
Non-ssl Traffic	Select Pass to allow the non-SSL traffic to pass through the application and select Drop to drop the non- SSL traffic before passing through the application.

4. Click **Save**.
5. Click **Deploy**. The Select nodes to deploy the monitoring session page appears.
6. Select the GigaVUE V Series Nodes you want to deploy and select an interface for each GigaVUE V Series Node. Then, click **Deploy**.

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 ($\text{numBuckets} \times \text{bucketInterval}$ 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.

Field	Description
	 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.
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	failedSBITransactionVlanID
pc-mm	100	200
pc-sm	300	400
ccsm-5gc	500	
ccsm-cnhss	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	
sc-sepp	2000	
others	2100	
unclassified	2200	

4. Click **Choose File** in **File Name** field to upload the CSV file into GigaVUE-FM.

5. Click **Validate** to validate the CSV file.
6. 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](#).

5G Cloud 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 HTTP2 protocol. These HTTP2 transactions are mirrored using some specific network functions, which are in JSON-encoded format.

The 5G Cloud application supports synthesizing of L2, L3, L4, and HTTP2 headers from the JSON-encoded data of the 5G core network. Currently, Ericsson VTAP 5G core input format is supported in DPDK based platforms and in non-DPDK based platforms, Casa vTAP, Oracle SCP, Nokia SCP Inbound, and Nokia SCP Inbound/Outbound input formats are supported.

In non-DPDK mode, V Series Node supports processing of the HTTP2/HTTP2S input formats and extracts the actual packet in the 5G core, and then sends it out on the egress tunnel to probe using either VXLAN or L2GRE or UDP-GRE Tunnel. To achieve this, V Series supports two new applications:

- GVHTTP2
- 5G-Cloud

Supported Platforms

The application is supported on the following platforms:

- VMware ESXi
- OpenStack
- Third Party Orchestration


To know more about the configuration of GVHTTP2 and 5G-Cloud application for Casa vTAP, Oracle SCP, and Nokia SCP, refer to the following sections:

- [5G-Cloud Casa vTAP Support](#)
- [5G-Cloud Oracle and Nokia SCP Support](#)
- [5G-Cloud Ericsson SCP Support](#)
- [Reference - Configuration Table](#)

5G-Cloud Casa vTAP Support

V Series Node supports receiving data from 5G Casa network functions in the Casa vTAP input format. This augments the currently existing functionality of receiving data from Ericsson network functions in Ericsson vTAP format and receiving data from Nokia network functions in Nokia format. 5G-Cloud Casa vTAP is a non-service-based interface architecture in which input from the Casa vTAP system is processed, decoded, and sent out through an egress tunnel.

Casa systems vTAP network function mirrors the traffic in the 5G core and sends it out in gRPC encoded protobuf format. To convert the encoded protobufs to actual network traffic in 5G core, V Series Node supports two applications, GVHTTP2 and the 5G-Cloud.



GVHTTP2 Application

The first application, GVHTTP2, acts as an HTTP2 server. It receives the gRPC data over HTTP2 and establishes a TCP connection with the Casa vTAP system by acknowledging it. The application will extract the HTTP2 header and HTTP2 body, which are gRPC encoded, then encapsulate the data into a VXLAN header and send it out over the VXLAN tunnel to the next application, 5G-Cloud.

5G-Cloud Application

The 5G-Cloud application receives the VXLAN packet, decapsulates, and decodes the gRPC protobuf content. The application will extract the PFCP packets (N4 traffic) and convert the traffic into PCAPng format. It then sends the extracted packets to the NetScout tool probe over the UDP-GRE/L2GRE/VXLAN tunnels.

NOTE: Since the NetScout probe only understands Ericson vTAP format, which is UDP-GRE encapsulated in PCAPng format, the same is followed in the Casa vTAP system. As part of the 6.7 release, only N4 non-SBI traffic (PFCP) is supported.

Configuration of 5G-Cloud Casa vTAP


To configure the GVHTTP2 and 5G-Cloud applications in GigaVUE-FM, refer to the following sections:

- [Configuration of GVHTTP2 Application for Casa vTAP](#)
- [Configuration of 5G-Cloud Application for Casa vTAP](#)

Configuration of GVHTTP2 Application for Casa vTAP

In GigaVUE-FM, you must do the following to add the GVHTTP2 application in the Monitoring Session of a Monitoring Domain:

1. Create an ingress REP (Raw End Point) to receive the gRPC data over TCP on top of HTTP2. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the GVHTTP2 application in the Monitoring Session.
3. Create a link between ingress REP and the GVHTTP2 application.
4. Create egress REP.
5. Create a link between the GVHTTP2 application and egress REP.



You can add a GVHTTP2 application to:

- **New Monitoring Session** - Add the GVHTTP2 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 add a GVHTTP2 application:

1. In the canvas, drag and drop the GVHTTP2 application and select **Details**. The GVHTTP2 quick view appears.

GVHTTP2

Application

Alias*

HTTP2 Listening IP*

HTTP2 Listening Port*

Max Stream*

Worker Thread*

Mode*

IS TUNNEL

Type*

Source IP*

Source Port*

Destination IP*

Destination Port*

VNI ID*

IS SETTINGS

TLS* ☐ Enable ☒ Disable

ADVANCED SETTINGS

CSV* ☒ Enable ☐ Disable

PCAP* ☐ Enable ☒ Disable

Log Directory*

Log Level

2. On the application quick view, enter or select the required information as described in the following table:

Field	Description
Application	The name GVHTTP2 appears by default.
Alias	Enter the alias name as gvhttp2 .

HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port. Enter a value between 1 and 65535.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specifies the number of worker threads. Enter a value between 1 and 128. The default value is one.
Mode	Select Casa vTAP as mode from the drop-down menu.
TX Tunnel (Specifies the tunnel configuration from GVHTTP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the source port to bind for sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the destination IP address to which packets will be sent.
Destination Port	Specify the destination port where packets will be sent.
VNI id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TLS Setting	
TLS	Select a TLS/SSL Monitor Mode from one of the following options: ▪ Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. ▪ Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on.
Advanced Setting	
Log Directory	Specify the path to store the log files.
Log Level	In the log level drop-down list, select the severity log level of the events from the following options: • PCAP • CSV • Info • Detail • Full-Parse

Downloading Logs for GVHTTP2 application

You can view the log files of a V Series Node or download them as CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic > VIRTUAL >** select your cloud platform.
2. Select the required Monitoring Session and click the **APPS LOGGING** link.

3. The **GVHTTP2 Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Configuration of 5G-Cloud Application for Casa vTAP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain:

1. Create ingress REP to extract the HTTP2 body (gRPC content). Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the 5G-Cloud application in the Monitoring Session.
3. Create a link between ingress REP and the 5G-Cloud application.
4. Create egress REP.
5. Create a link between the 5G-Cloud application and egress REP.




You can add a 5G-Cloud application to:

- **New Monitoring Session** - Add the 5G-Cloud 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 add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud 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-Cloud appears by default.
Alias	Enter the alias name as cloud5g .
Vendor Integrations	From the drop-down list, select Casa Systems vTAP .
RX Tunnel (Specifies the tunnel configuration from GVHTTP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Listening IP	Specify the tunnel's local listen IP address (TX tunnel destination IP address used in GVHTTP2) to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is 0.
TX Tunnel	
Type	Select the tunnel type from the drop-down list. The available tunnel types are: • UDPGRE • VXLAN • L2GRE
Tool IP	Specify the remote IP address to which the packet will be sent.
Destination Port	Specify the tunnel destination port where the packet will be sent.

	Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind for sending the packet. Enter a value between 1 and 65535.
VNI ID (Applicable only when the selected tunnel type is VXLAN)	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key (Applicable only when the selected tunnel type is L2GRE)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU. Note: For V Series Node version 6.8.00 and above, the range should be between 1400 and 8800. The default value is 8800. For V Series Node version below 6.8.00, range should be between 1500 and 8800. The default value is 8800.
Log Directory	Specify the path to store the log files.
Log Level	In the tunnel log level drop-down list, select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug

Downloading Logs for 5G-Cloud application

You can view the log files of a V Series Node or download them as CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic > VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g** application, click the  menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
3. The **CLOUD5G Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Limitations

- The 5G-Cloud application does not validate the VNI ID configured by the user in the GVHTTP2 application when processing the VXLAN packets.
- The VNI ID configured by the user will not be validated while processing the packets received on the ingress interface. So, the 5G-Cloud application will not drop the VXLAN received on the ingress interface if the VNI ID received in the packet does not match the VNI ID configured by the user.
- After deploying the Monitoring Session from GigaVUE-FM, users can only change the following configuration parameters for the GVHTTP2 application:
 - Server (Listen IP Address and Port Number)
 - VXLAN (Source IP and Port Number)
 - VXLAN (Destination IP and Port Number)
- After deploying the Monitoring Session from GigaVUE-FM, users can only change the following configuration parameters for the 5G-Cloud application:
 - Ingress VXLAN (Listen IP, Listen Port Number, and Source Port Number)
 - Egress UDP-GRE/L2GRE/VXLAN Tunnels (Destination IP Address, Destination Port, Source IP, and Source Port)

Rules and Notes

- On a V Series Node, only one application (GVHTTP2 or 5G-Cloud or AMX) can be deployed across Monitoring Sessions.
- GVHTTP2, 5G-Cloud, and AMX applications should not be combined with general GigaSMART applications in a Monitoring Session or across Monitoring Sessions.
- For GVHTTP2, 5G-Cloud, and AMX applications, the source and destination links should be Raw Endpoints.
- Support for GVHTTP2, 5G-Cloud, and AMX applications is limited to 'Customer Orchestrated Source' tapping type.
- The 5G-Cloud and GVHTTP2 applications require the "VBL-ADD-ADV5GC" or "VBL-ADD-5GC" license.
- To deploy a Monitoring Session, ensure that the 'HTTP2 Listening IP address' and 'TX Tunnel Source IP address' in GVHTTP2 corresponds to the IP addresses of one of the V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the 'RX Tunnel Listening IP address' and 'TX Tunnel Source IP address' in 5G-Cloud correspond to the IP addresses of one of the V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the TX Tunnel Destination IP address used in GVHTTP2 and the RX Tunnel Listening IP address used in 5G-Cloud are the same.

