



GigaVUE Cloud Suite for Azure - Deployment Guide

GigaVUE Cloud Suite

Product Version: 6.14.xx

Document Version: 2.0

(See Change Notes for document updates.)

Copyright © 2026 Gigamon Inc. All rights reserved.

Information in this document is subject to change without notice. The software described in this document is furnished under a license agreement or nondisclosure agreement. No part of this publication may be reproduced, transcribed, translated into any language, stored in a retrieval system, or transmitted in any form or any means without the written permission of Gigamon Inc.

Trademark Attributions

Gigamon and the Gigamon logo are trademarks of Gigamon in the United States and/or other countries. Gigamon trademarks can be found at www.gigamon.com/legal-trademarks. All other trademarks are the trademarks of their respective owners.

Gigamon Inc.
3300 Olcott Street
Santa Clara, CA 95054
408.831.4000

Change Notes

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

Product Version	Document Version	Date Updated	Change Notes
6.14.01/6.14.02	1.0	07/10/2026	The original release of this document with 6.14.02 GA. This release incorporates changes in the following topics for the listed features: - Port Optimization Subnet and Security Group for Amazon VPC
6.14.00	1.0	07/10/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Cloud Suite for Azure - Deployment Guide	1
Change Notes	3
Contents	4
GV-Cloud-Azure	11
Get Started	11
Overview of GigaVUE Cloud Suite for Azure	11
GigaVUE Fabric Components	13
GigaVUE-FM Orchestrated Fabric Deployment	13
Third Party Orchestrated Fabric Deployment	16
Key Concepts	18
Monitoring Domain	19
Monitoring Session	19
Inline V Series (Azure)	19
Secure Communication between GigaVUE Fabric Components	22
Precryption™	24
Supported Kernel Version	29
Secure Tunnels	30
Prefiltering	31
Customer Orchestrated Source - Use Case	32
Check for Required IAM Permissions in Azure	32
Traffic Acquisition using Azure Virtual Network TAP	34
Deployment Overview	36
Deployment Planning	36
Identify the Deployment Method	37
Identify the Azure Regions	37
Define GigaVUE-FM Deployment Location	37
Identify the Traffic Acquisition Method	38
Identify the Azure Subscriptions Involved	38
Identify the VNets and Subnets Involved	38
Identify the Required Azure RBAC Permissions	39
Identify VPN Connectivity Requirements	39
Identify the UCT-V Deployment Details	39
Identify the UCT-V Controller Deployment Details	39
Deployment Prerequisites	39
Points to Note for GigaVUE Cloud Suite for Azure	40
Prerequisites for GigaVUE Cloud Suite for Azure	40

New Deployments	43
Existing Deployments	44
Migration to Reduced Control-Plane Ports	44
Network Firewall Requirement	62
GigaVUE FM	71
UCT-V Controller	74
UCT-V	75
GigaVUE V Series Node	76
Giga VUE V Series Proxy(Optional)	78
UCT-C Controller - deployed in Kubernetes worker mode	78
UCT-V Controller	79
GigaVUE V Series Node	79
GigaVUE V Series Proxy(Optional)	80
Licensing for GigaVUE Cloud Suite for Azure	80
Bundle Replacement Policy	83
View active Volume-Based License	85
View Inactive Volume-Based License	86
Activate Volume-Based Licenses	87
Supported Operating Systems for UCT-V	88
VPN Connectivity	88
Enable Subscription for GigaVUE Cloud Suite for Azure	89
Configure Tokens	90
Permissions and Privileges (Azure)	93
Assign a Custom Role using CLI	94
Assign a Custom Role using Azure Portal	95
Install GigaVUE-FM on Azure	101
Install GigaVUE-FM Using Azure VM Dashboard	101
Install GigaVUE-FM Using Azure Marketplace	101
Upgrade GigaVUE-FM in Azure	103
Upgrade GigaVUE-FM from the UI	103
Upgrade GigaVUE-FM using Snapshot in Azure	106
Upgrade GigaVUE Fabric Components in GigaVUE-FM for Azure	108
Prerequisite	109
Upgrade UCT-V Controller	109
Upgrade GigaVUE V Series Node and GigaVUE V Series Proxy	111
Deploy	114
Deploy GigaVUE Cloud Suite for Azure using GigaVUE-FM	
Orchestration	114
Acquire Traffic using UCT-V as Traffic Acquisition Method	115
Post Deployment Check:	124
Acquire Traffic using vTAP as Traffic Acquisition Method	140
Acquire Traffic using Customer Orchestrated Source as Traffic	
Acquisition Method	151
Acquire Traffic using Inline V Series Solution	160

Deploy GigaVUE Cloud Suite for Azure using Third Party	
Orchestration	181
Generic Mode	181
Integrated Mode	182
Deploy GigaVUE Fabric Components using Azure - Generic Mode	182
Deployment GigaVUE Fabric Components using Azure -	
Integrated Mode	182
GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud	184
GigaVUE Cloud Suite Cloud Suite in Azure Government	184
Linux UCT-V Installation	185
Linux UCT-V Installation Scenarios	188
Post Deployment Check:	192
Windows UCT-V Installation	193
Rules and Notes:	193
Windows UCT-V Installation Scenarios	196
Prerequisites	204
Create GigaVUE V Series Node with Management and Data NIC	
Attached using CLI	205
Create GigaVUE V Series Node with Management and Data NIC	
Attached using Azure GUI	205
Configure UCT-V Controller in Azure	206
Register UCT-V Controller during Virtual Machine Launch	206
Register UCT-V Controller after Virtual Machine Launch	207
Configure UCT-V in Azure	209
Configure GigaVUE V Series Node and GigaVUE V Series Proxy in	
Azure	210
Register GigaVUE V Series Node and GigaVUE V Series Proxy during	
Virtual Machine Launch	211
Register GigaVUE V Series Node and GigaVUE V Series Proxy after	
Virtual Machine Launch	212
What to do Next	213
GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud	213
GigaVUE Cloud Suite Cloud Suite in Azure Government	213
Prerequisites	219
Create GigaVUE V Series Node with Management and Data NIC	
Attached using CLI	220
Create GigaVUE V Series Node with Management and Data NIC	
Attached using Azure GUI	220
Configure GigaVUE V Series Node and GigaVUE V Series Proxy in	
Azure	221
Register GigaVUE V Series Node and GigaVUE V Series Proxy during	
Virtual Machine Launch	222
Register GigaVUE V Series Node and GigaVUE V Series Proxy after	
Virtual Machine Launch	222

What to do Next	224
Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform	224
License	225
Supportability	225
Rules and notes	226
Prerequisites	226
Get Started	226
Deployment Using Terraform	228
Install the Gigamon Terraform Provider Without Internet Access	229
Single-file configuration	231
Module-based configuration	232
Run the Terraform Deployment File	233
Best Practices	237
Verification and Troubleshooting	238
Configure	240
Configure Monitoring Session	240
Create a Monitoring Session (Azure)	240
Enable Prefiltering	243
Enable Precryption	244
Validate Precryption connection	245
Limitations	245
Apply Threshold Template	246
Enable User Defined Applications	246
Enable Distributed De-duplication	247
Tool Exclusion	247
Create Ingress and Egress Tunnels (Azure)	248
Create Raw Endpoint (Azure)	257
Create a New Map (Azure)	258
Add Applications to Monitoring Session (Azure)	263
What's Next	270
What's Next	272
Limitations	277
What's Next	277
Configure Application Visualization for Virtual Environment	280
Configure Filtering and Metadata Export for Selected Applications in	
Application Visualization	282
Filter Traffic for Selected Applications	282
Export Metadata for Selected Applications	283
View Application Statistics for Application Visualization	283
Supported Protocols and Attributes	284
Metadata	289
Supported RegExp Syntax	289
Limitations	290

Create Rules for User Defined Application 290

Configure Application Filtering Intelligence for Virtual Environment	292
View Application Statistics for Application Filtering	297
Configure Application Metadata Intelligence for Virtual Environment	298
Extend Your AMI and AMX Deployment	306
View Application Statistics for Application Metadata	307
Application Metadata Exporter Deployment Options	308
On-Premises	308
Hardware (AMI)	308
Hardware (Control Plane Metadata)	309
Private Cloud (VMware)	310
Public Cloud	311
Application Metadata Exporter - Export Scenarios	311
Export Application Metadata in JSON format over HTTPS/Kafka	311
Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	312
Supported Platforms	313
Important Notes	313
Export of 3G/4G/5G Control Plane Metadata by AMX	314
Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	315
Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	316
Export of GigaVUE Enriched Metadata for Cloud Workloads (Virtual Machines)	316
Export of GigaVUE Enriched Metadata for Cloud Workloads Kubernetes containers - Azure Kubernetes Service (AKS)	317
Identity Enrichment using Zscaler Private Access (ZPA)	318
Rules and Notes	320

Configure Metadata Export Filter for AMX 320

Rules and Notes	321
-----------------------	-----

**Deploy Application Metadata Exporter in Azure using
FM Orchestration 321**

Predeployment Checklist	322
Install GigaVUE-FM on Azure	324
Upload GigaVUE V Series Node Image into GigaVUE-FM	324
(Optional) Configure Metadata Export Filter for AMX	324
(Optional) Integrate Private CA	324
Rules and Notes	324
Generate CSR	325
Upload CA Certificate	325

Create Monitoring Domain to Acquire Traffic from Customer	
Orchestrated Source	326
Deploy V Series Node	326
Create a Monitoring Session to Configure Application Metadata Exporter	331
Configure Ingestor in AMX to Process Traffic	333
Scenario 1 - Export Application Metadata in JSON format over	
HTTPS/Kafka	333
Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks	
and Export over HTTPS/Kafka	333
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads	
and Export over HTTPS/Kafka	334
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and	
export over HTTPS/Kafka	335
Integrate AMX with Third Party Tool	335
Set up the Third Party Tool for AMX Integration	335
Configure Exporter to Export Attributes to Third Party tools	336
Configure Cloud Tool Exporter	336
Deploy the Monitoring Session	336
Cloud Tool Exporter - Field References	336
Configure Kafka Exporter	337
Deploy the Monitoring Session	337
Kafka Exporter - Field References	338
Verify Your Deployment	339
Deploy Application Metadata Exporter in Azure using	
 Third Party Orchestration	339
Predeployment Checklist:	340
(Optional) Install GigaVUE-FM on Azure	342
Install GigaVUE-FM Using Azure VM Dashboard	342
Install GigaVUE-FM Using Azure Marketplace	342
(Optional) Configure Metadata Export Filter for AMX	344
(Optional) Configure Role-Based Access for Third Party Orchestration	345
Create a Custom User	345
Create a Custom Role and Assign the User	345
Create a Custom User Group and Assign the Custom Role and the	
User	346
Create Monitoring Domain to Acquire Traffic from Customer	
Orchestrated Source	346
Create Tokens to Secure REST API and Third Party Orchestration Access	
to GigaVUE-FM	347
Deploy V Series Node and Register with GigaVUE-FM	347
Verify GigaVUE V Series Node Registration	350
Create a Monitoring Session to Configure Application Metadata Exporter	350
Configure Ingestor in AMX to Process Traffic	352

Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka	352
Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka	352
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka	353
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka	354
Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access (ZPA)	354
Prerequisites	354
Integrate ZPA with Gigamon AMX	355
Configure Ingestor for AMX	356
Integrate AMX with Third Party Tool	357
Set up the Third Party Tool for AMX Integration	357
Configure Exporter to Export Attributes to Third Party tools	358
Configure Cloud Tool Exporter	358
Deploy the Monitoring Session	358
Cloud Tool Exporter - Field References	358
Configure Kafka Exporter	359
Deploy the Monitoring Session	359
Kafka Exporter - Field References	360
Verify Your Deployment	361
View Application Metadata Exporter Statistics	361
Attributes for GigaVUE Enriched Metadata for Mobile Networks	365
Control Plane Attributes	366
User Plane Attributes	368
Attributes for GigaVUE Enriched Metadata for Cloud Workloads	371
Azure	372
Azure	373
VMware vCenter	374
Azure Kubernetes Service (AKS)	375
Zscaler Private Access (ZPA) Attributes	377
Supported Element Types for NetFlow Integrator Element Mapping	378
Migrate Application Intelligence Session to Monitoring Session	382
Post Migration Notes for Application Intelligence	383
What's Next	390

Upload SSL Key Pairs	404
Create SSL Service	406
Prerequisite	406
Key Mapping	406
SSL Key Pair Store	407
Add SSL Decrypt to Monitoring Session	407
Interface Mapping (Azure)	409
Deploy Monitoring Session (Azure)	410
Monitor and Troubleshoot	412
Monitor Cloud Health	413
View Health Status	413
View Monitoring Session Statistics (Azure)	416
Visualize the Network Topology (Azure)	418
Administer GigaVUE Cloud Suite for Azure	419
Manage Monitoring Domain	420
Configure Proxy Server	423
Configure Azure Settings	425
Role Based Access Control	426
About Events	428
About Audit Logs	430
High Availability for Cloud Infrastructure	431
Analytics for Virtual Resources	452
Virtual Inventory Statistics and Cloud Applications Dashboard	452
Analytics for Inline V Series Solution (Azure)	456
Debuggability and Troubleshooting	459
Sysdumps	459
FAQs - Secure Communication between GigaVUE Fabric	
Components (Azure)	460

GV-Cloud-Azure

Get Started

Overview of GigaVUE Cloud Suite for Azure


GigaVUE Cloud Suite for Azure gives you clear visibility into your workloads in Azure. It helps your security and monitoring tools find and respond to threats quickly using detailed, real-time network data.

This cloud-native solution runs entirely within your Azure Virtual Networks (VNETs). It gathers traffic from all areas of your cloud environment—including Azure’s native traffic mirroring. Then it processes that traffic to create useful metadata. This helps you monitor activity, spot issues early, and keep your cloud environment secure.

All parts of GigaVUE Cloud Suite for Azure live in the cloud. They collect traffic using UCT-V, small agent-like components installed on each Virtual Machine (VM). As your VMs change, Gigamon automatically adjusts to keep pace.

