



GigaVUE Cloud Suite Deployment Guide - Google Cloud Platform (GCP)

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

Contents

GigaVUE Cloud Suite Deployment Guide - Google Cloud Platform (GCP)	1
Contents	3
GV-Cloud-GCP	9
Get Started	9
Overview of GigaVUE Cloud Suite for GCP	9
GigaVUE Fabric Components and Deployment Options	11
travenKey Concepts	15
Precryption Technology on Single Node	26
Precryption Technology on Multi-Node	27
Supported Kernel Version	28
Note	29
Supported Platforms	30
Deployment Prerequisites	32
Reduced Control-Plane Ports	34
New Deployments	35
Existing Deployments	35
Migration to Reduced Control-Plane Ports	36
GigaVUE-FM	36
UCT-V Controller	36
UCT-V	37
GigaVUE V Series Node	37
GigaVUE V Series Proxy (Optional)	37
UCT-V OVS Controller	38
UCT-V OVS Module	38
GigaVUE FM	38
UCT-V Controller	39
UCT-V	40
GigaVUE V Series Node	41
GigaVUE V Series Proxy (Optional)	42
UCT-V OVS Controller	43
UCT-V OVS Module	43
Legacy Control-Plane Ports	44
GigaVUE-FM	44
GigaVUE V Series Node	44

GigaVUE V Series Proxy (optional)	45
Reduced Control-Plane Ports	45
New Deployments	46
Existing Deployments	47
Migration to Reduced Control-Plane Ports	47
GigaVUE-FM	47
UCT-V Controller	48
UCT-V	48
GigaVUE V Series Node	48
GigaVUE V Series Proxy (Optional)	49
UCT-V OVS Controller	49
UCT-V OVS Module	49
GigaVUE FM	49
UCT-V Controller	50
UCT-V	51
GigaVUE V Series Node	52
GigaVUE V Series Proxy (Optional)	53
UCT-V OVS Controller	54
UCT-V OVS Module	54
Legacy Control-Plane Ports	55
Single Network Interface Configuration	62
Multiple Network Interface Configuration	62
Loopback Network Interface Configuration	63
Linux Network Firewall Requirements	63
Install Linux UCT-Vs	64
Prerequisites	64
Install Linux UCT-Vs using Installation Script	64
Linux UCT-V Installation Scenarios	67
Install Linux UCT-Vs using Manual Configuration	67
Install UCT-V from Ubuntu/Debian Package	67
Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS	70
Post Deployment Check:	72
Register Linux UCT-V	72
Windows Network Firewall Requirements	73
Install Windows UCT-Vs	73
Rules and Notes:	73
Install Windows UCT-Vs using Installation Script	74
Windows UCT-V Installation Scenarios	76
Install Windows UCT-Vs using Manual Configuration	77
Register Windows UCT-V	79
Upgrade UCT-V through GigaVUE-FM (Recommended Method)	80
Rules and Notes:	80
Upload the UCT-V Images	81

Upgrade the UCT-V	81
Upgrade UCT-V Manually	83
Deploy	85
	86
Deploy Fabric Components using Generic Mode	86
Acquire Traffic using UCT-V - By Creating Monitoring Domain	86
Post Deployment Check:	96
Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain	116
Configure GigaVUE Fabric Components using GCP	121
Configure Packet Mirroring for GCP	129
Configure Secure Communication between Fabric Components in FMHA	131
Configure Secure Tunnel for Third Party Orchestration	132
Configure Prefiltering	138
Create Precryption Template for UCT-V	140
Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform	145
License	145
Supportability	145
Rules and notes	146
Prerequisites	147
Get Started	147
Deployment Using Terraform	148
Install the Gigamon Terraform Provider Without Internet Access	150
Single-file configuration	151
Module-based configuration	153
Run the Terraform Deployment File	154
Best Practices	158
Verification and Troubleshooting	159
Configure	161
Configure UCT-V Features	161
Configure Prefiltering	161
Create Precryption Template for UCT-V	163
Configure Monitoring Session	168
Create a Monitoring Session (Third Party Orchestration)	168
Configure Application Filtering Intelligence for Virtual Environment	187
View Application Statistics for Application Filtering	192
Application Metadata Exporter Deployment Options	193
On-Premises	194
Hardware (AMI)	194
Hardware (Control Plane Metadata)	194
Private Cloud (VMware)	195
Public Cloud	196
Application Metadata Exporter - Export Scenarios	196

Export Application Metadata in JSON format over HTTPS/Kafka	196
Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	197
Supported Platforms	198
Important Notes	198
Export of 3G/4G/5G Control Plane Metadata by AMX	199
Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	199
Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	201
Export of GigaVUE Enriched Metadata for Cloud Workloads (Virtual Machines)	201
Export of GigaVUE Enriched Metadata for Cloud Workloads Kubernetes containers - Azure Kubernetes Service (AKS)	202
Rules and Notes	203
Configure Application Metadata Exporter	204
Switching to outer_ip for Mobility Enrichment	212
Application Monitoring Options	212
Deploy Application Metadata Exporter in GCP using Third Party Orchestration	213
Predeployment Checklist:	213
(Optional) Install GigaVUE-FM using ISO Image File on VMware ESXi Host	215
Access GigaVUE-FM	216
(Optional) Configure Metadata Export Filter for AMX	217
(Optional) Configure Role-Based Access for Third Party Orchestration	217
Create a Custom User	217
Create a Custom Role and Assign the User	218
Create a Custom User Group and Assign the Custom Role and the User	218
Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source	219
Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM	219
Deploy V Series Node and Register with GigaVUE-FM	220
Register GigaVUE V Series Node and GigaVUE V Series Proxy during Instance Launch	221
Register GigaVUE V Series Node and GigaVUE V SeriesProxy after Instance Launch	221
Create a Monitoring Session to Configure Application Metadata Exporter	222
Configure Ingestor in AMX to Process Traffic	224
Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka	224

Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	225
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	225
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	226
Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)	227
Prerequisites	227
Integrate ZPA with Gigamon AMX	228
Configure Ingestor for AMX	228
Integrate AMX with Third Party Tool	230
Set up the Third Party Tool for AMX Integration	230
Configure Exporter to Export Attributes to Third Party tools	230
Configure Cloud Tool Exporter	230
Deploy the Monitoring Session	230
Cloud Tool Exporter - Field References	231
Configure Kafka Exporter	231
Deploy the Monitoring Session	232
Kafka Exporter - Field References	232
Extend Your AMI and AMX Deployment	233
Verify Your Deployment	234
View Application Metadata Exporter Statistics	234
Attributes for GigaVUE Enriched Metadata for Mobile Networks	237
Control Plane Attributes	238
User Plane Attributes	240
Attributes for GigaVUE Enriched Metadata for Cloud Workloads	243
AWS	244
Azure	245
VMware vCenter	247
Azure Kubernetes Service (AKS)	247
Zscaler Private Access (ZPA) Attributes	249
Supported Element Types for NetFlow Integrator Element Mapping	250
Configure Application Metadata Intelligence for Virtual Environment	255
Extend Your AMI and AMX Deployment	262
View Application Statistics for Application Metadata	263
Configure Application Visualization for Virtual Environment	263

Configure Filtering and Metadata Export for Selected Applications in Application Visualization	265
Filter Traffic for Selected Applications	265
Export Metadata for Selected Applications	266
View Application Statistics for Application Visualization	266
Feature Overview	268
Configure De-duplication Application	268
Distributed De-duplication	269
Limitations	270
What's Next	270
What's Next	271
Create NetFlow Session for Virtual Environment	272
Limitation	276
Configure Header Stripping Application	276
Create Custom Port Template	278
What's Next	279
Enhanced Load Balancing	281
What's Next	285
What's Next	288
What's Next	290
Supported Protocols, Algorithms, and Ciphers for SSL Decrypt	292
Configure SSL Decrypt	294
Upload SSL Key Pairs	295
Create SSL Service	296
Prerequisite	296
Key Mapping	296
SSL Key Pair Store	297
Add SSL Decrypt to Monitoring Session	297
Interface Mapping	301
Tool Exclusion	301
Configure Precryption in UCT-V	301
Monitor and Troubleshoot	303
Monitor Cloud Health	303
Configuration Health Monitoring	303
Traffic Health Monitoring	304
Apply Threshold Template to Monitoring Session	308
Apply Threshold Template to Applications	309
Clear Thresholds for Applications	309
Clear Thresholds across the Monitoring Session	309
View Health Status	310
View Monitoring Session Statistics	312
Visualize the Network Topology (Third Party Orchestration)	312
Administer GigaVUE Cloud Suite for Third Party Orchestration	314

Configure Certificate Settings	314
Interface Mapping	314
Configure Third Party Orchestration Settings	315
About Audit Logs	317
High Availability for Cloud Infrastructure	319
Debuggability and Troubleshooting	340
Sysdumps	340
Additional Sources of Information	342
Documentation	342
Documentation Feedback	345
Contact Technical Support	346
Contact Sales	346
The VUE Community	347

GV-Cloud-GCP

Get Started

Overview of GigaVUE Cloud Suite for GCP

GigaVUE Cloud Suite for GCP delivers a cloud-native visibility and analytics solution for Google Cloud Platform (GCP).

It enables full network observability so you can:

- Migrate workloads to the cloud
- Optimize your existing cloud environment

The suite removes network blind spots. It helps strengthen security and compliance and lets you troubleshoot performance issues faster.

Traffic Acquisition

GigaVUE Cloud Suite for GCP captures network traffic from workloads running in GCP. It uses a simple and efficient process:

- Copies and mirrors traffic from selected workloads
- Sends the traffic to GigaVUE V Series Nodes

- Filters and processes the packets
- Forwards the processed traffic to your security and monitoring tools

The suite supports multiple traffic acquisition methods in GCP, including:

- GigaVUE Universal Cloud Tap (UCT)
- GCP Packet Mirroring

Unified Data Visibility

The suite gives you a single, clear view of all data in motion. It works across single, hybrid, and multi-cloud environments.

GigaVUE Cloud Suite for GCP:

- Collects traffic from the workloads you choose to monitor
- Optimizes traffic before sending it to tools
- Delivers relevant traffic to the right monitoring and security tools

This end-to-end visibility helps your tools analyze raw packet data, application activity, and enriched metadata so you can detect threats faster and resolve issues sooner.

Enriched Metadata

GigaVUE Enriched Metadata (GEM) converts raw packet data into actionable insights. It extracts and enriches metadata and sends the resulting information to security and analytics tools.

This helps speed up threat detection and performance monitoring by linking platform context with packet-level visibility.

Key Capabilities

GigaVUE Cloud Suite for GCP offers the following capabilities:

- Captures traffic in GCP using UCT or Packet Mirroring
- Deploys deep observability nodes that aggregate, select, and distribute traffic to tools
- Uses GigaSMART applications such as Packet De-duplication, Application Filtering Intelligence, and Application Metadata Intelligence to improve tool efficiency
- Scales dynamically as VM and container instances change

- Extends existing security, monitoring, and compliance workflows into GCP environments
- Supports deployment through third-party orchestration for GCP environments

GigaVUE Cloud Suite for GCP helps you extend deep observability into Google Cloud so you can improve security, reduce blind spots, and make better operational decisions across your cloud environment.

GigaVUE Fabric Components and Deployment Options

You can deploy GigaVUE Fabric Components in your OCI cloud environment and configure to support the visibility policies you have defined.

GigaVUE Cloud Suite for OCI includes the following fabric components:

- GigaVUE-FM
- UCT-V
- UCT-V Controller
- GigaVUE V Series Node
- GigaVUE V Series Proxy

You can use the following approach to deploy GigaVUE Fabric Components:

Third Party Orchestrated Fabric Deployment

GigaVUE fabric components support flexible deployment options beyond the standard GigaVUE-FM deployment method. You can use third-party orchestration tools, such as Terraform or custom scripts, to deploy these components, enabling seamless integration with existing infrastructure management workflows.

When deploying fabric components through an external orchestration system,

- You can provide the necessary parameters for self-registration with GigaVUE-FM.
- Once registered, GigaVUE-FM serves as the central point for configuring Monitoring Sessions and associated services
- It ensuring unified management and control of the deployed fabric.

Benefits

- Leverage existing orchestration workflows for automated deployment.
- Ensure seamless integration with infrastructure management tools.
- Maintain centralized monitoring and service configuration through GigaVUE-FM.
- Enhances flexibility while maintaining control over network monitoring infrastructure.

Deployment Methods with Third Party Orchestration

- [User Driven Fabric Deployment](#)
- [Automated Fabric Deployment](#)

User Driven Fabric Deployment

This method uses a combination of the cloud platform's UI or CLI.

To use the user-driven fabric deployment method:

1. Deploy GigaVUE-FM

Deploy GigaVUE-FM in your OCI environment by using the OCI Console after importing the required custom image.

2. Create a Monitoring Domain

Log in to GigaVUE-FM and create a Monitoring Domain and Connection to define the boundaries for your fabric components.

3. Deploy the fabric components

Use the OCI Console to deploy the following components into the required VCN/subnets within the Monitoring Domain:

- UCT-V Controller: Deploy this if you are using UCT-V to acquire traffic.
- GigaVUE V Series Nodes: Deploy these for traffic processing.
- GigaVUE V Series Proxy: Deploy this if required for your environment.

With these steps, you can efficiently deploy and configure the fabric components in your OCI environment.

Automated Fabric Deployment

When using Infrastructure as Code (IaC) methodology, you can automate the deployment of a fabric component using scripts written in tools such as Terraform, Ansible, or CloudFormation.

To use the automated fabric deployment method,

1. Upload and import the images

Upload the required GigaVUE QCOW2 images to an OCI Object Storage bucket and import them as custom images in OCI Compute before you start deployment.

2. Deploy GigaVUE-FM

Deploy GigaVUE-FM in your OCI environment using the OCI Console.

NOTE: Deploy GigaVUE-FM before deploying any other fabric component because other components depend on it for registration and management.

3. Create a Monitoring Domain

Log in to GigaVUE-FM and create a Monitoring Domain and Connection to define the boundaries for your fabric components.

4. Deploy the fabric components

Use the OCI Console to deploy the following components into the required VCN/subnets within the Monitoring Domain:

- **UCT-V Controller:** Deploy this if you are using UCT-V to acquire traffic.
- **UCT-V:** Deploy Linux or Windows UCT-V instances, if applicable.
- **GigaVUE V Series Nodes:** Deploy these for traffic processing.
- **GigaVUE V Series Proxy:** Deploy this if required for your environment.

5. Add additional VNICs for GigaVUE V Series

For GigaVUE V Series, attach the required secondary VNICs after instance launch and reboot the instance so OCI can apply the IP configuration correctly inside the guest OS.

6. Verify registration and health

In GigaVUE-FM, verify that UCT-V Controller, UCT-V, GigaVUE V Series Nodes, V Series Proxy, and GigaVUE-FM appear under the expected Monitoring Domain and Connection, and confirm that their status is Healthy.

With these steps, you can efficiently deploy and configure the fabric components. Refer to [Third Party Orchestrated Fabric Deployment](#) topic for more details about fabric components.

Third Party Orchestrated Fabric Deployment

You can use your orchestration system, instead of GigaVUE-FM to deploy GigaVUE fabric components. The Third Party orchestration feature allows you to register these fabric components, after deployment, with GigaVUE-FM using the information you provide. After registering, you can set up monitoring sessions and other services in GigaVUE-FM.

Components in Third Party Orchestration

The GigaVUE Cloud Suite for Third Party Orchestration consists of the following components:

- GigaVUE-FM
- UCT-Vs
- UCT-V Controllers
- GigaVUE V Series Proxy
- GigaVUE V Series Nodes

GigaVUE-FM is a key component of the GigaVUE Cloud Suite Cloud solution. GigaVUE-FM is a web-based fabric management interface that provides a single-pane-of-glass to manage and view both physical and virtual network traffic.

What You Need to Do

If you prefer to use the Third Party orchestration deployment option, you need to:

- Install and start GigaVUE-FM on a supported cloud or enterprise platform.
- Launch the fabric components in your environment.
- Register the fabric components with GigaVUE-FM.
- Delete the fabric components from the platform when you no longer need them.

Getting the Images of Components

- Download the component images from the [Gigamon Customer Portal](#)
- Get public cloud images from the respective marketplace.

NOTE: Contact Gigamon Technical Support team if your platform doesn't support the available Gigamon images.

For information about installing GigaVUE-FM in your enterprise data center, refer to the *GigaVUE-FM Installation and Upgrade Guide*.

Components for Third Party Orchestration

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 (earlier known as	UCT-V is a standalone service installed in your Virtual Machine (VM). This service

Component	Description
G-vTAP Agent)	mirrors the selected traffic from the VMs to the GigaVUE® V Series node. The UCT-V is offered as a Debian (.deb), Red Hat Package Manager (.rpm), or Windows package. For details, refer to Install UCT-Vs .
Next generation UCT-V (earlier known as Next Generation G-vTAP Agent)	Next generation UCT-V is a lightweight solution that: - Acquires traffic from VMs and improves UCT-V mirroring capability with a lightweight solution. - Reduces the traffic flow from the UCT-V to the V Series node with tap-level prefiltering, lowering the V Series load. - Activates UCT-V on Linux systems with a Kernel version above 5.4. - Allows you to prefilter traffic at UCT-Vs before sending it to V Series nodes. - Allows you to create a prefiltering policy template to apply on a monitoring session.
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 Proxy (optional)	- Manages multiple V Series nodes and orchestrates the traffic flow from GigaVUE V Series nodes to the monitoring tools. - GigaVUE-FM uses one or more GigaVUE V Series Proxies to communicate with the GigaVUE V Series nodes.
GigaVUE® V Series Node	GigaVUE® V Series Node is a visibility node that: - Aggregates mirrored traffic from multiple UCT-Vs. - Applies filters and manipulates the packets using GigaSMART applications. - 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.

travenKey Concepts

This section provides the key concepts required to understand GCP.

For details, refer to the following section:

- [Configure Monitoring Session](#)
- [Traffic Acquisition](#)
- [Precryption Technology on Single Node](#)
- [Configure Prefiltering](#)
- [Secure Communication between GigaVUE Fabric Components](#)
- [Supported Platforms](#)
- [Traffic Acquisition](#)

Monitoring Domain

Your OCI cloud infrastructure may encompass numerous resources spread across multiple OCI accounts and Virtual Private Clouds (VPCs). When implementing GigaVUE Cloud Suite for OCI to monitor specific workloads, it's crucial to limit access to only the necessary parts of your cloud environment.

To clearly define the scope of deployment, GigaVUE Cloud Suite for OCI follows the concept of Monitoring Domain. This domain establishes a boundary within your Oracle cloud environment where you implement the visibility solution.

Monitoring Domain

You can create a Monitoring Domain in OCI using:

- OCI accounts
- Regions
- VPCs

GigaVUE Cloud Suite for OCI interacts only with resources within the specified accounts and VPCs listed in a Monitoring Domain.

Connection


To control access better, GigaVUE Cloud Suite for OCI uses a Connection.

A connection,

- Includes a set of credentials and permissions that precisely limit what GigaVUE Cloud Suite for OCI can access.
- Helps you deploy GigaVUE Cloud Suite for OCI to the appropriate resources with minimal credentials and permissions, enhancing security and control.

Monitoring Domains and Connections work together to provide secure, targeted monitoring. You define what GigaVUE Cloud Suite for OCI can see and do.

The following diagram provides more details about the concepts of Monitoring Domain and Connection:



© 2017-2025 Gigamon Inc.. All rights reserved.