NOTE: For information on statistics details of Casa vTAP for GVHTTP2 and 5G-Cloud applications, refer to [Appendix - Casa vTap Statistics](#).

5G-Cloud Oracle and Nokia SCP Support

GigaVUE V Series Node supports receiving data from 5G network functions in Oracle SCP and Nokia SCP input formats. In both Oracle and Nokia SCP, 3GPP standard method is followed, which has all the standard network components implemented.

Oracle SCP

Oracle systems input format is HTTP2/HTTP2S POST over TCP with TLS enabled. In Oracle SCP, the 5G-Cloud application synthesizes the packets and egresses them to the tool port after it has received all the 4-leg transactions. Oracle SCP processing in the 5G-Cloud application supports converting the discovery and the 4-legged (consumer, SCP, and producer) transactions to packets as per 3GPP Model C and Model D.

Model C

The Consumer knows the producer and sends messages through SCP. Model C comprises of 4-legged transactions between the consumer, SCP, and the producer. The 5G-Cloud application converts the 4-legged transactions to packets and sends them to the probe.

Model D

SCP sends discovery messages to discover the producer after receiving the messages from the consumer. Model D comprises of 2-legged transactions between the consumer and the SCP, 4-legged transactions between the SCP, SCP discovery, and the NRF to find the producer, and the 2-legged transactions between the producer and the SCP. The 5G-Cloud application converts the 4-legged transactions between the consumer, SCP, and the producer to packets and sends them to the probe. The 5G-Cloud application also converts the discovery transactions between the SCP and the NRF and sends them to the probe.

5G- Cloud application will ensure that both the discovery transaction and the 4-legged transaction are transmitted in order, to the connected tool or probe.


Nokia SCP

In Nokia SCP, traffic is mirrored between control functions using the HTTP2 protocol. This traffic is mirrored from a service called SCP (Service Control Proxy), a centralized point through which all the communications between all the control plane functions pass. SCP encrypts all the messages that are going across the control plane network functions and forwards that information in JSON-encoded format to GVHTTP2.

Both Model C and Model D transactions can be received from the SCP. However, Model D cannot be identified as Model D, because a separate message id is used for the discovery messages. In the case of Model D, SCP will include the Model D discovery transactions that

happened with NRF along with the network function communication. Therefore, 6 transactions will be forwarded from SCP for a single request and response pair between any two network functions. The Model C transactions will be processed in 2 modes:

- **Transparent (2 legged):** The 5G-Cloud application will synthesize a direct transaction between the network functions. The 4-legged transactions received between the consumer, SCP, and the producer are synthesized to look like 2-legged transactions between the consumer and the producer.
- **Inbound and outbound (4 Legged):** The 5G-Cloud application will synthesize indirect transactions, meaning all transactions between any network function will go via SCP. All the transactions received in Model C or Model D are synthesized and sent to the probes.



GVHTTP2 Application

The first application, GVHTTP2, acts as an HTTP2 server, receives the data over HTTP2 frame and establishes a TCP connection with multiple Oracle SCP/DD or Nokia SCP pods. It receives the HTTP2 headers and stores them in a temporary buffer. Once it receives all the headers and body of that stream, it acknowledges back to the client. The application will extract the HTTP2 header and HTTP2 body, encapsulate the data into a VXLAN header, and send it out over the VXLAN tunnel to the next application, 5G-Cloud.

5G-Cloud Application

The 5G-Cloud application receives and decapsulates the VXLAN packet. It extracts the HTTP2 body, parses the SCP/DD JSON mirrored message, and associates the mirrored message transaction ID to track the complete transaction. It then synthesizes TCP Flow and HTTP2 Session representing full 5G CNF / SCP / 5G CNF transaction data and sends it out over L2GRE/VXLAN tunnel to the probe.

The 5G-Cloud application runs in Oracle SCP, Nokia SCP Transparent, and Nokia SCP Inbound-Outbound modes.

- In Nokia SCP Inbound-Outbound mode, both Inbound (between Consumer and SCP) and Outbound (between SCP and Producer) traffic is synthesized and sent to the connected Tool port. In Nokia SCP Transparent mode, only Inbound (between Consumer and SCP) traffic is processed and sent to the Tool Port.
- In the Nokia SCP Transparent mode, the 5G-Cloud Application will drop the Outbound traffic from the SCP (if received in the input).
- In the Nokia SCP Transparent mode, the Producer replaces the SCP in transactions between the Consumer and the SCP.
- In the Nokia SCP Transparent mode, authority header value will be replaced (if the replace authority configuration is enabled).

Configuration of 5G-Cloud Oracle/Nokia SCP


To configure the GVHTTP2 and 5G-Cloud applications in GigaVUE-FM, refer to the following sections:

- [Configuration of GVHTTP2 Application for Oracle SCP/Nokia SCP](#)
- [Configuration of 5G-Cloud Application for Oracle SCP/Nokia SCP](#)

Configuration of GVHTTP2 Application for Oracle SCP/Nokia SCP

In GigaVUE-FM, you must do the following to add the GVHTTP2 application in the Monitoring Session of a Monitoring Domain:

1. Create an ingress REP (Raw End Point) to receive the HTTP2/HTTP2S data over TCP. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the GVHTTP2 application in the Monitoring Session.
3. Create a link between ingress REP and the GVHTTP2 application.
4. Create egress REP.
5. Create a link between the GVHTTP2 application and egress REP.




You can add a GVHTTP2 application to:

- New Monitoring Session - You can add the GVHTTP2 application after creating a new Monitoring Session and when the canvas appears.

- Existing Monitoring Session - Click **Edit** on the existing Monitoring Session, and the GigaVUE-FM canvas appears.

To add a GVHTTP2 application:

- In the canvas, drag and drop the GVHTTP2 application and select **Details**. The GVHTTP2 quick view appears.



- On the application quick view, enter or select the required information as described in the following table:

Field	Description
Application	The name GVHTTP2 appears by default.
Alias	Enter the alias name as gvhttp2 .
HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specifies the number of worker threads. Enter a value between 1 and 128. The default value is four.
Mode	Select Oracle SCP or Nokia SCP as mode from the drop-down menu.
TX Tunnel (Specifies the tunnel configuration from GVHTPP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.

Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the source port to bind for sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the destination IP address to which packets will be sent.
Destination Port	Specify the destination port where packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TLS Setting	
TLS	Select a TLS/SSL Monitor Mode from one of the following options: ▪ Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. ▪ Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on.
Server Private key*	Specify the server private key only when TLS is enabled.
Server Certificate*	Specify server certificate only when TLS is enabled.
Advanced Setting	
CSV	Enable or disable the CSV log level.
PCAP	Enable or disable the PCAP log level.
Log Directory	Specify the path to store the log files.
Log Level	In the log level drop-down list, select the severity log level of the events from the following options. The default value level is Notification. • Notification • Information • Debug

Download Logs - GVHTTP2 application

You can view the log files of a V Series Node or download them as .CSV or .txt file.

To download the log files to a local environment:


1. Go to **Traffic > VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and click the **APPS LOGGING** link.
3. The **GVHTTP2 Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Configuration of 5G-Cloud Application for Oracle SCP/Nokia SCP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain:

1. Create ingress REP to extract the HTTP2 body. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the 5G-Cloud application in the Monitoring Session.
3. Create a link between ingress REP and the 5G-Cloud application.
4. Create egress REP.
5. Create a link between the 5G-Cloud application and egress REP.

GigaVUE V Series in Sydney



The GigaVUE V Series map rules in Sydney take virtual packets matching the criteria specified by the map rules, encapsulate them in UDP packets, and forward them across the Internet to the receiving GigaVUE HC Series in Melbourne.

The map rules include the destination IP address and UDP port (the address of the IP interface associated with the network port in Melbourne).

GigaVUE HC Series with GigaSMART in Melbourne

IP interface associated with network port 1/1/x3 in Melbourne decapsulates traffic arriving from Sydney and makes it available for standard GigaVUE packet distribution to local tool ports. The decapsulation operation bound to IP interface associated with network port 3 as part of a map specifies the listening UDP port matching the one configured in Sydney.

Adding 5G-Cloud Application in Oracle SCP/Nokia SCP

Prerequisites

The pre-requisite to add a 5G-Cloud Application in Oracle SCP/Nokia SCP is:


- You must upload CSV files containing a valid FQDN, NF Instance, and User Agent ID's and a valid IPv4/IPv6 addresses. To upload the CSV file, refer to [Adding CSV file for IP Mapping](#).

You can add a 5G-Cloud application to:

- New Monitoring Session - You can add the 5G-Cloud application after creating a new Monitoring Session and when the canvas appears.
- Existing session - Click **Edit** on the existing Monitoring Session, and the GigaVUE-FM canvas appears.

To add a 5G-Cloud application:

- In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



The screenshot displays the 'Edit Monitoring Session' interface. On the left, a canvas shows a network diagram with a central '5G-Cloud' application icon connected to two other nodes. On the right, the '5G-Cloud cloud5g' configuration panel is open, showing various settings for the application.