Benefits

- **Improves tool capacity:** Offloads security and monitoring work from your tools. This improves performance and reduces cost.
- **Fully automates the infrastructure:** Finds new or moved workloads on its own. It also launches visibility components and updates traffic rules without manual setup.
- **Simplifies operation:** Provides one dashboard for monitoring and control—across hybrid and cloud networks.
- **Helps accelerate cloud migrations:** Connects on-premises and cloud setups with a single visibility system. This makes migrations faster and easier.



GigaVUE Fabric Components

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

GigaVUE Cloud Suite for Azure includes the following fabric components:

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

You can use one of the following two different approaches to deploy GigaVUE Fabric Components:

- [GigaVUE-FM Orchestrated Fabric Deployment](#)
- [Third Party Orchestrated Fabric Deployment](#)

GigaVUE-FM Orchestrated Fabric Deployment

GigaVUE Fabric Components are the set of resources used to acquire, transport, process and distribute traffic to support your monitoring goals.

Using the GigaVUE-FM GUI, like UCT-V Controllers, GigaVUE V Series Nodes, and GigaVUE V Series Proxy, you can define Monitoring Domains and then use the UI to deploy the fabric components into your Azure environment.

GigaVUE Cloud Suite for Azure includes the following fabric components:

- [GigaVUE-FM](#)
- [UCT-V](#)
- [UCT-V Controller](#)
- [GigaVUE V Series Node](#)
- [GigaVUE V Series Proxy](#)

GigaVUE-FM

GigaVUE-FM fabric manager provides unified access, centralized administration, and high-level visibility for all GigaVUE traffic visibility nodes in the enterprise or data center, allowing a global perspective which is not possible from individual nodes.

In addition to centralized management and monitoring GigaVUE-FM helps you with configuration of the physical and virtual traffic policies for the visibility fabric thereby allowing administrators to map and direct network traffic to the tools and analytics infrastructure.

You have the flexibility of installing GigaVUE-FM across various supported platforms. Additionally, you can effectively manage deployments in any of the cloud platform as long as there exists IP connectivity for seamless operation.

UCT-V

Universal Cloud Tap - Virtual Machine(**UCT-V**) (earlier known as G-vTAP Agent) is a standalone service that is installed in the VM instance. UCT-V mirrors the selected traffic from the instances (virtual machines) to the GigaVUE V Series Node. The UCT-V is offered as a Debian (.deb), Redhat Package Manager (.rpm) package, ZIP and MSI .

Next generation UCT-V is a lightweight solution that acquires traffic from Virtual Machines and in-turn improves the performance of the UCT-V mirroring capability. The solution has a prefiltering capability at the tap level that reduces the traffic flow from the UCT-V to GigaVUE V Series Node and in-turn reduces the load on the GigaVUE V Series Node. Next generation UCT-V gets activated on Windows and also on Linux systems with a Kernel version above 4.18.

Prefiltering helps you reduce the costs significantly. It allows you to filter the traffic at UCT-Vs before sending it to the GigaVUE V Series Node. For prefiltering the traffic, GigaVUE-FM allows you to create a prefiltering policy template and the template can be applied to a monitoring session.