In the above diagram:


- You create Monitoring Domain A for VPC 1 in Account A
- You assign a Connection that grants access to only VPC 1
- GigaVUE-FM uses this connection to find and monitor resources like:
 - Workload VMs
 - Subnets
 - Security groups
 - Key pairs

It then deploys GigaVUE components to collect and process traffic. The data is forwarded to your monitoring tools. GigaVUE-FM operates only within VPC 1, following the rules of the Monitoring Domain and its associated connection.

The Monitoring Domain defines a clear boundary within which GigaVUE-FM operates. This ensures both security and performance goals are achieved. In this case, you do not want GigaVUE-FM to access any VPCs other than VPC 1, GigaVUE-FM complies with your intent. It restricts its operations strictly to the boundary defined by the Monitoring Domain. Additionally, the credentials and permissions you provide through the connection associated with the Monitoring Domain further restrict the actions of GigaVUE-FM.

The Monitoring Domain is a logical concept. Its definition is based on concepts defined by the underlying cloud platform. On OCI, you can define any boundary using accounts and VPCs.

The following diagram illustrates a Monitoring Domain encompassing resources in VPCs in two different OCI accounts:



In this diagram:

- Monitoring Domain B includes VPC 2 in Account A and VPC 3 in Account B
- You assign Connection 1 to access VPC 2 and Connection 2 to access VPC 3

GigaVUE-FM accesses only VPC 2 and VPC 3. It does not see or touch VPC 1 or VPC 4. The domain and connections define strict access limits.

For more information on creating a Monitoring Domain, see [Create a Monitoring Domain](#).

Monitoring Session

GigaVUE Cloud Suite for OCI offers a sophisticated approach to network visibility through customizable Monitoring Sessions. These sessions allow you to tailor your visibility strategy with precision, focusing on specific aspects of your OCI environment:

Workload Selection

You can select target workloads to define the scope of your monitoring:

- Select VMs or pods based on various criteria
- Utilize filters such as VPC name, subnet, IP address, VM tags, or pod labels
- Align selections with your predefined Monitoring Domain boundaries

Traffic Acquisition

You can specify the types of traffic to capture from selected sources:

- Options vary depending on the chosen acquisition technology
- Customize traffic selection to match your monitoring requirements

Traffic Processing

You can leverage a comprehensive suite of traffic processing features:

- Apply various GigaSMART Applications to refine and optimize captured traffic.
- Enhance security and efficiency of your monitoring solutions
- Tailor the data sent to analysis tools for maximum relevance


Configuring these elements within a Monitoring Session allows you to create a finely-tuned visibility policy that aligns perfectly with your specific monitoring objectives in OCI environments.

After you define a Monitoring Session, GigaVUE-FM

- Automatically configures the required traffic acquisition points
- Processes the traffic.
- Sends the captured traffic to GigaVUE V Series Nodes.
- Forwards the process data to to your analysis tools.

The following diagram provides a high-level overview of an example Monitoring Session:

Example: Visibility Policy Defined as Monitoring Session A



The above diagram display an example to monitor VMs tagged with project = "Saturn". For this, you

- Define a Monitoring Session with this tag filter
- Choose UCT-V to acquire traffic
- Apply packet slicing in the V Series Nodes
- Forward traffic to your analysis tool at IP 1.2.3.4

When deployed, GigaVUE-FM:

- Discovers all matching VMs in your Monitoring Domain
- Configures UCT-V on these VMs
- Starts forwarding sliced packets to your tool

If new Saturn-tagged VMs are launched, GigaVUE-FM detects them automatically. It applies the same monitoring policy without manual action.

Benefits

- Focus monitoring on only the workloads you need
- Control how traffic is captured and processed


- Automatically update monitoring as your environment changes

For more information, see [Configure Monitoring Session](#).

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

- 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.

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](#)

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.

- [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

- 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](#).

Note

- See the [Configure Precryption in UCT-V](#) section for details on how to enable Precryption™ in VM environments.
- See how [Secure Tunnels](#) feature can enable secure delivery of precrypted data.

Secure Tunneling to GigaVUE V Series Nodes

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.

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



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)

Refer to [Configure Secure Tunnel](#) for instructions on how to configure Secure tunnels.

Traffic Acquisition

You can acquire traffic from multiple virtual machines and container pod instances using UCT-V

The acquired traffic is forwarded to the GigaVUE V Series Node to conduct core intelligence and additional GigaSMART processing.

Traffic-Flow-Across-Components

GigaVUE Cloud Suite for GCP mirrors traffic from workload VMs, sends the mirrored packets to a GigaVUE V Series Node for processing, and then forwards the optimized traffic to the monitoring tools for analysis. Production application traffic remains on the normal data path, while visibility traffic follows a separate path through the Gigamon visibility fabric.

Traffic Flow Accross Components

1. **Workload VMs generate production traffic**

Applications running on the monitored GCP VM instances generate production traffic on their network interfaces. This application traffic continues to flow on the normal path without interruption.

2. **UCT-V mirrors the selected traffic from the workload VMs**

UCT-V is installed on the VM instances that must be monitored. UCT-V mirrors the selected traffic from the source interface and can be configured to capture ingress, egress, or both directions of traffic.

3. **UCT-V encapsulates the mirrored traffic and sends it to the GigaVUE V Series Node**

UCT-V encapsulates the mirrored packets by using GRE or VXLAN tunneling and forwards the tunneled traffic to the GigaVUE V Series Node for further processing.

4. **UCT-V Controller manages the UCT-V agents**

UCT-V Controller manages multiple UCT-V agents, and GigaVUE-FM uses one or more UCT-V Controllers to communicate with the UCT-Vs and orchestrate the mirrored traffic flow to the GigaVUE V Series Nodes.

5. **GigaVUE-FM sends the monitoring configuration to the deployed components**

In Third Party Orchestration, you create a monitoring domain and select the traffic acquisition method. If you choose UCT-V, GigaVUE-FM uses the monitoring domain configuration to manage the UCT-V Controller and the traffic flow to the GigaVUE V Series Nodes.

6. **The GigaVUE V Series Node decapsulates and processes the traffic**

The GigaVUE V Series Node receives the tunneled mirrored traffic and applies the configured monitoring logic, such as filtering, mapping, and transformation, before preparing the traffic for tool consumption.

7. The GigaVUE V Series Node forwards the optimized traffic to the monitoring tools

After processing, the GigaVUE V Series Node forwards the resulting traffic or metadata stream to the connected monitoring and analysis tools.

NOTE: If you use GCP packet mirroring instead of UCT-V, the traffic acquisition method is Customer Orchestrated Source. In that case, the mirrored traffic is directed to the GigaVUE V Series Nodes through the GCP internal load balancer, and you do not need to configure UCT-V or UCT-V Controller.

Deployment Prerequisites

Before you deploy GigaVUE Cloud Suite for GCP, make sure that the required GCP resources, machine types, firewall rules, and registration details are in place. In GCP, the deployment uses Third Party Orchestration, and the prerequisite set changes slightly depending on whether you use UCT-V or GCP packet mirroring.

Complete the following prerequisites for all GCP deployments:

- Use Third Party Orchestration to deploy and register the GigaVUE fabric components in GCP.
- Select the required machine types:
- GigaVUE V Series Node: c2-standard-4 for 2 interfaces and c2-standard-8 for 3 interfaces.
- GigaVUE V Series Proxy: e2-micro.
- UCT-V Controller: e2-micro.
- Set the tunnel MTU correctly. With the default GCP MTU of 1460, use 1410 for IPv4 tunnels and 1390 for IPv6 tunnels.
- Configure the required GCP firewall rules before you register the fabric components.
- Create the Monitoring Domain first. Use the Monitoring Domain name as groupName and the Connection name as subGroupName in the registration data.
- Create tokens in GigaVUE-FM for component registration. If SSH access is required, add the SSH key during instance launch.

Additional Prerequisites for UCT-V Deployments

If the traffic acquisition method is UCT-V, deploy and register the following components:

- UCT-V Controller.
- UCT-V.

- GigaVUE V Series Node.
- GigaVUE V Series Proxy, if a proxy path is required.

Additional Prerequisites for GCP Packet Mirroring

If you use GCP packet mirroring, complete these additional prerequisites:

- Select Customer Orchestrated Source as the Traffic Acquisition Method.
- Enable Uniform Traffic Policy in the Monitoring Domain.
- Use a minimum of three NICs on the GigaVUE V Series Node and select the c2-standard-8 machine type.
- Create a GCP instance template.
- Create a managed instance group with autoscaling.
- Create a TCP or UDP internal load balancer with packet forwarding enabled.

NOTE: The internal load balancer forwards raw traffic. For packet mirroring deployments, the Monitoring Session must start with a Raw End Point.

With these prerequisites in place, you can deploy and configure the fabric components in your GCP environment.

Supported Compute Requirements for GCP

In GCP, flavors set the vCPU, memory, and storage requirements for an image. Gigamon recommends that you create a flavor on your choice that matches or exceeds the supported requirements listed in the following table.

The following table lists the minimum requirements for deploying the fabric components:

Fabric Component	Machine type
GigaVUE V Series Node	• c2-standard-4 for 2 interfaces • c2-standard-8 for 3 interfaces
GigaVUE V Series Proxy	e2-micro
UCT-V Controller	e2-micro

NOTE: A single UCT-V Controller can manage up to 500 UCT-Vs. For more than 500 UCT-Vs, you must add an additional UCT-V Controller to scale up accordingly.

The instance size of the GigaVUE V Series Node is configured and packaged as part of the qcow2 image file.

Network Requirements

The following table lists the recommended requirements to setup the network topology.

Network	Purpose
Management	Identify the subnets that GigaVUE-FM uses to communicate with the GigaVUE V Series nodes and controllers.
Data	Identify the subnets that receives the mirrored tunnel traffic from the monitored instances. In data network, if a tool subnet is selected then the V Series node egress traffic on to the destinations or tools.

NOTE: If you are using IPv6 in the tenant network, then it is recommended to use SLAAC or stateless DHCPv6 for dynamic address assignment.

Network Firewall Requirements

Following are the Network Firewall Requirements for Gigamon fabrics for Nutanix deployments.

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	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.

GigaVUE-FM

Direction	Type	Protocol	Port	CIDR	Purpose
Inbound	HTTPS	TCP	443	Anywhere Any IP	Allows GigaVUE® V Series Nodes, GigaVUE V Series Proxy, and GigaVUE-FM administrators to communicate with GigaVUE-FM
Inbound	SSH	TCP	22	Anywhere Any IP	Allows GigaVUE® V Series Nodes, GigaVUE V Series Proxy, and GigaVUE-FM administrators to communicate with GigaVUE-FM
Outbound (optional)	Custom TCP Rule	TCP	8890	GigaVUE V Series Proxy IP	Allows GigaVUE-FM to communicate with GigaVUE V Series Proxy
Outbound	Custom TCP Rule	TCP	8889	GigaVUE V Series Node IP	Allows GigaVUE-FM to communicate with GigaVUE V Series Node

GigaVUE V Series Node

Direction	Type	Protocol	Port	CIDR	Purpose
Inbound	Custom TCP Rule	TCP	9903	GigaVUE V Series Proxy IP	Allows GigaVUE V Series Proxy to communicate with GigaVUE® V Series Nodes
Inbound	UDP	UDPGRE	4754	Ingress Tunnel	Allows to UDPGRE tunnel to communicate and

Direction	Type	Protocol	Port	CIDR	Purpose
					tunnel traffic to GigaVUE V Series Nodes
Outbound	Custom TCP Rule	TCP	5671	GigaVUE-FM IP	Allows GigaVUE® V Series Node to communicate and tunnel traffic to the Tool
Outbound	Custom UDP Rule	- UDP (VXLAN) - IP Protocol (L2GRE)	- VXLAN (default 4789) - L2GRE (IP 47)	Tool IP	Allows GigaVUE® V Series Node to communicate and tunnel traffic to the Tool
Outbound (optional)	Custom ICMP Rule	ICMP	- echo request - echo reply	Tool IP	Allows GigaVUE® V Series Node to health check the tunnel destination traffic.

GigaVUE V Series Proxy (optional)

Direction	Type	Protocol	Port	CIDR	Purpose
Inbound	Custom TCP Rule	TCP	8890	GigaVUE-FM IP	Allows GigaVUE-FM to communicate with GigaVUE V Series Proxy
Outbound	Custom TCP Rule	TCP	8889	GigaVUE V Series Node IP	Allows GigaVUE-FM to communicate with 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	Allows UCT-V to

				IP	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	Allows GigaVUE V Series

				Proxy	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	Allows CLI access to

			Subnet	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	TCP	8891	GigaVUE V	Allows GigaVUE V Series Node to

(This is the port used for Third Party Orchestration)			Series Proxy IP	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.

Installing GigaVUE-FM on GCP

This section describes how to install GigaVUE-FM on GCP. It consists of the following sections:

- [Before You Install](#) describes the pre-requisites and the preparatory steps required for installing GigaVUE-FM.
- [Install GigaVUE-FM on GCP](#) describes the steps to install and deploy GigaVUE-FM on GCP.

Default Login Credentials

Using the default credentials, you can log in to the GigaVUE V Series Node and GigaVUE V Series Proxy.

Product	Login credentials
GigaVUE V Series Node	Using ssh, you can log in to the GigaVUE V Series Node. The default username and password are: • Username: gigamon • Password: Gigamon123!
GigaVUE V Series proxy	Using ssh, you can log in to the GigaVUE V Series proxy. The default username and password are: • Username: gigamon • Password: Gigamon123!

Install UCT-V

UCT-V is the primary Gigamon monitoring module that you install on your Virtual Machines (VMs). UCT-V mirrors the selected traffic from a source interface to a destination mirror interface. UCT-V encapsulates the mirrored traffic using GRE or VXLAN tunneling and then forwards to the GigaVUE V Series Node.

NOTE: Install UCT-V only when the UCT-V is your traffic acquisition method.

The Workflow

- A UCT-V can consist of multiple source interface and a single destination interface.
- UCT-V collects the network packets from the source interface and sends to the destination interface.
- From the destination interface, the packets traverse through the L2GRE, VXLAN tunnel interface, or Secure Tunnels to the GigaVUE V Series Node.

You can configure a source interface with one or more Network Interfaces. While configuring a source interface, specify the traffic direction to monitor: ingress, egress, or both.

NOTE: For environments with both Windows and Linux or just windows UCT-V, VXLAN tunnels in the UCT-V Controller specification is required.

Supported Operating System for UCT-V¹ is 6.5.00, 6.6.00, 6.7.00, 6.8.00, 6.9.00, 6.10.00, 6.11.00, 6.12.00, 6.13.00

The table below lists the validated and supported versions of the Operating Systems for UCT-V.

Operating System	Supported Versions
Ubuntu/Debian	Versions 16.04 through 22.04
CentOS	Versions 7.5 through 9.0
RHEL	Versions 7.5 through 9.4
Windows Server	Versions 2012 through 2022 Note: Ensure the send buffer size of the network adapters is set to 128 MB for optimal performance and to minimize traffic disruption.
Rocky OS	Versions 8.4 through 8.8

GigaVUE-FM version 6.14 supports UCT-V version 6.14 as well as (n-2) versions. We recommend to use the latest version of UCT-V with GigaVUE-FM for better compatibility.

¹From Software version 6.4.00, G-vTAP is renamed to UCT-V.

Linux UCT-V Installation

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

Refer to the following sections:

- [Single Network Interface Configuration](#)
- [Multiple Network Interface Configuration](#)
- [Loopback Network Interface Configuration](#)
- [Linux Network Firewall Requirements](#)
- [Install Linux UCT-Vs](#)
- [Register Linux UCT-V](#)

Single Network Interface Configuration

A single network interface card (NIC) serves as both the source and destination interface. UCT-V, with a single network interface configuration, enables you to monitor both ingress and egress traffic from the same NIC. The system uses the same interface to send monitored traffic.

Example

Consider a single interface eth0 in the monitoring instance. In the UCT-V configuration, you can configure eth0 as both source and destination, and also specify monitoring for both ingress and egress traffic. The monitored traffic from eth0 is mirrored and sent using the same eth0 interface.

NOTE: Using a single NIC as the source and destination can lead to increased latency when sending traffic.

Sample Configuration

Example of the UCT-V configuration file for a single NIC configuration:

Grant permission to monitor ingress and egress traffic at iface

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

Multiple Network Interface Configuration

In a multiple NIC configuration, UCT-V enables you to configure two NICs, one for the source and another for the destination.

Example

Consider two NICs, eth0 and eth1, in the monitoring instance.

In the UCT-V configuration, you can configure:

- eth0 as the source interface, and specify to monitor egress traffic.
- eth1 as the destination interface.

Then, the mirrored traffic from eth0 is sent to eth1. From eth1, the traffic is sent to the GigaVUE V Series Node.

Sample: Example of the UCT-V configuration file for a dual NIC configuration:

Grant permission to monitor ingress and egress traffic at iface

```
# 'eth0' to monitor and 'eth1' to transmit the mirrored packets
# eth0 mirror-src-ingress mirror-src-egress
# eth1 mirror-dst
```

Loopback Network Interface Configuration

UCT-V supports the ability to tap and mirror the loopback interface. You can tap the loopback interfaces on the workload that carries application-level traffic inside the Virtual Machine. The loopback interface is always configured as bidirectional traffic, regardless of the configurations provided in the configuration file.

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

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

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. For details, refer to [Network Firewall Requirement for GigaVUE Cloud Suite](#).

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 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.

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 token and place it in the **/etc/gigamon-cloud.conf** configuration file. For more information, refer to .

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:

- a. Download the UCT-V **6.14.00** Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance, contact [Contact 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-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 [Contact 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	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 Single Network Interface 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 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.

Options	Use Command	Description
		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. For details, refer to Single Network Interface 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:

1. Download the UCT-V6.14.00 Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance contact [Contact 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-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-V6.14.00 RPM (.rpm) package from the [Gigamon Customer Portal](#).
For assistance contact [Contact 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
```
- After installing UCT-V, refer to [Deploy Fabric Components using Generic Mode](#) for platform specific information to configure UCT-V using Third Party Orchestration.

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/
```

Register Linux UCT-V

It is mandatory to create a cloud configuration file and add the token to authenticate the UCT-V package with GigaVUE-FM. The token is required only for initial registration before generating the certificate. You can use the token only once and do not need to maintain. Refer below to register UCT-V in your virtual machine.

Third Party Orchestration: The third-party orchestration feature allows you to deploy UCT-V using your own orchestration system. UCT-V uses the information of the user to register with GigaVUE-FM. You can register UCT-V to GigaVUE-FM using Third Party Orchestration with the following two modes:

- Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
- Integrated Mode - Deploy GigaVUE Fabric Components using Integrated Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration

NOTE: If you have already configured gigamon-cloud.conf file in the /tmp directory, you can directly use the **uctv-wizard configure** command (sudo uctv-wizard configure). This action automatically fetches the configuration file and completes the registration process.

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 Windows UCT-Vs](#)
- [Register Windows UCT-V](#)

Windows Network Firewall Requirements

If your environment uses network firewall rules or security groups, you must open specific ports for the virtual machine. For details, refer to [Network Firewall Requirement for GigaVUE Cloud Suite](#).



Notes:

- After installing UCT-V, ensure the following TCP ports are configured:
 - Port 8301 (Inbound)
 - Port 8300 (Outbound)
- You can configure the ports using the following PowerShell commands. Make sure to run PowerShell as **Administrator**:
 1. `New-NetFirewallRule -DisplayName "GigaVUE UCT-V (http01_challenge_port)" -Group "Virtual Tap" -Direction "Inbound" -Program "C:\Program Files (x86)\Uctv\step.exe" -LocalPort "8301" -Protocol "TCP" -Action`
 2. `New-NetFirewallRule -DisplayName "GigaVUE UCT-V (pki_ra_port)" -Group "Virtual Tap" -Direction "Outbound" -Program "C:\Program Files (x86)\Uctv\uctvd.exe" -LocalPort "8300" -Protocol "TCP" -Action Allow`

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 [Configure Tokens](#).