Field	Description
Application	5G-Cloud
Alias*	cloud5g
Mode*	cloud5g
RX TUNNEL	
Type*	10000
Listening IP*	192.168.20.12
Listening Port*	12432
Source Port*	12431
VNI id	121
TX TUNNEL	
Type*	10000
Dest IP*	192.168.10.16
Destination Port*	4789
Source IP*	192.168.10.12
Source Port*	12433
L2GRE Key*	122
ADVANCED SETTINGS	
Tunnel MTU*	8000
Log Directory*	/var/log
Log Level	Debug
SCP CONFIG	
FQDN Mapping*	cloud5g
NF Instance Mapping*	cloud5g

- On the application quick view, enter or select the required information as described in the following table:

Field	Description
Application	The name 5g-Cloud appears by default.
Alias	Enter the alias name as cloud5g .
Mode	From the drop-down list, select the required mode from the following options: - Oracle SCP - Nokia SCP Transparent - Nokia SCP Inbound/Outbound
RX Tunnel (Specifies the tunnel configuration from GVHTPP2 to 5G-Cloud)	

Type	Specify the tunnel type. VXLAN is selected by default.
Listening IP	Specify the tunnel's local listen IP address (TX tunnel destination IP address used in GVHTTP2) to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TX Tunnel	
Type	Specify the tunnel type.
Tool IP	Specify the remote IP address to which the packet will be sent.
Destination Port	Specify the tunnel destination port where the packet will be sent. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind for sending the packet. Enter a value between 1 and 65535.
VNI Id (Applicable only when the selected tunnel type is VXLAN)	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key (Applicable only when the selected tunnel type is L2GRE)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU. Note: For V Series Node version 6.8.00 and above, the range should be between 1400 and 8800. The default value is 8800. For V Series Node version below 6.8.00, range should be between 1500 and 8800. The default value is 8800.
Log Directory	Specify the path to store the log files.
Log Level	In the tunnel log level drop-down list, select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information

	• Debug
SCP Config	
FQDN Mapping	Specify the alias name created for the uploaded FQDN table CSV file. You can also add new FQDN by utilizing the ADD button. Note: Header details: gigamon@test-5gc-new:/var/log/cloud5g_tabledir\$ cat FQDN.txt FQDN,IngressIP1,IngressIP2,IngressPort,EgrIP,NFType,NFLocation (The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional) Examples - 5G-Cloud FQDN Entries: Oracle SCP: FQDN.txt:scp1.pp2.scp.5gc.mnc001.mcc208.3gppnetwork.org,10.64.104.30,10.64.104.36,9443,10.64.104.45,SCP,plolp Nokia SCP: s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,172.23.13.187,,8080,,SCP,plolp
NF Instance Mapping	Specify the alias name created for the uploaded NFID table CSV file. You can also add new NF Instance by utilizing the ADD button. Note: Header details: gigamon@test-5gc-new:/var/log/cloud5g_tabledir\$ cat NFID.txt NFID,IngressIP1,IngressIP2,IngressPort,EgrIP,NFType,NFLocation (The fields NFID and ingressIP1 are mandatory. All other fields are optional) Examples - 5G-Cloud NFID Entries: • Oracle SCP: NFID.txt:04fba086-6ac1-4472-9ee3-cf10e57849f8,10.64.107.85,10.64.107.79,8080,,SMF,plolp • Nokia SCP: 8b94c7ee-d750-4bcd-95be-634b31a5fa0d,172.23.14.4,,8080,,AMF,jeolp
User Agent Mapping	Specify the alias name created for the uploaded User Agent table CSV file. You can also add new User Agent by utilizing the ADD button. Note: Header details: gigamon@test-5gc-new:/var/log/cloud5g_tabledir\$ cat UA.txt UserAgent,IngressIP1,NFType,NFLocation (The fields useragentID and ingressIPAddress are mandatory. All other fields are optional) Examples - 5G-Cloud User Agent Entries: • Oracle SCP: pp5gs2nd15gcam01,10.64.106.67,AMF,jeolp • Nokia SCP: pp5gs2nd15gcam01,10.64.106.67,AMF,plolp pp5gs2nd15gcam02,10.64.106.107,AMF,plolp ofrpp25gcsmsf01,10.64.106.138,SMF,jeolp
SCP Advanced Config	
No. of TCP Flows	Specify the total number of concurrent TCP flows. Enter a value between 128 and 2048. The default value is 1024.
No. of Transaction Flows	Specify the total number of Transaction Stream Flows allocated for tracking active Transaction Streams. Enter a value between 128 and 5500. The default value is 2048.

TCP Flow Timeout	Specify the TCP flow timeout ranging between 0 to 7200. The default value is 900.
SCP Transaction Timeout	Specify the SCP transaction timeout value ranging between 1 and 300. The default value is 10. Note: The recommended timeout value is between 10 and 20 seconds. An increase in the transaction timeout value will impact the total number of transactions that can be processed in parallel.
Minimum TCP Client Port	Specify the minimum value to be used for TCP client port ranging between 1023 and 65535. The default value is 32768.
Maximum TCP Client Port	Specify the maximum value to be used for client port ranging between 1023 and 65535. The default value is 36863.
Header Index	Enable or disable the header indexing flag from the available options. The default value is disable.
Header Compression Code Note: HTTP/2 requests and responses typically contain a large number of redundant headers. HTTP/2 uses header compression to compress headers significantly.	Enable or disable the header compression code from the available options. The default value is disable.
Gigamon Header Note: Helps in the debugging process to identify all transactions.	Enable or disable the Gigamon Header from the available options. The default value is Enable.
Nokia Use 3gpp Target	Specify the value which upon detecting producer IP/FQDN, treats the 3gpp Target API Root to be predictive of the Producer IP if value is non-zero as either enable or disable. The default value

API Root (Applicable only when the selected mode is Nokia SCP Transparent)	is disable.
Nokia Replace Authority (Applicable only when the selected mode is Nokia SCP Transparent)	Specify the value which upon detecting producer IP/FQDN, treats the 3gpp Target API Root to be predictive of the Producer IP if the value is non-zero as either enable or disable. The default value is disable.
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
5G-Cloud Log Level	Select the required 5G-Cloud CSV log level from the drop down list. The default value is None. • All • Flow • Message • Transaction • None

Download Logs - 5G-Cloud application

You can view the log files of a V Series Node or download them as .CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic > VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and click the **APPS LOGGING** link.
3. The **CLOUD5G Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.

4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

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. Select any of the following from the **Type** as per the requirement:
 - **5G-Cloud FQDN** - Adding the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping. Adding FQDN is mandatory.
 - **5G-Cloud NF Instance** - Adding the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.
 - **5G-Cloud UA** - Adding the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.
4. Click **Choose File** in **FileName** field to upload the CSV file into GigaVUE-FM.
5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

NOTE: For Oracle SCP, configuration of FQDN Mapping, NF Instance Mapping, and User Agent Mapping are mandatory. For Nokia SCP Transparent and Nokia SCP Inbound/Outbound, configuration of FQDN Mapping is mandatory.

Limitations

- The following configuration parameters for the GVHTTP2 application are non-editable and can be configured only during the initial configuration:
 - Mode
 - TX Tunnel – txType
 - Worker thread
 - RX VNI Id
 - Max Streams
 - TLS configuration
 - Log Level
 - Log Directory
- The following configuration parameters for the 5G-Cloud application are non-editable and can be configured only during the initial configuration:
 - Mode
 - TX Tunnel – txType

- RX Tunnel - rxType
- RX VNI Id
- Tool MTU
- Log Directory
- Log Level
- Header Index
- Header Compression Code
- Packet capture log level
- SCP log level
- Gigamon Header
- FQDN Mapping
- NF Instance Mapping
- User Agent Mapping

Rules and Notes

- The 5G-Cloud and GVHTTP2 applications require the "VBL-250T-ADD-5GC" license.
- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to **Procedure to deploy V Series Node in VMware ESXi with SR-IOV Adapte** topic in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration guide for more details.
- On a V Series Node, only one application (GVHTTP2 or 5G-Cloud or AMX) can be deployed across Monitoring Sessions.
- GVHTTP2, 5G-Cloud, and AMX applications should not be combined with general GigaSMART applications in a Monitoring Session or across Monitoring Sessions.
- For GVHTTP2, 5G-Cloud, and AMX applications, the source and destination links should be Raw Endpoints.
- Support GVHTTP2, 5G-Cloud, and AMX applications is limited to 'Customer Orchestrated Source' tapping type.
- Configuration of GVHTTP2 and 5G-Cloud in V Series Node will fail in versions less than v6.7.00.
- To deploy a Monitoring Session, ensure that the 'HTTP2 Listening IP address' and 'TX Tunnel Source IP address' in GVHTTP2 correspond to the IP addresses of one of the GigaVUE V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the 'RX Tunnel Listening IP address' and 'TX Tunnel Source IP address' in 5G-Cloud correspond to the IP addresses of one of the GigaVUE V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the TX Tunnel Destination IP address used in GVHTTP2, and the RX Tunnel Listening IP address used in 5G-Cloud are the same.
- When you have already deployed the 5G-Cloud application with all the table files uploaded, you should undeploy and redeploy the Monitoring Session again with the updated files to update the table.

- In the GVHTTP2 application, when the TLS/SSL configuration is enabled, the private key and certificate path must be configured.
- Before upgrading GigaVUE V Series Node, user should undeploy the Monitoring Session where the GVHTTP2 application is deployed and take the backup of the SSL certificate and private key files. Following the upgrade, user should manually install/copy the private key and the certificate file in GigaVUE V Series Node and redeploy the GVHTTP2 application.
- To upload the table configuration file in CSV format, go to Inventory -> Virtual ->Settings -> 5G-Apps, enter the alias, type the CSV file name, and upload it.
- To download the files, go to Inventory -> Virtual ->Settings -> 5G-Apps and select the alias name to be downloaded.
- From release version 6.9.00, the maximum Transaction Flow value for GigaVUE V Series Nodes with memory less than 32 GB is 5500. For V Series Nodes classified as extra-large with memory greater than 32 GB, the maximum Transaction Flow value is 26000.
- From release version 6.9.00, to make use of the multithreading support for egress/ingress tunnels in GVHTTP2/5G-Cloud applications, you should update the thread count after the GigaVUE V Series Node is upgraded.
- The configuration of the 5G-Cloud application should be performed on extra-large VMs.
 - Configuration specifications for extra large VM:
 - vCPUs: 12-20
 - RAM: 32 GB
 - Disk: 160 GB
 - Supported Ports: Management Ports, Tunneling Ports, and 2 Data Ports

FHA Dashboards for 5G-Cloud Applications

After configuring the GVHTTP2 and 5G-Cloud applications, you can monitor the statistics for Oracle SCP and Nokia SCP by the reports displayed in the Dashboard.

To access the details, refer to **FHA Dashboards for 5G-Cloud Applications** topic in GigaVUE Fabric Management Guide.

5G-Cloud Ericsson SCP Support

The 5G Core follows a service-based architecture where control plane network functions communicate using the HTTP2 protocol. In Ericsson SCP Mode, HTTP2 transactions are mirrored through the Service Communication Proxy (SCP) and encapsulated using UDP-GRE for further processing.


Ericsson SCP

In **Ericsson vTAP**, traffic is tapped from the Packet Core Controller (PCC), including AMF and SMF, in Layer 7 HTTP SBI JSON format. The 5G SBI application supports this PCC traffic, ensuring efficient processing and monitoring. Refer to [5G-Service Based Interface](#)

Application.