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-Venables 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
```

UCT-V Controller

UCT-V Controller (earlier known as G-vTAP Controller) manages multiple UCT-Vs and orchestrates the flow of mirrored traffic to GigaVUE V Series Nodes. GigaVUE-FM uses one or more UCT-V Controllers to communicate with the UCT-Vs. A UCT-V Controller can only manage UCT-Vs that has the same version.

For example, the UCT-V Controller6.14.00 can only manage UCT-Vs 6.14.00. If you have the previous version of UCT-V still deployed in the Virtual Network, you must configure

both UCT-V Controller 6.14.00 and the previous version. While configuring the UCT-V Controllers, you can also specify the tunnel type to be used for carrying the mirrored traffic from the UCT-Vs to the GigaVUE V Series Nodes.

NOTE: 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](#).

GigaVUE V Series Node

GigaVUE® V Series Node is a visibility node that aggregates mirrored traffic. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to on premise device or tools. GigaVUE Cloud Suite for Azure uses the TLS-PCAPNG, ERSPAN, L2GRE, UDPGRE and, VXLAN tunnels to deliver traffic to tool endpoints.

NOTE: 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](#).

GigaVUE V Series Proxy

GigaVUE V Series Proxy manages multiple GigaVUE V Series nodes and orchestrates the traffic flow of traffic of control and management communication from GigaVUE V Series nodes to the GigaVUE-FM. GigaVUE-FM uses one or more GigaVUE V Series Proxies to communicate with the GigaVUE V Series nodes.

NOTE: 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](#).

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 G-vTAP Agent)	UCT-V is a standalone service installed in your Virtual Machine (VM). This service 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.

Component	Description
	- 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.

For detailed information about configuring and deploying fabric components using different traffic acquisition methods, refer to [Deploy GigaVUE Cloud Suite for Azure using Third Party Orchestration](#).

Key Concepts

This section provides the key concepts required to understand GigaVUE Cloud Suite for Azure.

For details, refer to the following section:

- [Monitoring Domain](#)
- [Monitoring Session](#)
- [Inline V Series \(Azure\)](#)
- [Secure Communication between GigaVUE Fabric Components](#)
- [Precryption™](#)
- [Secure Tunnels](#)
- [Prefiltering](#)
- [Customer Orchestrated Source - Use Case](#)

- [Check for Required IAM Permissions in Azure](#)
- [Traffic Acquisition using Azure Virtual Network TAP](#)

Monitoring Domain

Monitoring domain helps you establish connection in between GigaVUE-FM and Azure platform. Once the connection is established, you can use GigaVUE-FM to launch the GigaVUE V Series Nodes, GigaVUE V Series Proxy and UCT-V Controller.

Monitoring Session

Monitoring sessions are the rules created in GigaVUE-FM to collect inventory data from all target instances in your cloud environment. You can design your monitoring session to include or exclude the instances 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 to your monitoring session. Similarly, when an instance is removed, it updates the monitoring sessions.

For more information on creating a monitoring session, see [Configure Monitoring Session](#).

Inline V Series (Azure)

The Inline V Series solution provides an advanced, scalable, agentless traffic acquisition mechanism that integrates seamlessly into your network. By deploying V Series Nodes in inline mode, you can mirror and process traffic efficiently while ensuring the reinjection of production traffic without disruption.

In AWS and Azure environments, the Inline V Series solution leverages Gateway Load Balancers (GWLB) to enable efficient traffic handling and visibility. This feature ensures low-latency performance, making it ideal for continuous traffic inspection and monitoring. Designed for simplicity and operational efficiency, the Inline V Series allows you to gain deep insights into network activity while maintaining high performance in demanding network environments.

You can use this solution for forwarding inline traffic and traffic processing. When traffic reaches the Inline V Series Node, a copy of the packet is taken as out-of-band traffic. You can forward the copied traffic to a GigaVUE V Series Node for additional processing or directly to monitoring tools. During boot-up, the Inline V Series Node initializes with the default Inline application.

A Monitoring Session is required to:

- Tap the inline traffic
- Create a copy for out-of-band forwarding
- Send the traffic to the desired tools.

Inline V Series Deployment Types

Single Tier Deployment

You can use this deployment model when traffic has to be tapped, filtered, and directly sent to tools without any processing.

Multi-Tier Deployment

Use this model when you need to process traffic through GigaVUE V Series applications before forwarding it to the tools. The first tier taps the traffic, and the second tier processes it using the GigaVUE V Series applications and forwards it to the tools.


Architecture of Inline V Series Solution in Azure

Components required for configuring Inline V Series Solution in Azure:

- Application VNet
- Appliance VNet
- Public Load balancer
- Gateway Load balancer
- Inline V Series Node

Application VNet consists of multiple workload VMs, Public Load Balancer, Public IP Load Balancer, and Application Server in the Backend pool. The appliance VNet consists of Gateway Load Balancer, Inline V Series Node. Any traffic reaching the Gateway Load Balancer will be routed to the Inline V Series Node.

The below architecture diagram explains how the Inline V Series solution works:



Traffic from the internet to the application server:

1. The traffic from the internet is sent to the Public Load Balancer configured in Application VNet using an Public IP LB configuration.
2. This traffic is routed the Gateway Load balancer.
3. The Gateway Load Balancer in the Appliance VNet forwards the traffic to the Inline V Series Nodes. The following actions are performed in the Inline V Series Node:
 - Once the traffic reaches the Inline V Series Nodes, a copy of the packet is taken as out of band traffic.
 - The Out of Band traffic is forwarded to the GigaVUE V Series Node for further processing or it can be forwarded to the tools.
 - The Inline V Series swaps the IP address and the Mac of the packets, where the source and destination are interchanged. As a result the Inline V Series Node becomes the source and Gateway Load Balancer becomes the destination.

NOTE: Packets sent from the Gateway Load Balancer will be VXLAN encapsulated and forwarded to the Inline V Series Nodes.

4. The inline traffic is sent back to the Gateway Load Balancer.
5. The Gateway Load Balancer forwards the inline traffic to the application servers in the Application VNet.

For a detailed workflow on acquiring traffic through the Inline V Series, refer to [Acquire Traffic using Inline V Series Solution](#) .


Limitation

This solution can be implemented only to tap the North-South traffic.

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.

Precription™

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

Gigamon Precription™ 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 Precription Technology Works](#)
- [Why Gigamon Precription](#)
- [Key Features](#)
- [Key Benefits](#)
- [Precription Technology on Single Node](#)
- [Precription Technology on Multi-Node](#)

¹**Disclaimer:** The Precription 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.

- [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	• Azure • Azure • GCP (via Third Party Orchestration)
Private Cloud	• OpenStack • VMware ESXi (via Third Party Orchestration only)

Platform Type	Platform
	- 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)
4.18.0-348.12.2.el8_5.x86_64	RHEL release 8.5 (Ootpa)

Kernel-Version	Operating System
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 Tunnels


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

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

How Does it Work!

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

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)

For information about how to configure secure tunnels, refer to the section [Configure Secure Tunnel](#).

Prefiltering

Prefiltering allows you to filter the traffic at UCT-Vs before sending it to the GigaVUE V Series Nodes. For prefiltering the traffic, GigaVUE-FM allows you to create a prefiltering policy template. You can apply the policy template to a monitoring session.

You can define a policy template with rules and filter values. A policy template once created can be applied to multiple monitoring sessions. However a monitoring session can use only one template.

Each monitoring session can have a maximum of 16 rules.

You can also edit a specific policy template with required rules and filter values for a particular monitoring session while editing a monitoring session. However, the customized changes are not saved in the template.

Some of the points that must be remembered for prefiltering in Next Generation UCT-Vs are:

- Prefiltering is supported only in Next Generation UCT-Vs. It is not supported for classic mirroring mechanism.
- Prefiltering is supported for both Linux and Windows UCT-Vs .
- For single monitoring session only one prefiltering policy is applicable. All the agents 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 the prefiltering policy cannot be applied. The policy is set to PassAll and prefiltering is not performed.
- When multiple monitoring sessions utilize a single agent to capture traffic, and one session uses a prefilter while the other does not, then the prefiltering policy is not applied. In this scenario, the policy defaults to PassAll, resulting in the omission of any prefiltering.

For details, refer to [Create Prefiltering Policy Template](#)

Customer Orchestrated Source - Use Case

Customer Orchestrated Source is a traffic acquisition method that allows to tunnel traffic directly to the GigaVUE V Series Nodes. In cases where UCT-V or VPC Mirroring cannot be configured due to firewall or other restrictions, you can use this method and tunnel the traffic to GigaVUE V Series Node, where the traffic is processed.

When using Customer Orchestrated Source, you can directly configure tunnels or raw endpoints in the monitoring session, where you can use other applications like Slicing, Masking, Application Metadata, and Application Filtering., to process the tunneled traffic.

For details on how to configure Tunnels and Raw End Points in the Monitoring Session, refer to [Create Ingress and Egress Tunnels \(Azure\)](#) and [Create Raw Endpoint \(Azure\)](#).

You can configure an Ingress tunnel in the Monitoring Session with the GigaVUE V Series Node IP address as the destination IP address, then the traffic is directly tunneled to that GigaVUE V Series Node.

Check for Required IAM Permissions in Azure

GigaVUE-FM allows you to validate whether the policy attached to the GigaVUE-FM using "Managed Identity" or "Application ID with client secret" has the required IAM permissions and notifies the users about the missing permissions. You can check

permissions while creating a Monitoring Domain and deploying GigaVUE Fabric Components using GigaVUE-FM by clicking the **Check Permissions** button on the Create Monitoring Domain page and Azure Fabric Launch page. The GigaVUE-FM displays the minimum required IAM permissions.

IMPORTANT: "Microsoft.Authorization/roleAssignments/read" permission is required for validating the required permissions. Ensure to include "Microsoft.Authorization/roleAssignments/read" permission in your IAM policy.

Prerequisites to deploy GigaVUE Cloud Suite for Azure:

- IAM permissions: Check whether the minimum required permissions are granted for the instance where the GigaVUE-FM is deployed. For details, refer to [Permissions and Privileges \(Azure\)](#).
- Access to public cloud endpoints: Check for access to the Azure cloud endpoint APIs.
- Subscription to the GigaVUE Cloud Suite for Azure: Before deploying the solution, you must subscribe to the GigaVUE Cloud Suite components from the Azure marketplace. For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).
- Security Group: Checks whether the required ports are configured in the security group. For more information on the security groups, see [Network Security Groups](#)

After you press the **Check Permissions** button, GigaVUE-FM verifies the minimum required permissions. Any missing permissions are highlighted with the respective message against the permission in a dialog box. You can use the displayed IAM Policy JSON as a reference and update the policy that is attached to the GigaVUE-FM.

Points to Note

1. When using Managed Identity (MSI), the IAM policy modified in Azure Portal takes a long duration to reflect in GigaVUE-FM. For details, refer to the [Limitation of using managed identities for authorization](#) section in Azure Documentation.
2. The Check Permissions feature is not supported when the **Traffic Acquisition Method** is set to **vTAP**.

The following table lists the different available status and their descriptions.

Access Status	Description
Allowed	This status is displayed if permission is configured correctly.
Denied	This status is displayed if permission is missing. For Example: If a permission is not configured in the IAM policy or if the permission access is explicitly denied in Azure, then the status is displayed as Denied.

Access Status	Description
Failed	This status is displayed if GigaVUE-FM fails to validate a permission. The reason and the probable cause are also displayed.
Not Executed	This status is displayed if a higher level of permission is denied or not configured, then GigaVUE-FM cannot validate a permission. For Example: If a subscription level permission is in denied or failed state then the resource level permission cannot be validated.
Undeterminable	The "Microsoft.Authorization/roleAssignments/read" permission is required to validate the required permissions. If this permission is not configured, the status of several other permissions cannot be determined.

Refer to the following section for more detailed information:

- [Check Permissions while Creating a Monitoring Domain](#)
- [Check Permissions while Configuring GigaVUE Fabric Components using GigaVUE-FM](#)
- [View Permission Status Reports](#)

View Permission Status Reports

The permission status reports consist of previously run **Check permissions** reports. They are auto purged once every 30 days. You can change the purge interval from the **Advanced Settings** page. For details, refer to [Configure Azure Settings](#).

You can view the Permission Status Report in the following two ways:

- In the Monitoring Domain page, click **Actions > View Permission Status Report**.
- In the Monitoring Domain page, you can navigate to **Settings** and then click **Permission Status Report**

On the **Permission Status Report** page, you can use the Filter button to filter the reports based on File Name, Type, and Date.


To view or delete individual reports, select the report and click **Actions** button.

Traffic Acquisition using Azure Virtual Network TAP

NOTE: Microsoft Azure vTAPs are currently in Public Preview in select regions. Gigamon has fully tested our orchestration and automation against the current APIs to ensure compatibility. Refer to Microsoft's website: [Virtual network TAP](#) for the latest Azure vTAP release information and regional availability.

Azure Virtual Network TAP allows traffic mirroring directly from virtual machine network interfaces to designated target network interfaces. The mirrored traffic, a deep copy of inbound and outbound network packets, can be forwarded to a destination IP endpoint

or an internal load balancer within the same or peered virtual networks. GigaVUE V Series Nodes receive traffic directly from source VMs using vTAP, simplifying traffic acquisition and visibility.



In the above diagram, the traffic from the source VMs are mirrored and forwarded to the GigaVUE V Series Node. GigaVUE-FM creates VTAP source configurations for each source VM NIC and a VTAP destination configuration for the GigaVUE V Series Node NIC. The source VMs and GigaVUE V Series Nodes can reside in different VNets, provided the VNets are peered. Multiple NICs can be configured for the same source VM and the traffic can be tapped and forwarded to GigaVUE V Series Node.

For more details on Azure virtual network TAP, refer to the [Virtual network TAP](#) Microsoft Azure documentation.

Rules and Notes

- Destination VM and Source VM must be in the same region.
- If workload VMs are present in multiple resource groups or Virtual Network (VNet), then Virtual Network peering has to be enabled between workload VNets and VNet where the GigaVUE V Series Node is deployed.


DISCLAIMER: Keep in mind that these guidelines are inherent to Azure, subject to change, and beyond Gigamon's purview. Refer to the Azure documentation for the most up-to-date instructions.

Limitation

- IPv6 tunnels are not supported by Azure VTAP.
- The Check Permissions feature is not supported when the **Traffic Acquisition Method** is set to **VTAP**.

Deployment Overview

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



Deployment Planning

This section focuses on the different decisions and information that you need to know before deploying GigaVUE Cloud Suite for Azure. It assumes you have read the previous chapters and have a good understanding of the Cloud Suite components and concepts, along with Azure components, services, and architectures.

Refer to the following sections for more detailed information:

- [Identify the Deployment Method](#)
- [Identify the Azure Regions](#)
- [Define GigaVUE-FM Deployment Location](#)
- [Identify the Traffic Acquisition Method](#)
- [Identify the Azure Subscriptions Involved](#)
- [Identify the VNets and Subnets Involved](#)
- [Identify the Required Azure RBAC Permissions](#)
- [Identify VPN Connectivity Requirements](#)
- [Identify the UCT-V Deployment Details](#)

- [Identify the UCT-V Controller Deployment Details](#)

Identify the Deployment Method

GigaVUE Cloud Suite for Azure supports two options for deploying the Fabric Components:

- GigaVUE-FM Orchestration
- Third Party Orchestration (Azure portal or CLI)

In either scenario, you need to first launch GigaVUE-FM and apply the required license. You can deploy GigaVUE-FM from the Azure Marketplace.

GigaVUE-FM Orchestration

In this method, you use the GigaVUE-FM user interface to deploy the GigaVUE Fabric Components. After launching GigaVUE-FM and applying the required license, you can deploy GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller using the GigaVUE-FM interface. Once deployed, GigaVUE-FM monitors the state of the deployed fabric components, and automatically replaces or relaunches them if any error occurs.

The GigaVUE-FM Orchestration option simplifies the GigaVUE Cloud Suite for Azure deployment but requires Azure API access and permissions to deploy and manage visibility fabric VMs.

Third Party Orchestration

In this method, you can deploy the fabric Components using the Azure portal, Azure CLI, or third-party tools. GigaVUE-FM does not launch the fabric components directly; instead, the components register with GigaVUE-FM after they are launched.

Identify the Azure Regions

Identify the Azure regions where your workloads are running and where you need traffic visibility. Consider the following:

- GigaVUE V Series Nodes must be deployed in the same region as the workloads they monitor.
- If workloads span multiple regions, plan for separate Monitoring Domains per region.
- GigaVUE-FM can manage multiple regions from a single instance.

Define GigaVUE-FM Deployment Location

GigaVUE-FM can be deployed in one of the following locations:

- **Inside Azure** (VNet) — GigaVUE-FM is launched as a VM within an Azure VNet. This is the recommended approach when the GigaVUE-FM Orchestration deployment method is used.

- **Outside Azure** (On-premises or external cloud) — GigaVUE-FM is deployed outside Azure and connects to the Azure environment. In this case, VPN connectivity between GigaVUE-FM and the Azure environment is required.

Identify the Traffic Acquisition Method

GigaVUE Cloud Suite for Azure supports the following traffic acquisition methods:

- **UCT-V (Universal Cloud Tap - Virtual)** — A lightweight agent installed on each target VM to mirror traffic to the GigaVUE V Series Node. Supports Linux and Windows VMs.
- **vTAP (Virtual TAP)** — An agentless method that uses Azure Network Packet Capture or Azure Gateway Load Balancer (GWLB) for traffic acquisition.
- **Customer Orchestrated Source** — Traffic is acquired from customer-deployed sources and forwarded to GigaVUE V Series Nodes.
- **Inline V Series** — Deploys GigaVUE V Series Nodes inline with the traffic path for active inspection.

The choice of traffic acquisition method determines the deployment option and the fabric components required.

Identify the Azure Subscriptions Involved

Most GigaVUE Cloud Suite for Azure deployments provide network visibility across multiple Azure subscriptions.

Consider the following:

- Identify the Azure subscriptions containing workloads to be monitored.
- Verify whether cross-subscription access is required. If so, create Azure Credentials in GigaVUE-FM for each subscription.
- Identify the subscription where GigaVUE-FM and fabric components will be deployed.

Identify the VNets and Subnets Involved

Plan the VNets and subnets where the GigaVUE Fabric Components will be deployed. Identify:

- The VNets containing the target workloads (monitored VMs).
- The subnets for GigaVUE V Series Nodes, UCT-V Controllers, and GigaVUE V Series Proxies.
- Whether VNet peering is required for cross-VNet traffic visibility.
- The management subnet for GigaVUE-FM communication.

Identify the Required Azure RBAC Permissions

GigaVUE-FM requires specific Azure Role-Based Access Control (RBAC) permissions to discover workloads, deploy fabric components, and manage the monitoring environment. For the list of required permissions, refer to [Permissions and Privileges](#) in Azure.

Identify VPN Connectivity Requirements

If GigaVUE-FM is deployed outside the Azure environment, a VPN connection between the GigaVUE-FM host and the Azure VNet is required. Identify:

- The VPN type (Azure VPN Gateway, ExpressRoute, or site-to-site VPN).
- The IP address ranges to be routed between GigaVUE-FM and the Azure VNet.
- Network firewall rules required to allow management traffic between GigaVUE-FM and the fabric components.

Identify the UCT-V Deployment Details

If UCT-V is the selected traffic acquisition method, plan the UCT-V deployment:

- Identify all target VMs (Linux and Windows) that require traffic mirroring.
- Choose the UCT-V installation method (manual installation or automated using Azure VM extensions or scripts).

NOTE: Each target VM requires one UCT-V instance, regardless of the number of monitoring tools.

- Ensure that the UCT-V can communicate with the UCT-V Controller over the required ports.

Identify the UCT-V Controller Deployment Details

The UCT-V Controller acts as the control plane proxy between GigaVUE-FM and the UCT-Vs. Plan the UCT-V Controller deployment:

- A single UCT-V Controller can manage up to 1000 UCT-Vs.
- Deploy the UCT-V Controller in the same VNet (or a peered VNet) as the target VMs it manages.
- Identify the VM size and subnet for the UCT-V Controller.
- Plan for redundant controllers if high availability is required.

Deployment Prerequisites

The following sections explain the prerequisites for successful GigaVUE Cloud Suite for Azure deployment:

- [Points to Note for GigaVUE Cloud Suite for Azure](#)
- [Prerequisites for GigaVUE Cloud Suite for Azure](#)
- [Network Firewall Requirement](#)
- [Licensing for GigaVUE Cloud Suite for Azure](#)
- [Supported Operating Systems for UCT-V](#)
- [VPN Connectivity](#)
- [Enable Subscription for GigaVUE Cloud Suite for Azure](#)
- [Configure Tokens](#)
- [Permissions and Privileges \(Azure\)](#)

Points to Note for GigaVUE Cloud Suite for Azure

IMPORTANT: If you are using a Cloud Solution Provider (CSP) in Azure, we require your CSP tenant ID and company name to be included in our Azure publishing portal, contact Gigamon Sales.

- When tool is deployed outside Azure, ensure there is connectivity between GigaVUE V Series Node tool interface and the tool. You can create connectivity by configuring a Network Address Translation (NAT) gateway.
- When GigaVUE-FM and GigaVUE V Series Nodes are deployed in different cloud platforms, then the GigaVUE-FM public IP address must be added to the **Data Notification Interface** as the Target Address in the Event Notifications page. For details, refer to [Configuration Settings](#).
- Fragmentation in the network should be avoided from UCT-V to GigaVUE V Series Node and from GigaVUE V Series Node to tool by setting appropriate MTU for the interfaces as there is a chance of fragment packets getting reordered in the network before it is received in GigaVUE V Series Node and the tool. If the tool VM MTU is less than that of the GigaVUE V Series Node, then the GigaVUE V Series Node fragments the packets.

Prerequisites for GigaVUE Cloud Suite for Azure

To enable the flow of traffic between the components and the monitoring tools, you must create the following requirements:

- [Resource Group](#)
- [Virtual Network](#)
- [Subnets for VNet](#)
- [Network Interfaces \(NICs\) for VMs](#)
- [Network Security Groups](#)
- [Virtual Network Peering](#)
- [Access control \(IAM\)](#)

- [Default Login Credentials](#)
- [GigaVUE-FM Version Compatibility](#)
- [Recommended Instance Types](#)

Resource Group

The resource group is a container that holds all the resources for a solution.

To create a resource group in Azure, refer to [Create a resource group](#) topic in the Azure Documentation.

Virtual Network

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. VNet enables many types of Azure resources, such as Azure Virtual Machines (VM), to securely communicate with each other, the internet, and on-premises networks.

You can only configure the GigaVUE fabric components in a Centralized VNet only. In case of a shared VNet, you must select a VNet as your Centralized VNet for GigaVUE fabric configuration.

To create a virtual network in Azure, refer to [Create a virtual network](#) topic in the Azure Documentation.

Subnets for VNet

The following table lists the two recommended subnets that your VNet must have to configure the GigaVUE Cloud Suite Cloud components in Azure.

You can add subnets when creating a VNet or add subnets on an existing VNet. For details, refer to [Add a subnet](#) topic in the Azure Documentation.

Subnet	Description
Management Subnet	Subnet that the GigaVUE-FM uses to communicate with the GigaVUE V Series Nodes and Proxy.
Data Subnet	A data subnet can accept incoming mirrored traffic from agents to the GigaVUE V Series Nodes or be used to egress traffic to a tool from the GigaVUE V Series Nodes. There can be multiple data subnets. ▪ Ingress is VXLAN from agents ▪ Egress is either VXLAN tunnel to tools or to GigaVUE HC Series tunnel port, or raw packets through a NAT when using NetFlow. Note: If you are using a single subnet, then the Management subnet will also be used as a Data Subnet.

Subnet	Description
Tool Subnet	A tool subnet can accept egress traffic to a tool from the GigaVUE V Series Nodes. There can be only one tool subnet. Egress is either VXLAN tunnel to tools or to GigaVUE HC Series tunnel port, or raw packets through a NAT when using NetFlow.

Network Interfaces (NICs) for VMs

When using UCT-V as the traffic acquisition method, for the UCT-Vs to mirror the traffic from the VMs, you must configure one or more Network Interfaces (NICs) on the VMs.

- **Single NIC**—If there is only one interface configured on the VM with the UCT-V, the UCT-V sends the mirrored traffic out using the same interface.
- **Multiple NICs**—If there are two or more interfaces configured on the VM with the UCT-V, the UCT-V monitors any number of interfaces but has an option to send the mirrored traffic out using any one of the interfaces or using a separate, non-monitored interface.

Network Security Groups

A network security group defines the virtual firewall rules for your VM to control inbound and outbound traffic. When you launch GigaVUE-FM, GigaVUE V Series Proxy, GigaVUE V Series Nodes, and UCT-V Controllers in your VNet, you add rules that control the inbound traffic to VMs, and a separate set of rules that control the outbound traffic.

To create a network security group and add in Azure, refer to [Create a network security group](#) topic in the Azure Documentation.

We recommend to create a separate security group for each component using the rules and port numbers.

In your Azure portal, select a network security group from the list. In the Settings section select the Inbound and Outbound security rules to the following rules.

Following are the Network Firewall Requirements:

Reduced Control-Plane Ports

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

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

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

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



Important

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

Choose the section that applies to your deployment scenario:

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

New Deployments

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

The reduced configuration uses only the three ports listed below:

Direction	Port	Purpose
Inbound/Outbound	443	HTTPS/REST management and control-plane communication, including GigaVUE-FM UI/API access, component registration, and certificate-related requests.

Inbound/Outbound	5671	RMQ/RA control-plane messaging and statistics or health updates between GigaVUE-FM and fabric components.
Inbound/Outbound	9902	RPC and control-plane communication, especially between UCT-V Controllers and UCT-V agent.

Existing Deployments

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

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

Migration to Reduced Control-Plane Ports

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

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

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

Firewall Rule Changes for Migrating to Reduced Control-Plane Ports

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

GigaVUE-FM

Action	Direction	Port
Remove	Inbound	9600

Remove	Outbound	8890
Remove	Outbound	8889
Remove	Outbound	80
Remove	Outbound	9990
Add	Inbound	443
Add	Inbound	5671
Add	Outbound	443

UCT-V Controller

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

UCT-V

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

GigaVUE V Series Node

Action	Direction	Port
Remove	Inbound	8889
Remove	Inbound	80
Remove	Outbound	8891
Remove	Outbound	8892
Remove	Outbound	9600

Remove	Outbound	8300
Add	Inbound	443
Add	Outbound	443
Add	Outbound	5671

GigaVUE V Series Proxy (Optional)

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

UCT-V OVS Controller

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

UCT-V OVS Module

Action	Direction	Port
Add	Inbound	9901

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

GigaVUE FM

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

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR2	Purpose
Inbound	TCP	443	9600	GigaVUE V Series Node IP	Allows GigaVUE-FM to receive certificate

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



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

UCT-V Controller

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

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

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



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

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

UCT-V

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

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	5671	NA	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller for

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



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

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

GigaVUE V Series Node

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

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8889	GigaVUE-FM IP	Allows GigaVUE V Series Node to communicate control and management plane traffic with GigaVUE-FM.

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



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

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

GigaVUE V Series Proxy (Optional)

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

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	8890	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic with GigaVUE V Series Proxy.
Inbound	TCP	443	8891	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive PKI requests from GigaVUE V Series Node.
Inbound	TCP	443	8892	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive

					registration requests and heartbeat messages from GigaVUE V Series Node.
Inbound	TCP	5671	8300	GigaVUE V Series Node	Allows GigaVUE V Series Proxy to receive certificate requests from GigaVUE V Series Node for the configured params and provides the certificate using those parameters.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	443	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from GigaVUE V Series Proxy.



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

UCT-V OVS Controller

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

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

UCT-V OVS Module

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

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

Legacy Control-Plane Ports

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

When to Use Legacy Ports

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

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

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

GigaVUE FM

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

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	443	Administrator Subnet	Allows GigaVUE-FM to accept Management connection using REST API. Allows users to access GigaVUE-FM UI securely through an HTTPS connection.
Inbound	TCP	22	Administrator Subnet	Allows CLI access to user-initiated

				management and diagnostics.
Inbound (This is the port used for Third Party Orchestration)	TCP	443	UCT-V Controller IP	Allows GigaVUE-FM to receive registration requests from UCT-V Controller using REST API.
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	443	UCT-V Controller IP	Allows GigaVUE-FM to receive registration requests from UCT-C Controller 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	5671	UCT-V Controller IP	Allows GigaVUE-FM to receive statistics from UCT-V Controllers.
Inbound	TCP	9600	UCT-V Controller	Allows GigaVUE-FM to receive certificate requests from UCT-V Controller.
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	TCP	5671	UCT-V Controller IP	Allows GigaVUE-FM to receive statistics from UCT-C Controllers.
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	TCP	9900	UCT-V Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic with UCT-V Controller.
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	8443	UCT-C Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic to UCT-C Controller.
Outbound	TCP	80	UCT-V Controller IP	Allows GigaVUE-FM to send ACME challenge requests to UCT-V Controller.
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 Controller

The following table defines the network communication parameters—protocols, ports, and CIDRs—required for UCT-V Controller to interact with GigaVUE-FM and UCT-V components, supporting registration, diagnostics, certificate exchange, and control-plane operations including third-party orchestration..

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	9900	GigaVUE-FM IP	Allows UCT-V Controller to communicate control and management plane traffic with GigaVUE-FM.
Inbound	TCP	9900	UCT-V or Subnet IP	Allows UCT-V Controller to receive traffic health updates from UCT-V.
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 UCT-V Controller to receive the ACME challenge requests from GigaVUE-FM.
Inbound	TCP	8300	UCT-VSubnet	Allows UCT-V Controller to receive the certificate requests from the UCT-V.
Inbound (This is the port used for Third Party Orchestration)	TCP	8892	UCT-V Subnet	Allows UCT-V Controller to receive the registration requests and heartbeat from UCT-V.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound (This is the port used for Third Party Orchestration)	TCP	443	GigaVUE-FM IP	Allows UCT-V Controller to send the registration requests to GigaVUE-FM using REST API.
Outbound	TCP	5671	GigaVUE-FM IP	Allows UCT-V Controller to send traffic health updates to GigaVUE-FM.
Outbound (This is the port used for Third Party Orchestration)	TCP	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from the UCT-V Controller.
Outbound	TCP	9902	UCT-V Subnet	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs for UCT-Vs with version greater than 6.10.00.
Outbound	TCP	8301	UCT-V Subnet	Allows ACME validation flow from UCT-V Controller to UCT-V.

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.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	9902	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller.
Inbound	TCP	8301	UCT-V Controller IP	Allows UCT-V to receive the ACME challenge requests from the UCT-V Controller.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	UDP (VXLAN)	VXLAN (default 4789)	GigaVUE V Series Node IP	Allows UCT-V to tunnel VXLAN traffic to GigaVUE V Series Nodes.
Outbound	IP Protocol (L2GRE)	L2GRE (IP 47)	GigaVUE V Series Node IP	Allows UCT-V to tunnel L2GRE traffic to GigaVUE V Series Nodes.
Outbound (Optional - This port is used only for Secure Tunnels)	TCP	11443	GigaVUE V Series Node IP	Allows UCT-V to securely transfer the traffic to the GigaVUE V Series Node.
Outbound	TCP	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	8892	UCT-V Controller IP	Allows UCT-V to receive the registration requests and heartbeat to UCT-V Controller.
Outbound	TCP	8300	UCT-V Controller IP	Allows UCT-V to receive ACME validation flow from UCT-V Controller.

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	UDP (VXLAN)	VXLAN (default 4789)	UCT-V Subnet IP	Allows GigaVUE V Series Nodes to receive VXLAN tunnel traffic to UCT-V.
Inbound	IP Protocol (L2GRE)	L2GRE	UCT-V Subnet IP	Allows GigaVUE V Series Nodes to receive L2GRE tunnel traffic to UCT-V.
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 Secure Tunnels)	TCP	11443	UCT-V subnet	Allows to securely transfer the traffic to GigaVUE V Series Nodes.
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.

UCT-C Controller - deployed in Kubernetes worker mode

The following table outlines UCT-C Controller's network communication parameters in Kubernetes worker mode, specifying TCP ports and CIDRs required for management, statistics exchange, and secure connectivity with GigaVUE-FM.

UCT-C Controller deployed inside Kubernetes worker node				
Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	8443 (configurable)	GigaVUE-FM IP	Allows GigaVUE-FM to communicate with UCT-C Controller.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	5671	Any IP address	Allows UCT-C Controller to send statistics to GigaVUE-FM.
Outbound	TCP	443	GigaVUE-FM IP	Allows UCT-C Controller to communicate with GigaVUE-FM.

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.

UCT-V Controller

The following table specifies the communication parameters required for third-party orchestration, detailing the TCP ports and CIDRs used by UCT-V Controller to manage registration and control-plane traffic with UCT-V components.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound (This is the port used for Third Party Orchestration)	TCP	8891	UCT-V or Subnet IP	Allows UCT-V Controller to receive the registration requests from UCT-V.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	9901	UCT-V Controller IP	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs.

GigaVUE V Series Node

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

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

GigaVUE V Series Proxy(Optional)

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

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

Virtual Network Peering

If workloads VMs are present in multiple resource groups or Virtual Network (VNet), then Virtual Network peering has to be enabled between workload VNets and VNet where the GigaVUE V Series Node is deployed. Virtual network peering enables you to seamlessly connect two or more Virtual Networks in Azure. Virtual Network Peering is only applicable when multiple Virtual Networks are used in a design. For details, refer to [Virtual Network Peering](#) topic in Azure documentation.

Access control (IAM)

You must have full resource access to the control the GigaVUE Cloud Suite cloud components. For details, refer to [Check access for a user](#) topic in the Azure documentation.

Default Login Credentials

You can login to the GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller by using the default credentials.

Product	Login credentials
GigaVUE V Series Node	You can login to the GigaVUE V Series Node by using ssh. The default username and password is not configured.
GigaVUE V Series proxy	You can login to the GigaVUE V Series Node by using ssh. The default username and password is not configured.
UCT-V Controller	You can login to the GigaVUE V Series Node by using ssh. The default username and password is not configured.

GigaVUE-FM Version Compatibility

GigaVUE-FM version 6.14.00 supports the latest version (6.14.00) of GigaVUE V Series Node, GigaVUE V Series Proxy, UCT-V Controller, and UCT-V, as well as (n-2) versions. For better compatibility, we recommend to use the latest version of fabric components with GigaVUE-FM.

Recommended Instance Types

NOTE: Additional instance types are supported. You can choose the instance type that best fits your deployment needs. If you're unsure which instance to select, contact Support, Sales, or Professional Services for deployment optimization.

Product	Instance Type	vCPU	RAM	Disk Size
GigaVUE V Series Node	Standard_D4s_v4	4 vCPU	16GB	10GB
	Standard_D8S_V4	8 vCPU	32GB	10GB
GigaVUE V Series Proxy	Standard_B1s	1 vCPU	1GB	
UCT-V Controller	Standard_B4ms	4 vCPU	16GB	4GB

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.

Network Firewall Requirement

Reduced Control-Plane Ports

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

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

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

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



Important

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

Choose the section that applies to your deployment scenario:

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

New Deployments

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

The reduced configuration uses only the three ports listed below:

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

Existing Deployments

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

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

Migration to Reduced Control-Plane Ports

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

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

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

Firewall Rule Changes for Migrating to Reduced Control-Plane Ports

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

GigaVUE-FM

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

UCT-V Controller

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

UCT-V

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

GigaVUE V Series Node

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

GigaVUE V Series Proxy (Optional)

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

UCT-V OVS Controller

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

UCT-V OVS Module

Action	Direction	Port
Add	Inbound	9901

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

GigaVUE FM

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

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



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

UCT-V Controller

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

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



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

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

UCT-V

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

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



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

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

GigaVUE V Series Node

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

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



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

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

GigaVUE V Series Proxy (Optional)

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

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



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

UCT-V OVS Controller

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

Direction	Protocol	Port	Replaces Legacy Port	Source CIDR	Purpose
Inbound	TCP	443	9900	GigaVUE-FM IP	Allows GigaVUE-FM to communicate control and management plane traffic. The controller supports both ports for compatibility and falls

					back to 9900 when required.
Direction	Protocol	Port	Replaces Legacy Port	Destination CIDR	Purpose
Outbound	TCP	9901	NA	UCT-V OVS Module IP	Allows to communicate with UCT-V OVS Modules (RPC).

UCT-V OVS Module

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

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

Legacy Control-Plane Ports

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

When to Use Legacy Ports

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

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

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

GigaVUE FM

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

components.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	443	Administrator Subnet	Allows GigaVUE-FM to accept Management connection using REST API. Allows users to access GigaVUE-FM UI securely through an HTTPS connection.
Inbound	TCP	22	Administrator Subnet	Allows CLI access to user-initiated management and diagnostics.
Inbound (This is the port used for Third Party Orchestration)	TCP	443	UCT-V Controller IP	Allows GigaVUE-FM to receive registration requests from UCT-V Controller using REST API.
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	443	UCT-V Controller IP	Allows GigaVUE-FM to receive registration requests from UCT-C Controller 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	5671	UCT-V Controller IP	Allows GigaVUE-FM to receive statistics from UCT-V Controllers.
Inbound	TCP	9600	UCT-V Controller	Allows GigaVUE-FM to receive certificate requests from UCT-V Controller.
Inbound	TCP	9600	GigaVUE V Series	Allows GigaVUE-FM to

			Proxy	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	TCP	5671	UCT-V Controller IP	Allows GigaVUE-FM to receive statistics from UCT-C Controllers.
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	TCP	9900	UCT-V Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic with UCT-V Controller.
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	8443	UCT-C Controller IP	Allows GigaVUE-FM to communicate control and management plane traffic to UCT-C Controller.
Outbound	TCP	80	UCT-V Controller IP	Allows GigaVUE-FM to send ACME challenge requests to UCT-V Controller.

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 Controller

The following table defines the network communication parameters—protocols, ports, and CIDRs—required for UCT-V Controller to interact with GigaVUE-FM and UCT-V components, supporting registration, diagnostics, certificate exchange, and control-plane operations including third-party orchestration..

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	9900	GigaVUE-FM IP	Allows UCT-V Controller to communicate control and management plane traffic with GigaVUE-FM.
Inbound	TCP	9900	UCT-V or Subnet IP	Allows UCT-V Controller to receive traffic health updates from UCT-V.
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 UCT-V Controller to receive the ACME challenge requests from GigaVUE-FM.
Inbound	TCP	8300	UCT-VSubnet	Allows UCT-V Controller to receive the certificate requests from the UCT-V.
Inbound (This is the port used for Third Party Orchestration)	TCP	8892	UCT-V Subnet	Allows UCT-V Controller to receive the registration requests and heartbeat from UCT-V.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound (This is the port used for Third Party Orchestration)	TCP	443	GigaVUE-FM IP	Allows UCT-V Controller to send the registration requests to GigaVUE-FM using REST API.

Outbound	TCP	5671	GigaVUE-FM IP	Allows UCT-V Controller to send traffic health updates to GigaVUE-FM.
Outbound (This is the port used for Third Party Orchestration)	TCP	9600	GigaVUE-FM IP	Allows GigaVUE-FM to receive certificate requests from the UCT-V Controller.
Outbound	TCP	9902	UCT-V Subnet	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs for UCT-Vs with version greater than 6.10.00.
Outbound	TCP	8301	UCT-V Subnet	Allows ACME validation flow from UCT-V Controller to UCT-V.

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.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound	TCP	9902	UCT-V Controller IP	Allows UCT-V to receive control and management plane traffic from UCT-V Controller.
Inbound	TCP	8301	UCT-V Controller IP	Allows UCT-V to receive the ACME challenge requests from the UCT-V Controller.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	UDP (VXLAN)	VXLAN (default 4789)	GigaVUE V Series Node IP	Allows UCT-V to tunnel VXLAN traffic to GigaVUE V Series Nodes.
Outbound	IP Protocol (L2GRE)	L2GRE (IP 47)	GigaVUE V Series Node IP	Allows UCT-V to tunnel L2GRE traffic to GigaVUE V Series Nodes.
Outbound (Optional - This port is used only for Secure Tunnels)	TCP	11443	GigaVUE V Series Node IP	Allows UCT-V to securely transfer the traffic to the GigaVUE V Series Node.

Outbound	TCP	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	8892	UCT-V Controller IP	Allows UCT-V to receive the registration requests and heartbeat to UCT-V Controller.
Outbound	TCP	8300	UCT-V Controller IP	Allows UCT-V to receive ACME validation flow from UCT-V Controller.

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	UDP (VXLAN)	VXLAN (default 4789)	UCT-V Subnet IP	Allows GigaVUE V Series Nodes to receive VXLAN tunnel traffic to UCT-V.
Inbound	IP Protocol (L2GRE)	L2GRE	UCT-V Subnet IP	Allows GigaVUE V Series Nodes to receive L2GRE tunnel traffic to UCT-V.
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 Secure Tunnels)	TCP	11443	UCT-V subnet	Allows to securely transfer the traffic to GigaVUE V Series Nodes.
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.

UCT-C Controller - deployed in Kubernetes worker mode

The following table outlines UCT-C Controller's network communication parameters in Kubernetes worker mode, specifying TCP ports and CIDRs required for management, statistics exchange, and secure connectivity with GigaVUE-FM.

UCT-C Controller deployed inside Kubernetes worker node				
Direction	Protocol	Port	Source CIDR	Purpose

Inbound	TCP	8443 (configurable)	GigaVUE-FM IP	Allows GigaVUE-FM to communicate with UCT-C Controller.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	5671	Any IP address	Allows UCT-C Controller to send statistics to GigaVUE-FM.
Outbound	TCP	443	GigaVUE-FM IP	Allows UCT-C Controller to communicate with GigaVUE-FM.

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.

UCT-V Controller

The following table specifies the communication parameters required for third-party orchestration, detailing the TCP ports and CIDRs used by UCT-V Controller to manage registration and control-plane traffic with UCT-V components.

Direction	Protocol	Port	Source CIDR	Purpose
Inbound (This is the port used for Third Party Orchestration)	TCP	8891	UCT-V or Subnet IP	Allows UCT-V Controller to receive the registration requests from UCT-V.
Direction	Protocol	Port	Destination CIDR	Purpose
Outbound	TCP	9901	UCT-V Controller IP	Allows UCT-V Controller to communicate control and management plane traffic with UCT-Vs.

GigaVUE V Series Node

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

Direction	Protocol	Port	Source CIDR	Purpose
Outbound (This is the port used	TCP	8891	GigaVUE V Series Proxy IP	Allows GigaVUE V Series Node to send registration requests and

for Third Party Orchestration)				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.

Licensing for GigaVUE Cloud Suite for Azure

You can license the GigaVUE Cloud Suite for Azure using the following method:

- [Bundle Replacement Policy](#)

For purchasing licenses with the Volume-Based License (VBL) option, contact our Sales. Refer to [Contact Sales](#).

For more detailed information on Volume-Based Licensing and instructions on how to generate and apply license, refer to the following topics:

- [Bundle Replacement Policy](#)
- [Activate Volume-Based Licenses](#)
- [Manage Volume-Based Licenses](#)

Default Trial Licenses

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

This license includes the following applications:

- ERSPAN
- GENEVE
- Slicing
- Masking
- Trailer
- Tunneling

- Load Balancing
- Enhanced Load Balancing
- Flow map
- Header Stripping
- Header Addition
- De-duplication
- NetFlow
- Application Packet Filtering
- Application Filtering Intelligence
- Application Metadata Intelligence
- Application Metadata Exporter
- Inline SSL
- SSL Decrypt
- Precryption

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

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

License Requirements for Monitoring Sessions

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

- CoreVUE Base Bundle (minimum required)
- NetVUE
- SecureVUE Plus

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

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

Volume Based License (VBL)

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

applications that run on these nodes.

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

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

Related Information

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

Base Bundles

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

- CoreVUE
- NetVUE
- SecureVUE Plus

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

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

Feature	CoreVUE	NetVUE	SecureVUE Plus
Slicing	√	√	√
Masking	√	√	√
Advanced Load Balancing	√	√	√
Advanced Tunneling	√	√	√
Deduplication	-	√	√
NetFlow Generation	-	√	√
Adaptive Packet Filtering (APF)	-	-	√

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

Feature	CoreVUE	NetVUE	SecureVUE Plus
Application Filtering Intelligence (AFI)	-	-	√
Application Metadata Intelligence (AMI)	-	-	√
Application Metadata Exporter (AMX)	-	-	√
Inline Decryption	-	-	√
OOB Decryption	-	-	√
Precryption	-	-	√
Gigamon Enriched Metadata	-	-	√

Bundle Replacement Policy

Refer to the following notes:

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

Add-on Packages

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

The following add-on SKUs are available:

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

Rules for add-on packages:

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

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

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

How GigaVUE-FM Tracks Volume-Based License Usage

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

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

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

Activate Volume-Based Licenses

To activate Volume-Based Licenses,

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

Manage Volume-Based Licenses

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

View active Volume-Based License

To view active Volume-Based License (VBL):

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

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

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

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

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

View Inactive Volume-Based License

To view inactive Volume-Based License (VBL):

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

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

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

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

Activate Volume-Based Licenses

To activate Volume-Based Licenses,

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

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

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

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

Supported Operating Systems for UCT-V

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 6.14.02 supports UCT-V version 6.14.02 as well as (n-2) versions. We recommend to use the latest version of UCT-V with GigaVUE-FM for better compatibility.

VPN Connectivity

GigaVUE-FM requires Internet access to integrate with the public API endpoints to integrate with the GigaVUE Cloud Suite Cloud platform. If there is no Internet access, refer to [Configure Proxy Server](#).

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

Enable Subscription for GigaVUE Cloud Suite for Azure

For GigaVUE-FM to be able to launch the fabric images,

- You must accept the terms of the end user license agreements (EULAs)
- Enable programmatic access in the Azure portal or through Azure Portal Cloud Shell.

For details, refer to the following topics:

- [Enable Subscription using CLI](#)
- [Enable Subscription using Azure Portal](#)

NOTE: For accepting EULA, you need to have Owner role on the Subscription.

Enable Subscription using CLI

1. **BYOL FM:** The following example shows how to accept EULA for BYOL FM using Azure Portal Cloud Shell