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

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

- [Install Windows 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 [Contact Technical Support](#).
2. Install the downloaded MSI package as **Administrator**.
The UCT-V service starts automatically.

3. 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 Windows Network Firewall Requirements
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. • Enter N if you wish to restore it manually. For details, refer to Install Windows UCT-Vs using Manual

Options	Use Command	Description
		Configuration.
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. For details, refer to Install Windows UCT-Vs using Manual Configuration.
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

1. **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.
2. **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 [Contact Technical Support](#).
2. Install the downloaded MSI package as **Administrator**. The UCT-V service starts automatically.

3. Modify the file **C:\ProgramData\Uctv\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.

NOTE: After installing UCT-V, refer to [Deploy Fabric Components using Generic Mode](#) for platform specific information to configure UCT-V using Third Party Orchestration.

Register Windows UCT-V

It is mandatory to create a cloud configuration file and add the token to authenticate the UCT-V package with GigaVUE-FM. The token is required only for initial registration before generating the certificate. You can use the token only once and do not need to maintain. Refer below to register UCT-V in your virtual machine.

Third Party Orchestration: The third-party orchestration feature allows you to deploy UCT-V using your own orchestration system. UCT-V uses the information of user to register with GigaVUE-FM.

UCT-V can register with GigaVUE-FM using Third Party Orchestration in one of the following two modes:

- Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
- Integrated Mode - Deploy GigaVUE Fabric Components using Integrated Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration

NOTE: If you have already configured gigamon-cloud.conf file in the directory (C:\Users\<username>\AppData\Local), you can directly use the **uctv-wizard configure** command (sudo uctv-wizard configure). This action automatically fetches the configuration file and complete the registration process.

Uninstall UCT-V

This section describes how to uninstall Linux UCT-V and Windows UCT-V.

Uninstallation Method

- Linux:
 - Uninstall the UCT-V in Ubuntu/Debian, RPM, Red Hat Enterprise Linux, and CentOS packages
 - Use the following command: **sudo uctv-wizard uninstall**
- Windows
 - Uninstall the UCT-V in the MSI package.

- Use the following command: **CMD `uctv-wizard uninstall`**

NOTE: Uninstall command automatically stops the UCT-V service, removes the firewall rules, and uninstalls the UCT-V.

When an UCT-V is uninstalled, it moves to the Unknown status. If it remains in this state for more than 24 hours, it is considered a stale entry and is automatically removed from GigaVUE-FM every day at 12:30 AM (system time), unless it is part of an active or scheduled upgrade.

Upgrade or Reinstall UCT-V

You can upgrade UCT-V in your virtual machine using the following options:

- [Upgrade UCT-V through GigaVUE-FM \(Recommended Method\)](#)
- [Upgrade UCT-V Manually](#)

Refer to the below sections for detailed information and instructions.

Upgrade UCT-V through GigaVUE-FM (Recommended Method)

Upgrading UCT-V manually involves a series of steps to uninstall, install, and restart the service again. This upgrade method is applicable to both GigaVUE-FM Orchestration and Third Party orchestration. For a list of supported platforms, refer to .

NOTE: This method is complicated if you need to upgrade UCT-Vs for a large number of VMs. However, you can upgrade UCT-V in the workload VM without any hands-on involvement through GigaVUE-FM.

Refer to the sections below for details and instructions:

1. [Upload the UCT-V Images](#)
2. [Upgrade the UCT-V](#)

Rules and Notes:

- Upgrade is allowed only to versions 6.9.00 or later. Ensure that the UCT-V Controller version is compatible with the version to which you are upgrading.
- Do not trigger system upgrades or other upgrades at the same time as the UCT-V upgrade.
- You should have Infrastructure Management permission to upgrade the UCT-Vs.
- Currently, you can upgrade the UCT-Vs to n+2 versions and any number of patch releases through GigaVUE-FM.
- Before you proceed with the upgrade, ensure that the UCT-Vs are in a healthy state.

- Make sure that a UCT-V is performing only one active job at a time. If the selected UCT-V is part of another job, you cannot trigger the immediate job using the same UCT-V.
- You must upload a compatible image type to upgrade the UCT-V; otherwise, the UCT-V is rejected for the upgrade job.
- Upgrade through GigaVUE-FM is not applicable to OVS Modules. For OVS tapping, you should upgrade the UCT-Vs manually.

Upload the UCT-V Images

Perform the following steps to upload UCT-V image files in GigaVUE-FM:

1. Go to **Inventory > Virtual** and select your cloud platform.
The **Monitoring Domain** page appears.
2. Select the **UCT-V Upgrade** drop-down menu and select **Images**.
3. In the **Images** page, click **Upload**. The **Upload Internal Image Files** wizard appears.
4. Select **Choose File**, upload the UCT-V files from your local, and select **Ok**.



Notes:

- You can download the UCT-V image files from Gigamon software portal.
- You can upload a maximum of 15 UCT-V files at a time.
- The supported file formats are **.deb**, **.rpm**, and **.msi**.
- Ensure that you do not change the file names. GigaVUE-FM does not accept the image files with modified names.
- When the upload is in process, GigaVUE-FM does not allow uploading a file with similar type and version.

5. Verify if the uploaded UCT-V images is listed in the **Images** page.

Delete the file

You can delete one or multiple images.

1. In the **Images** page, select **Filter** to find the images based on Image Name, Version, and Image Type.
2. Select the required images.
3. From the Actions drop-down menu, select **Delete** or **Delete All**.

You can only delete those image files that are not associated with any tasks created for the upgrade process.

Upgrade the UCT-V

Follow the steps below to upgrade UCT-V in GigaVUE-FM:

1. In the **UCT-V Upgrade** drop-down menu, select **Dashboard** to view the UCT-V upgrade landing page.
In the Dashboard page, you can view the upgrade status of individual UCT-Vs and the stages of the upgrade process (Fetch, Install, Verify). The page also displays the overall progress of the upgrade.
2. Select the required UCT-Vs, and from the **Actions** drop-down menu, select **Upgrade**.
The **UCT-V Upgrade task** page appears.
3. Enter the task name.
4. In the **Image Version** drop-down menu, select the required version you want to upgrade to from the list of available image versions. You can choose to upgrade immediately or schedule a time for the upgrade to happen.
5. Select the required option in the **Time Selection** field. If you prefer to schedule the upgrade, enter the choice of your date and time in the respective fields.
Do not schedule the upgrade for a time in the past.
6. Select **Create**.
The image upgrade task is now created.



Note:

- You cannot edit the upgrade task once it is created.
- You can only reschedule the scheduled task but cannot edit the UCT-V selected for the particular task.
- In the event of the errors listed below, GigaVUE-FM displays a pop-up message with the list of UCT-Vs that are not compatible for upgrade. Select **Proceed** to ignore the unsupported UCT-Vs and upgrade the compatible ones, or select "**Edit**" to modify your changes. The errors include:
 - Controller version is not compatible with the upgrade version.
 - Inconsistency between the uploaded image file type and the selected UCT-V.

You can view the created task details (both immediate and scheduled) in the **UCT-V Upgrade > Jobs** section.



Notes:

- For better progress monitoring, it is recommended to split the upgrade task to a limited number, such as 50 or 100 UCT-Vs.
- When you create a new upgrade task for the same UCT-V, the status of any existing UCT-V changes to 'In Progress' until the latest task is completed. Once the upgrade for the existing tasks is successfully finished, you can create another task for that same UCT-V.

You can view the different stages of the upgrade process in UCT-V Upgrade Dashboard

page. Each stage is marked with  if it is successful and  in case of failure. If the upgrade is successful, GigaVUE-FM updates the upgrade status as **Success** for the selected UCT-V.



Notes:

- The default wait time for the **Upgrade Status** to get updated is 15 minutes.
- The default wait time for the **Image Version** to get updated is 5 minutes.
- In case of failure, you can upgrade the failed instance manually.

Upgrade UCT-V Manually

To upgrade UCT-V manually on a virtual machine, delete the existing UCT-V and install the new version of UCT-V.

NOTE: Before deleting the UCT-V, take a backup copy of the **/etc/uctv/uctv.conf** configuration file. This step avoids reconfiguring the source and destination interfaces.

1. Uninstall the existing UCT-V. Refer to the *Uninstall UCT-V* section in the respective GigaVUE Cloud Suite Deployment Guide.
2. Install the latest version of the new UCT-V. Refer to the Linux UCT-V Installation and the Windows UCT-V Installation topics in the respective GigaVUE Cloud Suite Deployment Guides.
3. Restart the UCT-V service.
 - Linux platform:

```
$ sudo service uctv restart
```
 - Windows platform: Restart from the Task Manager.

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.

Adding Certificate Authority

This section describes how to add CA Certificate chain Authority in GigaVUE-FM.

The CA Certificate chain List page allows you to add the root CA for the devices.

To upload the CA Certificate chain using GigaVUE-FM, follow these steps:

1. Go to **Inventory > Resources > Security > CA List**.
2. Select **Add**, to add a new Custom Authority.
The **Add Certificate Authority** page appears.
3. In the **Alias** field, enter the alias name of the CA Certificate chain Authority
4. Use one of the following options to enter the CA Certificate chain Authority:
 - **Copy and Paste:** In the **Certificate** field, enter the certificate.
 - **Install from URL:** In the **Path** field, enter the URL in the format:
<protocol>://<username>@<hostname/IP address>/<file path>/<file name>. In the **Password** field, enter the password.
 - **Install from Local Directory:** Select **Choose File** to browse and select a certificate from the local directory.
5. Select **Save**.

Edit SSL Configuration

You can add certificate authority and SSL keys to your fabric components after deploying it.

To add certificate authority and SSL keys when using secure tunnels,

1. Go to **Inventory > VIRTUAL > Third Party Orchestration**.
2. Select **Monitoring Domain**.
The **Monitoring Domain** page appears.
3. From the page, select the desired monitoring domain to add the SSL key.
4. From the **Actions** drop down list, select **Edit SSL Configuration**.
An **Edit SSL Configuration** window appears.
5. From the **UCT-V Agent Tunnel CA** drop down list, select **CA**.
6. From the **V Series Node SSL key** drop down list, select the SSL key.
7. Select **Save**.

Deploy

Deploy Fabric Components using Generic Mode

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.

Acquire Traffic using UCT-V - By Creating Monitoring Domain

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 GCP](#)
- [Configure Role-Based Access for Third Party Orchestration](#)
- [Install UCT-Vs](#)
- [Create Monitoring Domain](#)
- [Configure and Deploy GigaVUE Fabric Components using GCP - Generic Mode](#)
- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)

Install GigaVUE-FM on GCP

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

This section describes how to install GigaVUE-FM on GCP. It consists of the following sections:

- [Before You Install](#) describes the pre-requisites and the preparatory steps required for installing GigaVUE-FM.
- [Install GigaVUE-FM on GCP](#) describes the steps to install and deploy GigaVUE-FM on GCP.

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-FMUser 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.

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](#)

NOTE: If you want to avoid downloading and installing the UCT-Vs every time there is a new instance to be monitored, you can save the UCT-V running on an instance as a private AMI.

To save the UCT-V as an AMI from your EC2 console, right click on the instance and navigate to **Image and Templates > Create Image**.

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 Linux UCT-Vs](#)
- [Register Linux UCT-V](#)

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:

- a. Download the UCT-V6.14.00 Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance, contact [Contact 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-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 [Contact 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	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 [Contact 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-V6.14.00 RPM (.rpm) package from the [Gigamon Customer Portal](#).
For assistance contact [Contact 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/
```

Register Linux UCT-V

It is mandatory to create a cloud configuration file and add the token to authenticate the UCT-V package with GigaVUE-FM. The token is required only for initial registration before generating the certificate. You can use the token only once and do not need to maintain.

You can register UCT-V in your virtual machine using:

- **Third Party Orchestration:** The third-party orchestration feature allows you to deploy UCT-V using your own orchestration system. UCT-V uses the information of the user to register with GigaVUE-FM. You can register UCT-V to GigaVUE-FM using Third Party Orchestration with the following two modes:
 - a. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
 - b. Integrated Mode - Deploy GigaVUE Fabric Components using Integrated Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration

NOTE: If you have already configured gigamon-cloud.conf file in the /tmp directory, you can directly use the **uctv-wizard configure** command (sudo uctv-wizard configure). This action automatically fetches the configuration file and completes the registration process.

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 Linux UCT-Vs](#)
- [Register Linux UCT-V](#)

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 Linux 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 [Configure Tokens](#).
- 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 Linux 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 [Contact 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 [Contact 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.

Register Windows UCT-V

It is mandatory to create a cloud configuration file and add the token to authenticate the UCT-V package with GigaVUE-FM. The token is required only for initial registration before generating the certificate. You can use the token only once and do not need to maintain.

You can register UCT-V in your virtual machine using:

- **Third Party Orchestration:** The third-party orchestration feature allows you to deploy UCT-V using your own orchestration system. UCT-V uses the information of user to register with GigaVUE-FM. UCT-V can register with GigaVUE-FM using Third Party Orchestration in one of the following two modes:
 - a. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
 - b. Integrated Mode - Deploy GigaVUE Fabric Components using Integrated Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration

NOTE: If you have already configured `gigamon-cloud.conf` file in the directory (C:\Users\<username>\AppData\Local), you can directly use the **uctv-wizard configure** command (`sudo uctv-wizard configure`). This action automatically fetches the configuration file and complete the registration process.

Create Monitoring Domain

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.

New Monitoring Domain

Form elements with * are required.

Cancel Save

You can think of a Monitoring Domain as a geographic region in your overall deployment. The next step is to create a Monitoring Session to set filter rules for your application traffic to their destination.

Monitoring Domain Name*

Connection Name*

Traffic Acquisition Method*

Traffic Acquisition Tunnel MTU*

Enable IPv6 Preference ☐ Disabled

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: • 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 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 using GCP - Generic Mode

This section provides step-by-step information on how to register GigaVUE fabric components using Google Cloud Platform (GCP) or a configuration file.

Minimum Requirements

The following table lists the minimum requirements for deploying the fabric components:

Fabric Component	Machine type
GigaVUE V Series Node	- c2-standard-4 for 2 interfaces - c2-standard-8 for 3 interfaces
GigaVUE V Series Proxy	e2-micro
UCT-V Controller	e2-micro

Deployment Guideines

Keep in mind the following when deploying the fabric components using GCP:

- Set the correct MTU for tunnels: With a default MTU of 1460 in GCP, use 1410 for IPv4 tunnels and 1390 for IPv6. To update the Traffic Acquisition Tunnel MTU:
 - Select the monitoring domain.
 - Select the **Edit Monitoring Domain** option.
 - Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

For details, refer to [Traffic Acquisition Tunnel MTU](#)

- For successful registration of fabric components, configure firewall rules to open ports.
 - For details on how to configure firewall rules, refer to [Use VPC firewall rules](#) topic in GCP documentation.
 - For details on required ports, refer to [Network Firewall Requirement](#).
- When you deploy the fabric components using third party orchestration, you cannot delete the Monitoring Domain or change the MTU without unregistering the registered fabric components.
- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to [Create Monitoring Domain](#).
- When launching an instance, to access it using a private key, add the SSH key. The default password is gigamon.
- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to [Configure Tokens for Third Party Orchestration](#).
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section.

In your GCP, you can configure the following GigaVUE fabric components:

Configure UCT-V Controller in GCP

IMPORTANT: You must enable the basic authentication to launch the GigaVUE fabric components for version 6.9 and lower. For more instructions on the steps to enable the basic authentication, refer to [Authentication Type](#).

You can configure more than one UCT-V Controller in a monitoring domain.

To register UCT-V Controller in GCP, use one of the following methods:

Register UCT-V Controller during Instance Launch

In your GCP, you can launch the UCT-V Controller and to register UCT-V Controller using custom metadata.

Perform the following steps:

1. On the VM instances page of GCP, select **Create instances** . For details, refer to [Create and Start instance](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field enter the below mentioned text in the following format and deploy the instance. The UCT-V Controller uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the GigaVUE-FM>
      sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
      remotePort: 443
```

Register UCT-V Controller after Instance Launch

You can register UCT-V Controller after launching a Virtual Machine using a configuration file.

Perform the following steps:

1. Log in to the UCT-V Controller.
2. Download the cloud configuration file (for example, md_gigamon-cloud.conf) using the Download Cloud Configuration File option in the Monitoring Domain Actions tab, and use the details for registering UCT-V Controller

```
Registration:
  groupName: <Monitoring Domain Name>
  subGroupName: <Connection Name>
  token: <Token>
  remoteAddress: <IP address of the GigaVUE-FM>
  sourceIP: <IP address of UCT-V Controller> (Optional. Required only in multi-NIC
  deployments to specify which UCT-V Controller interface to use for communicating with
  GigaVUE-FM.)
  remotePort: 443
```

3. Restart the UCT-V Controller service.
`$ sudo service uctv-cntlr restart`

The deployed UCT-V Controller registers with the GigaVUE-FM. After successful registration, the UCT-V Controller 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, then GigaVUE-FM tries to reach the UCT-V Controller. If that fails as well then GigaVUE-FM unregisters the UCT-V Controller and it is removed from GigaVUE-FM.

The UCT-V Controller deployed in GCP appears on the Third Party Orchestration Monitoring Domain page of GigaVUE-FM.

☐ Monitoring Domain	Connection	Fabric	Management IP	Fabric Version	Status
☐ MD1					
	pubsub-vpc				✔ Connected
		G-vTapController	34.219.250.141	1.7-304	✔ Ok
		Gigamon-VSeriesProxy-1	34.211.211.49	2.1.0	✔ Ok
		Gigamon-VSeriesNode-1	172.26.24.188	2.2.0	✔ Ok

Assign Static IP address for UCT-V Controller

By default, the UCT-V Controller gets assigned an IP address using DHCP.

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
  ethernet:
    <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
  ethernet:
    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 UCT-V Controller service.

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

The deployed UCT-V Controller registers with the GigaVUE-FM. After successful registration the UCT-V Controller 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 UCT-V Controller. If that fails as well then GigaVUE-FM unregisters the UCT-V Controller and removes from GigaVUE-FM.

Configure UCT-V in GCP

Linux and Windows platforms support Deployment of GigaVUE fabric components through a third-party orchestrator.

NOTE: You can configure more than one UCT-V Controller for a UCT-V, so that if one UCT-V Controller goes down, the UCT-V registration takes place through another active Controller.

Install the Windows UCT-V

Perform the following steps:

1. Deploy Windows server in GCP. For details, refer to [Create a Windows Server VM instance in Compute Engine](#) topic in Google documentation.
2. Follow the instruction in the *Connect to the VM instance by using RDP* section of [Set up Chrome Remote Desktop for Windows on Compute Engine](#) topic in the GCP documentation.
3. Download UCT-V build in your desktop and copy it to the RDP session.
4. Turn off the Windows Firewall Defender.
5. Install the Windows UCT-V. For details, refer to [Windows UCT-V Installation](#) for step-by-step instructions on how to install Windows UCT-V.

To register UCT-V in GCP, use one of the following methods.

Register UCT-V during Instance Launch

You need to register UCT-V through the registered UCT-V Controller. It communicates through PORT 8891.

NOTE: Registering UCT-V during Virtual Machine Launch is not applicable to Windows UCT-V. You can register your Windows UCT-V after launching the Virtual machine, using a configuration file.