In **Ericsson SCP** mode, the Service Communication Proxy (SCP) mirrors 5G SBI traffic within the Cloud Core Signaling Controller, capturing communication between multiple network functions via SCP.

The 5G-Cloud application in GigaVUE V Series Nodes is enhanced to receive and process traffic from the Ericsson SCP and forward it to the monitoring tools. This setup supports both configuration and visualization through GigaVUE-FM integration.



Traffic from the Ericsson 5G Core is encapsulated using UDP-GRE tunnels, with specific keys 3, 17, and 19 to identify different traffic flows from SCP. GigaVUE-FM serves as the centralized management platform orchestrating the visibility fabric, which includes the V Series Nodes. The V Series Nodes inspect, aggregate, and process the incoming packets.


After processing, traffic is encapsulated using L2GRE or VXLAN tunneling protocols and forwarded to either physical or virtual probes as follows:

- **Physical Probe:** If IP connectivity is not available, the traffic is sent to a GigaVUE HC/GigaVUE TA Series device, which serves as a traffic aggregator and distributor. The device decapsulates and distributes the traffic to the probe.
- **Virtual Probe:** If IP connectivity is available, the processed traffic is forwarded directly to virtual probe for detailed traffic inspection and analytics.

This end-to-end solution ensures comprehensive visibility into 5G network traffic, facilitating efficient monitoring and performance analysis across both physical and cloud-based infrastructures.

How Ericsson SCP Solution works

The following image shows the block diagram of the data flow in the V Series Nodes.



The encapsulated traffic enters the system, and based on the key values, is passed through a PCAP-NG, 5G SBI, and LB application to capture and monitor the data.

For traffic mirroring and flow processing in Ericsson VTap 5G Core network with key value 1, refer to [5G-Service Based Interface Application](#).

When traffic is processed through keys 3, 17, and 19, the **PCAPng** application reads and validates various blocks within received PCAPng files before forwarding them to the Load Balancing application. It processes data based on the packet type, which includes a combination of the following blocks.

1. Mandatory Blocks:
 - Section Header Block (SHB)
2. Optional Blocks:
 - Interface Description Block (IDB)
 - Enhanced Packet Block (EPB)
 - Simple Packet Block
 - Name Resolution Block
 - Interface Statistics Block

Also, the PCAPng application classifies domains by using Ericsson vTAP metadata and a configured domain table. Each Ericsson domain (for example, pc-mm, pc-sm, ccsn-5gc, ccrc-nrf, nrfagent-nrf, ccpc-pcf, sc-scp, sc-sepp) maps to a VLAN ID. The V Series applies this VLAN ID on the egress tunnel so downstream probes can distinguish traffic by Ericsson domain.

The domain table can also include others and unclassified to handle traffic that does not match known domains or cannot be classified. Refer to [PCAPng Application](#) for more details.

The **Load Balancing** application efficiently distributes traffic with different flow IDs across multiple cores, enabling parallel processing of packets with distinct stream IDs. It ensures that request and response frames with the same stream ID are processed sequentially within the same core, preventing race conditions where these frames might otherwise be handled in parallel across different cores. Refer to [Load Balancing](#).

The extracted data is encapsulated within a VXLAN and routed to the 5G-Cloud application. The 5G-Cloud application processes VXLAN packets by first decapsulating them to extract the HTTP2 headers and body. It then parses the EEVTAP TCP mirrored message and maps the transaction ID to track the full transaction.

After synthesizing the TCP flow and HTTP2 session to represent the complete 5G CNF/SCP/5G CNF transaction data, it forwards the data over an L2GRE/VXLAN tunnel to the probe.

You should enable VLAN Tagging on the egress tunnel (VXLAN or L2GRE) used with the domain-classification PCAPng instance so the tunnel applies the VLAN ID from the Domain VLAN Mapping table for each Ericsson domain. Refer to [Create Ingress and Egress Tunnels](#).



Notes:

- Ericsson SCP will process only Model C traffic.
- You can deploy traffic ingress applications together in a single V Series Node and Monitoring Session. However, for the 5G-Cloud and egress applications, you must deploy a separate Monitoring Session and V Series Node.
- 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.
- Ensure that all V Series Nodes within a single Monitoring Domain are running the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.

Configuration of 5G-Cloud Ericsson SCP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain in the V Series:

V Series Node 1:

S.No	Steps	Refer to
1	Configure UDP-GRE Ingress TEP to receive the HTTP2 messages. Note: You should configure individual UDP-GRE tunnels for each of the following key values: - Key Value 1: Ericsson vTAP - Key Values 3 and 17: Ericsson SCP	Create Ingress and Egress Tunnel (VMware vCenter) Create Ingress and Egress Tunnels (OpenStack) Create Ingress and Egress Tunnel (Third Party Orchestration)
2	Configure PCAPng application and link ingress TEP and PCAPng application instances in the Monitoring Session. NOTE: For Domain tagging configuration, refer to Configure the PCAPng Application for Domain Tagging section in PCAPng Application .	PCAPng Application
3	Configure LB application and link LB application and PCAPng application instances in the Monitoring Session.	Load Balancing
4	Create a VXLAN egress TEP. NOTE: Enable the Domain Tagging option when creating an egress tunnel to tag packets with the Ericsson domain specific VLAN IDs.	Create Ingress and Egress Tunnel (VMware vCenter) Create Ingress and Egress Tunnels (OpenStack) Create Ingress and Egress Tunnel (Third Party Orchestration)
5	Create a link between egress TEP and LB application.	




Note: When upgrading GigaVUE-FM with a Monitoring Session in V Series Node 1 to version 6.12 from an earlier release, follow these steps to ensure a successful migration:

- Undeploy the existing V Series Node 1 Monitoring Session.
- Delete any Egress VXLAN TEPs with the Multi-Tunnel feature enabled.
- Recreate the Egress VXLAN TEP with the required configuration and enable the Multi-Tunnel flag.
- Establish a Link connection between the LB application and the Egress VXLAN TEP.
- Deploy the Monitoring Session.

V Series Node 2:

S.No	Steps	Refer to
1	Create a ingress REP to extract the HTTP2 body.	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (OpenStack) Create Raw Endpoint (Third Party Orchestration)
2	Create a link between the VXLAN egress TEP and	

	ingress REP.	
3	Add the 5G-Cloud application in the Monitoring Session.	5G-Cloud Application
4	Create a link between ingress REP and the 5G-Cloud application.	
5	Create egress REP.	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (OpenStack) Create Raw Endpoint (Third Party Orchestration)
6	Create a link between the 5G-Cloud application and egress REP.	



Add 5G-Cloud Application in Ericsson SCP

Pre-requisite:


- You must upload CSV files containing a valid FQDN, NF Instance, and User Agent IDs and a valid IPv4/IPv6 address. Refer to [Add CSV file for IP Mapping](#).

You can add a 5G-Cloud application to:

- New Monitoring Session - Add the 5G-Cloud 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 add a 5G-Cloud application:

- In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



- On the application quick view, enter or select the required information as described in the [5G Cloud Ericsson SCP Configuration - Field References](#).

NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

Download Logs - 5G-Cloud application

You can view the log files of a V Series Node or download them as .CSV or .txt files.

To download the log files to a local environment:

- Go to **Traffic > VIRTUAL** > select your cloud platform.
- Select the required Monitoring Session and click **Details** on the 5G-Cloud application. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
- Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
- Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add CSV file for IP Mapping

To add the CSV file for IP mapping:

- Go to **Inventory > VIRTUAL** > select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
- 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:
 - **5G-Cloud FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.
 - **5G-Cloud NF Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.
 - **5G-Cloud UA** - Add the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.
4. Click **Choose File** in the **FileName** field to upload the CSV file to GigaVUE-FM.



Notes - Header details of the CSV file:

- **FQDN Mapping:** gigamon@vseries:/var/log/cloud5g_tabledir\$ cat FQDN.txt (The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional)
 - Header details:
FQDNid,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation
 - Example:
s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,170.00.13.187,,8080,,SCP,plolp
- **NF Instance Mapping:** gigamon@vseries:/var/log/cloud5g_tabledir\$ pwd NFID.txt (The fields NFID and ingressIP1 are mandatory. All other fields are optional)
 - Header details:
NFID,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation
 - Example: 8b94c7xx-d700-4xcd-90be-634x11a5xx0d,100.23.14.4,,8080,,AMF,plolp
- **User Agent Mapping:** gigamon@vseries:/var/log/cloud5g_tabledir\$ pwd UA.txt (The fields useragentID and ingressIPAddress are mandatory. All other fields are optional)
 - Header details: UserAgent,IngressIP1,NFType,NFLocation
 - Example: xx5gs20015gcay01,192.64.10.67,AMF,plolp

The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.

5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

FHA Dashboards for 5G-Cloud Applications

After configuring the 5G-Cloud application, you can monitor the statistics for Ericsson SCP by the reports displayed in the Dashboard. To access the details, refer to [FHA Dashboards for 5G-Cloud Applications](#).

5G Cloud Ericsson SCP Configuration - Field References

Field	Description
Application	The name 5g-Cloud appears by default.
Alias	Enter the alias name as cloud5g .

Mode	From the drop-down list, select Ericsson SCP Transparent .
RX Tunnel	
Type	Specify the tunnel type. The default is VXLAN.
Listening IP	Specify the tunnel's local listen IP address to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
RX Thread	Specify the RX value to receive the packet. Enter a value between 1 and 16. The default value is 8.
TX Tunnel	
Type	Specify the tunnel type. The available tunnel types are: • VXLAN • L2GRE
Tool IP	Specify the remote IP address to send the packet.
Destination Port	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending the packet. Enter a value between 1 and 65535.
VNI Id (Applicable only when the selected tunnel type is VXLAN)	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key (Applicable only when the selected tunnel type is L2GRE)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU.
Log Directory	Specify the path to store the log files.
Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug
SCP Config	
FQDN Mapping	Specify the alias name created for the uploaded FQDN table CSV file.
NF Instance Mapping	Specify the alias name created for the uploaded NFID table CSV file.