```
az vm image terms accept --urn gigamon-inc:gigamon-gigavue-cloud-suite-
v2:gfm-azure-v6.14.xx:6.14.02
{
  "accepted": true,
  "id": "<Enter Subscription ID>",
  "licenseTextLink": "<Provide License text file link>",
  "marketplaceTermsLink": "<Provide Market Place Terms text file link>",
  "name": "gfm-azure",
  "plan": "gfm-azure",
  "privacyPolicyLink": "https://www.gigamon.com/privacy-policy.html",
  "product": "gigamon-gigavue-cloud-suite",
  "publisher": "gigamon-inc",
  "retrieveDatetime": "2023-05-02T20:09:36.1347592Z",
  "signature":
    "SZL3CYR5MMU5QC5FEBIDHLMOYE7DD4CBSMLOVRMCKAAUD5CKLG4RIWPALULYWCFWCENMFF7
    7RCXM4CM2B24WV3PGEFWW7UL4VMI3BVI",
  "systemData": {
    "createdAt": "2023-05-02T20:09:38.101210+00:00",
    "createdBy": "6447eb55-9d09-481b-89bc-52e96bb52823",
    "createdByType": "ManagedIdentity",
    "lastModifiedAt": "2023-05-02T20:09:38.101210+00:00",
    "lastModifiedBy": "6447eb55-9d09-481b-89bc-52e96bb52823",
    "lastModifiedByType": "ManagedIdentity"
  },
  "type": "Microsoft.MarketplaceOrdering/offertypes"
}
```

2. **Fabric Images (need to accept on all 3):** The following examples show how to accept EULA for different fabric components using Azure Portal Cloud Shell

For UCT-V Controller