In your GCP, to launch the instance and register the UCT-V using Custom Metadata, perform the following steps:

1. On the VM instances page of GCP, select **Create instances**. For details, refer to [Create and Start instance](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field, enter the below mentioned text in the following format and deploy the instance. The UCT-V uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the UCT-V Controller 1>,<IP address of the UCT-V
Controller 2>
      sourceIP: <IP address of UCT-V> (Optional Field)
```

Register UCT-V after Instance Launch

To register UCT-V after launching a Virtual Machine using a configuration file, perform the following steps:

1. Install the UCT-V in the Linux or Windows platform.
2. Log in to the UCT-V.
3. Create a local configuration file and enter the following user data.



- **/etc/gigamon-cloud.conf** is the local configuration file in Linux platform.
- **C:\ProgramData\uctv\gigamon-cloud.conf** is the local configuration file in Windows platform.
- When creating **C:\ProgramData\uctv\gigamon-cloud.conf** file, ensure that the file name extension is **.conf**. To view the file name extension in Windows, perform the following steps:
 - a. Go to File Explorer and open the File Location.
 - b. On the top navigation bar, select **View**.
 - c. In the **View** tab, enable the **File name extensions** check box.

Registration:

```
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the UCT-V Controller 1>,<IP address of the UCT-V
Controller 2>
sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
```

4. Restart the UCT-V service.
 - Linux platform:
`$ sudo service uctv restart`
 - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in GCP

NOTE: It is not mandatory to register GigaVUE V Series Nodes via V Series proxy however, if there is a large number of nodes connected to GigaVUE-FM or if the user does not wish to reveal the IP addresses of the nodes, then you can register your nodes using GigaVUE V Series Proxy. In this case, GigaVUE-FM communicates with GigaVUE V Series Proxy to manage the GigaVUE V Series Nodes.

To register GigaVUE V Series Node and GigaVUE V Series Proxy in GCP, use one of the following methods.

IMPORTANT: You must enable the basic authentication to launch the GigaVUE fabric components for version 6.9 and lower. For more instructions on the steps to enable the basic authentication, refer to [Authentication Type](#).

Register GigaVUE V Series Node and GigaVUE V Series Proxy during Instance Launch

1. On the VM instances page of GCP, select **Create an instance** . For details, refer to [Create VMs with multiple network interfaces](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field, enter the below mentioned text in the following format and deploy the instance. The UCT-V uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
      remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

Register GigaVUE V Series Node and GigaVUE V SeriesProxy after Instance Launch

You can register GigaVUE V Series Node and GigaVUE V Series Proxy after launching the virtual machine using a configuration file.

To register,

1. Log in to the GigaVUE V Series Node or Proxy.
2. Create a local configuration file (**/etc/gigamon-cloud.conf**) and enter the following user data.

```
Registration:
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the proxy>
remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

3. Restart the GigaVUE V Series node or proxy service.

- V Series node:
\$ sudo service vseries-node restart
- V Series proxy:
\$ sudo service vps restart

The deployed GigaVUE V Series proxy registers with the GigaVUE-FM. After successful registration the GigaVUE V Series proxy 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 proxy. If that fails as well, then GigaVUE-FM unregisters the GigaVUE V Series proxy and it is removed from GigaVUE-FM.

NOTE: When the GigaVUE V Series Node is stopped or terminated from the GCP, it does not send any unregistration request and GigaVUE-FM unregisters the V Series Node soon after.

What to do Next

You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab.

Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain

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

Install GigaVUE-FM on GCP

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

This section describes how to install GigaVUE-FM on GCP. It consists of the following sections:

- [Before You Install](#) describes the pre-requisites and the preparatory steps required for installing GigaVUE-FM.
- [Install GigaVUE-FM on GCP](#) describes the steps to install and deploy GigaVUE-FM on GCP.

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.

Create Monitoring Domain

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 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 using GCP - Generic Mode

Configure GigaVUE Fabric Components using GCP

This section provides step-by-step information on how to register GigaVUE fabric components using Google Cloud Platform (GCP) or a configuration file.

Minimum Requirements

The following table lists the minimum requirements for deploying the fabric components:

Fabric Component	Machine type
GigaVUE V Series Node	- c2-standard-4 for 2 interfaces - c2-standard-8 for 3 interfaces
GigaVUE V Series Proxy	e2-micro
UCT-V Controller	e2-micro

Deployment Guideines

Keep in mind the following when deploying the fabric components using GCP:

- Set the correct MTU for tunnels: With a default MTU of 1460 in GCP, use 1410 for IPv4 tunnels and 1390 for IPv6. To update the Traffic Acquisition Tunnel MTU:
 - Select the monitoring domain.
 - Select the **Edit Monitoring Domain** option.
 - Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

For details, refer to [Traffic Acquisition Tunnel MTU](#)

- For successful registration of fabric components, configure firewall rules to open ports.
 - For details on how to configure firewall rules, refer to [Use VPC firewall rules](#) topic in GCP documentation.
 - For details on required ports, refer to [Network Firewall Requirement](#).
- When you deploy the fabric components using third party orchestration, you cannot delete the Monitoring Domain or change the MTU without unregistering the registered fabric components.
- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to [Create Monitoring Domain](#).
- When launching an instance, to access it using a private key, add the SSH key. The default password is gigamon.

- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to [Configure Tokens for Third Party Orchestration](#) .
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section.

In your GCP, you can configure the following GigaVUE fabric components:

Configure UCT-V Controller in GCP

IMPORTANT: You must enable the basic authentication to launch the GigaVUE fabric components for version 6.9 and lower. For more instructions on the steps to enable the basic authentication, refer to [Authentication Type](#).

You can configure more than one UCT-V Controller in a monitoring domain.

To register UCT-V Controller in GCP, use one of the following methods:

Register UCT-V Controller during Instance Launch

In your GCP, you can launch the UCT-V Controller and to register UCT-V Controller using custom metadata.

Perform the following steps:

1. On the VM instances page of GCP, select **Create instances** . For details, refer to [Create and Start instance](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field enter the below mentioned text in the following format and deploy the instance. The UCT-V Controller uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the GigaVUE-FM>
      sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
      remotePort: 443
```

Register UCT-V Controller after Instance Launch

You can register UCT-V Controller after launching a Virtual Machine using a configuration file.

Perform the following steps:

1. Log in to the UCT-V Controller.
2. Create a local configuration file (**/etc/gigamon-cloud.conf**) and enter the following user data:

```
Registration:
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the GigaVUE-FM>
sourceIP: <IP address of UCT-V Controller> (Optional. Required only in multi-NIC
deployments to specify which UCT-V Controller interface to use for communicating with
GigaVUE-FM.)
remotePort: 443
```

3. Restart the UCT-V Controller service.
\$ sudo service uctv-cntlr restart

The deployed UCT-V Controller registers with the GigaVUE-FM. After successful registration, the UCT-V Controller 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, then GigaVUE-FM tries to reach the UCT-V Controller. If that fails as well then GigaVUE-FM unregisters the UCT-V Controller and it is removed from GigaVUE-FM.

The UCT-V Controller deployed in GCP appears on the Third Party Orchestration Monitoring Domain page of GigaVUE-FM.

☐ Monitoring Domain	Connection	Fabric	Management IP	Fabric Version	Status
☐ MD1					
	pultra-vpc				✓ Connected
		G-vTapController	34.219.250.141	1.7-304	✓ Ok
		Gigamon-VSeriesProxy-1	34.211.211.49	2.1.0	✓ Ok
		Gigamon-VSeriesNode-1	172.16.14.188	2.2.0	✓ Ok

Assign Static IP address for UCT-V Controller

By default, the UCT-V Controller gets assigned an IP address using DHCP.

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
  ethernet:
    <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
  ethernet:
    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 UCT-V Controller service.

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

The deployed UCT-V Controller registers with the GigaVUE-FM. After successful registration the UCT-V Controller 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 UCT-V Controller. If that fails as well then GigaVUE-FM unregisters the UCT-V Controller and removes from GigaVUE-FM.

Configure UCT-V in GCP

Linux and Windows platforms support Deployment of GigaVUE fabric components through a third-party orchestrator.

NOTE: You can configure more than one UCT-V Controller for a UCT-V, so that if one UCT-V Controller goes down, the UCT-V registration takes place through another active Controller.

Install the Windows UCT-V

Perform the following steps:

1. Deploy Windows server in GCP. For details, refer to [Create a Windows Server VM instance in Compute Engine](#) topic in Google documentation.
2. Follow the instruction in the *Connect to the VM instance by using RDP* section of [Set up Chrome Remote Desktop for Windows on Compute Engine](#) topic in the GCP documentation.
3. Download UCT-V build in your desktop and copy it to the RDP session.
4. Turn off the Windows Firewall Defender.
5. Install the Windows UCT-V. For details, refer to [Windows UCT-V Installation](#) for step-by-step instructions on how to install Windows UCT-V.

To register UCT-V in GCP, use one of the following methods.

Register UCT-V during Instance Launch

You need to register UCT-V through the registered UCT-V Controller. It communicates through PORT 8891.

NOTE: Registering UCT-V during Virtual Machine Launch is not applicable to Windows UCT-V. You can register your Windows UCT-V after launching the Virtual machine, using a configuration file.

In your GCP, to launch the instance and register the UCT-V using Custom Metadata, perform the following steps:

1. On the VM instances page of GCP, select **Create instances**. For details, refer to [Create and Start instance](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field, enter the below mentioned text in the following format and deploy the instance. The UCT-V uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the UCT-V Controller 1>,<IP address of the UCT-V
Controller 2>
      sourceIP: <IP address of UCT-V> (Optional Field)
```

Register UCT-V after Instance Launch

To register UCT-V after launching a Virtual Machine using a configuration file, perform the following steps:

1. Install the UCT-V in the Linux or Windows platform.
2. Log in to the UCT-V.
3. Create a local configuration file and enter the following user data.



- **/etc/gigamon-cloud.conf** is the local configuration file in Linux platform.
- **C:\ProgramData\uctv\gigamon-cloud.conf** is the local configuration file in Windows platform.
- When creating **C:\ProgramData\uctv\gigamon-cloud.conf** file, ensure that the file name extension is **.conf**. To view the file name extension in Windows, perform the following steps:
 - a. Go to File Explorer and open the File Location.
 - b. On the top navigation bar, select **View**.
 - c. In the **View** tab, enable the **File name extensions** check box.

Registration:

```
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the UCT-V Controller 1>,<IP address of the UCT-V
Controller 2>
sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
```

4. Restart the UCT-V service.
 - Linux platform:
`$ sudo service uctv restart`
 - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in GCP

NOTE: It is not mandatory to register GigaVUE V Series Nodes via V Series proxy however, if there is a large number of nodes connected to GigaVUE-FM or if the user does not wish to reveal the IP addresses of the nodes, then you can register your nodes using GigaVUE V Series Proxy. In this case, GigaVUE-FM communicates with GigaVUE V Series Proxy to manage the GigaVUE V Series Nodes.

To register GigaVUE V Series Node and GigaVUE V Series Proxy in GCP, use one of the following methods.

IMPORTANT: You must enable the basic authentication to launch the GigaVUE fabric components for version 6.9 and lower. For more instructions on the steps to enable the basic authentication, refer to [Authentication Type](#).

Register GigaVUE V Series Node and GigaVUE V Series Proxy during Instance Launch

1. On the VM instances page of GCP, select **Create an instance** . For details, refer to [Create VMs with multiple network interfaces](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field, enter the below mentioned text in the following format and deploy the instance. The UCT-V uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
      remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

Register GigaVUE V Series Node and GigaVUE V SeriesProxy after Instance Launch

You can register GigaVUE V Series Node and GigaVUE V Series Proxy after launching the virtual machine using a configuration file.

To register,

1. Log in to the GigaVUE V Series Node or Proxy.
2. Create a local configuration file (**/etc/gigamon-cloud.conf**) and enter the following user data.

```
Registration:
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the proxy>
remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

3. Restart the GigaVUE V Series node or proxy service.

- V Series node:
\$ sudo service vseries-node restart
- V Series proxy:
\$ sudo service vps restart

The deployed GigaVUE V Series proxy registers with the GigaVUE-FM. After successful registration the GigaVUE V Series proxy 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 proxy. If that fails as well, then GigaVUE-FM unregisters the GigaVUE V Series proxy and it is removed from GigaVUE-FM.

NOTE: When the GigaVUE V Series Node is stopped or terminated from the GCP, it does not send any unregistration request and GigaVUE-FM unregisters the V Series Node soon after.

What to do Next

You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab.

Configure Packet Mirroring for GCP

Packet Mirroring clones the traffic of specified instances in your Virtual Private Cloud (VPC) network and forwards it for examination. Packet Mirroring captures all traffic and packet data, including payloads and headers. You can configure the capture for both egress and ingress traffic, only ingress traffic, or only egress traffic.

NOTE: When deploying GigaVUE V Series Nodes for configuring Application Intelligence Session, do not use Packet Mirroring. As Application Intelligence uses stateful traffic, you may experience packet drop due to GCP platform limitation.

Refer to the following topics for detailed information.

- [Configure Packet Mirroring in GCP](#)
- [Configure Monitoring Session with Packet Mirroring](#)

Rules and Notes

- Load Balancer forwards raw traffic. When configuring the monitoring session, use the Raw End Point as the first component that receives traffic.
- You must configure three NICs for GigaVUE V Series Node because REP and TEP cannot share the same interface.

A typical GCP deployment to support the internal load balancer and packet mirroring requires the following components:

- GigaVUE-FM
- GigaVUE V Series Node
- GCP Internal Load Balancer (uniformly distributes traffic from GCP target VMs to GigaVUE V Series Nodes)

Configure Packet Mirroring in GCP

This section provides instructions on how to configure packet mirroring in GCP.

1. Create an instance template in GCP to deploy the GigaVUE V Series Node in GCP using Third Party Orchestration. For details, refer to [Register GigaVUE V Series Node and GigaVUE V Series Proxy during Instance Launch](#).



- When using packet mirroring, configure a minimum of 3 NICs and apply c2-standard-8 (8 vCPU, 32GB memory) as Machine Type.
- Enable **IP Forwarding** when creating the instance template in GCP.

2. Create Instance Group in GCP with autoscaling in Managed Instance Group. For details, refer [Create a MIG with autoscaling enabled](#) topic in Google Cloud Platform documentation.
3. Configure TCP or UDP internal Load balancer with packet forwarding enabled and ensure that the GigaVUE V Series Nodes data NICs are used to receive traffic. For details, refer to [Create a load balancer for Packet Mirroring](#) section in Google Cloud Platform documentation.
4. Configure packet mirroring in GCP. For instructions, refer to [Use Packet Mirroring](#) topic in Google Cloud documentation.

After configuring packet mirroring in GCP, edit the Monitoring Domain in GigaVUE-FM and configure the Monitoring Session.

Configure Monitoring Session with Packet Mirroring

To configure monitoring session with packet mirroring enabled in GCP, perform the following steps:

Step 1: Edit the monitoring domain and update the following details:

1. Go to **Inventory > VIRTUAL > Third Party Orchestration**, and then click **Monitoring Domain**. The **Monitoring Domain** page appears.
2. Select the Monitoring Domain that contains the GigaVUE V Series Node deployed with packet mirroring.
3. Select **Actions > Edit**.
4. On the **Monitoring Domain Configuration** page, select **Customer Orchestrated Source** as the Traffic Acquisition method.
5. Enable the **Uniform Traffic Policy** check box. When enabling this option, same monitoring session configuration is applied to all GigaVUE V SeriesNodes.
6. Select **Save**.

Step 2: Create a monitoring session with the following instructions:

1. On the left page of GigaVUE-FM, select **Traffic > Virtual > Orchestrated Flows** and select **Third Party Orchestration**. The **Monitoring Sessions** page appears.
2. Select **New** to open the **Create a New Monitoring Session** page. For details, refer to [Create a Monitoring Session](#).
3. In the **Edit Monitoring Session** page, add Raw End Point as the first component and Tunnel End Point as the final component. For details, refer to [Create Raw Endpoint](#) and [Create Ingress and Egress Tunnel](#).
4. Add your application to the monitoring session to connect your components.
5. Select the **Deploy** button in the edit monitoring session page to deploy the monitoring session after adding the Raw End Point.
6. In the **Select nodes to deploy the Monitoring Session** dialog box, from the drop-down menu, select the interface for REP and TEP.

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.
7. Copy the script in the standby instance in the **/home/admin directory** and execute it using the following command:

```
sh /home/admin/runstepca_fmha
```

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

- 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 Keys 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 certificates 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 Custom Authority Certificate (CA)	You must upload a Custom Certificate to UCT-V Controller to establish a connection with the GigaVUE V Series Node. To upload the CA using GigaVUE-FM, follow the steps given below: - Go to Inventory > Resources > Security > CA List. - Select New to add a new Custom Authority. The Add Custom Authority page appears. - Enter or select the following information. - Alias - Alias name of the CA. - File Upload - Choose the certificate from the desired location. - Select Save. For more information, refer to the section Adding Certificate Authority.
2.	Upload a SSL Key	You must add a SSL key to GigaVUE V Series Node. To add SSL Key, follow the steps in the section Supported Protocols, Algorithms, and Ciphers for SSL Decrypt

S. No	Task	Refer to
3	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.
4.	Select the SSL Key and CA certificate, after deploying the fabric components.	You must select the added SSL Key and CA Authority 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 and CA Authority in GigaVUE V Series Node.

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 Certificate Authority (CA) Certificate	You must upload a Custom Certificate to UCT-V Controller to establish a connection with the GigaVUE V Series Node. To upload the CA using GigaVUE-FM, follow the steps given below: 1. Go to Inventory > Resources > Security > CA List. 2. Select New to add a new Custom Authority. The Add Custom Authority page appears. 3. Enter or select the following information. • Alias - Alias name of the CA. • File Upload - Choose the certificate from the desired location. 4. Select Save. For more information, refer to the section Adding Certificate Authority.
2.	Upload an SSL Key	You must add an SSL key to GigaVUE V Series node. To add an SSL Key, follow the steps in the section Upload SSL Keys.
3	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 precrypted traffic.
4	Select the SSL Key and CA certificate, after deploying the fabric components.	You must select the added SSL Key and CA Authority 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 and CA Authority in GigaVUE V Series Node.
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. 2. In the canvas, select New > New Tunnel, drag and drop a new tunnel template to the workspace. The Add Tunnel

S. No	Task	Refer to
		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 Souce 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 Version1.3. - 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 the second GigaVUE V Series Node (Destination IP). 4. Select Save.
6	Select the added SSL Key after deploying the fabric components in the second GigaVUE V Series Node	You must select the added SSL Key in the second GigaVUE V Series Node. Select the 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 GigaVUE V Series Node with tunnel type as TLS-PCAPNG

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.	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: • 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 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. • IP Version - The version of the Internet Protocol. IPv4 and IPv6 are supported. • Remote Tunnel IP - Enter the interface IP address of the first GigaVUE V Series Node (Destination IP). 4. Select Save.

Viewing Status of Secure Tunnel

GigaVUE-FM allows you to view the status of secure tunnel connection in UCT-V. You can verify whether the tunnel is connected to the tool or GigaVUE V Series Node through the status.

To verify the status of secure tunnel,

1. Go to **Inventory > VIRTUAL > AWS**, and then select **Monitoring Domain**.
2. In the **Monitoring Domain** page, view the status.

Tunnel status displays the status of the tunnel. The green color represents that the tunnel is connected and the red represents that the tunnel is not connected.

For configuring secure tunnel, refer to the **Configure Secure Tunnel** section.

Configure Prefiltering

For prefiltering the traffic, GigaVUE-FM enables you to create a prefiltering policy template and apply it to a Monitoring Session,

Considerations

- You can create and reuse a policy template across multiple monitoring sessions.
- Each monitoring session supports up to 16 prefiltering rules.
- You can edit the policy template during monitoring session configuration to customize rules and filter values.
- GigaVUE-FM does not save changes made to the template during monitoring session edits. These changes apply only to that specific session.

Rules and Notes

- Prefiltering is supported on UCT-V for Windows systems and Linux systems running Kernel version 4.18 or later.
- Prefiltering is supported for both Linux and Windows UCT-Vs .
- For single monitoring session only one prefiltering policy is applicable. All UCT-Vs in that monitoring sessions are configured with respective prefiltering policy .
- For multiple monitoring session using the same agent to acquire the traffic, if a monitoring session uses a prefilter and the other monitoring session does not use a prefilter, then you cannot apply the prefiltering policy. The policy is set to PassAll and prefiltering is not performed.

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 for a Monitoring Session, refer to [Configure Monitoring Session Options \(AWS\)](#).

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**.

Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform

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

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

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

License

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

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

Supportability

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

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

Supported Deployment Models

GigaVUE-FM Orchestration

- VMware ESXi-Traffic acquisition using Platform Mirroring.

Third-Party Orchestration (Generic Mode)

Supported platforms

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

Supported traffic acquisition methods

- UCT-V
- Customer-Orchestrated Source

Supported GigaSMART Applications

The Gigamon Terraform provider supports the following GigaSMART applications:

- AMX
- Dedup
- Header Stripping
- Load Balancing
- Masking
- Slicing

Rules and notes

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

Prerequisites

Before starting a Terraform deployment, ensure the following:

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

Get Started

Install Terraform

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

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

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