User Agent Mapping	Specify the alias name created for the uploaded User Agent table CSV file.
TCP Server Ports	Specify the TCP server port or port range to allow TCP communication endpoints.
SCP Advanced Config	
No. of TCP Flows	Specify the total number of concurrent TCP flows. Enter a value between 128 and 2048. The default value is 1024.
No. of Transaction Flows	Specify the total number of Transaction Stream Flows allocated for tracking active Transaction Streams. Enter a value between: 128 and 26000 for V Series Node of extra large form factor (total memory $\geq$ 32GB). 128 and 5500 for V Series Node of large form factor (total memory < 32GB). The default value is 2048.
TCP Flow Timeout	Specify the TCP flow timeout ranging between 0 to 7200. The default value is 900.
SCP Transaction Timeout	Specify the SCP transaction timeout value ranging between 1 and 300. The default value is 10. Note: The recommended timeout value is between 10 and 20 seconds. An increase in the transaction timeout value will impact the total number of transactions that can be processed in parallel.
Minimum TCP Client Port	Specify the minimum value to be used for TCP client port ranging between 1023 and 65535. The default value is 32768.
Maximum TCP Client Port	Specify the maximum value to be used for client port ranging between 1023 and 65535. The default value is 36863.
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. Enter a value between 1 and 16. The default value is 8.
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. Enter a value between 100 and 8000. The default value is 1000.
Header Index	Enable or disable the header indexing. The default value is disabled.
Header Compression Code	Enable or disable the header compression code. The default value is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default value is Enable.
Packet Capture Level	Select the packet capture level from the drop-down list: - All - Receive - Transmit - None
5G-Cloud Log Level	Select the required 5G-Cloud CSV log level from the drop down list. The default value is None. - All - Flow - Message - Transaction

	None
HTTP2 Monitored Flows	
Number of Streams	Specify the number of monitored stream flows. Enter a value between 1024 and 16384. The default is 1024.
Request Timeout	Specify the HTTP/2 request timeout in seconds. Enter a value between 1 and 300. The default is 15 seconds.
Response Timeout	Specify the HTTP/2 response timeout in seconds. Enter a value between 1 and 300. The default is 2 seconds.
TCP Monitored Flows	
Number of Flows	Specify the number of monitored TCP flows. Enter a value between: 1024 and 32768 for V Series Node of extra large form factor (total memory > 32GB). 1024 and 2048 for V Series Node of lesser form factor. The default value is 2048.
Flow Timeout	Specify the timeout for TCP flows in seconds. Enter a value between 1 and 7200. The default value is 60 seconds.
Flow Reassembly Timeout	Specify the timeout for TCP flow reassembly in seconds. Enter a value between 1 and 7200. The default value is 500 seconds.

5G Cloud Nokia HEP3 Support

Nokia HEP3 (Homer Encapsulation Protocol Version 3) method is supported by Nokia for network functions across 3G,4G, 5G and the IMS core. HEP3 efficiently organizes and encodes data for various network protocols and enhances the functionality of Nokia's core network solutions by accommodating diverse packet requirements.

The 5G-Cloud application supports processing the Nokia HEP3 format from the **IMS** network functions and 3G, 4G network functions. Non-SBI traffic is received as a TCP stream with HEP3 encoding by the 5G Cloud application. The 5G-Cloud application receives encapsulated traffic from the IMS vTAP in HEP3 over TCP, using either clear text or mTLS. Non-SBI control traffic is decoded from HEP3, emphasizing data handling without IP translation.

Nokia Cloud Mobility Manager (**CMM**), including AMF and MME, generates SBI and Non-SBI traffic over HTTP2 using clear text / TLS. The gvHTTP2 application receives this traffic and forwards it to the 5G-Cloud application. The 5G-Cloud synthesizes SBI payloads into L2GRE/VXLAN packets, enabling seamless monitoring across 5G core interfaces.

In **Nokia SCP** mode, the traffic is mirrored between the 5G control functions using HTTP2 protocol, which is mirrored from a service called SCP (Service Communication Proxy). The 5G-Cloud application receives the JSON encoded data from the VXLAN ingress TEPs. Refer to [Nokia SCP](#).

The system supports the following Nokia source functions and traffic types:

- 5G Cloud Nokia HEP3 IMS Integration (Non-SBI)
- 5G Cloud Nokia HEP3 CMM Integration (SBI and Non-SBI)

Supportability Table

The supportability table provides a quick reference to the supported Nokia functions:

Source Functions	Traffic Type	Protocols	Reference
3G, 4G, IMS	Non-SBI	HEP3 over TCP with Clear Text / mTLS	5G Cloud Nokia HEP3 IMS Support
CMM (AMF, MME)	SBI and Non-SBI	HEP3 over HTTP2 over TCP with Clear Text / TLS (via gvHTTP2)	5G Cloud Nokia HEP3 CMM Support


5G Cloud Nokia IMS Solution (Non-SBI)

Nokia HEP3 (Homer Encapsulation Protocol Version 3) is a binary encapsulation format that Nokia uses for control-plane traffic in the IMS core and across 3G and 4G network functions. HEP3 efficiently encodes protocol metadata and payloads so that Nokia core network solutions can handle diverse packet types.

In this solution, the 5G-Cloud application processes non-SBI IMS control traffic that Nokia vTAPs export as HEP3 over TCP, in clear text or protected with mTLS. 5G-Cloud also supports non-SBI HEP3 traffic from 3G and 4G network functions. It terminates the TCP session, decodes the HEP3 payloads, and exposes the non-SBI control traffic for monitoring without performing IP address translation.

How Nokia HEP3 IMS Solution works

The following image shows the ingress and egress data flow through the V Series Node for Nokia HEP3 IMS solution.



Traffic from client vTAP systems (in clear text or secured with mTLS) travels over TCP to the V Series node. The 5G-Cloud application receives this TCP traffic, terminates the TCP connections, decodes the HEP3 payloads, and synthesizes the data for monitoring.

GigaVUE-FM serves as the centralized management platform that orchestrates the visibility fabric, including the 5G-Cloud V Series nodes. The V Series nodes inspect, aggregate, and process the incoming traffic. After processing, the nodes encapsulate the traffic using L2GRE or VXLAN tunneling protocols and forward it to probes as follows:

Physical probe- If IP connectivity is not available, the node sends traffic to a GigaVUE HC or GigaVUE TA Series device, which aggregates and distributes traffic to the probe after decapsulation.

Virtual probe- If IP connectivity is available, the node forwards the processed traffic directly to the virtual probe for detailed inspection and analytics.

This end-to-end solution provides comprehensive visibility into 5G IMS control-plane traffic and supports efficient monitoring and performance analysis across both physical and cloud-based infrastructures.


Configure Nokia IMS HEP3 Solution (Non-SBI)

In GigaVUE-FM, you must do the following to add the 5G-Cloud IMS integration to either a new or existing monitoring session of a Monitoring Domain:

Pre-requisite:

Before starting, ensure that:

- You have deployed the V Series nodes and added them to a monitoring domain in GigaVUE-FM.
- IMS vTAPs can reach the V Series node over TCP on the HEP3 collector IP address and port.
- You must upload CSV files containing a valid FQDN and a valid IPv4/IPv6 address. Refer to [Add CSV file for IP Mapping](#).



1. Create New or Use Existing Monitoring Session

In GigaVUE-FM, go to the monitoring domain that will carry IMS HEP3 traffic. You can add the 5G-Cloud application, to a new or existing monitoring session of a Monitoring Domain.

- **New Monitoring Session** - Add the 5G-Cloud 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 Monitoring Session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.


2. Create an ingress REP to receive the data over TCP

- On the canvas, create an ingress raw endpoint (REP) on the V Series node to receive HEP3/TCP from the IMS vTAPs.
- Configure the listening IP address and TCP port that the vTAPs use as the HEP3 collector.
- Refer to [Create Raw Endpoint \(VMware vCenter\)](#), [Create Raw Endpoint \(OpenStack\)](#), or [Create Raw Endpoint \(Third Party Orchestration\)](#) for platform-specific steps.

3. Configure Nokia HEP3 in 5G-Cloud Application

To add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



- On the application quick view, enter or select the required information as described in the [Reference - Configuration Table](#).

NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

Reference - Configuration Table

Field	Description
Application	The name 5g-Cloud appears by default.
Alias	Enter the required alias name. (for example: NokiaHep3)
Vendor Integrations	From the drop-down list, select Nokia IMS .
SBI Operational Mode	The Operational Mode is grayed out when you select Nokia CMM or Nokia IMS.
Ingress Tunnel	
Type	Specify the tunnel type. The default is TCP.
Listening IP	Specify the tunnel's local listen IP address to receive packets.
Listening Port	Specify the tunnel's local listening port to bind for receiving packets. The application listens to the traffic on the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port to send packets. Enter a value between 1 and 65535.
Egress Tunnel	
Type	Specify the tunnel type. Available tunnel types are: VXLAN

	L2GRE
Tool IP	Specify the remote IP address to which you want to send the packets.
Destination Port	Specify the tunnel destination port to send the packets. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending packets. Enter a value between 1 and 65535.
Tool MTU	Specify the tool port MTU: - For V Series Node version 6.8.00 and above, enter a value between 1400 and 8800. The default value is 8800. - For V Series Node version below 6.8.00, enter a value between 1500 and 8800. The default value is 8800.
VNI ID (Applicable only when the selected tunnel type is VXLAN)	Specify the ID for VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key (Applicable only when the selected tunnel type is L2GRE)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Nokia IMS / CMM Configuration	
Number of Ingress TCP Connections	Specify the number of concurrent TCP connections VTAP can establish. Enter a value between 128 and 2048. The default value is 1024.
Ingress TCP Timeout	Specify the timeout value for an Ingress TCP connection in seconds. Enter a value between 30 and 3600 seconds. The default is 60 seconds. NOTE: If no packets are received within the configured time, the TCP connection will terminate due to a timeout.
Number of Egress TCP Flows	Specify the total number of TCP flows allocated for tracking non-SBI traffic. - For V Series Node (extra-large form factor, total memory $\geq$ 32 GB): 1024 – 16000000 - For V Series Node (large form factor, total memory < 32 GB): 1024 – 16384. - Default value: 4096.
Egress TCP Flow Timeout	Specify the timeout value for an Egress TCP flow in seconds. Enter a value between 30 and 7200 seconds. The default is 900 seconds. NOTE: If no packets are received within the configured time, the TCP flow will terminate due to a timeout.
Number of Receiver Threads	Specify the number of receiver threads for processing incoming packets. Enter a value between 1 and 128. The default value is 8.
Number of Egress SCTP Flows	Specify the total number of SCTP flows allocated to track non-SBI traffic. Enter a value between 128 and 2000000. The default value is 1024.