```
az vm image terms accept --urn gigamon-inc:gigamon-gigavue-cloud-suite-
v2:uctv-cntlr-v6.14.xx:6.14.02
{
  "accepted": true,
  .....
  "type": "Microsoft.MarketplaceOrdering/offertypes"
}
```

For GigaVUE V Series Node

```
az vm image terms accept --urn gigamon-inc:gigamon-gigavue-cloud-suite-
v2:vseries-node-v6.14.xx:6.14.02
{
  "accepted": true,
  .....
  "type": "Microsoft.MarketplaceOrdering/offertypes"
}
```

For GigaVUE V Series Proxy

```
az vm image terms accept --urn gigamon-inc:gigamon-gigavue-cloud-suite-
v2:vseries-proxy-v6.14.xx:6.14.02
{
  "accepted": true,
  .....
  "type": "Microsoft.MarketplaceOrdering/offertypes"
}
```

Enable Subscription using Azure Portal

Enable the subscription for GigaVUE-FM and its fabric components like GigaVUE V Series Node, UCT-V Controller, and GigaVUE V Series Proxy.

Perform the following steps to accept the terms using Azure Portal:

1. Go to Market Place and search Gigamon.
2. From the search results, select **Gigamon GigaVUE Cloud Suite for Azure**.
3. From the **Plan** drop-down menu, select the required image.
4. Select the **"Want to deploy programmatically? Get started"** link.
5. Review the terms of service and the subscription name and then select **Enable**.

Configure Tokens

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

This feature verifies the identity of a user for accessing the GigaVUE-FM REST APIs by generating tokens.

GigaVUE-FM allows you to generate a token only if you are an authenticated user and based on your privileges in accessing the GigaVUE-FM. You can copy the generated tokens from the GUI, which can be used to access the REST APIs. Token inherits the Role-Based Access (RBAC) privilege (read or write) of the user groups assigned to a particular user.

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

Prerequisite

You must create user groups in GigaVUE-FM. For details, refer to

Rules and Notes

- Authentication using a token is an additional mechanism to access GigaVUE-FM REST APIs, and it does not replace the existing GigaVUE-FM authentication mechanism.
- Only authenticated users can create tokens.
- The token expires or becomes invalid under the following circumstances:
 - Based on the configured value for expiry:
 - Default value: 30 days
 - Maximum value: 105 days
 - Deleting a related user group that exists as part of the token leads to deletion of the corresponding token.
 - A password change for the user(local) deletes the corresponding token.
 - A change in the authentication type deletes all the tokens.
- During the back up and restoration of GigaVUE-FM, previously generated tokens are not available.
- In FMHA role changeover, active GigaVUE-FM tokens are active.
- For basic authentication, activities such as creating, revoking, and reviewing of Token APIs are restricted.
- For expired or invalid tokens, you notice the error code 401 on GigaVUE-FM REST API access.

This section explains about the following:

- [Create Token](#)
- [Revoke Tokens](#)
- [Export Token](#)
- [Configure Tokens](#)

Create Token

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

To create a token, follow these steps:

1. Go to , select **Authentication > GigaVUE-FM User Management**. The **User Management** page appears.
2. In the **User Management** page, select **Tokens**.

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

3. Select **New Token**.
4. Enter a name for the new token in the **Name** field.
5. Enter the days until the token is valid in the **Expiry** field.
6. Select the user group for which you are privileged to access GigaVUE-FM from the **User Group** drop-down list.
7. Select **OK** to generate a new token.

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

Copy and Paste a Token

1. Select the token that you want to copy.
2. Select **Actions>Copy Token**.

The token is copied.

3. Paste the copied token in the required areas.

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

Revoke Tokens

You can revoke tokens that other users create.

Prerequisite: Write access in GigaVUE-FM Security Management.

To revoke tokens,

1. Go to , select **Authentication > GigaVUE-FM User Management**.
2. In the **User Management** page, select **Tokens**.
3. From the drop-down, select **Token Management**. You can view the created tokens.
4. Select the token that you want to revoke.
5. Select **Action> Revoke**.

Export Token

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

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

Permissions and Privileges (Azure)

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. This is used for GigaVUE-FM to integrate with Azure APIs and to automate the fabric deployment and management.

IMPORTANT: "Microsoft.Authorization/roleAssignments/read" permission is required for validating the required permissions. Ensure to include "Microsoft.Authorization/roleAssignments/read" permission in your IAM policy.

Prerequisite

- Pre-defined custom roles or
- Create new custom roles that you can attach to the resource group or subscription level.

For details, refer to [Custom Roles](#) topic.

GigaVUE-FM supports two types of authentications with Azure. Refer to the following sections for more detailed information on how to enable each type of authentication for GigaVUE-FM and how to assign the above created custom roles for GigaVUE-FM:

- [Managed Identity \(recommended\)](#)
- [Application ID with client secret](#)

Managed Identity (recommended)

NOTE: The following information applies only to **System assigned managed identity**.

Managed Identity (MSI) is a feature of Azure Active Directory. When you enable MSI on an Azure service, Azure automatically creates an identity for the service VM in the Azure AD tenant used by your Azure subscription.

Managed Identity (MSI) is only available when GigaVUE-FM is launched inside Azure. If GigaVUE-FM is launched in one VNet and the GigaVUE V Series Nodes are deployed in a different VNet, then Virtual Network Peering must be configured. For details, refer to [Virtual Network Peering](#).

IMPORTANT: When using Managed Identity (MSI), the IAM policy modified in Azure portal can take up to **1 hour** to propagate and appear in GigaVUE-FM. Refer to the [Limitation of using managed identities for authorization](#) section in Azure Documentation for more detailed information.

To work with MSI,

1. Enable MSI on the VM running in GigaVUE-FM. It can be done in using Azure portal or CLI.
 - Azure Portal: Refer to [Configure managed identities using the Azure portal](#) in the Azure documentation for detailed instructions.
 - Azure CLI:
 - Enable system-assigned managed identity : **az vm identity assign -g <Resource group where GigaVUE-FM is deployed> -n <myVmName>**

For more information, refer to [Configure managed identities for Azure resources using Azure CLI](#) topic in the Azure Documentation.

2. Assign permissions to this VM on all the resources where you need GigaVUE-FM to manage.

After enabling MSI, you can assign custom roles to GigaVUE-FM at a resource group level or subscription level.



NOTE: Use a system-assigned managed identity in Azure when a single resource needs to authenticate to other services, and you want the identity's lifecycle tied to the resource's. This means the identity is created and deleted along with the resource.

Assign a Custom Role using CLI

1. Assign a custom role at resource group level where you will deploy the fabric:

```
az vm identity assign -g <Resource group where GigaVUE-FM is deployed>
-role <Custom Role> -n <GigaVUE-FM name> --scope <resource group id>
```
2. Assign a custom role at the subscription level to view the complete account details:

```
az vm identity assign -g <Resource group where GigaVUE-FM is deployed>
-role <Custom Role> -n <GigaVUE-FM name> --scope <subscription id>
```

If you want to update the Role, you can edit the JSON file, and then update the Role in Azure using the following CLI command:

```
az role definition update --role-definition <Custom Role>.json
```

You can run these commands in the Azure Portal in a cloud shell (icon in the upper right of the portal as seen here): .

Assign a Custom Role using Azure Portal

You can assign roles to GigaVUE-FM using Azure Portal for Resource Group Level or Subscription Level. For details, refer to [Assign Azure roles](#) topic in Azure Documentation.

Application ID with client secret

GigaVUE-FM supports application id with client secret authentication. When using GigaVUE-FM to connect to Azure, it uses a service principal. A service principal is an account for a non-human such as an application to connect to Azure. When GigaVUE-FM is launched outside Azure, Application ID with client secret is preferred.

To create a service principal in Azure, refer to the following topics in the Azure Documentation:

- [Create an Azure service principal with the Azure CLI](#)
- [Create an Azure service principal with Azure PowerShell](#)
- [Create an Azure service principal with Azure Portal](#)



GigaVUE-FM must be able to access the URLs listed in the [Allow the Azure portal URLs on your firewall or proxy server](#) in order to connect to Azure.

Following are the required endpoints for Azure:

- login.microsoftonline.com
- management.azure.com
- login.microsoftonline.us (for Azure Gov cloud deployments)
- management.usgovcloudapi.net (for Azure Gov cloud deployments)

After creating service principal in Azure, you can add custom roles. For details, refer to [Assign a Custom Role using CLI](#) or [Assign a Custom Role using Azure Portal](#).

The key fields required for GigaVUE-FM to connect to Azure are Subscription ID, Tenant ID, Application ID, and Application Secret.

- When creating the service principal using the Azure CLI, the output of that command will display the "appId" and "password" fields. These two are the Application ID and Application Secret fields that are required for GigaVUE-FM to connect to Azure. Copy them.
- Now, using the Azure CLI again, do an 'account show' command and copy the Subscription ID and the Tenant ID of your subscription.

The Subscription ID, Tenant ID, Application ID, and Application Secret will be used when creating credentials in GigaVUE-FM.

Custom Roles

The 'built-in' roles provided by Microsoft are open to all resources. You can create a custom role if required. For more information, refer to [Azure custom roles](#) topic in the Azure Documentation.

You can use the following command to create custom roles in CLI:

```
az role definition create --role-definition {roleDefinition}
```

{roleDefinition} - You can provide the input either as a JSON object or by specifying the path to a JSON file (e.g., ~/roles/vmoperator.json).

The following examples provides the minimum permissions that are required for GigaVUE-FM to deploy the fabric components and/or inventory the UCT-V. The permissions can be applied at the resource group level or subscription level.

You can use the following snippet in the example JSON file mentioned below to assign your custom role at either resource group level or subscription level.

For Subscription level: Assigning roles at the subscription level grants access to all resource groups and resources within that subscription.

```
"assignableScopes": [  
  "/subscriptions/<Subscription ID>/"  
],
```

For Resource group level: Assigning roles at the resource group level provides granular control. You can limit access to specific sets of resources by assigning roles only to selected resource groups.

```
"assignableScopes": [  
  "/subscriptions/<Subscription ID>/resourceGroups/<resourceGroup name>"  
],
```

Example 1: Create Custom Role for GigaVUE-FM to deploy visibility fabric components and inventory UCT-V

```
{
  "Name": "CustomRoleFabricDeploymentAndInventory",
  "description": "The minimum requirements for FM to deploy Fabric Components and inventory UCT-V",
  "assignableScopes": [
    "/subscriptions/<Subscription ID>"
  ],
  "permissions": [
    {
      "actions": [
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Compute/virtualMachines/write",
        "Microsoft.Compute/virtualMachines/delete",
        "Microsoft.Compute/virtualMachines/start/action",
        "Microsoft.Compute/virtualMachines/powerOff/action",
        "Microsoft.Compute/virtualMachines/restart/action",
        "Microsoft.Compute/virtualMachines/instanceView/read",
        "Microsoft.Compute/locations/vmSizes/read",
        "Microsoft.Compute/images/read",
        "Microsoft.Compute/disks/read",
        "Microsoft.Compute/disks/write",
        "Microsoft.Compute/disks/delete",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Network/networkInterfaces/write",
        "Microsoft.Network/virtualNetworks/subnets/join/action",
        "Microsoft.Network/virtualNetworks/subnets/read",
        "Microsoft.Network/networkInterfaces/join/action",
        "Microsoft.Network/networkInterfaces/delete",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/publicIPAddresses/write",
        "Microsoft.Network/publicIPAddresses/delete",
        "Microsoft.Network/publicIPAddresses/join/action",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/virtualNetworks/virtualMachines/read",
        "Microsoft.Network/networkSecurityGroups/read",
        "Microsoft.Network/networkSecurityGroups/join/action",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/publicIPAddresses/write",
        "Microsoft.Network/publicIPAddresses/delete",
        "Microsoft.Network/publicIPAddresses/join/action",
        "Microsoft.Resources/subscriptions/locations/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read",
        "Microsoft.Resources/subscriptions/resourcegroups/resources/read"
      ],
      "notActions": [],
      "dataActions": [],
      "notDataActions": []
    }
  ]
}
```

Example 2: Create Custom Role for GigaVUE-FM to only inventory UCT-V

```
{
```

```
"Name":"CustomRoleInventoryUCT-V",
"description":" Minimum requirements for FM to inventory UCT-V ",
"assignableScopes":[
  "/subscriptions/<Subscription ID>"
],
"permissions":[
  {
    "actions":[
      "Microsoft.Compute/virtualMachines/read",
      "Microsoft.Compute/virtualMachines/instanceView/read",
      "Microsoft.Compute/images/read",
      "Microsoft.Compute/disks/read",
      "Microsoft.Network/networkInterfaces/read",
      "Microsoft.Network/virtualNetworks/subnets/read",
      "Microsoft.Network/publicIPAddresses/read",
      "Microsoft.Network/virtualNetworks/read",
      "Microsoft.Network/virtualNetworks/virtualMachines/read",
      "Microsoft.Network/networkSecurityGroups/read",
      "Microsoft.Network/publicIPAddresses/read",
      "Microsoft.Resources/subscriptions/locations/read",
      "Microsoft.Resources/subscriptions/resourceGroups/read",
      "Microsoft.Resources/subscriptions/resourcegroups/resources/read"
    ],
    "notActions":[],
    "dataActions":[],
    "notDataActions":[]
  }
]
}
```

Example 3: Create Custom Role for GigaVUE-FM to deploy visibility fabric components, inventory VMs and configure vTAPs in Azure

```
{
  "Name":"CustomRolevTAP",
  "description":"Minimum requirements for GigaVUE-FM to deploy visibility fabric components,
inventory VMs and configure vTAPs in Azure",
  "assignableScopes":[
    "/subscriptions/<Subscription ID>"
  ],
  "permissions":[
    {
      "actions":[
        "Microsoft.Network/virtualNetworkTaps/read",
        "Microsoft.Network/virtualNetworkTaps/delete",
        "Microsoft.Network/virtualNetworkTaps/write",
        "Microsoft.Network/virtualNetworkTaps/join/action",
        "Microsoft.Network/networkInterfaces/tapConfigurations/read",
        "Microsoft.Network/networkInterfaces/tapConfigurations/write",
        "Microsoft.Network/networkInterfaces/tapConfigurations/delete",
        "Microsoft.Network/networkInterfaces/ipconfigurations/join/action",
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Compute/virtualMachines/write",
        "Microsoft.Compute/virtualMachines/delete",
        "Microsoft.Compute/virtualMachines/start/action",
        "Microsoft.Compute/virtualMachines/powerOff/action",
```

```

        "Microsoft.Compute/virtualMachines/restart/action",
        "Microsoft.Compute/virtualMachines/instanceView/read",
        "Microsoft.Compute/locations/vmSizes/read",
        "Microsoft.Compute/images/read",
        "Microsoft.Compute/disks/read",
        "Microsoft.Compute/disks/write",
        "Microsoft.Compute/disks/delete",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Network/networkInterfaces/write",
        "Microsoft.Network/virtualNetworks/subnets/join/action",
        "Microsoft.Network/virtualNetworks/subnets/read",
        "Microsoft.Network/networkInterfaces/join/action",
        "Microsoft.Network/networkInterfaces/delete",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/publicIPAddresses/write",
        "Microsoft.Network/publicIPAddresses/delete",
        "Microsoft.Network/publicIPAddresses/join/action",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/virtualNetworks/virtualMachines/read",
        "Microsoft.Network/networkSecurityGroups/read",
        "Microsoft.Network/networkSecurityGroups/join/action",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/publicIPAddresses/write",
        "Microsoft.Network/publicIPAddresses/delete",
        "Microsoft.Network/publicIPAddresses/join/action",
        "Microsoft.Resources/subscriptions/locations/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read",
        "Microsoft.Resources/subscriptions/resourceGroups/resources/read"
    ],
    "notActions":[],
    "dataActions":[],
    "notDataActions":[]
  }
]
}

```

Example 4: Create Custom Role for GigaVUE-FM to only inventory VMs and configure vTAPs in Azure

```

{
  "Name":"CustomRoleInventoryvTAP",
  "description":"Minimum requirements for GigaVUE-FM to only inventory VMs and configure vTAPs in Azure",
  "assignableScopes":[
    "/subscriptions/<Subscription ID>"
  ],
  "permissions":[
    {
      "actions":[
        "Microsoft.Network/virtualNetworkTaps/read",
        "Microsoft.Network/virtualNetworkTaps/delete",
        "Microsoft.Network/virtualNetworkTaps/write",
        "Microsoft.Network/virtualNetworkTaps/join/action",
        "Microsoft.Network/networkInterfaces/tapConfigurations/read",
        "Microsoft.Network/networkInterfaces/tapConfigurations/write",
        "Microsoft.Network/networkInterfaces/tapConfigurations/delete",

```

```

        "Microsoft.Network/networkInterfaces/ipconfigurations/join/action",
        "Microsoft.Compute/virtualMachines/read",
        "Microsoft.Compute/virtualMachines/instanceView/read",
        "Microsoft.Compute/images/read",
        "Microsoft.Compute/disks/read",
        "Microsoft.Network/networkInterfaces/read",
        "Microsoft.Network/virtualNetworks/subnets/read",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/virtualNetworks/virtualMachines/read",
        "Microsoft.Network/networkSecurityGroups/read",
        "Microsoft.Network/publicIPAddresses/read",
        "Microsoft.Resources/subscriptions/locations/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read",
        "Microsoft.Resources/subscriptions/resourceGroups/resources/read"
    ],
    "notActions":[],
    "dataActions":[],
    "notDataActions":[]
  }
}

```

Example 5: Create Custom Role for GigaVUE-FM to configure Inline V Series in Azure