```
terraform version
```

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

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

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

Create a working directory

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

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

Deployment Using Terraform

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

Third-Party Orchestration

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

Deploy Third Party Orchestration by using Terraform

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

Supported platforms (Generic Mode):

- AWS

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

This deployment supports the following traffic acquisition methods:

- UCT-V
- Customer-Orchestrated Source

Before you begin

Complete the following setup before starting the Terraform deployment.

1. **Install GigaVUE-FM.**

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

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

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

3. **Verify third-party networking configuration**

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

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

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

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

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

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

Install the Gigamon Terraform Provider Without Internet Access

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

Step 1: Download and install the provider binary

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

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

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

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

3. **Extract the zip file:**

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

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

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

5. **Make the binary executable:**

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

Step 2: Update the provider block

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

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

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

Step 3: Continue with the standard deployment flow

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

Prepare the Third Party Orchestration Deployment Files

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

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

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

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

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

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

1. [Single-file 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 UCT-V Features

Refer to the following sections for more detailed information:

- [Configure Prefiltering](#)
- [Create Precryption Template for UCT-V](#)
- [Configure Secure Tunnel \(AWS\)](#)

Configure Prefiltering

For prefiltering the traffic, GigaVUE-FM enables you to create a prefiltering policy template and apply it to a Monitoring Session,

Considerations

- You can create and reuse a policy template across multiple monitoring sessions.
- Each monitoring session supports up to 16 prefiltering rules.
- You can edit the policy template during monitoring session configuration to customize rules and filter values.
- GigaVUE-FM does not save changes made to the template during monitoring session edits. These changes apply only to that specific session.

Rules and Notes

- Prefiltering is supported on UCT-V for Windows systems and Linux systems running Kernel version 4.18 or later.
- Prefiltering is supported for both Linux and Windows UCT-Vs.

- For single monitoring session only one prefiltering policy is applicable. All UCT-Vs in that monitoring sessions are configured with respective prefiltering policy .
- For multiple monitoring session using the same agent to acquire the traffic, if a monitoring session uses a prefilter and the other monitoring session does not use a prefilter, then you cannot apply the prefiltering policy. The policy is set to PassAll and prefiltering is not performed.

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 for a Monitoring Session, refer to [Configure Monitoring Session Options \(AWS\)](#).

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

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

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

- [Create a Monitoring Session](#)
- [Create Tunnel Endpoint](#)
- [Create Raw Endpoint](#)
- [Create Map](#)
- [Add Applications to Monitoring Session](#)
- [Interface Mapping](#)
- [Deploy Monitoring Session](#)
- [View Monitoring Session Statistics](#)
- [Visualize the Network Topology](#)

Create a Monitoring Session (Third Party Orchestration)

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 (Third Party Orchestration)

The table outlines core tabs in the Monitoring Session interface, each enabling targeted configuration, traffic analysis, and infrastructure visibility to support efficient session management.

Tab	Description
Overview	You can view the high level information of the selected Monitoring Session such as, connections, tunnel details, health status, deployment status, and information related to Application Intelligence statistics. You can also view the statistics of the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis. You can filter the statistics based on the elements associated with the Monitoring Session. For more information, refer to View Monitoring Session Statistics .
Sources	Displays the sources and target details monitored by the Monitoring Session. You can view and edit the connection details of the Monitoring Session. You can view the deployment

Tab	Description
	status, number of targets, and targets source health. Note: In the case of OVS Mirroring, the Sources tab also displays the Hypervisor details along with the Instances.
Traffic Processing	You can view, add, and configure applications, tunnel endpoints, raw endpoints, and maps. You can view the statistical data for individual applications and also apply threshold template, enable user defined applications, and enable or disable distributed De-duplication. Refer to Configure Monitoring Session Options for more detailed information.
V Series Nodes	You can view the V Series nodes associated with the Monitoring Session. In the split view, you can view details such as name of the V Series Node, health status, deployment status, Host VPC, version, and Management IP. You can also change the interfaces mapped to an individual GigaVUE V Series Node. Refer to Interface Mapping section for details.
Topology	Displays the fabric and monitored instances based on the connections configured in your network. You can select a specific connection to explore its associated subnets and instances in the topology view, offering a clear visualization of the monitored network elements. Refer to Visualize the Network Topology (Third Party Orchestration) .

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

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

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

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

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

Monitoring Session Options (Third Party Orchestration)

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 Raw Endpoint (Third Party Orchestration)

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.

NOTE: The maximum number of links that can egress from any endpoint in V Series is four.

To add Raw Endpoint to the Monitoring Session,

1. Drag and drop **New Raw Endpoint** from the **New** expand menu to the graphical workspace.
2. On the new raw endpoint icon, click 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 click **Save**.

4. Perform the following steps to deploy the Monitoring Session after adding the Raw Endpoint:
 - a. On the **TRAFFIC PROCESSING** page, select **Actions>Deploy**. The **Deploy Monitoring Session** dialog box appears.
 - b. Select the desired V Series Nodes to deploy the Monitoring Session.
 - c. Select the interfaces for each of the REPs and the TEPs deployed in the Monitoring Session from the drop-down menu for the selected individual V Series Nodes.
 - d. Select **Deploy**.
5. Select **Export** to download all or selected V Series Nodes in CSV and XLSX formats.

vendithaCreate Ingress and Egress Tunnel (Third Party Orchestration)

Traffic from the GigaVUE V Series Node is distributed to tunnel endpoints in a monitoring session. You can create a tunnel endpoint using a standard L2GRE, VXLAN, UDPGRE, or ERSPAN tunnel.

Create a new tunnel endpoint

To create,

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

The GigaVUE-FM Monitoring Session canvas page appears.


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

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

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

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

5. Select **Save**.



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

Apply a threshold template to Tunnel End Points

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

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

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

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

Table 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 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.

Field	Description	
Out	Choose Out (Encapsulation) for creating an egress tunnel from the GigaVUE V Series Node to the destination endpoint.	
	Remote Tunnel IP	For egress tunnel, the Remote Tunnel IP is the IP address of the tunnel destination endpoint.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.
	Time to Live	Enter the value of the time interval for which the session needs to be available. The value ranges from 1 to 255. The default value is 64.
	DSCP	Differentiated Services Code Point (DSCP) is a value that network devices use to identify traffic to be handled with higher or lower priority. The values ranges from 0 to 63 with 0 being the highest priority and 63 being the lowest priority.
	Flow Label	Unique value, which is used to identify packets that belong to the same flow. A flow is a sequence of packets that need to be treated as a single entity that may require special handling. The accepted value is between 0 and 1048575.
	VXLAN Network Identifier	Unique value which is used to identify the VXLAN. The value ranges from 1 to 16777215.
	Multi Tunnel	Enable the multi-tunnel flag to create multiple tunnels for flow distribution to the 5G-Cloud application. Refer to 5G-Cloud Ericson SCP Support. Applicable Platforms: OpenStack, Third Party Orchestration, VMware ESXi  Notes: - You can configure either a single-step or multi-step 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,

Field	Description	
		 then recreate the LB and VXLAN TEP with the new configuration, and re-establish the link between them.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
	Domain Tagging	Enable this option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. NOTE: This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
UDPGRE		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	Choose In (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Key	Identifier used to differentiate different UDPGRE/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		

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

Field	Description	
		the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. NOTE: This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
ERSPAN		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Flow ID	The ERSPAN flow ID is a numerical identifier that distinguishes different ERSPAN sessions or flows. The value ranges from 1 to 1023.
TLS-PCAPNG		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. Note: In the scenario where secure tunnels need to be established between a GigaVUE V Series and a GigaVUE HC Series, you can utilize the Configure Physical Tunnel option provided in the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels on your physical device . For details, refer to Secure Tunnels section.		

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

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

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

Create a New Map (Third Party Orchestration)

Terms to know before creating a map:

Parameter	Description
Rules	A rule (R) contains specific filtering criteria that the packets must match. The filtering criteria lets you determine the targets and the (egress or ingress) direction of tapping the network traffic.
Priority	Priority determines the order in which the rules are executed. The priority value can range from 1 to 5, with 1 being the highest and 5 is the lowest priority.
Pass	The traffic from the virtual machine is passed to the destination.
Drop	The traffic from the virtual machine is dropped when passing through the map.
Traffic Filter Maps	A set of maps that are used 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 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: • 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. • Use the GigamonNode Tag to exclude any Gigamon devices from the target. • When using VM Name filter for selecting the Virtual Machines using Inclusion and Exclusion Maps, wild- cards in VM names are not supported. You must use the prefix of the Virtual Machine name.
Group	A group is a collection of maps that are pre-defined and saved in the map library for reuse.

Rules and Notes:

- Directional rules do not work on single NIC VMs that are running a Windows UCT-V.
- Loopback captures bidirectional traffic from both ingress and egress. To prevent duplicate tapping, only egress tapping is permitted.
- If 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.

To create a new map:

1. Drag and drop **New Map** from the **New** expand menu to the graphical workspace. The **Map** quick view appears.



2. On the new Map quick view, select the **General** tab and enter the required information as described below.
 - a. Enter the **Name** and **Description** of the new map.
 - b. Enable the **Application Filtering** option if you wish to use Application Filtering Intelligence.

Enabling this option allows you to filter traffic based on Application name or family. Refer to [Application Filtering Intelligence](#).

NOTE: Pass and Drop rule selection with Automatic Target Selection (ATS) differ with the Map type as follows:

- Traffic Map—Only Pass rules for ATS
- Inclusion Map—Only Pass rules for ATS
- Exclusion Map—Only Drop rules for ATS

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

Through the map, you can drop or pass packets 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. To add ATS rules for an Inclusion/Exclusion map, you must select at least one rule condition. For details, refer to [Example- Create a New Map using Inclusion and Exclusion Maps](#).

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

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

Example- Create a New Map using Inclusion and Exclusion Maps

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

1. Drag and drop a new map template to the workspace. The New map quick view appears.
2. In the **GENERAL** tab, enter the name as Map 1 and enter the description. 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** is 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.

Add Applications to Monitoring Session

GigaVUE Cloud Suite with GigaVUE V Series Node supports the following GigaSMART applications in the GigaVUE-FM canvas:

You can also configure the following GigaSMART operations from the **Traffic > Solutions > Application Intelligence**:

- Application Metadata Intelligence
- Application Filtering Intelligence

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>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr></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	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows								
VMware	Large (8 vCPU and 16GB RAM)	200k								
Nutanix	Large (8 vCPU and 16GB RAM)	200k								
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.									

5. 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 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 Deployment Options](#)

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.


Boot - Device boot , uptime init
 FS - First Packet for the flow switched
 LS - Last Packet for the flow switched
 Export - When the flow is exported
 EXP_TS - abs ts at export [present in netflow v9 header]
 UPTIME - Time since this device is booted [present in netflow v9 header]
 FIRST_SWITCHED - System uptime at which the first packet of this flow was switched
 LAST_SWITCHED - System uptime at which the last packet of this flow was switched



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

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

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



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

The metadata enrichment enhances service provider analytics, by generating metadata on 5G/4G/3G network traffic. The AMX correlates the user plane metadata produced by AMI with the control plane metadata produced by the GTP/5G correlation mobility

application to produce an enriched metadata feed for the mobile networks. This data feed helps with use cases like service personalization, planning, and many others by containing information about the:

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


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

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

- GigaVUE-HC3
- GigaVUE-HC1-Plus

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

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



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

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

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

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

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

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

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


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

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

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

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


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



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

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

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



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

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

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

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

Rules and Notes

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


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

Configure Application Metadata Exporter

NOTE: Refer to [Prerequisites for Application Metadata Exporter](#) topic for all information related to the prerequisites.

In GigaVUE-FM, you must do the following to configure the Application Metadata Exporter application in the Monitoring Session.

S.No	Steps	Refer to
1	Create an ingress REP	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (VMware NSX-T) Create Raw Endpoint (AWS) Create Raw Endpoint (Azure) Create Raw Endpoint (Third Party Orchestration)
2	Create a link between the ingress REP and Application Metadata Exporter application	N/A
3	Create an egress REP	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (VMware NSX-T) Create Raw Endpoint (AWS) Create Raw Endpoint (Azure) Create Raw Endpoint (Third Party Orchestration)
4	Create a link between the Application Metadata Exporter application and egress REP	N/A



To add AMX application:

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Session** page appears.



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.
- You can use the Third-Party Orchestration method when deploying fabric components—such as the AMX V Series Node—on your own.
- When deploying AMX using the Third-Party Orchestration method, Monitoring Session must also be configured in the same work flow. To deploy through Third Party Orchestration, go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**.
- Before deploying AMX, ensure to change the **Traffic Acquisition Method** to **Customer Orchestrated Source** in the Monitoring Domain. Refer to [Create Monitoring Domain](#).

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. In the canvas, click the  icon on the left side of the page to view the traffic processing elements. Drag and drop **AMX** from **APPLICATIONS** to the graphical workspace. The application quick view appears.
4. Enter the **Alias** for the application.

5. In the **Ingestor** section, enter or select the following details. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.

Fields	Description
Name	Enter name for the Ingestor.
Port	Enter the port number to which the Application Metadata or the NetFlow/IPFIX or the Control Plane metadata must be ingested.
Type	Select any one of the following: AMI - Select this option if the input is AMI. Mobility Control Plane - Select this option if the input is Control Plane Metadata. NetFlow/IPFIX - Select this option if the input type is NetFlow/IPFIX. Note: If your selected ingestor type is NetFlow/IPFIX, Metadata Enrichment will be disabled by default. NetFlow/IPFIX input type is not supported for Metadata Enrichment.

6. Enter or select the following details in the **Metadata Enrichment** section:



You can use the **Actions** button to add multiple Metadata Enrichment. Keep in mind the following when configuring multiple Metadata Enrichment:

- You can only configure either Mobility or Workload enrichment
- You can only configure one Metadata enrichment with either **Mobility** or **Workload** as the **Type**
- You can only configure a maximum of 5 Metadata enrichment with **Others** as the **Type**.

Fields	Description
Enrichment Name	Enter a unique name for each enrichment.
Enable	Use this option to enable the enrichment.
Type	Select the type from the drop-down menu. - o Mobility - o Workload - o Others
Mobility	
Attribute Fields	Select the attributes from the list. You can use the Select All option to select all the available attributes. Refer to Attributes for GigaVUE Enriched Metadata for Mobile Networks for more detailed information on the list of available attributes and their description.
Workload	
Platforms	Select the platform in which your Workload Virtual Machines are present. • VMware vCenter • AWS • Azure • Kubernetes
Attribute Fields	Select the attributes from the list. You can use the Select All option to select all the available attributes. Refer to Attributes for GigaVUE Enriched Metadata for Cloud Workloads for more detailed information on the list of available attributes and their description. Advanced Options The advanced options allow you to configure additional details like interval and delay. Click Add. Enter the following details: - o Enter the Key. - o Enter the Value.
Source Information	Enter the details of the source. Name: Enter a unique name for each Source Information. The name should be unique across the ingestor and the source information.

Fields	Description
	Key: The default keys for each platform are listed as follows. Click + to add more keys. AWS: Refer to Manage access keys for IAM users section in AWS documentation for more detailed information on how to create an access key ID and a secret access key. - aws_access_key_id - aws_secret_access_key - aws_region Azure: Refer to Register a Microsoft Entra app and create a service principal section in Azure documentation for more detailed information on how to configure client ID, client secret, tenant ID, and subscription ID in Azure. - azure_client_id - azure_client_secret - azure_tenant_id - azure_subscription_id VMware: - url - The URL of VMware vCenter. - username - Username of the VMware vCenter - password - vCenter password used to connect to the vCenter - self_signed_certificate Kubernetes: - Upload kubeconfig YAML file from your local. Refer to Kubernetes for Prerequisites details. Value: Enter the value for the keys. This field is editable after saving the changes. Secure Keys: Use this option to mask the value. After saving the changes, Secure Keys and Key fields are disabled and the value is masked. Click  to add another Source Information and click  to remove a Source Information. You can add multiple Source Information if your workloads are across different vCenters, AWS Accounts, or Azure Subscriptions. You can create an individual Source Information for each of the vCenters, AWS Accounts, or Azure Subscriptions.
	Others: Use this setting to perform custom mapping for advanced use cases. For details contact Gigamon Support.
Attribute Fields	Click Add and enter the attribute.
Settings	Click Add and enter the details.

7. Enter or select the following details in **Cloud Tools Export** section:

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. Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads. NetFlow/IPFIX - Select this option to export NetFlow/IPFIX packets to tools such as HTTPS or Kafka.
Account ID	Enter the account ID number of the selected Cloud Tool.
API Key	Enter the API key of the Cloud Tool.
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.
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Format	The default output format is JSON.
Zip	Enable the box to compress the output file. Note: Enable this field when using New Relic as the cloud tool.
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.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Export Retries	The number of times the application tries to export the entries to Cloud Tool. The recommended minimum value is 4 and the maximum is 10.
Maximum Entries	The number of JSON entries in a file. The maximum number of allowed entries is 20000 and the minimum is 10, however 1000 is the default value.

Fields	Description
Backoff Reset Window	With a non-zero value, the exporter keeps trying to reach the tool as many times as the retry. The default setting is 0, and the Backoff Reset Window is disabled at this value.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
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.

8. Enter or select the following details in the **Kafka Exports** section:

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. Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata. 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.
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.
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Format	The default output format is JSON.
Zip	Enable the box to compress the output file.
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.
Export Retries	The number of times the application tries to export the entries to Kafka. The recommended minimum value is 4 and the maximum is 10.
Maximum Entries	The number of JSON entries in a file. The maximum number of allowed entries is 20000 and the minimum is 10, however 1000 is the default value.
Backoff Reset Window	With a non-zero value, the exporter keeps trying to reach the tool as many times as the retry. The default setting is 0, and the Backoff Reset Window is disabled at this value.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value.
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

9. Click **Save** to deploy the Monitoring Session. The **Select nodes to deploy the Monitoring Session** dialog box appears. Select the GigaVUE V Series Node for which you wish to deploy the Monitoring Session.
10. 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 all the V Series Nodes.

Switching to outer_ip for Mobility Enrichment

In 5G Control Plane (CP) records, only the lower 64 bits of the inner IP address are captured. To overcome this limitation, you can switch to outer_ip. Follow the steps below based on your configuration:

If AMX is already configured and the Monitoring Session is deployed:

1. Undeploy the Monitoring Session in GigaVUE-FM.
2. Create a new .yaml file in the GigaVUE V Series Node: `vi /etc/amx_settings.yaml`
3. Modify the enrichment mode:

```
# Default mobility enrichment mode: inner_ip (default) or outer_ip
mobility_outer_ip_version: <Enter version as v6 or v4>
mobility_enrichment_mode: outer_ip
```
4. Deploy the Monitoring Session.

During Initial AMX Configuration:

1. Create a new .yaml file in the GigaVUE V Series Node: `vi /etc/amx_settings.yaml`
2. Modify the enrichment mode:

```
# Default mobility enrichment mode: inner_ip (default) or outer_ip
mobility_outer_ip_version: <Enter version as v6 or v4>
mobility_enrichment_mode: outer_ip
```
3. Save the file and proceed with the deployment.

Application Monitoring Options

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. Refer to [View Application Metadata Exporter Statistics](#) for more detailed information.

Deploy Application Metadata Exporter in GCP 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	

Requirements	Small Scale Deployment
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:
 - o Enter a name for the virtual machine.
 - o Select the **Guest OS family** as Linux.
 - o Select the **Guest OS version** as Other Linux (64-bit).
 - o 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. 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.

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

NOTE: It is not mandatory to register GigaVUE V Series Nodes via V Series proxy however, if there is a large number of nodes connected to GigaVUE-FM or if the user does not wish to reveal the IP addresses of the nodes, then you can register your nodes using GigaVUE V Series Proxy. In this case, GigaVUE-FM communicates with GigaVUE V Series Proxy to manage the GigaVUE V Series Nodes.

To register GigaVUE V Series Node and GigaVUE V Series Proxy in GCP, use one of the following methods.

IMPORTANT: You must enable the basic authentication to launch the GigaVUE fabric components for version 6.9 and lower. For more instructions on the steps to enable the basic authentication, refer to [Authentication Type](#).

Register GigaVUE V Series Node and GigaVUE V Series Proxy during Instance Launch

1. On the VM instances page of GCP, select **Create an instance** . For details, refer to [Create VMs with multiple network interfaces](#) topic in the GCP documentation.
2. Under the **Metadata** tab, enter the **key** as **user-data** and in the **value** field, enter the below mentioned text in the following format and deploy the instance. The UCT-V uses this Custom Metadata to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#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>
      remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
      remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

Register GigaVUE V Series Node and GigaVUE V SeriesProxy after Instance Launch

You can register GigaVUE V Series Node and GigaVUE V Series Proxy after launching the virtual machine using a configuration file.

To register,

1. Log in to the GigaVUE V Series Node or Proxy.
2. Create a local configuration file (**/etc/gigamon-cloud.conf**) and enter the following user data.