Egress SCTP Flow Timeout	Specifies the timeout value for an Egress SCTP flow in seconds. Enter a value between 30 and 7200. The default value is 900 seconds. NOTE: If no packets are received within the configured time, the SCTP flow will terminate due to a due to timeout.
mTLS Configuration	
mTLS	Use the toggle button to enable or disable mTLS encryption and decryption in Monitor Mode.
mTLS key	- If mTLS encryption is enabled, select the mTLS key alias for the SSL certificate from the drop-down list. - To create a key alias, click Add next to the text box. For further details, refer to Configure SSL Decrypt.
SBI / IMS Mappings	
Refer to Add CSV file for IP Mapping section for more details.	
FQDN Mapping	Use this option to provide mappings between Fully Qualified Domain Names and IP addresses for all Network Functions in the network. Select an alias: Choose the alias name associated with the uploaded FQDN table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one.
Network Function Instance Mapping (Optional)	Use this option to provide mappings between NF Instance IDs and IP addresses for all Network Functions in the network. Select an alias: Choose the alias name associated with the uploaded NFID table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one.
User Agent Mapping (Optional)	Use this option to provide mappings between User Agents and IP addresses for all Network Functions in the network. Select an alias: Choose the alias name associated with the uploaded User Agent table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one.
Service Mapping (Optional)	Use this option to map ephemeral attribute values to standard service attributes and define service-specific actions. Select an alias: Choose the alias name associated with the uploaded Service Mapping table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one. None: Select None to clear any previously selected Service Map entry. Refer to Reference - Service Map Field Definitions section for example details.
SBI Advanced Configuration	
NOTE: The below configurations apply only to SBI traffic and will not be used for non-SBI traffic.	
TCP Server Ports	Specify the TCP server port or port range to enable TCP communication between endpoints. - Valid port range: 1-65535. - Maximum number of ports: 1024.

	NOTE: TCP communication is allowed only on the configured ports.
Number of TCP Flows	Specify the total number of TCP Flows allocated for tracking active TCP Flows. - Valid range: 128-2048 - Default value: 1024
Number of Transaction Flows	Specify the total number of Transaction Stream flows allocated for tracking active Transaction. Refer to Recommend Form Factor for VMware vCenter (Instance Types) for more details. Streams: - Large Form Factor VMs (total memory < 32GB): 128-5500. - Extra Large Form Factor VMs (total memory ≥ 32GB): 128-26000 - Default value: 2048
TCP Flow Timeout	Specify the TCP flow timeout value. - Valid range: 0-7200 - Default value: 900
SCP Transaction Timeout	Specify the Transaction timeout value. - Valid range: 1-300 - Default value: 10
Minimum TCP Client Port	Specify the starting port number for Synthesize Flow. - Valid range: 1023-65535 - Default value: 32768
Maximum TCP Client Port	Specify the end port number for the Synthesize Flow. - Valid range: 1023-65535 - Default value: 36863
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. - Valid range: 1-16 - Default value: 8
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. - Valid range: 100-8000 - Default value: 1000
Header Index	Enable or disable the header indexing. The default option is disabled.
Header Compression Code	Enable or disable the header compression code. The default option is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default option is Enable.
Diagnostic Settings	
Log Directory	Specify the path to store the log files.
Application Log Level	Select the severity log level of the events from the following options: - None - Fatal - Error - Notification - Information

	• Debug
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
Packet Capture Timestamp	Select the required timestamp used to save the Egress PCAPs. The available timestamps are: • Observed Timestamp • System Timestamp NOTE: Timestamp option is disabled if the Packet capture level is set to None.
Message Trace Log Level	Select the 5G-Cloud CSV log level from the drop-down list. • All • Flow • Message • Transaction • None (Default Value)

1. Create links between the ingress REP, 5G-Cloud, and egress REP

- On the canvas, draw a link from the ingress raw endpoint to the 5G-Cloud application.
- Create an egress raw endpoint (REP) for the target probe or GigaVUE HC/TA device (using the same raw endpoint procedures as in Step 2).
- Draw a link from the 5G-Cloud application to the egress raw endpoint.

2. Deploy the Monitoring Session

- After adding all the required elements to the canvas, to deploy the session do the following:
- 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.

3. 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:
 - a. **Success:** Not deployed on one or more instances due to V Series Node failure.
 - b. **Failure:** Not deployed on all V Series Nodes or Instances.

4. Download Logs - 5G-Cloud application

You can view the log files of a V Series Node or download them as .CSV or .txt files.

To download the log files to a local environment:

- Go to **Traffic > VIRTUAL** > select your cloud platform.
- Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g** application, click the  menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
- Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
- Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add 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.
 - o **5G-Cloud FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
FQDNID	abc.xyz.com
IngressIP1	1.1.1.1
IngressIP2	2.2.2.2
IngressPort	100
EgressIP	3.3.3.3
NFType	SMF
NFLocation	US
IMSIP	1001:0008:85a2:0000:0000:8fle:0340:7110
IMSIntfName	100 200
IMSMask	96
IMSType	nef

**Notes:**

- The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.
- The fields imsipaddress, imsinterfacename, imsmask and imstype are applicable only when the Integrated solution Vendor is "Nokia HEP3".

- o **5G-Cloud NF Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
NFID	amf.local
ingressIPAddress1	1.1.1.1
ingressIPAddress2	2.2.2.2
IngressPort	100
egressIPAddress	3.3.3.3
NFType	AMF
NFLocation	US

NOTE: The fields NFID and ingressIP1 are mandatory. All other fields are optional.

- **5G-Cloud UA** - Add the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
useragentid	nef.nb.5gc.mc005.mcc580.yournetwork
ingressIPAddress	1.1.1.1
NFType	AMF
NFLocation	US

NOTE: The fields useragentID and ingressIPAddress are mandatory. All other fields are optional.

- **5G-Cloud Service Map** - Add the CSV file containing valid ephemeral port details and a valid IPv4/IPv6 address for IP mapping. Refer to [Reference - Service Map Field Definitions](#) for header details and examples.
4. Click **Choose File** in the **File Name** field to upload the CSV file into GigaVUE-FM.
 5. Click **Validate** to validate the CSV file.
 6. Click **Save** to add the CSV file.

Recommended VM Specifications

The configuration of the 5G-Cloud Nokia HEP3 application should be performed on extra-large and large VMs.

- Configuration specifications for Extra Large VM:
 - vCPUs: 26
 - Memory: 32 GB
 - Disk Space: 200 GB

NOTE: Nokia HEP3 configuration on extra-large VMs is supported only through the Third-Party Orchestration method.

- 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.

Reference - Service Map Field Definitions

Service mapping normalizes SIP traffic by converting vendor-specific or ephemeral ports to standard SIP service ports (such as 5060). GigaVUE-FM matches incoming SIP control messages against configurable rules and, when a rule applies, rewrites the source and/or destination port before forwarding the traffic. This ensures that monitoring tools that rely on standard SIP ports continue to work without any changes.

Refer below for field specific details:

NOTE: All fields except sipMessageType, srcPortMap, and dstPortMap are mandatory.

1. **srcIP** - Source IP address of the SIP traffic, as carried in HEP3. You can specify a single IPv4 / IPv6 address, a prefix, or ipany to match any address.

Example:

- ipany – matches SIP traffic from any source IP.
- 1.1.1.1 – matches only traffic coming from 1.1.1.1.

2. **dstIP** - Destination IP address of the SIP traffic. You can specify a single IPv4 / IPv6 address, a prefix, or ipany to match any address.

Example:

- ipany – any destination IP.
- 2.2.2.2 – match traffic going to 2.2.2.2.

3. **srcPort** - Source port number of the SIP traffic. You can use a specific port to target a known source, or portany to match all source ports.

Example:

- portany – match SIP traffic regardless of the source port.
- 5060 – match only when the source port is 5060.

4. **dstPort** - Destination port number of the SIP traffic.

Example:

- portany – any destination port.
- 5060 – only traffic with destination port 5060.

5. **protocolFamily** - Identifies the IP family of the traffic.

Example:

- IPv4 SIP traffic.
- IPv6 SIP traffic.

6. **transportProtocol** - Identifies the L4 transport protocol (TCP or UDP).

NOTE: This field accepts only tcp or udp; it does not support integer values.

7. **applicationProtocol** - Encodes the application protocol type carried in the HEP3 encapsulation.

Example:

- 1 – match SIP messages.

8. **interfaceName** - Identifies the Nokia interface number on which the traffic is received.

Example:

- 90 – match traffic received on Nokia HEP3 vTAP interface 90 (IMS Non-SBI).

9. **Direction** - Traffic direction from the node's point of view.

Example:

- ingress – apply this rule only to traffic arriving at the node.
- egress – apply this rule only to traffic sent out from the node.

10. **sipMessageType** - Tells whether the SIP message is a request or a response.

Example:

- request – map ports only for SIP requests.
- response – map ports only for SIP responses.

11. **action** - Specifies what the 5G Cloud application does with matching traffic.

Example:

- portmap – change source and/or destination ports as specified in srcPortMap / dstPortMap.
- discard – drop matching traffic.
- ignore – match but do not modify the ports.

12. **srcPortMap** - Applies only to server response frames. Use a valid port number in response frames (sipMessageType = response) and set it to portany for request frames.

Example:

- 5060 – set the source port to 5060 for matching SIP responses.
- portany – keep the original source port.

13. **dstPortMap** - Applies onto client-to-server request frames. Use a valid port number in request frames (sipMessageType = request) and set it to portany for response frames.

Example:

- 7777 – set the destination port to 7777 for matching SIP requests.
- portany – keep the original destination port.

FHA Dashboards for 5G-Cloud Applications

After configuring the 5G-Cloud application, you can monitor the statistics for Nokia HEP3 by the reports displayed in the Dashboard. To access the details, refer to [FHA Dashboards for 5G-Cloud Applications](#).