```

{
  "Name":"CustomRoleForInline",
  "description":"Minimum requirements for FM in inline tapping",
  "assignableScopes":[
    "/subscriptions/<Subscription ID>"
  ],
  "permissions":[
    {
      "actions":[
        "Microsoft.Resources/subscriptions/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/loadBalancers/read",
        "Microsoft.Network/loadBalancers/backendAddressPools/read",
        "Microsoft.Network/loadBalancers/backendAddressPools/backendPoolAddresses/read",
        "Microsoft.Compute/virtualMachineScaleSets/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/networkInterfaces/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/networkInterfaces/ipConfigurations/read",
        "Microsoft.Compute/virtualMachines/read"
      ],
      "notActions":[],
      "dataActions":[],
      "notDataActions":[]
    }
  ]
}

```


NOTE: Add the `Microsoft.Compute/sshPublicKeys/read` permission only when you want GigaVUE-FM to pull SSH keys from Azure. This permission is optional and not required for all deployments.

To add a role assignment, refer to [Steps to assign an Azure role](#).

DISCLAIMER: These are general guidelines for enabling a deployment in Azure. Since the Azure interface is subject to change and is outside Gigamon's purview, please see Azure documentation for instructions on using Azure.

Install GigaVUE-FM on Azure

You can launch GigaVUE-FM from the Azure VM dashboard or Azure Marketplace.

Install GigaVUE-FM Using Azure VM Dashboard

To install,

1. Go to **Azure VM Dashboard > Virtual Machines** and select **Create** to create an Azure Virtual Machine.

For details, refer to [Create a Linux virtual machine in the Azure](#) in Azure Documentation.

2. Enter the details as mentioned in [Table 1: GigaVUE-FM Installation Steps](#).

Install GigaVUE-FM Using Azure Marketplace

You can install GigaVUE-FM using the Azure Marketplace.

To install,

1. Go to Azure Marketplace and search for Gigamon.

The latest version of Gigamon GigaVUE Cloud Suite for Azure appears.

2. Open the latest version of GigaVUE-FM.
3. Review and accept the terms for Gigamon GigaVUE Cloud Suite for Azure.

For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).

4. For details, refer to [Create a Linux virtual machine in the Azure](#) in Azure Documentation.
5. Enter the details as mentioned in [Table 1: GigaVUE-FM Installation Steps](#).

The following table describes the important fields.

Table 1: GigaVUE-FM Installation Steps

Field	Description
Basics	
Subscription	Select your subscription.
Resource Group	Select an existing resource group or create a new resource group. For more information, refer to Create a resource group topic in the Azure Documentation.
System-assigned managed identity	Use a system-assigned managed identity when a resource needs to authenticate to other services, and you want the identity to be created and deleted with the resource. NOTE: If you update any role it would take more than an hour to reflect in GigaVUE-FM, however, if you use APP registration it would take between 5-10 minutes to update in GigaVUE-FM.
Virtual machine name	Enter a name for the VM.
Region	Select a region for Azure VM.
Availability Zone	Choose your availability zone
Security Type	To enable UEFI secure boot, select Trusted launch virtual machines from the drop-down list. Click Configure security features and ensure that the Enable secure boot check box is enabled.
Image	Select the latest GigaVUE-FM images. Note: You cannot select multiple images for a VM. Refer to Configure GigaVUE Fabric Components in Azure for more details on configuring GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller in Azure.
Size	The recommended instance types are as follows: - GigaVUE-FM - Standard_D4s_v3 - UCT-V Controller - Standard_B4ms - V Series Node - Standard_D4s_v4 - V Series Proxy - Standard_Bms
Authentication Type	We support only SSH public key authentication type - SSH public key - Enter the administrator username for the VM. - Enter the SSH public key pair name. - Password - Enter the administrator username for the VM. - Enter the administrator password.
Disks	
Disk Size	The required disk size for GigaVUE-FM is 2 x 40GB .

Field	Description
Networking	
Virtual Network	Select an existing VNet or create a new VNet. For more information, refer to Create a virtual network topic in the Azure Documentation. On selecting an existing VNet, the Subnet and the Public IP values are auto-populated.
Configure network security group	Select an existing network security group or create a new network security group. For more information, refer to Network Security Groups . Configure the Network Security Group to allow GigaVUE-FM to communicate with the rest of the components.

NOTE: Verify the summary before proceeding to create. It will take several minutes for the VM to initialize. After the initialization is completed, you can verify the VM through the Web interface.

After the deployment, navigate to the VM overview page, copy the **Public IP address**, and paste it in a new web browser tab.

If GigaVUE-FM is deployed in Azure, use **admin123A!!** as the password for the **admin** user to login to GigaVUE-FM. You must change the default password after logging in to GigaVUE-FM.

Upgrade GigaVUE-FM in Azure

This chapter describes how to upgrade the GigaVUE-FM instance on Azure.

Refer to the following sections for details:

- [Upgrade GigaVUE-FM from the UI](#)
- [Upgrade GigaVUE-FM using Snapshot in Azure](#)

Upgrade GigaVUE-FM from the UI

This section describes the steps to upgrade GigaVUE-FM from the UI. You can upgrade by using an image that is located on an external image server, or you can use GigaVUE-FM as the image server.

Keep in mind the following rules and notes before you upgrade GigaVUE-FM from the UI:

- In place GigaVUE-FM upgrade is supported from GigaVUE-FM version 6.6 to a later releases.
- When using the GigaVUE-FM UI to upgrade GigaVUE-FM, you can only upgrade to the currently available version or to the next version. You cannot downgrade.

- When using the Firefox browser, clear the cache before upgrading to prevent issues with the browser.
- When upgrading GigaVUE-FM to the latest image file, the upgrade operation will fail if GigaVUE-FM already has three image files. Delete the unused image files before proceeding to download a new image file.

NOTE: You can enable auto-purge of GigaVUE-FM images using an API. Contact Gigamon Technical Support for assistance.

Upgrade from External Image Server

This section provides the steps for upgrading the GigaVUE-FM from an image stored on an external server. The image can be transferred from the server to the GigaVUE-FM using SCP.

To upgrade with an image stored on an external image server, do the following:

1. Upload the image to the external image server to make it available to GigaVUE-FM.
To obtain software images, register on the [Gigamon Customer Portal](#) and download the software.
2. Add the image server to GigaVUE-FM. This stores the credentials, image file name, and IP address of the server on GigaVUE-FM.

To add the image server:

- a. Go to **Settings**.
- b. Select **System** on the left navigation panel and go to **Images > External Servers**. The External Servers page appears.
- c. Click **Add**. The Add External Server page appears.
- d. On the Images Server page, specify the following:
 - An alias to help identify the image server.
 - The host IP address of the server.
 - The protocol to use for the download: SCP
 - The user name and password.
- e. Click **Save**.

The External Server page displays the newly added external server. From the **Admin** drop-down list in the top right corner of the window, select **Upgrade** to open the GigaVUE-FM Image Upgrade page.

To monitor the progress and status of the upgrade, click  "GigaVUE Administration Guide".

3. On the GigaVUE-FM Image Upgrade page, click on the Image Server field and select the server added in [Step 2](#).

4. In the Image File Path, enter the image path and filename on the external file server.
5. Upgrade any deployed GigaVUE Cloud Suite-VMs.

After upgrading GigaVUE-FM, you must also upgrade any deployed GigaVUE Cloud Suite-VMs. Otherwise, maps may not work and the GigaVUE Cloud Suite-VMs will be unreachable. For information about upgrading GigaVUE Cloud Suite-VM, refer to the “*Bulk Upgrading GigaVUE-VM Nodes*” section in the *GigaVUE Cloud Suite Deployment Guide - VMware*.

Upgrade with GigaVUE-FM as the Image Server

This section provides the steps for upgrading GigaVUE-FM when GigaVUE-FM is used as the file server.

To upgrade a GigaVUE-FM using internal image files, perform the following:

1. Download the images from the Gigamon website and store them in the instance where GigaVUE-FM is installed.

To obtain software images, register on the [Gigamon Customer Portal](#) and download the software.

2. Upload the images file to GigaVUE-FM.
 - a. Go to **Settings** and select to **System > Images > Internal Image Files**.
 - b. On the Internal Image File page, click **Upload**.
 - c. Click **Browse** to locate the image file.
 - d. Click **OK** to upload the file. The page displays the progress of the upload.

After the upload completes, you can see the GigaVUE-FM image to use for the upgrade on the Internal Images Files page.

3. Click the **Admin** drop-down list on the top right of the window and select **Upgrade**.
4. On the GigaVUE-FM Image Upgrade page, click in the **Image Server** field and select **Internal Image Server**.

Software Upgrade

Image File

Select the file server where you have uploaded the image file.

Image Server:

Version:

Image File:

- From the **Version** drop-down list, select the version to which you are upgrading.

NOTE: You can only upgrade to another instance of the current version or the immediate next version. Downgrading to a lower version is not supported through the UI.

- Click **Upgrade**.

NOTE: When you change the IP address of the GigaVUE-FM instance using the jump-start configuration, the internal database and the in-memory caches of the GigaVUE-FM instance are not updated. The Database continues to have the IP address of the old GigaVUE-FM, and the image upgrade using the internal server option does not work. To fix this, you must restart the GigaVUE-FM instance after upgrade.

NOTE: If you are using FireFox, clear the cache before upgrading to prevent issues with the browser.

Upgrade GigaVUE-FM using Snapshot in Azure

This chapter describes how to upgrade the GigaVUE-FM instance deployed on Azure using snapshot:

- [Upgrade the GigaVUE-FM using the Snapshot method](#)
- [Stop GigaVUE-FM Instance](#)
- [Create Snapshot of the GigaVUE-FM Instance](#)
- [Upgrade GigaVUE-FM Instance](#)

Upgrade the GigaVUE-FM using the Snapshot method

To upgrade the GigaVUE-FM instance successfully, you must perform the following steps:

1. Stop the existing version of the GigaVUE-FM instance.
2. Create a snapshot of the data disk of the current version of the GigaVUE-FM instance.
3. Deploy the latest version of the GigaVUE-FM instance from the marketplace.
4. In the disk tab, select **Create and attach a new disk**, and attach the Snapshot.
5. Launch the GigaVUE-FM and wait for it to initialize.

Stop GigaVUE-FM Instance

Before upgrading the GigaVUE-FM instance, the existing version of the GigaVUE-FM instance must be stopped.

NOTE: Do not terminate the GigaVUE-FM instance.

To stop the GigaVUE-FM instance:

1. Login to the Azure portal and select the **Virtual Machines** that have the GigaVUE-FM deployed
2. Click the Stop button in the portal to shutdown GigaVUE-FM
3. Wait for few seconds for Azure to completely stop the virtual machine

Create Snapshot of the GigaVUE-FM Instance

You must create a snapshot of the blob of the existing version of the GigaVUE-FM instance. Snapshots capture data that are written to your Azure Virtual Disks at the time the snapshot is taken. This excludes any data that are cached by any applications or the operating system.

To create a snapshot:

1. Click the Data Disk link in the existing version of the GigaVUE-FM properties in the Azure portal.
2. Click Create Snapshot in the screen. The Create Snapshot (using SSD) dialog box is displayed.
3. In the Create Snapshot dialog box, enter the following information:

Table 1: Fields for Creating a Snapshot

Field	Description
Name	The name of the snapshot
Subscription	Subscription ID
Resource-group	The name of the resource group
Location	Location of the resource group
Account-type	Select the account type to use to store the snapshot.

4. Click **Create**. It will take several minutes for the snapshot to be created.

Upgrade GigaVUE-FM Instance

To upgrade the GigaVUE-FM instance:

1. In the Azure Marketplace, search for the GigaVUE-FM.
2. Select the latest version GigaVUE-FM of and then click **Create**.
3. On the **Basic** page, configure basic information for the GigaVUE-FM and then click **Next**.
4. In the **disk** tab, select **Create and attach a new disk**.
5. In the **Create a new disk page**, enter the name.
6. Select the **Source type** as **Snapshot** and then select the **Snapshot**. For more information, see [Create Snapshot of the GigaVUE-FM Instance](#).
7. Enable the option, **Delete disk with VM** and click **Ok**
8. On the **Networking Page**, configure the network settings and then click **Next**.
9. On the **Management page**, enable System assigned managed Identity and then click **Next**.
10. Configure the **Advanced page** as required and then click **Review and Create**.

Upgrade GigaVUE Fabric Components in GigaVUE-FM for Azure

This chapter describes how to upgrade GigaVUE V Series Proxy and GigaVUE V Series Node. For more detailed information about UCT-V Controller, GigaVUE V Series Proxy and Node Version, refer to the *GigaVUE-FM Version Compatibility* section in the [Prerequisites for GigaVUE Cloud Suite for Azure](#).



IMPORTANT NOTE:



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

- Create Token in GigaVUE-FM for UCT-V Installation and update it in the configuration file.
- Create Tokens for deploying the Fabric Components using Third Party Orchestration. For more details, refer to [Token-based Authentication](#).
- Open the required ports in the cloud platform. For more details, refer to [Network Firewall Requirement for GigaVUE Cloud Suite](#).

For more details, refer for more information:

- [Prerequisite](#)
- [Upgrade UCT-V Controller](#)
- [Upgrade GigaVUE V Series Node and GigaVUE V Series Proxy](#)

Prerequisite

Before you upgrade the GigaVUE V Series Proxy and GigaVUE V Series Node, you must upgrade GigaVUE-FM to software version 5.13.01 or above.

Upgrade UCT-V Controller

NOTE: UCT-V Controllers cannot be upgraded. Only a new version that is compatible with the UCT-V's version can be added or removed in the **Azure Fabric Launch Configuration** page.

To change the UCT-V Controller version follow the steps given below:

To change UCT-V Controller version between different major versions

NOTE: You can only add UCT-V Controllers which has different major versions. For example, you can only add UCT-V Controller version 1.8-x if your existing version is 1.7-x.

- In the **Azure Fabric Launch Configuration** page, under **Controller Versions**, click **Add**.
- From the **Image** drop-down list, select a UCT-V Controller image that matches with the version number of UCT-Vs installed in the instances.
- From the **Size** drop-down list, select a size for the UCT-V Controller. The default size is Standard_B1s.

- d. In **Number of Instances**, specify the number of UCT-V Controllers to launch. The minimum number you can specify is 1.

You cannot change the IP Address Type and the Additional Subnets details, provided at the time of UCT-V Controller configuration.

After installing the new version of UCT-V Controller, follow these steps:

1. Install UCT-V with the version same as the UCT-V Controller.
2. Delete the UCT-V Controller with older version.

Change UCT-V Controller version with in the same major version

NOTE: This is only applicable, if you wish to change your UCT-V Controller version from one minor version to another with in the same major version. For example, from 1.8-2 to 1.8-3.

- a. From the **Image** drop-down list, select a UCT-V Controller image with in the same major version.
- b. Specify the **Number of Instances**. The minimum number you can specify is 1.
- c. Select the **Subnet** from the drop-down.



- You cannot modify the rest of the fields.
- After installing the new version of UCT-V Controller, install the UCT-V with the same version.

Upgrade GigaVUE V Series Node and GigaVUE V Series Proxy

GigaVUE-FM lets you upgrade GigaVUE V Series Proxy and GigaVUE V Series Node at a time.

You can upgrade the GigaVUE V Series Proxy and Node using the following options:

- Launch and replace the complete set of nodes and proxys at a time.


For example, if you have 1 GigaVUE V Series Proxy and 10 GigaVUE V Series Nodes in your VNet, you can upgrade all of them at once. First, the new version of GigaVUE V Series controller is launched. Next, the new version of GigaVUE V Series nodes are launched. Then, the old version of V Series controller and nodes are deleted from the VNet.

NOTES:

- When the new version of node and proxy is launched, the old version still exists in the VNet until they are deleted. Make sure the instance type determined during the configuration can accommodate the total number of new and old instances present in the VNet. If the instance type cannot support so many instances, you can choose to upgrade in multiple batches.
- If an error happens while upgrading the complete set of proxys and nodes present in the VNet, the new version of the fabric is immediately deleted and the old version of the fabric is retained as before.
- If you have deployed your nodes using Public IP address while creating the monitoring domain, then select the same number of Public IP addresses defined in your Max Instances when upgrading your nodes. For details, refer to *Create Monitoring Domain* in GigaVUE Cloud Suite Deployment Guide - Azure.
- Launch and replace the nodes and proxy in multiple batches.
For example, if you need to upgrade 18 GigaVUE V Series Nodes, specify how many you want to upgrade per batch.

To upgrade the GigaVUE V Series Proxy and GigaVUE V Series Node,

1. Go to **Inventory > VIRTUAL > Azure**, and then select **Monitoring Domain**. The **Monitoring Domain** page appears.
2. On the Monitoring Domain page, select the connection name check box and select **Actions**



3. Select **Upgrade Fabric** from the drop-down list. The Fabric Nodes Upgrade page is displayed.

Fabric Nodes Upgrade

V Series Proxy

Upgrade	☒
Current Version	2.3.0
Image	gigamon-gigavue-vseries-proxy-2.3.2-284364
Change Size	☐
Batch Size	1

V Series Node

Upgrade	☒
Current Version	2.3.0
Image	gigamon-gigavue-vseries-node-2.3.2-284421
Change Size	☐
Batch Size	1
Public IPs	104.42.58.54 104.42.183.63

UpgradeCancel

4. Select the **Upgrade** check box, upgrade the GigaVUE V Series Node/Proxy.

5. From the **Image** drop-down list, select the latest version of the GigaVUE V Series Proxy/Nodes.
6. Select the **Change Size** check box to change the flavor of the node/proxy, only if required.
7. Specify the batch size in the **Batch Size** box to upgrade the GigaVUE V Series Node/Proxy.

For example, if you have 7 GigaVUE V Series Nodes, specify 7 as the batch size and upgrade all of them. Alternatively, you can specify 3 as the batch size, and launch and replace 3 V Series nodes in each batch. In the last batch, the remaining 1 V Series node is launched.

8. From the Public IPs drop-down list, select the IP addresses equal to the Max Instances defined when creating a monitoring domain.

NOTE: This is only applicable for nodes deployed using Public IP, when creating a monitoring domain.

9. Select **Upgrade**.

The upgrade process takes a while depending on the number of GigaVUE V Series Proxys and Nodes upgrading in your Azure environment. First, the new version of the GigaVUE V Series Proxy is launched. Next, the new version of GigaVUE V Series Nodes is launched. Then, the older version of both is deleted from the project. The monitoring session is deployed automatically.

To view the detailed upgrade status, select **Upgrade in progress** or **Upgrade successful**, the **V Series Node Upgrade Status** dialog box appears.

Fabric Nodes Upgrade Status

Monitoring Domain: md

Start Time

2021-10-11 20:58:56

End Time

2021-10-11 21:04:03

Status

Fabric upgrade completed successfully

	Proxies	Nodes
Total	1	1
Upgraded	1	1
Upgrading	0	0
Remaining	0	0
Failures	0	0

Clear

Close

- Click **Clear** to delete the monitoring domain upgrade status history of successfully upgraded nodes.

Deploy

Deploy GigaVUE Cloud Suite for Azure using GigaVUE-FM Orchestration

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

Refer to the following deployment options for more detailed information:

- [Acquire Traffic using UCT-V as Traffic Acquisition Method](#)
- [Acquire Traffic using vTAP as Traffic Acquisition Method](#)
- [Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method](#)
- [Acquire Traffic using Inline V Series Solution](#)

Acquire Traffic using UCT-V as Traffic Acquisition Method

This section outlines the workflow to acquire traffic with UCT-V and deploy GigaVUE Fabric Components using Azure. Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Obtain GigaVUE-FM Image](#)
- [Acquire Traffic using UCT-V as Traffic Acquisition Method](#)
- [Check the Required Permissions and Privileges](#)
- [Install UCT-Vs](#)
- [Create the Azure Credentials](#)
- [Create Monitoring Domain](#)
- [Configure GigaVUE Fabric Components in GigaVUE-FM](#)
- [Acquire Traffic using UCT-V as Traffic Acquisition Method](#)

Obtain GigaVUE-FM Image

The image for the GigaVUE Cloud Suite Cloud is available in both the Azure Public Cloud and in the Azure Government portal.

GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud

GigaVUE Cloud Suite Cloud is available in the Azure Marketplace with the Volume Based License options.

GigaVUE Cloud Suite Cloud Suite in Azure Government

Azure Government is an isolated Azure region that contains specific regulatory and compliance requirements of the US government agencies.

To monitor the VMs that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the Azure Government (US) Region, the Azure Government solution provides the same robust features in Azure Government as in the Azure public cloud.

Install GigaVUE-FM on Azure

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

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Check the Required Permissions and Privileges

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. Refer to [Permissions and Privileges \(Azure\)](#) for more details.

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](#)
- [Acquire Traffic using UCT-V as Traffic Acquisition Method](#)

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

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

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

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

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

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

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

sudo uctv-wizard

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

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

Options	Use Command	Description
pre-check	sudo uctv-wizard pre-check	Checks the status of the required packages and firewall requirements. - If any package is missing, it displays an appropriate message with the missing package details. - If installation includes all the packages, it displays a success message indicating that UCT-V is ready for configuration.
pkg-install	sudo uctv-wizard pkg-install	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-V6.14.00 Debian (.deb) package from the [Gigamon Customer Portal](#). For assistance contact [Gigamon Technical Support](#).
- Copy this package to your instance.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5. Save the file.

6. Restart the UCT-V service.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
$ sudo service uctv restart
```