```
Registration:
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the proxy>
remotePort: 443
```



NOTE: You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **remotePort** value as 443 and the **remoteIP** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using GigaVUE V Series proxy then, enter the **remotePort** value as 8891 and **remoteIP** as <IP address of the Proxy>.

3. Restart the GigaVUE V Series node or proxy service.

- V Series node:
\$ sudo service vseries-node restart
- V Series proxy:
\$ sudo service vps restart

The deployed GigaVUE V Series proxy registers with the GigaVUE-FM. After successful registration the GigaVUE V Series proxy 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 proxy. If that fails as well, then GigaVUE-FM unregisters the GigaVUE V Series proxy and it is removed from GigaVUE-FM.

NOTE: When the GigaVUE V Series Node is stopped or terminated from the GCP, it does not send any unregistration request and GigaVUE-FM unregisters the V Series Node soon after.

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.
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.

Configure Exporter to Export Attributes to Third Party tools

AMX supports the following types of exporters:

- [Configure Cloud Tool Exporter](#)
- [Configure Kafka Exporter](#)

Configure Cloud Tool Exporter

To configure Cloud Tool Export settings in AMX:

1. On the AMX quick view in the **Cloud Tool Exports** section, enter or select the required details as described in [Cloud Tool Exporter - Field References](#).
2. Click **Save**.

Deploy the Monitoring Session

1. Click **Deploy** from the **Actions** menu.
2. The **Select nodes to deploy the Monitoring Session** dialog box appears. Select the GigaVUE V Series Node for which you wish to map the interfaces and deploy the Monitoring Session.
3. Select the interfaces for the ingress and egress REPs configured in the Monitoring Session. The interface naming includes the type and IP address of the interface, for example, DATA (ens192) IPv4: 10.114.42.135 IPv6: fcab:0:1:2040:a0a1:874e:7a66:40a2.
4. Select **Deploy**. Once deployed, the status is displayed as **Success** in the Monitoring Sessions **Sources** tab. The session is successfully deployed on the selected V Series Node.

Cloud Tool Exporter - Field References

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable Secure Endpoint option to hide the API keys from the URL. Example: - Normal Endpoint: http://examples.com/mykey/key123 - Masked Endpoint: http://examples.com/mykey/****
Headers	Enter the custom header value of the cloud tool instance. Example: Add Content-Type: application/json in a separate header field. Ensure that the Secure Keys option is enabled.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to transform NetFlow/IPFIX records to JSON format.
Labels	Click Add. Enter the following details: - Enter the Key. - Enter the Value. Note: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Zip	Enable the box to compress the output file. Note: Enable this field when using New Relic as the cloud tool.
Interval (sec)	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
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

Fields	Description
	value is 5000.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Kafka Broker. Both IPv4 and IPv6 are supported. This field is optional.
Producer Configurations	Click Add to enter the authentication details if a Kafka broker needs authentication. For Example: • security.protocol=SASL_SSL • sasl.mechanism=PLAIN or sasl.mechanisms=PLAIN • sasl.username=username • sasl.password=password
Message Configurations	Use this field to specify custom Kafka message settings, such as partition key selection, that GigaVUE-FM passes to the Kafka exporter within the Application Metadata Exporter application. Valid values are: • Key - partition.key and Value - flow_id, five_tuple, three_tuple, round_robin (default) • Key - partition.keyRotateInterval and Value - <seconds> (default: 3600; range: 60–86400). This key is rotated every <value> seconds, and this rotated key is used as a fallback when the selected key option (flow_id, three_tuple, or five_tuple) cannot be determined for a record. NOTE: If the five_tuple option is selected, but port attributes are not found in the record, the partition selection will be determined using the three_tuple.

Extend Your AMI and AMX Deployment

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

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

Verify Your Deployment

You can view the Ingestor, Exporter and Enrichment statistics for the deployed Monitoring Session in the **TRAFFIC PROCESSING** tab. For more details, refer to [View Application Metadata Exporter Statistics](#).

To troubleshoot any issues in Application Metadata Exporter deployment, refer to [Troubleshoot Application Metadata Exporter Deployment](#).

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.

Parameters	Description
Packets Dropped	Displays the packets dropped when the traffic passes from AMI to AMX ingestor.
Octets Received	Displays the octets received from the AMI to AMX ingestor.
Octets Dropped	Displays the octets dropped when the traffic passes from AMI to AMX ingestor.
Records Added	Displays the number of workload information pulled from the workload environment.
Records Dropped	Displays the number of the workload details that were dropped from the workload environment.
Request Success	Displays if the workload information was fetched from the cloud environment. The default interval at which the AMX fetches the workload information is 60 seconds.
Request Authentication Errors	Displays if the workload information was not fetched from the cloud environment due to any authentication errors like wrong password.
Request Timedout Errors	Displays if the workload information was not fetched from the cloud environment even after the default interval of 60 seconds.
Request Errors	Displays when the workload information could not be fetched from the cloud environment due to issues other than authentication failures or timeout conditions.
Decode Success	Displays number of successful decoding operations performed by a system when processing NETFLOW/IPFIX packets.
Decode Fail	Displays the number of failed attempts to decode NETFLOW/IPFIX packets.
Elements Skipped	Displays number of data elements that were not processed or were omitted during the NETFLOW/IPFIX ingestion process.
Pkts Rx Rolling 15 Sec	Rolling average of packets received per second over the past 15 seconds.
Pkts Rx Rolling 30 Sec	Rolling average of packets received per second over the past 30 seconds.
Pkts Rx Rolling 60 Sec	Rolling average of packets received per second over the past 60 seconds.
Mbytes Rx	Total number of megabytes ingested.
Mbytes Rx Rolling 15 Sec	Rolling average of total megabytes received over the past 15 seconds.
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

core_ipv6_addr	SMF identifier for N11 and GGSN for Gn,Gp
ctrl_tun_core_teid	Control Plane Core Tunnel Endpoint Identifier
user_tun_up_link_ipv4_addr	User Plane Uplink Tunnel IPv4 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_ipv6_addr	User Plane Uplink Tunnel IPv6 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_teid	User Plane Uplink Tunnel Endpoint Identifier
user_tun_down_link_ipv4_addr	User Plane Downlink Tunnel IPv4 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_ipv6_addr	User Plane Downlink Tunnel IPv6 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_teid	User Plane Downlink Tunnel Endpoint Identifier
if_name	3GPP Control Plane Interface Name
name	Name of the event on the 3GPP control plane interface
cause_code	Outcome of the event on the 3GPP control plane interface. Mostly set to 16 to indicate a successful event
sm_context_ref	Session Management Context
ebi	EPS(Evolved Packet System) Bearer ID
lbi	Linked Bearer Identity
pdu_session_id	PDU(Packet Data Unit) Session ID
apn	Access Point Name
dnn	Data Network Name
ue_ipv4_addr	User Equipment IPv4 Address
ue_ipv6_addr	User Equipment IPv6 Address
ue_ipv4v6_addr	User Equipment IPv4v6 Address
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.

Attribute	Description
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)	✓	✓	✓

¹The numbers within the enclosed brackets represent the element ID in IPFIX/NetFlow v10 and the field type ID in NetFlow v9.

	- 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) - 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)	✗	✓	✓
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)	✗	✓	✓
packets_total	Running packet counter for a permanent flow • NetFlow v9 - IN_PERMANENT_PKTS (86) • NetFlow v10/IPFIX - packetTotalCount (86)	✗	✓	✓
flow_direction	Flow direction • NetFlow v9 - DIRECTION (61) • NetFlow v10/IPFIX - flowDirection(61)	✗	✓	✓
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)	✗	✓	✓
icmp_type	Type of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpTypeXXX (176, 178), icmpTypeCodeXXX (32, 139)	✗	✓	✓
icmp_code	Code of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpCodeXXX (177, 179) icmpTypeCodeXXX (32, 139)	✗	✓	✓
flow_end_reason	The reason for flow termination. Values are listed in the flowEndReason registry.	✗	✓	✓
application_id	Specifies an Application ID - NetFlow v9 - APPLICATION_TAG (95) - NetFlow v10/IPFIX - applicationId (95)	✗	✓	✓

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	✓

Application Metadata Intelligence

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

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

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

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

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

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

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

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

Configure Application Metadata Intelligence for Virtual Environment

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



Notes:

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

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

view appears.

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

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

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

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

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

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

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

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

- Click **Save**.

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

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

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

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

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

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

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

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

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

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

Extend Your AMI and AMX Deployment

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

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

View Application Statistics for Application Metadata

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

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

Application 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><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>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	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows								
VMware	Large (8 vCPU and 16GB RAM)	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.

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.



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 than the existing de-duplication.

NOTE: From version 6.9, Traffic Distribution option is renamed to Distributed De-duplication.

From version 6.5, distributed de-duplication is supported on Azure, AWS, GCP, Nutanix, VMware NSX-T, VMware ESXi, and OpenStack. An enhanced configuration profile for the load balancer will be set by default with no option for modification. The default profile will use source and destination IP addresses, source and destination ports as the configuration for calculating the hash value for traffic distribution.

Important:

1. False traffic health alarms could be raised due to distribution of traffic across GigaVUE V Series Nodes.
2. Statistics are displayed for all the applications. Distributed De-duplication requires additional entities which will be listed in statistics page.

Limitations

- When using distributed de-duplication application, if **Prefer IPv6** option is enabled when configuring GigaVUE V Series Node using GigaVUE-FM, IPv6 tunnel will be created only between GigaVUE V Series Node and UCT-V and the rest of the tunnels created are IPv4 tunnels. Refer to Configure GigaVUE Fabric Components section in the respective cloud deployment guide for more detailed information on how to deploy GigaVUE V Series Node with IPv6 tunnels.
- When using distributed de-duplication application, if **Enable IPv6 Preference** option is enabled when configuring GigaVUE V Series Node using Third Party Orchestration, IPv6 tunnel will be created only between GigaVUE V Series Node and UCT-V and the rest of the tunnels created are IPv4 tunnels. Refer to *Create Monitoring Domain* section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration Guide for more detailed information on how to create a monitoring domain to register GigaVUE fabric components with IPv6 tunnels.

What's Next

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

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


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

GENEVE Decapsulation

The GENEVE Decapsulation application is used to acquire and strip GENEVE headers. To route the traffic through the third-party network appliances seamlessly, the AWS gateway load balancer with a VPC adds GENEVE header to packets as they are forwarded to a third-party network appliance. Each appliance is expected to terminate the GENEVE tunnel and process the GENEVE encapsulated traffic. When the GigaVUE-FM directs the acquisition of the customer traffic, the packets are encapsulated and forwarded as GENEVE tunnels that are terminated in GigaVUE V Series nodes.

To add a GENEVE application:

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



3. Enter an alias for the GENEVE application.
4. Click **Save**.

What's Next

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

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

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

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>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	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows								
VMware	Large (8 vCPU and 16GB RAM)	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

Parameter	Description
	to be filtered from the traffic. Attributes for the selected application is displayed in the Attribute column. You can select the required attributes.
NETWORK & TRANSPORT PARAMETERS: Select the Network and the transport packet attributes with the respective parameters	
Data Link	Select any one of the parameters such as Source MAC Address, Destination MAC Address and VLAN.
Interface	Select any one of the parameter such as Input Physical, Output Physical and Input Name.
IP	Select the parameter as Version if required.
IPv4	Select the required attributes. By default, Source Address, Destination Address, and Protocol are enabled.
IPv6	Select the required attributes. By default, Source Address, Destination Address, and Next Header are enabled.
Transport	Select the required attributes. By default, Source Port, Destination Port are enabled.
Counter	Select the Bytes, and Packets.
Timestamp	Select the required timestamp such as System Uptime First, Flow Start, System Uptime Last, and Flow End.
Flow	Select the parameter as End Reason if required.
GTP-U	Select the required parameters such as QFI and TEID.
Outer IPv4	Select any one of the parameter such as Source or Destination.
Outer IPv6	Select any one of the parameter such as Source or Destination.

5. Click **Save**.

Header Stripping

Header Stripping application efficiently examines the packets for specified headers like GTP, ISL, ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-tag, VXLAN, FM6000Ts, GENEVE, and generic and remove them before sending the packet to the appropriate security and analysis tools. Each packet is examined for the packet forwarding addition, and it also ensures that the headers are removed before sending the packet to the tools. This application is useful when working with tools that either cannot recognize these headers or must engage in additional processing to adjust for them.

Furthermore, the presence of protocols like GTP, ISL, ERSPAN, MPLS, MPLS+VLAN, VLAN, VN-tag, VXLAN, GENEVE, and FM6000Ts in the packet can restrict or limit the ability to apply filtering and flow-based load balancing to the traffic as it is forwarded to specific tools. To address each of these challenges, Header Stripping of these protocols is required.

List of Protocols that are supported for stripping:

- GTP
- ISL
- ESPRAN
- MPLS
- MPLS+VLAN
- VLAN
- VN-tag
- VXLAN
- FM6000Ts
- GENEVE
- Generic



Note:

Header Stripping for the GENEVE protocol is currently supported for the following platforms:

- VMware vCenter
- VMware NSX-T
- Nutanix

Limitation

The GENEVE protocols currently supported for header stripping are the Transparent Ethernet Bridge (0x6558), IPv4 (0x0800), and IPv6 (0x86DD).

Configure Header Stripping Application

To configure the Header Stripping application in GigaVUE-FM, follow the steps given below:

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

3. In the application quick view, enter or select the required information as described in the following table:

Field	Description
Alias	Enter the alias name for the application.
Protocol	Select the type of protocol.
VLAN: Use this option to strip the VLAN header from the packets. You can either strip the outer VLAN header or the entire VLAN header. When choosing VLAN as your protocol for stripping, enter the following details:	
VLAN Header	The VLAN Header that should be stripped. The supported minimum value is 0, and the maximum value is 16777215. The default value is 0.
VXLAN: Use this option to strip VXLAN (Virtual eXtensible Local Area Network) headers. You can strip either matching VXLAN headers or all VXLAN headers. When choosing VXLAN as your protocol for stripping, enter the following details:	
VXLAN ID	The VXLAN ID that should be stripped. The default value is outer.
FM6000Ts: Use this option to strip FM6000Ts time stamp headers. Packets entering the application from other devices may contain FM6000 timestamps. FM6000 is an Intel chip used for time stamping. FM6000 has a hardware timestamp in the packet. When choosing FM6000Ts as your protocol for stripping, enter the following details:	
Time Stamp Format	The format of the time stamp you wish to strip. The only supported format for the time stamp is None .
ESPRAN: Use this option to strip ERSPAN Type II and Type III headers. When choosing ESPRAN as your protocol for stripping, enter the following details:	
ESPRAN FlowID	Specify an ERSPAN flow ID between 0 to 1023. A flow ID of zero is a wildcard value that matches all flow IDs.
GENEVE: Use this option to strip the GENEVE header from the packets. The default L4 Destination Port value is 6081. To change the default L4 destination port value, follow the steps given below:	
1. Expand the Custom Port section. 2. You can either select an existing template with the port details or provide them directly in the pop-up window. - o Using Template - If you wish to use an existing template, select the template from the Template drop-down menu. To create a new template, refer to Create Custom Port Template.	
Note: You cannot modify an existing template by editing the values directly in the Monitoring Session. However, you can modify the values in the Monitoring Session and save it as a new template.	
o Without Template - You can directly enter your port details in the L4 Destination Port field. Click the Save as New Template button to save this existing configuration as a template. You can view the newly created template in the Custom Port Template page.	

Generic: Using this option to strip any header without worrying about the header level. When choosing generic as your protocol for stripping, enter the following details:	
Ah1	The anchor header (AH1), after which the header to be stripped occurs.
Offset	Based on the selected Offset, enter the following details: • Offset Range: If you wish to use offset range as your offset, then enter the following details: a. Offset Range Value: Specify the offset of the header occurrence from the above anchor header. The minimum supported value is one, and the maximum supported value is 1500. b. Header Count: Specify the number of headers the application should remove from the offset. The minimum supported value is one, and the maximum is 32. c. Custom Len: The length (in bytes) of the header that should be stripped. d. Ah2: The next possible standard header that occurs immediately after the header. • Start / End: If you wish to use start or end as your offset, then enter the following details: a. Header Count: Specify the number of headers that the application should remove from the offset.. The minimum supported value is one, and the maximum is 32. b. Custom Len: The length (in bytes) of the header that should be stripped. The minimum supported value is one, and the maximum supported value is 1500. c. Ah2: The next possible standard header that occurs immediately after the header.

4. Click **Save**.

Create Custom Port Template

GigaVUE-FM enables you to specify custom L4 destination port values when using the Header Stripping application. You can create a template and utilize it to define custom port values when configuring the Header Stripping application.

To create a custom port template for the Header Stripping application:

1. Go to **Inventory > Resources > Custom Port Template**. The **Custom Port Template** page appears.
2. Click **New** to create a new template. The **New Custom Port Template** dialog box appears.
3. Enter a name for the template.
4. Enter valid port numbers between the range 1025 to 65535 in the **L4 Destination Port** field. Enter only comma separated integer values. You can provide a maximum of 10 ports. For example: 1034,1098,1039
5. Click **Save**.

You can view the created template on the **Custom Port Template** page. If a template is associated with a Monitoring Session, you can view the name of the Monitoring Session to which the template is associated in the Affected Entity column.

You can use the **Actions** buttons to perform the following:

- **Edit:** You can use this option to edit a template. If a template is associated with a Monitoring Session, editing it will be reflected in all the Monitoring Sessions associated with that template.
- **Delete:** You can use this option to delete a template. You cannot delete a template associated to a Monitoring Session.

What's Next

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

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

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


Load Balancing

Load balancing application performs stateless distribution or Enhanced Load Balancing of the packets between different endpoints. Stateless load balancing distributes the processed traffic to multiple tool ports or tunnel endpoints based on hash values generated from predefined protocol fields in the packet.

When Enhanced Load Balancing is enabled and an endpoint fails, the traffic is redistributed for the failed endpoint. When the failed endpoint recovers, the redistributed traffic is restored to the recovered endpoint. The traffic across other endpoints remain undisturbed during this process.

To add a load balancing application:

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



X Application Save

Application	Load Balancing
Alias	lb
Hash Fields	ipOnly
Field Location	outer
Load balancing groups	

+ - Application Endpoint ID Weight

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

Metric	Description
Alias	Enter a name for the load balancing application
Stateless	Select this option to enable Stateless Load Balancing
Enhanced Load Balancing	Select this option to enable Enhanced Load Balancing and select the ELB profile.
Hash Field	• ipOnly : The source IP and destination IP addresses. • ipAndPort: The source IP and destination IP addresses, and Layer 4 source port and destination port numbers. • fiveTuple: The source IP and destination IP addresses, source port and destination port numbers, and protocol field in the IP header. • gtpuTeid The GTP-u tunnel identifier (ID). NOTE: There is no inner or outer field location for GTPU-TEID. • greFlowid: The flow identifier location. NOTE: The default field location value is outer. When configuring 5G-Cloud Ericsson SCP Transparent, greFlowid hash field should be selected. Refer to 5G-Cloud Ericsson SCP Support.
Field Location	• Outer: The first occurrence of header or field. For example, IP Only outer is the first IP header in the packet, which could be IPv4 or IPv6. • Inner: The second occurrence of header or field. The supported IP encapsulation types are: IP-in-IP, VXLAN, GTP, GRE, and ERSPAN.
Load balancing groups	Add or remove an application with the Endpoint ID and Weight value (1-100). A load balancing group can have minimum of two endpoints. Endpoint with higher weight receives more traffic.

4. Click **Save**.

NOTE: When you configure the Load Balancing Application in enhanced mode, you can associate it with only a single enhanced load balancing profile. However, you have the flexibility of changing the association to different profile as needed.

Enhanced Load Balancing

Enhanced Load Balancing redistributes the traffic from a failed endpoint to another endpoint which is a part of the ELB group. When the failed endpoint recovers, the redistributed traffic is restored to the recovered endpoint. The traffic across other endpoints remain undisturbed during this process. Hashing is performed on the packets based on outer and/or inner headers.

To enable to enhanced load balancing application:

1. Create an Enhanced Load Balancing Profile.
 - a. In the Monitoring domain page, click Settings, then click **Enhanced Load Balancing**
 - b. Click **New** and then create a profile.

Fields	Values
Hash fields	The various hash options are - ip - ip-src - ip-dst - l4-port - l4-portsrc - l4-portdst - gtpoteid
Position	Select either inner/outer location of the hash field to be matched with the incoming packet..
Hash-mask	When the hash-fields ip-src/ip-dst are defined along with the IP, then ip-src/ip-dst hash-mask will overwrite the IP hash mask. For load balancing, you must apply this mask before hashing the IP. The default options and their default values are: - IPv4 : 255.255.255.255 - IPv6 :FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF:FFFF

2. Go to monitoring session page and
 - a. Drag and drop **Load Balancing** from **APPLICATIONS** to the graphical workspace.
 - b. Click the load balancing application and select **Details**. The Application quick view appears.

× **Load Balancing: lb** Save

DETAILS THRESHOLDS HEALTH STATUS STATISTICS

Application Load Balancing

Alias* lb

☐ Stateless
☒ Enhanced Load balancing

ELB profiles* Select... ▼

Load balancing groups

Application Endpoint ID	2	Weight	1	⊕ ⊖
Application Endpoint ID	3	Weight	1	⊕ ⊖

⏪

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

Metric	Description
Alias	Enter a name for the load balancing application
Enhanced Load Balancing	Select this option to enable Enhanced Load Balancing and select the ELB profile.
Load balancing groups	Add or remove an application with the Endpoint ID and Weight value (1-100). A load balancing group can have minimum of two endpoints.

3. Click **Save**.

NOTE: When you configure the Load Balancing Application in enhanced mode, you can associate it with only a single enhanced load balancing profile. However, you have the flexibility of changing the association to different profile as needed.

What's Next

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

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

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

Masking

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

To add a Masking application:

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

masking



X

Application

Save

Application

Masking

Alias

masking

Protocol

none

Offset

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 grid of application icons, where the 'Slicing' icon (a blue hexagon with a red gear) is selected. The main panel has a header bar with a close button 'X' and a 'Save' button. Below the header, the application details are displayed in a table-like format.

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.

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
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_	RSA	RSA	CAMELLIA_128_	128	SHA

Cipher Name	Kx	Au	Enc	Bits	Mac
SHA			CBC		
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 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**.



Field	Value
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**.

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.

Configure Precryption in UCT-V

GigaVUE-FM allows you to turn on or off the Precryption feature for a monitoring session.

To enable or disable the Precryption feature in UCT-V, refer to Create monitoring session.

Rules and Notes

- To avoid packet fragmentation, change the option precryption-path-mtu in the UCT-V configuration file (**/etc/uctv/uctv.conf**) within the range 1400-9000 based on the platform path MTU.
- Protocol version IPv4 and IPv6 are supported.
- Using IPv6 tunnels requires GigaVUE-FM and the fabric components version 6.6.00 or above.

To create a new monitoring session with Precryption, follow these steps:

1. On the left pane in GigaVUE-FM, select **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
The **Monitoring Sessions** page appears.
2. Select **New** to open the **Create a New Monitoring Session** page.
3. Enter the appropriate information for the monitoring session:
 - a. In the **Alias** field, enter the name of the monitoring session.
 - b. In the **Monitoring Domain** field, enter the name of the monitoring domain that you want to select.
 - c. In the **Connection** field, enter the desired connection(s) to include as part of the monitoring domain. You can select the connections required for the monitoring domain.
4. Select **Next**. The **Edit Monitoring Session** page appears with the new canvas.
5. Select **Options** button. The Monitoring Session options appear.
6. Select the **Precryption** tab.
7. Enable **Precryption**.
8. Select **Save**. The **Edit Monitoring Session** page appears. You can proceed to create map, tunnels, and add applications.

NOTE: We recommend enabling the secure tunnel feature whenever the Precryption feature is enabled. Secure tunnel helps to securely transfer the cloud-captured packets or precrypted data to a GigaVUE V Series Node. For more information, refer to Secure Tunnel .

Validate Precryption connection

To validate the Precryption connection, follow these steps:

- Navigate to the **Monitoring Session** dashboard and check the Precryption option. The **yes** status indicates an active state.
- Select **Status** to view the rules configured.

Limitations

During precryption, the agent generates a TCP message and captures the payload in clear text. It probes the SSL connect and accept APIs to extract Layer 3 and Layer 4 (L3/L4) details from the packet. When the agent receives the SSL data on a specific interface, it sets the default gateway's MAC address as the destination MAC address for the TCP packet. If the gateway is misconfigured, the agent sets the destination MAC address to all zeros.

Monitor and Troubleshoot

Monitor Cloud Health

GigaVUE-FM allows you to monitor the traffic and configuration health status of the monitoring session and its individual components. This section provides detailed information on how to view the traffic and configuration health status of the monitoring session and its individual components. Refer to the following topics for more detailed information on configuration health, traffic health and how to view the health status:

- [Configuration Health Monitoring](#)
- [Traffic Health Monitoring](#)
- [View Health Status](#)

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 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 5. Tx Dropped 6. Rx Dropped 7. Tx Errors 8. Rx Errors	1. Difference 2. Derivative	1. Over 2. Under
Map	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Slicing	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Masking	1. Tx Packets 2. Rx Packets 3. Packets Dropped	1. Difference 2. Derivative	1. Over 2. Under
Dedup	1. Tx Packets	1. Difference	1. Over

	2. Rx Packets 3. Packets Dropped	2. Derivative	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 6. Ingestor - Rx Octets 7. Ingestor - Octets Dropped 8. Ingestor - Records Dropped 9. Workload - Records Dropped 10. Workload - Req Auth Errors 11. Workload - Req Timedout Errors 12. Workload - Req Errors 13. Exporter - Avg File Size	1. Difference 2. Derivative	1. Over 2. Under

	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.

- 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.

- 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:

- On the left pane in GigaVUE-FM, select **Traffic > Virtual > Orchestrated Flows** and select your cloud platform. The **Monitoring Session** page appears.
- In the **TRAFFIC PROCESSING** tab, select **Options>Thresholds** menu.
- From the **Select Template** drop-down list, select the template you wish to apply across the Monitoring Session.
- 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

Operational events for Enrichment:

Refer below for message format and messages that indicate common issues that can occur during the operations:

Format: <Operation Type>_<Message>

Operation Types: GETSERVICES, GETPODS, GETNODES, GETENDPOINTS, WATCHALL

Message	Description
K8S_AUTHORIZATION_FAILURE	The request was denied due to insufficient permissions Example: GETPODS_K8S_AUTHORIZATION_FAILURE
K8S_AUTHENTICATION_FAILURE	Authentication failed. Verify your credentials Example: GETPODS_K8S_AUTHENTICATION_FAILURE
K8S_UNHANDLED_ERROR	An unspecified error occurred. Check the error description Example: GETPODS_K8S_UNHANDLED_ERROR

View Health Status for Individual GigaVUE V Series Nodes

You can also view the health status of an individual GigaVUE V Series Node. To view the configuration health status and traffic health status of the V Series Nodes:

1. On the Monitoring Session page, select the required Monitoring Session from the list view.
2. In the **Overview** tab, view the health status of the required GigaVUE V Series Node from the chart options.

View Application Health Status for Individual V Series Nodes

To view the application configuration and traffic health status of the GigaVUE V Series Nodes:

1. On the Monitoring Session page, select the required Monitoring Session from the list view.
2. On the Overview tab, select the GigaVUE V Series Node from the All V Series Nodes drop-down menu.

The list view displays the list of applications for the selected GigaVUE V Series Node and the health status of each application.

View Monitoring Session Statistics

The Monitoring Session **OVERVIEW** page lets you analyze the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis.

You can view the high level information of the selected Monitoring Session such as, connections, tunnel details, health status, deployment status, and information related to Application Intelligence statistics. You can view the detailed statistics of an individual traffic processing element in the **TRAFFIC PROCESSING** tab.

You can view the statistics by applying different filters as per the requirements of analyzing the data. GigaVUE-FM allows you to perform the following actions on the Monitoring Session Statistics page:

- You can view the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis.
- You can filter the traffic and view the statistics based on factors such as **Incoming, Outgoing, Ratio (Out/In), Incoming Packets, Outgoing Packets, Ratio (Out/In) Packets**. You can select the options from the drop-down list box in the **TOTAL TRAFFIC** section of the **OVERVIEW** page.
- You can also view the statistics of the Monitoring Session deployed in the individual V Series Nodes. To view the statistics of the individual GigaVUE V Series Node, select the name of the **V Series Node** for which you want to view the statistics from the GigaVUE V Series Node drop-down list on the bottom left corner of the **OVERVIEW** page.

Visualize the Network Topology (Third Party Orchestration)

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 Third Party Orchestration

You can perform the following administrative tasks in GigaVUE-FM for GigaVUE Cloud Suite for AWS:

- [Configure Third Party Orchestration Settings](#)
- [Role Based Access Control](#)

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**.

Interface Mapping

You can remap interfaces for individual GigaVUE V Series Nodes within a Monitoring Session.

Note: When using Raw and Tunnel In, Interface Mapping is mandatory before you deploy the Monitoring Session.

To perform interface mapping,

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
The **Monitoring Sessions** landing page appears.
2. Navigate to the **V SERIES NODES** tab and select **Interface Mapping**.
The **Deploy Monitoring Session** dialog box appears.
3. Select the GigaVUE V Series Nodes to which you wish to map the interface.

4. From the drop-down menu of the GigaVUE V Series Nodes, select the interfaces for the following deployed in the Monitoring Session:
 - REPs (Raw Endpoints)
 - TEPs (Tunnel Endpoints)
5. Select **Deploy**.


NOTE: The updated mappings take effect when deployed.

Configure Third Party Orchestration Settings

This topic allows you to set the maximum number of connections, adjust refresh intervals for instance and non-instance inventory, and define the maximum batch size for monitoring session updates.

To configure these settings,

1. In GigaVUE-FM, go to **Inventory > VIRTUAL > Third-Party Orchestration**.
2. Select **Settings**.
3. In the **Settings** page, select the **Advanced** tab and update the required values.



Settings	Description
Refresh interval for instance target selection inventory (secs)	Specifies the frequency for updating the state of the instances.
Refresh interval for fabric deployment inventory (secs)	Specifies the frequency for deploying the fabric components
Number of UCT-Vs per V Series Node	Specifies the maximum number of UCT-Vs that you can assign to the GigaVUE V Series Node.

Settings	Description
	Points to Note: - At the time of GigaVUE V Series Node deployment, GigaVUE-FM maps the UCT-Vs equally to all the available GigaVUE V Series Node. For example, there are 10 UCT-Vs, two GigaVUE V Series Nodes, and 10 f UCT-Vs per V Series Node. Even though the first GigaVUE V Series Node can accommodate all the ten UCT-Vs, it is shared equally between the two GigaVUE V Series Node to balance the traffic load. - When a new GigaVUE V Series Node is added to the deployment, GigaVUE-FM does not re-balance the UCT-Vs to the new GigaVUE V Series Node unless the number of UCT-Vs mapped to the GigaVUE V Series Node is greater than the Number of UCT-Vs per V Series Node.
Refresh interval for UCT-V inventory (secs)	Specifies the frequency for discovering the UCT-Vs available. Note: When you upgrade to version 6.5 or above, GigaVUE-FM resets the UCT-V Refresh Interval field to its default. To retain your custom value, you must manually reconfigure it.
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.

Settings	Description
Traffic distribution tunnel MTU	Specifies the MTU value for the traffic distribution tunnel.
Use UCT-V conf file	Enable this option to allow interface mirroring to follow the configuration defined in the file. Disable it to mirror traffic from all physical interfaces. Note: - When changing the UCT-V conf file option from enabled to disabled, ensure to undeploy the Monitoring Session and delete the Monitoring Domain. Once changed, you should create a new Monitoring Domain and configure the Monitoring Session. - When changing the UCT-V conf file option from disabled to enabled, do the following: - Edit the uctv.conf file Windows: C:\ProgramData\Uctv\uctv.conf Linux: /etc/uctv/uctv.conf - Delete the skipConf file from the backup folder Windows: C:\ProgramData\Uctv\bak\skipConf Linux: /var/lib/uctv/bak/skipConf - Restart the UCT-V Windows: Restart from the Task Manager Linux: sudo service uctv restart
Reboot threshold limit for UCT-V Controller down	Specifies the number of times GigaVUE-FM tries to reach UCT-V Controller. When the UCT-V Controller moves to down state, GigaVUE-FM retries every 60 seconds.

About Audit Logs

Audit logs track the changes and activities that occur in the virtual nodes due to user actions. You can filter the logs to view specific information.

Access Audit Logs

Navigate to **Dashboard > SYSTEM > Audit Logs**. The **All Audit Logs** page appears.

All Audit Logs

Filter Manage

Filter : none

Time	User	Operation Type	Entity Type	Source	Device IP	Hostname	Status	Description	Tags	
2020-1...	admin	login fmUser ad...	User	fm			SUCCESS			
2020-1...	admin	logout fmUser a...	User	fm			SUCCESS			
2020-1...	admin	login fmUser ad...	User	fm			SUCCESS			
2020-1...	admin	update fmUser ad...	User	fm			SUCCESS			

Go to page: 1 of 16 Total Records: 106

Parameters

The Audit Logs have the following parameters:

Parameters	Description
Time	Provides the timestamp on the log entries.
User	Provides the logged user information.
Operation Type	Provides specific entries that the system logs. For example, - Log in and Log out based on users. - Create/Delete/Edit tasks, GS operations, maps, virtual ports, and so on.
Source	Provides details about the usage either in GigaVUE-FM or on the node when the event occurred.
Status	Success or Failure of the event.
Description	For failed status provides a brief update on the reason..

NOTE: Verify if the GigaVUE-FM time is set correctly to ensure accuracy of the captured trending data.

Filtering the audit logs

You can filter to view specific type of logs based on the following criteria:

- **When:** Displays logs that occurred within a specified time range.
- **Who:** Displays logs related to a particular user or users.
- **What:** Displays logs for one or more operations, such as Create, Read, and Update.
- **Where:** Displays logs for GigaVUE-FM or devices.
- **Result:** Displays logs for success or failure.

To filter the audit logs,

1. Select **Filter**.

A quick view for Audit Log Filters displays.

2. Specify one or all of the following:

- **Start Date** and **End Date** to display logs within a specific time range.
- **Who** limits the scope of what displays on the Audit Logs page to a specific user or users.
- **What** narrows the logs to the types of operation that the log is related to. You can select multiple operations. Select **All Operations** to apply all operation types as part of the filter criteria.
- **Where** narrows the logs to particular of system that the log is related to. Select **All Systems** apply both GigaVUE-FM and device to the filter criteria. **Result** narrows the logs related to failures or successes. Select **All Results** to apply both success and failure to the filter criteria.

3. Select **OK** to apply the selected filters to the **Audit Logs** page.

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 GigaVUE Cloud Suite for Third Party Orchestration](#)
- [Deploy UCT-C Solution in Kubernetes](#)
- [Deploy GCB Controller Service and Pods](#)

Supportability Information

The GigaVUE-FM HA feature is supported on the following platforms:

Orchestration Method	Platforms/components Supported
GigaVUE-FM Orchestration	• AWS • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix
Third-party Orchestration	Generic Mode • AWS • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix • GCP • UCT-C • GCB Integrated Mode • AWS • Azure • OpenStack

Recommendations, Rules, Notes, and Limitations

Keep in mind the following recommendations, rules and notes when you configure the GigaVUE-FM HA feature:

Recommendations

- Do not configure GigaVUE-FM High Availability (HA) across multiple regions or multi-cloud environments. This setup is not recommended and may lead to instability.
- You can deploy HA across multiple availability zones within the same cloud region. Ensure that you deploy a GigaVUE-FM instance in each availability zone.
- Use the orchestrated upgrade procedure to upgrade the GigaVUE-FM instances.
- It is recommended to configure reverse DNS configuration while configuring FMHA with FQDN.
- When creating a GigaVUE-FM HA group, it is highly recommended to create the High Availability group using DNS name. GigaVUE-FM HA cannot manage the individual GigaVUE-FM instances if the IP address changes (if GigaVUE-FM IP(s) are dynamic).

Rules and Notes

- GigaVUE-FM cannot support more than 7 active-eligible nodes (this is due to the underlying constraint in MongoDB).
- The GigaVUE-FM instances in the HA group must be identical in terms of system configuration such as hard disk, memory, and network interfaces, which include domain server, NTP server, and name server. Domain server, NTP server, and name server should be configured in all GigaVUE-FM instances.
- The GigaVUE-FMs in the high availability group can be accessed using the IPv4/IPv6 address (DNS name) that is used to form the High Availability group.

NOTE: Do not access all the GigaVUE-FM instances at the same time, as this will impact the performance of the HA group.

- You can deploy the GigaVUE-FM HA virtual machines on a WAN link with a maximum latency of 200 ms.
- You cannot add a GigaVUE-FM Hardware Appliance and a GigaVUE-FM virtual machine in the same HA group.
- The **Reload All** Or **Reload any FM** commands will work only if there is an active GigaVUE-FM instance in the HA group.
- You must ensure that NTP is configured across all FM HA nodes.
- The GigaVUE-FM HA Cluster functions like standalone GigaVUE-FM's when a fail-over occurs during long-running operations, such as Fabric Launch, Monitoring Session Deployment, or Fabric Upgrade.
- When FMHA is configured with a reverse proxy, access to GigaVUE-FM requires non-SAML authentication, such as local authentication. Use the following sample URL to access GigaVUE-FM:
`https://<proxy_ip>/admin`
- If GigaVUE-FM has multiple IPv6 addresses, ensure FMHA uses the permanent IPv6 as the source IP. GigaVUE-FM excludes temporary IPv6 addresses from the IP set, which causes connection failures and FMHA formation issues.

Shrinking the size of the HA group:

- When reducing the size of a HA group that has both active-eligible and supporting instances, it is always recommended to remove the supporting instances first so that the active-eligible instances remain intact in the HA group. If active-eligible instances are removed first from the HA group, then the HA group will become unstable. This, however, depends on the number of active-eligible and supporting instances in the HA group. Consider a GigaVUE-FM HA group with 7 instances. The number of active and supporting instances in a HA group may be as follows:

HA group 1	With 3 active instances and 4 supporting instances	The recommendation is specifically applicable for HA group 1 in which there are only 3 active eligible instances in the group. Removing the active-eligible instances may cause the HA group to become unstable.
HA group 2	With 5 active instances and 2 supporting instances.	The recommendation is not applicable for HA group 2 as removing the active-eligible instances will not impact the stability of the HA group.

For a HA group to be stable, it must have a minimum of two active-eligible instances in the group.

- When reducing the size of HA group from 7 instance to 5 instance or 3 instance, the quorum is re-computed and adjusted accordingly. This is to ensure that the HA group remains intact with the available set of instances.
- When reducing the size of the HA cluster, during removal of an instance make sure to give time for the OpenSearch to rebalance and return to a complete (100%) healthy state before removing the next instance. If you remove the instances subsequently one after the other, the health state of the HA group will turn red as there will be no time for shard rebalancing amongst the available instances. Login to GigaVUE-FM CLI as a root user and execute the following command to check the ES cluster health:

```
curl -XGET "http://localhost:9200/_cluster/health?pretty"
```

Supporting GigaVUE-FM Instances:

- When a supporting instance is added to the HA group, its system information will not be available immediately and will be available only after a minute.
- When the HA group is upgraded, application services of the supporting instances will be turned on automatically for the upgrade and once GigaVUE-FM HA upgrade is completed it will be turned off automatically. Application services will be turned off automatically irrespective of the upgrade completion status, that is, either success or failure.



Notes:

- For supporting instances, system information will be populated into MongoDB as the Application Service is down by default (scheduler populates the system information every minute into MongoDB). Hence any changes in the system hostname or GigaVUE-FM version will be reflected in the GUI only after a minute.
- In case of suspended mode, MongoDB will go into write protected mode. Therefore, any system related updates performed in the supporting nodes will not be reflected in the GUI.
- If the supporting instance is not reachable, all system related information that is populated in the GUI reflects the previous known values and not the latest values for Application Status, Up Time, and other system related information.
- Whenever the application services are up and running, events will be logged in the Events page (with the severity level - Info)

Licensing Information

You must install a Prime license on the active GigaVUE-FM instance to configure a High Availability group.

If the Prime license expires or if you accidentally delete the license, the existing configurations will still be present in the GigaVUE-FMs that are part of the HA group, but you will not be able to perform any new configurations. Moreover, if you disassemble the HA group, you cannot reconfigure the HA group without installing a valid Prime license. You can re-use the same Prime license to form the HA group again.

Hostname Setups

The GigaVUE-FM instances are not required to be in the same subnet, but still must be able to communicate with each other.

In addition, ensure that the instances have unique host names. You must be able to ping a GigaVUE-FM instance from the other instances using the hostname or the IP address.

To add a GigaVUE-FM instance to a GigaVUE-FM HA group:

- The GigaVUE-FM instances in the HA group must be reachable to each other.
- If host names are used to configure the HA group, the host names of the GigaVUE-FM instances must be resolvable through a DNS server.

Configure GigaVUE-FM HA

To configure the GigaVUE-FM HA feature, you must have access to at least three authenticated GigaVUE-FM instances that reside on a trusted network. All GigaVUE-FM instances must run the same software version. The interfaces in the GigaVUE-FM instances must be up and must be assigned IPv4/IPv6 addresses. You can also choose to use DNS host names.

NOTE: You can configure the GigaVUE-FM HA feature only if you have administrative privileges. You must have Prime license installed for forming the GigaVUE-FM HA group.


Refer to the following sections:

- [Create and Add GigaVUE-FM Instances to HA Group](#)
- [Connect GigaVUE-FM High Availability to a Monitoring Domain using GigaVUE-FM Orchestration](#)
- [Connect GigaVUE-FM High Availability to a Monitoring Domain using Third Party Orchestration](#)
- [Integrate GigaVUE-FM High Availability with UCT-C](#)
- [Integrate GigaVUE-FM High Availability with Gigamon Containerized Broker](#)

Create and Add GigaVUE-FM Instances to HA Group

To add instances and to configure the HA group:

1. Go to  High Availability. The launch page displays the required prerequisites to create a HA group.



2. Click **Create**. In the **Group Name** field, enter a unique name for the HA group.

NOTE: Avoid using functionality names as the group alias when creating an FMHA group. Functionality names typically refer to actions or operations within the system and can cause confusion and conflicts. Choose a unique and descriptive name instead.

3. Choose any one of the following options to encrypt the communication between nodes:
 - **Pre-shared Secret Key** – Encrypts communication between GigaVUE-FM nodes using an application-managed shared secret key. Refer to [Encryption of Communication between FMHA Nodes](#).
 - **Certificate Based** – Encrypts communication between GigaVUE-FM nodes using manually uploaded certificates. Refer to [Encryption of Communication between FMHA Nodes](#).
4. In the **Add GigaVUE-FM Instance** section, the details of the GigaVUE-FM instance from which the HA group is created will be fetched automatically.

Edit the Primary Node

1. On the first instance, click the ellipses icon and select **Edit Node**.
2. Enter the following details: (These fields are optional)
 - Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
 - **Tunnel IP Address/DNS Name:** IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.
 - **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name:** If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
 - Enter the Entity ID and the IdP Metadata URL.
3. Click the **save** icon to save the configuration. Refer to [Additional Fields](#) for details about other fields displayed on the HA page.

High Availability

Create New

There must be atleast three active eligible nodes that are reachable to each other.

Group Name * FMHA1

Secure Communication ☒ Pre-Shared Key ☐ Certificate Based

Add GigaVUE-FM Instances

IP ADDRESS / DNS NAME	TUNNEL IP ADDRESS / D...	EXTERNAL IP ADDRESS / ...	ROLE	USERNAME	PASSWORD	ENTITY ID	IDP METADATA URL	
fcab0:1:2200:96ae:3c8...	—	—	Active Eligible	admin	*****	—	—	...

Add New Node

1. Click **Add New Node** to add the GigaVUE-FM instances. Enter the following details:.
 - **IP Address/DNS Name** - IP address/DNS name of the GigaVUE-FM instances (for communication between the GigaVUE-FMs and to configure the HA group). It is always recommended to use the DNS name while creating the GigaVUE-FM High Availability group.
 - (Optional) Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
 - **Tunnel IP Address/DNS Name:** IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.
 - **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name:** If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
 - When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.




Notes:

- When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.
- When using unique DNS for HA cluster configuration, ensure it always resolves to only one IP address—not multiple IPs.

- **Role** - Role of the GigaVUE-FM instance. Can be configured as below. Refer to [Active-Eligible instances](#) for details.
 - **Active Eligible**- provides resiliency and load distribution.
 - **Supporting Node**- provides scalability and performance.
 - **Username** - GigaVUE-FM GUI Username.
 - **Password** - GigaVUE-FM GUI password.
 - (Optional) **Entity ID** - Unique alphanumeric ID assigned to the GigaVUE-FM instance
 - (Optional) **IdP Metadata URL** - URL pointing to the IdP's XML configuration
2. Click **Save** to validate the details entered for the instances.

Use the **Add New Node** button to continue to add the GigaVUE-FM HA instances. You can click the **Edit Node** icon to change the external IP address. After adding the instances to the HA group, click **Submit**. The HA group is configured and will be displayed as follows. Refer to [Additional Fields](#) for details about other fields displayed on the HA page.



High Availability

Healthy
This HA group is Healthy. The members are protected from failure.

No supporting nodes in cluster ⓘ
Group Name: fmha
No of Nodes: 3
Encryption Authentication Method: Pre-Shared Key ⓘ
Tunnel Status: All tunnels are healthy

+ Add New Node Delete HA Group

STATUS	IP ADDRESS / D...	TUNNEL IP ADD...	ROLE	EXTERNAL IP A...	ENTITY ID	HOST NAME	APPLICATION STA...	SOFTWARE VERSI...	SYSTEM UPTIME	IDP METADATA ...	ⓘ
Standby		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...
Standby		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...
Active		—	Active Eligible			GigaVUE-FM-6...	Up		2 days, 21 hours...		...

Go to page: 1 of 1

Note: Please click Refresh to obtain the latest data.



Notes:

- When switching between the Pre-shared Secret Key and Certificate-Based options for encrypting communication between nodes, it's important to refresh the GigaVUE-FM to see the updated tunnel status. By clicking the link in the **Tunnel Status**, GigaVUE-FM displays the health status of the FMHA nodes.
- When creating a HA group, all the existing data on the first configured GigaVUE-FM instance will be preserved, allowing the system to continue using the data after it is fully set up. However, the data on the subsequent nodes will be lost during the setup.

The HA page also displays the following details:

Table 3: Additional Fields

Field	Description
Status and Reachability	Status of the GigaVUE-FM instances. Can be: - Active - Standby - Supporting Refer to Active-Eligible instances for details. Reachability: 🟢 - indicates that the instances are reachable. ⚠️ - indicates that the instances are unreachable.
Host Name	Host name of the GigaVUE-FM instance.
Application Status	Status of application services that are responsible for handling the user requests: Up: Indicates service is ready to handle the user requests

Field	Description
	• 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.

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

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 $n/2+1$).

		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_CNTLRL_EXT_IP_DNS
    value: "10.xxx.xx.xx" (IP address of the DNS server - external to the
    Kubernetes cluster)
  - name: FM_FQDN
    value: "fm.myorg.com" (FMs FQDN that is used for DNS lookups)
```