5G Cloud Nokia HEP3 CMM Solution (SBI and Non-SBI)

Nokia 5G network functions use the HEP3 binary format to encode and send packets. HEP3 supports both HTTP2 and TCP protocols, enabling efficient traffic transmission across core network functions. This solution supports traffic processing from Nokia Cloud Mobility Manager (CMM), and the Service Communication Proxy (SCP).


Cloud Mobility Manager (CMM) includes AMF and MME network functions. It generates both Service-Based Interface (SBI) and non-SBI traffic and sends it over HTTP2. The gvHTTP2 application receives CMM traffic and forwards it to the 5G Cloud application, which synthesizes SBI and non-SBI payloads into L2GRE/VXLAN packets before sending them to probes. This approach provides extended visibility across 5G core interfaces and enables operators to monitor both SBI and non-SBI traffic seamlessly.

How Nokia HEP3 CMM Solution works

Nokia vTAP systems in the mobility core send HEP3 data over both HTTP2 and TCP.

Nokia CMM Solution


The following image shows the ingress and egress data flow through the V Series node for the Nokia HEP3 CMM solution.



The Cloud Mobility Manager (CMM) function includes key control-plane elements such as AMF and MME. CMM generates both SBI and non-SBI traffic, making it central to mobility management in the 5G Core.

- CMM vTAPs export SBI and non-SBI control-plane traffic as HEP3 over HTTP2/TCP, in either clear text or TLS-protected sessions.

- The gvHTTP2 application on the V Series node listens on the configured HTTP2 IP address and port, terminates HTTP2 sessions from CMM, and extracts the encapsulated HEP3 payloads.
- gvHTTP2 sends the processed HEP3 payloads to the 5G-Cloud application over TCP by using a dedicated Nokia HEP3 over TCP tunnel between the gvHTTP2 and 5G-Cloud V Series nodes.
- 5G-Cloud receives HEP3/TCP from gvHTTP2, terminates the TCP connections, decodes the HEP3 payloads, and synthesizes SBI and non-SBI flows into a format that monitoring tools can consume without IP address translation.
- The V Series node encapsulates the synthesized traffic using L2GRE or VXLAN and then forwards it to:
 - Physical probe, through a GigaVUE HC or GigaVUE TA Series device that aggregates and decapsulates the traffic, or
 - Virtual probe, directly over IP when connectivity is available.



Configure Nokia CMM HEP3 Solution (SBI and Non-SBI)

In GigaVUE-FM, you must do the following to add the 5G-Cloud CMM integration to either a new or existing monitoring session of a Monitoring Domain:

Pre-requisite:

- You must upload CSV files containing a valid FQDN and a valid IPv4/IPv6 address. Refer to [Add CSV file for IP Mapping](#).

1. Create New or Use Existing Monitoring Session

In GigaVUE-FM, go to the monitoring domain that will carry IMS HEP3 traffic. You can add the 5G-Cloud application, to a new or existing monitoring session of a Monitoring Domain.

- **New Monitoring Session** - Add the 5G-Cloud 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 Monitoring Session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

2. Add gvHTTP2 and connect it to CMM vTAPs.


- On the canvas, drag and drop the gvHTTP2 application onto the V Series node.
- Select the gvHTTP2 node and click Details to open the quick view.
- Configuration of GvHTTP2 Application for CMM vTAP

Field	Description
Application	The name gvhttp2 appears by default.
Alias	Enter the alias name as gvhttp2 .
Vendor Integrations	From the drop-down list, select Nokia CMM .
HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specify the number of parallel threads that can be sued to process the requests. Enter a value between 1 and 16. The default value is 1.
TX Tunnel	
Type	Specify the tunnel type as Nokia HEP3 over TCP from the dropdown for Nokia CMM.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Destination Port	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Number of Egress TCP Connections	Specify the number of concurrent TCP connections GvHTTP2 can establish. Enter the value from 1 to 200. The default value is 4.
TLS Config	
TLS	Enable or disable TLS encryption or decryption using the toggle button.
Server Private key	Enter the path where the private key is stored.
Server Certificate	Enter the path where the certificate is stored.
Advanced Setting	
CSV	Enable this option to capture all stats counters for GvHTTP2.

PCAP	Enable this option to capture the packets at egress.
Log Directory	Specify the path to store the log files.
Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug

3. Add the 5G-Cloud application and connect it to gvHTTP2

- On the canvas, drag and drop the 5G-Cloud application onto the V Series node.
- Place 5G-Cloud downstream of gvHTTP2.
- Draw a link from gvHTTP2 to 5G-Cloud so that gvHTTP2 forwards HEP3/TCP to 5G-Cloud.
- Select the 5G-Cloud node and click Details to open the quick view.



NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

4. Configuration of 5G-Cloud Application for CMM vTAP

Field	Description
Application	The name 5g-Cloud appears by default.
Alias	Enter the required alias name. (for example: NokiaHep3)
Vendor Integrations	From the drop-down list, select Nokia CMM .
SBI Operational Mode	The Operational Mode is grayed out when you select Nokia CMM or Nokia IMS.
Ingress Tunnel	
Type	Specify the tunnel type. The default is TCP.
Listening IP	Specify the tunnel's local listen IP address to receive packets.
Listening Port	Specify the tunnel's local listening port to bind for receiving packets. The application listens to the traffic on the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port to send packets. Enter a value between 1 and 65535.
Egress Tunnel	
Type	Specify the tunnel type. Available tunnel types are: • VXLAN • L2GRE
Tool IP	Specify the remote IP address to which you want to send the packets.
Destination Port	Specify the tunnel destination port to send the packets. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending packets. Enter a value between 1 and 65535.
Tool MTU	Specify the tool port MTU: • For V Series Node version 6.8.00 and above, enter a value between 1400 and 8800. The default value is 8800. • For V Series Node version below 6.8.00, enter a value between 1500 and 8800. The default value is 8800.
VNI ID (Applicable only when the selected tunnel type is VXLAN)	Specify the ID for VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key (Applicable only when the selected tunnel type is L2GRE)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Nokia IMS / CMM Configuration	
Number of Ingress TCP Connections	Specify the number of concurrent TCP connections VTAP can establish. Enter a value between 128 and 2048. The default value is 1024.
Ingress TCP Timeout	Specify the timeout value for an Ingress TCP connection in seconds. Enter a value between 30 and 3600 seconds. The default is 60 seconds.

	NOTE: If no packets are received within the configured time, the TCP connection will terminate due to a timeout.
Number of Egress TCP Flows	Specify the total number of TCP flows allocated for tracking non-SBI traffic. - For V Series Node (extra-large form factor, total memory $\geq$ 32 GB): 1024 – 16000000 - For V Series Node (large form factor, total memory < 32 GB): 1024 – 16384. - Default value: 4096.
Egress TCP Flow Timeout	Specify the timeout value for an Egress TCP flow in seconds. Enter a value between 30 and 7200 seconds. The default is 900 seconds. NOTE: If no packets are received within the configured time, the TCP flow will terminate due to a timeout.
Number of Receiver Threads	Specify the number of receiver threads for processing incoming packets. Enter a value between 1 and 128. The default value is 8.
Number of Egress SCTP Flows	Specify the total number of SCTP flows allocated to track non-SBI traffic. Enter a value between 128 and 2000000. The default value is 1024.
Egress SCTP Flow Timeout	Specifies the timeout value for an Egress SCTP flow in seconds. Enter a value between 30 and 7200. The default value is 900 seconds. NOTE: If no packets are received within the configured time, the SCTP flow will terminate due to a due to timeout.
mTLS Configuration	
mTLS	Use the toggle button to enable or disable mTLS encryption and decryption in Monitor Mode.
mTLS key	- If mTLS encryption is enabled, select the mTLS key alias for the SSL certificate from the drop-down list. - To create a key alias, click Add next to the text box. For further details, refer to Configure SSL Decrypt.
SBI / IMS Mappings	
Refer to Add CSV file for IP Mapping section for more details.	
FQDN Mapping	Use this option to provide mappings between Fully Qualified Domain Names and IP addresses for all Network Functions in the network. Select an alias: Choose the alias name associated with the uploaded FQDN table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one.
Network Function Instance Mapping (Optional)	Use this option to provide mappings between NF Instance IDs and IP addresses for all Network Functions in the network. Select an alias: Choose the alias name associated with the uploaded NFID table (CSV file) from the drop-down list. Add a new alias: If no alias exists, click Add to create one.
User Agent Mapping (Optional)	Use this option to provide mappings between User Agents and IP addresses for all Network Functions in the network.

	• Select an alias: Choose the alias name associated with the uploaded User Agent table (CSV file) from the drop-down list. • Add a new alias: If no alias exists, click Add to create one.
Service Mapping (Optional)	Use this option to map ephemeral attribute values to standard service attributes and define service-specific actions. • Select an alias: Choose the alias name associated with the uploaded Service Mapping table (CSV file) from the drop-down list. • Add a new alias: If no alias exists, click Add to create one. • None: Select None to clear any previously selected Service Map entry.
SBI Advanced Configuration	
NOTE: The below configurations apply only to SBI traffic and will not be used for non-SBI traffic.	
TCP Server Ports	Specify the TCP server port or port range to enable TCP communication between endpoints. • Valid port range: 1-65535. • Maximum number of ports: 1024. NOTE: TCP communication is allowed only on the configured ports.
Number of TCP Flows	Specify the total number of TCP Flows allocated for tracking active TCP Flows. • Valid range: 128-2048 • Default value: 1024
Number of Transaction Flows	Specify the total number of Transaction Stream flows allocated for tracking active Transaction. Refer to Recommend Form Factor for VMware vCenter (Instance Types) for more details. Streams: • Large Form Factor VMs (total memory < 32GB): 128-5500. • Extra Large Form Factor VMs (total memory ≥ 32GB): 128-26000 • Default value: 2048
TCP Flow Timeout	Specify the TCP flow timeout value. • Valid range: 0-7200 • Default value: 900
SCP Transaction Timeout	Specify the Transaction timeout value. • Valid range: 1-300 • Default value: 10
Minimum TCP Client Port	Specify the starting port number for Synthesize Flow. • Valid range: 1023-65535 • Default value: 32768
Maximum TCP Client Port	Specify the end port number for the Synthesize Flow. • Valid range: 1023-65535 • Default value: 36863
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. • Valid range: 1-16

	• Default value: 8
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. • Valid range: 100-8000 • Default value: 1000
Header Index	Enable or disable the header indexing. The default option is disabled.
Header Compression Code	Enable or disable the header compression code. The default option is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default option is Enable.
Diagnostic Settings	
Log Directory	Specify the path to store the log files.
Application Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
Packet Capture Timestamp	Select the required timestamp used to save the Egress PCAPs. The available timestamps are: • Observed Timestamp • System Timestamp NOTE: Timestamp option is disabled if the Packet capture level is set to None.
Message Trace Log Level	Select the 5G-Cloud CSV log level from the drop-down list. • All • Flow • Message • Transaction • None (Default Value)

5. Create the egress raw endpoint and tool link

- On the canvas, create an egress raw endpoint (REP) on the V Series node that connects to your probe or GigaVUE HC/TA device.
- Draw a link from the 5G-Cloud application to the egress REP so that 5G-Cloud forwards encapsulated L2GRE/VXLAN traffic to the tool.