NOTE: The FQDN identifies either a standalone GigaVUE-FM or an FMHA cluster that includes multiple GigaVUE-FM instances. It returns the IP addresses (IPv4, IPv6, or both) for all GigaVUE-FMs associated with that FQDN.

For more details on the GigaVUE-FM High Availability configuration, refer to the *GigaVUE Administration Guide* 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.
7. 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 login to any of the GigaVUE-FM instances of the HA group, the dashboard page appears. Use the drop-down option in the header of GigaVUE-FM, which is available on specific pages, to view the details pertaining to that GigaVUE-FM. For example, the following pages in the GigaVUE-FM GUI have drop-down option in the header:

- FM Health
- IP Resolver
- [NTP Server for Clock Synchronization](#)

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:

1. Go to  High Availability.
2. Click the ellipsis on the standby GigaVUE-FM instance widget in the High Availability page as shown in [Figure 1Remove Standby GigaVUE-FM Instance](#).




Figure 1 Remove Standby GigaVUE-FM Instance

3. 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.

State	Number of GigaVUE-FM Instances	Description
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 deployment/upgrade during HA leader switchover	The newly elected leader resumes certificate handling and continues any in-progress fabric deployment or upgrade from the last known state. No manual intervention is required.

GigaVUE-FM High Availability Scenarios

The High Availability page displays the current state of the GigaVUE-FM HA group. When a failover occurs, the HA group state changes in the GUI.

The following table lists the GUI changes for the various scenarios:

Scenario	Changes in GUI
What happens to the High Availability page immediately after a failover?	The High Availability page may not update immediately or may not show all the GigaVUE-FM instances. NOTE: Refresh the page after a minute to view the new active GigaVUE-FM instance. However, the page will get updated automatically after 5 minutes.
What happens to an active GigaVUE-FM instance when it	One of the standby GigaVUE-FM instances changes to the active state. The GigaVUE-FM instance that was initially in the active state

Scenario	Changes in GUI
fails (either by itself or if failover is triggered manually)?	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)

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**.

Additional Sources of Information

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

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

Documentation

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

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

Table 1: Documentation Set for Gigamon Products

GigaVUE Cloud Suite 6.14 Hardware and Software Guides	
DID YOU KNOW? If you keep all PDFs for a release in common folder, you can easily search across the doc set by opening one of the files in Acrobat and choosing Edit > Advanced Search from the menu. This opens an interface that allows you to select a directory and search across all PDFs in a folder.	
Hardware	
how to unpack, assemble, rackmount, connect, and initially configure ports the respective GigaVUE Cloud Suite devices; reference information and specifications for the respective GigaVUE Cloud Suite devices	
GigaVUE-HC1 Hardware Installation Guide	
GigaVUE-HC3 Hardware Installation Guide	
GigaVUE-HC1-Plus Hardware Installation Guide	
GigaVUE-HCT Hardware Installation Guide	
GigaVUE-TA25 Hardware Installation Guide	
GigaVUE-TA25E Hardware Installation Guide	
GigaVUE-TA100 Hardware Installation Guide	
GigaVUE-TA200 Hardware Installation Guide	
GigaVUE-TA200E Hardware Installation Guide	

GigaVUE Cloud Suite 6.14 Hardware and Software Guides
GigaVUE-TA400 Hardware Installation Guide
GigaVUE-TA400E Hardware Installation Guide
GigaVUE-OS Installation Guide for DELL S4112F-ON
G-TAP A Series 2 Installation Guide
GigaVUE M Series Hardware Installation Guide
GigaVUE-FM Hardware Appliances Guide
Software Installation and Upgrade Guides
GigaVUE-FM Installation, Migration, and Upgrade Guide
GigaVUE-OS Upgrade Guide
GigaVUE V Series Migration Guide
Fabric Management and Administration Guides
GigaVUE Administration Guide covers both GigaVUE-OS and GigaVUE-FM
GigaVUE Fabric Management Guide how to install, deploy, and operate GigaVUE-FM; how to configure GigaSMART operations; covers both GigaVUE-FM and GigaVUE-OS features
GigaVUE Application Intelligence Solutions Guide
GigaVUE Inline Solutions Guide(NEW) (previously included in the GigaVUE Fabric Management Guide)
Cloud Guides how to configure the GigaVUE Cloud Suite components and set up traffic monitoring sessions for the cloud platforms
GigaVUE V Series Applications Guide
GigaVUE Cloud Suite Deployment Guide - AWS
GigaVUE Cloud Suite Deployment Guide - Azure
GigaVUE Cloud Suite Deployment Guide - OpenStack
GigaVUE Cloud Suite Deployment Guide - Nutanix
GigaVUE Cloud Suite Deployment Guide - VMware (ESXi)
GigaVUE Cloud Suite Deployment Guide - VMware (NSX-T)
GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
Universal Cloud TAP - Container Deployment Guide

GigaVUE Cloud Suite 6.14 Hardware and Software Guides	
Gigamon Containerized Broker Deployment Guide	
GigaVUE Cloud Suite Deployment Guide - AWS Secret Regions	
GigaVUE Cloud Suite Deployment Guide - Azure Secret Regions	
Reference Guides	
GigaVUE-OS CLI Reference Guide library of GigaVUE-OS CLI (Command Line Interface) commands used to configure and operate GigaVUE HC Series and GigaVUE TA Series devices	
GigaVUE-OS Security Hardening Guide	
GigaVUE Firewall and Security Guide	
GigaVUE Licensing Guide	
GigaVUE-OS Cabling Quick Reference Guide guidelines for the different types of cables used to connect Gigamon devices	
GigaVUE-OS Compatibility and Interoperability Matrix compatibility information and interoperability requirements for Gigamon devices	
GigaVUE-FM REST API Reference in GigaVUE-FM User's Guide samples uses of the GigaVUE-FM Application Program Interfaces (APIs)	
Factory Reset Guidelines for GigaVUE-FM and GigaVUE-OS Devices Sanitization guidelines for GigaVUE Fabric Management Guide and GigaVUE-OS devices.	
Release Notes	
GigaVUE-OS, GigaVUE-FM, GigaVUE-VM, G-TAP A Series, and GigaVUE Cloud Suite Release Notes new features, resolved issues, and known issues in this release ; important notes regarding installing and upgrading to this release Note: Release Notes are not included in the online documentation. Note: Registered Customers can log in to My Gigamon to download the Software and Release Notes from the Software and Docs page on to My Gigamon . Refer to How to Download Software and Release Notes from My Gigamon .	
In-Product Help	
GigaVUE-FM Online Help how to install, deploy, and operate GigaVUE-FM.	

How to Download Software and Release Notes from My Gigamon

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

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

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

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

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

Documentation Feedback

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

documentationfeedback@gigamon.com

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

Documentation Feedback Form		
About You	Your Name	
	Your Role	
	Your Company	
For Online Topics	Online doc link	<i>(URL for where the issue is)</i>
	Topic Heading	<i>(if it's a long topic, please provide the heading of the section where the issue is)</i>

For PDF Topics	Document Title	<i>(shown on the cover page or in page header)</i>
	Product Version	<i>(shown on the cover page)</i>
	Document Version	<i>(shown on the cover page)</i>
	Chapter Heading	<i>(shown in footer)</i>
	PDF page #	<i>(shown in footer)</i>
How can we improve?	Describe the issue	<i>Describe the error or issue in the documentation. (If it helps, attach an image to show the issue.)</i>
	How can we improve the content? Be as specific as possible.	
	Any other comments?	

Contact Technical Support

For information about Technical Support: Go to **Settings**  **> Support > Contact Support** in GigaVUE-FM.

You can also refer to <https://www.gigamon.com/support-and-services/contact-support> for Technical Support hours and contact information.

Email Technical Support at support@gigamon.com.

Contact Sales

Use the following information to contact Gigamon channel partner or Gigamon sales representatives.

Telephone: +1.408.831.4025

Sales: inside.sales@gigamon.com

Partners: www.gigamon.com/partners.html

Premium Support

Email Gigamon at inside.sales@gigamon.com for information on purchasing 24x7 Premium Support. Premium Support entitles you to round-the-clock phone support with a dedicated Support Engineer every day of the week.

The VÜE Community

[The VÜE Community](#) is a technical site where Gigamon users, partners, security and network professionals and Gigamon employees come together to share knowledge and expertise, ask questions, build their network and learn about best practices for Gigamon products.

Visit the VÜE Community site to:

- Find knowledge base articles and documentation
- Ask and answer questions and learn best practices from other members.
- Join special-interest groups to have focused collaboration around a technology, use-case, vertical market or beta release
- Take online learning lessons and tutorials to broaden your knowledge of Gigamon products.
- Open support tickets (Customers only)
- Download the latest product updates and documentation (Customers only)

The VÜE Community is a great way to get answers fast, learn from experts and collaborate directly with other members around your areas of interest.

Register today at community.gigamon.com

Questions? Contact our Community team at community@gigamon.com.