6. Deploy the monitoring session

- On the Monitoring Sessions page, select the session.
- From the Actions menu, select Deploy.
- Wait for deployment to complete on all V Series nodes.
- Confirm that the session Status shows Success.
- If the status shows Failure or partial success, open the SOURCES and V SERIES NODES tabs and click the Status link to view the deployment report.

Download Logs - 5G-Cloud / GvHTTP2 Application

You can view the log files of a V Series Node or download them as .CSV or .txt files.

To download the log files to a local environment:

1. Go to **Traffic > VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g / gvhttp2** application, click the  menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
3. Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add 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 **5G-Cloud FQDN**. You can also choose to download the template using the **Download Template** option and add the required entries.
4. Click **Choose File** in the **FileName** field to upload the CSV file into GigaVUE-FM. Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.



Notes - Header details of the CSV file:

FQDN Mapping: gigamon@vseries:/var/log/cloud5g_tabledir\$ cat FQDN.txt



- Header details:
FQDNid,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation,IMSIP,IMSIntfName,IMSMask,IMSType
- Example:
s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,170.00.13.187,,8080,,SCP,plolp,2000:4000:124:a0d9:ac1:4::,92,303 314 315,sbc

The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.

The fields imsipaddress, imsinterfacename, imsmask and imstype are applicable only when the Integrated solution Vendor is "Nokia HEP3".

5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

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.

For details, refer to the following topics:

- Configuration Health Monitoring
- Traffic Health Monitoring
- View Health Status
- View Monitoring Session Statistics
- Visualize the Network Topology

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 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
RawEnd Point	1. Tx Packets 2. Rx Packets 3. Tx Bytes 4. Rx Bytes	1. Difference 2. Derivative	1. Over 2. Under

	5. Tx Dropped 6. Rx Dropped 7. Tx Errors 8. Rx Errors		
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 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
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	1. Difference 2. Derivative	1. Over 2. Under

	6. Ingestor - Rx Octets 7. Ingestor - Octets Dropped 8. Ingestor - Records Dropped 9. Workload - Records Dropped 10. Workload - Req Auth Errors 11. Workload - Req Timeout Errors 12. Workload - Req Errors 13. Exporter - Avg File Size 14. Exporter - File Uploads 15. Exporter - File Uploads Errors 16. Enrichment - One Minute Percent		
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_CANCELED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_CANCELED

Message	Description
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 OVERVIEW page of Monitoring Session allows you to analyze the incoming and outgoing traffic across multiple time intervals. For example, hourly, daily, weekly, and monthly.

You can view the high-level information of the selected Monitoring Session, such as:

- Connections
- Tunnel details
- Health status
- Deployment status
- Information related to Application Intelligence statistics.

You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab.

View Statistics with Multiple Filters

You can apply different filters per the data analysis requirements to view the statistics.

GigaVUE-FM allows you to perform the following actions on the Monitoring Session Statistics page:

- Full Screen View
- Time Interval Options
- View with Traffic Data
- Per-Node Statistics
- Element-Level Insights

Option 1: View Full Screen

You can view the Statistics on full screen.

To view in full screen,

1. From the right side of the window, select the Actions drop-down list.
2. Select Full Screen.

Statistics appear in full screen.

Option 2: View using Time Interval

Select the options from the drop-down list box to view the incoming and outgoing traffic hourly, daily, weekly, and monthly.

Interval	Frequency of data points
Hourly	Every five minutes
Daily	Every one hour
Weekly	Every six hours
Monthly	Everyday

The most recent data point shown is slightly earlier than the current time (for example, viewing hourly stats at 11:30 will show the latest point at 11:25).

For each interval, the data displayed reflects the period starting from that data point. For instance, a data point marked 11:30 in daily view represents traffic from 11:30 to 12:30.

Option 3: View with Traffic Data

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.

Option 4: View Per-Node

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,

1. On the bottom-left corner of the OVERVIEW page, go to the GigaVUE V Series Node drop-down list.
2. Select the name of the V Series Node for which you want to view the statistics.

Option 5: Element-level statistics

You can view the statistics of the elements involved in the Monitoring Session.

To view the statistics,

1. Go to the Select Chart Options page.
2. Select the elements associated with the session.
3. On the graph, directly select the following options to view the statistics individually:
 - Incoming(Mbps)
 - Outgoing (Mbps)
 - Ratio (Out/In) (Mbps)

Raw EndPoint (REP)

A Raw EndPoint (REP) is part of the monitoring session but can also receive bypassed traffic that is not filtered by the map, so it records more packets than expected. For example, if the map has a rule as IPv4, but the REP can receive bypassed (non-IPv4) traffic, the recorded number of packets from the V Series node can be more than expected.

Visualize the Network Topology (VMware ESXi)

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 ESXi

You can perform the following administrative tasks in GigaVUE-FM for GigaVUE Cloud Suite for VMware ESXi.

Manage Monitoring Domain

You can perform the following actions in the Monitoring Domains and Fabric tabs of the Monitoring Domain page.

Actions	Description
Monitoring Domain 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
Edit Monitoring Domain	Used to edit a monitoring domain.
Deploy Fabric	Used to deploy GigaVUE V Series Nodes.
Upgrade Fabric	Used to upgrade GigaVUE V Series Nodes. For details, refer to Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi
Delete Monitoring Domain	Used to delete a Monitoring Domain.
Delete Fabric Nodes	Used to delete a GigaVUE V Series Node in the monitoring domain.
Connect	Used to start communication between GigaVUE-FM and the VMware vCenter.
Disconnect	Used to stop communication between GigaVUE-FM and the VMware vCenter.
Edit SSL Configuration	Used to add Certificate Authority and the SSL Keys when using the Secure Tunnels.

Actions	Description
Fabric tab	
Delete Fabric Nodes	Used to delete a specific GigaVUE V Series Node in the monitoring domain.
Rediscover	The changes made in vCenter for the GigaVUE V Series is reflected in GigaVUE-FM. Refer to Rediscover topic for more detailed information.
Power On	You can select an individual GigaVUE V Series Node and power it on. The status of the GigaVUE V Series Node is changed to Ok .
Power Off	You can select an individual GigaVUE V Series Node and power it off. If the GigaVUE V Series Node is turned off from GigaVUE-FM, then it is not considered as part of Cloud Health Monitoring and GigaVUE-FM does not try to turn it on. The status of the GigaVUE V Series Node is changed to Down .
Reboot	You can select an individual GigaVUE V Series Node and reboot it.
Rename	Used to rename the fabric nodes.
Generate Sysdump	You can select one or multiple GigaVUE V Series Nodes (Maximum 10) to generate the system files. The generation of sysdump takes a few minutes in a GigaVUE V Series Node. You can proceed with other tasks, and upon completion, the status appears in the GUI. These system files are helpful for troubleshooting. For more information, refer to Debuggability and Troubleshooting

More information

- To view and manage the generated sysdump files, select the GigaVUE V SeriesNode and select the **Sysdump** tab in the lower pane.
- To view the certificates associated with the fabric, select the fabric nodes and select the **Certificates** tab in the lower pane.

Configure VMware Settings

This section describes 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 vCenter (V Series)**
2. Select **Settings > Advanced Settings** to edit the VMware vCenter settings.

Advanced Settings

Edit

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

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.

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.
- o 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. Cancel Submit

Group Name *

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.:
 - o **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.
 - o (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 ①	EXTERNAL IP A...	ENTITY ID ①	HOST NAME	APPLICATION STA...	SOFTWARE VERSI...	SYSTEM UPTIME	IDP METADATA ... ①
✔ 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 Distribution	All GigaVUE-FMs are active-eligible instances.	Services will be equally 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_CNTLR_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* 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 1 Remove Standby GigaVUE-FM Instance](#).



The screenshot shows the 'High Availability' dashboard. At the top, a green checkmark indicates the HA group is 'Healthy'. Below this, a table lists the nodes in the group. The first node is in a 'Standby' state. A context menu is open for this node, showing options: 'Edit', 'Reboot', and 'Remove from group'. A red arrow points to the 'Remove from group' option. The table has columns for STATUS, IP ADDRESS, TUNNEL IP, ROLE, EXTERNAL IP, ENTITY ID, HOST NAME, APPLICATION STATUS, SOFTWARE VERSION, SYSTEM UPTIME, and IDP METADATA.

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 1 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 2 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	The newly elected leader resumes certificate handling and continues any

Reason for Failover	Description
deployment/upgrade during HA leader switchover	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 Analytics¹, 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

Dashboard	Displays	Visualizations	Displays
			as 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.
		<i>Dedup Packets Detected/Dedup</i>	Percentage of the de-

Dashboard	Displays	Visualizations	Displays
		<i>Packets Overload Percentage</i>	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 filters, based on which the visualizations will get updated:	<i>App Bytes</i>	Displays received traffic vs transmitted traffic, in Bytes.

Dashboard	Displays	Visualizations	Displays
	• 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.

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. 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**.
2. 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-ESXi)

This section addresses frequently asked questions about Secure Communication between GigaVUE Fabric Components and GigaVUE-FM. For more details, refer to the Secure Communication between GigaVUE Fabric Components section.

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.