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

```
$ sudo service uctv status
```



Notes:

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

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

Post Deployment Check:

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

1. Enter the command:

```
sudo uctvl uctv-show
```

2. Manually execute the following command:

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

Windows UCT-V Installation

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

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

- [Windows Network Firewall Requirements](#)
- [Install Linux UCT-Vs](#)
- [Acquire Traffic using UCT-V as Traffic Acquisition Method](#)

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 [Gigamon 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 Install Linux UCT-Vs using Manual Configuration.
adapter-restore	uctv-wizard adapter-restore	Using this command, you can restore the backup copy of the network adapter buffer size configuration saved in the in the uctv-wizard adapter-setup step. NOTE: You need to manually restart the network adapters for changes to take effect immediately. You can choose between the following: • If you wish to skip the prompts for restoring the buffer size of the compatible network adapters, enter the option as y.

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

**Notes:**

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

Windows UCT-V Installation Scenarios

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

Install Windows UCT-Vs using Manual Configuration

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

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

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



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

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

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

For IPv4:

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

For IPv6:

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

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

For IPv4:

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

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

For IPv6:

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

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

4. Save the file.

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

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

Registration:
token: <Enter the token created in GigaVUE-FM>

- a.

Create the Azure Credentials

You can monitor workloads across multiple Azure subscriptions within one monitoring domain. All the deployed GigaVUE fabric components are shared among many Azure subscriptions to reduce the cost. Earlier, each Azure subscription carried a set of GigaVUE fabric components.

- After launching GigaVUE-FM in Azure, the **Managed Identity** authentication credential is automatically added to the Azure Credential page as the default credential.
- You can only add the **Application ID with Client Secret** authentication credentials to the Azure Credential page.

To create Azure credentials,

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Credential**.
3. In the Azure Credential page, select **Add**.

The **Configure Credential** wizard appears.

Configure Credential

Save Cancel

Name* Credential Name

Authentication Type Application ID with Client Secret

Tenant ID* Tenant ID

Application ID* Application ID

Application Secret* Application Secret

Azure Environment Azure Environment...

Azure

AZURE_US_GOVERNMENT

4. Enter or select the appropriate information for the Azure credential:

a. **Name:** An alias used to identify the Azure credential.

b. **Authentication Type:**

Application ID with Client Secret: Connection with Azure with a service principal. Enter the values for the following fields.

- **Tenant ID:** A unique identifier of the Azure Active Directory instance.
- **Application ID:** A unique identifier of an application in Azure platform.
- **Application Secret:** a password or key to request tokens.

For details on how to create service principal and assign custom roles, refer to [Application ID with client secret](#)

c. **Azure Environment:** Select an Azure environment where your workloads are located. For example, Azure_US_Government.

5. Select **Save**.

You can view the list of available credentials in the Azure Credential page.

Create Monitoring Domain

Before configuring, you must establish a connection between GigaVUE-FM and your Azure environment. A Monitoring Domain in GigaVUE-FM allows you to define and manage this connection. Once established, GigaVUE-FM can deploy and manage UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes within your specified VNets and Resource Groups.

GigaVUE-FM connects to Azure using either an Application ID and Client Secret (Service Principal) or the Managed Service Identity (MSI) authentication method.

To create an Azure monitoring domain in GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**

2. Select **Monitoring Domain**.

The **Monitoring Domain** page appears.

3. In the Monitoring Domain page, select **New**.

The **Azure Monitoring Domain Configuration** wizard appears.

4. Enter or select the appropriate information for the Monitoring Domain:

- **Monitoring Domain:** An alias used to identify the monitoring domain.
- **Traffic Acquisition Method:** Select one of the following Tapping methods:
 - **UCT-V:** If you select UCT-V as the tapping method, the traffic is acquired from the UCT-Vs installed on your standard VMs in the Resource Group or in the Scale Sets. Then, the acquired traffic is forwarded to the GigaVUE V Series nodes. You must configure the UCT-V Controller to monitor the UCT-Vs.
- **Traffic Acquisition Tunnel MTU:** 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 default value is 1450.

- When using IPv4 tunnels, the maximum MTU value is 1450. Ensure that the UCT-V tunnel MTU is 50 bytes less than the UCT-V destination interface MTU size.
- When using IPv6 tunnels, the maximum MTU value is 1430. Ensure that the UCT-V tunnel MTU is 70 bytes less than the UCT-V destination interface MTU size.
- Enable the **Use FM to Launch Fabric** toggle, to deploy GigaVUE Fabric Components using GigaVUE-FM.
- Enable IPv6 Preference: Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes. This option appears only when **Use FM to Launch Fabric** is disabled and Traffic Acquisition Method is UCT-V.

- **Connections:**

Connections

- A Monitoring Domain can have multiple connections, however only one connection can have **Managed Service Identity** as the **Credential**.
- The connections in a monitoring domain can be a combination of multiple **Application ID with Client Secret** (Service Principal) accounts, or one **Managed Service Identity** and multiple **Application ID with Client Secret** (Service Principal) accounts.
- Each connection can have only one **Subscription ID**.
- **Name:** An alias used to identify the connection.
- **Credential:** Select an Azure credential. For details, refer to [Create Azure Credentials](#).
- **Subscription ID:** A unique alphanumeric string that identifies your Azure subscription.
- **Region:** Azure region for the monitoring domain. For example, West India.
- **Resource Groups:** Select the Resource Groups of the corresponding VMs to monitor.

**Notes:**

- This field is not applicable if you select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
- When you remove and re-add a resource group in IAM, it won't appear in GigaVUE-FM automatically. To refresh the list and make the resource group selectable, reselect the Subscription ID from the drop down. This action triggers the UI to reload the resource groups associated with the selected subscription.

5. Select **Save**.

The **Azure Fabric Launch Configuration** wizard appears.

**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 only view and delete the existing configuration for GigaVUE V Series Node 1. You cannot perform any other actions on the existing configuration for GigaVUE V Series Node 1 as the features are deprecated from GigaVUE-FM.

You can use the **Copy** button to copy the permissions to the clipboard. Also, you can use the **Download** button to download the permission in JSON format.

NOTE: After updating the IAM Policy, it takes around 5 minutes for the changes to reflect on the Check Permissions screen.

Configure GigaVUE Fabric Components in GigaVUE-FM

After configuring the Monitoring Domain, you reach the Azure Fabric Launch Configuration page. In the same page, you can configure all the GigaVUE fabric components.

Enter or select the required information as described in the following table.

Fields	Description
Connections	A connection that you created in the monitoring domain page. For details, refer to Create Monitoring Domain .
Centralized Virtual Network	Alias of the centralized VNet in which the UCT-V Controllers, V Series Proxies, and the GigaVUE V Series nodes are launched.
Authentication Type	Select SSH Public Key as the Authentication Type to connect with the Centralized VNet.
SSH Public Key	The SSH public key for the GigaVUE fabric components.
Resource Group	The Resource Groups created in Azure for communication between the controllers, nodes, and GigaVUE-FM.
Security Groups	The security group created for the GigaVUE fabric components.
Enable Custom Certificates	Enable this option to validate the custom certificate during SSL Communication. GigaVUE-FM validates the Custom certificate with the trust store. If the certificate is not available in Trust Store, communication does not happen, and an handshake error occurs. Note: If the certificate expires after the successful deployment of the fabric components, then the fabric components moves to failed state.

Fields	Description
Certificate	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes, GigaVUE V Series Proxy, and UCT-V Controllers. For details, refer to Secure Communication between GigaVUE Fabric Components .
Prefer IPv6	Enables IPv6 to deploy all the Fabric Controllers, and the tunnel between hypervisor to GigaVUE V Series Nodes using IPv6 address. If the IPv6 address is unavailable, it uses an IPv4 address. Note: This option can be enabled only when deploying a new GigaVUE V Series Node. If you wish to enable this option after deploying the GigaVUE V Series Node, then you must delete the existing GigaVUE V Series Node and deploy it again with this option enabled.

Select **Yes** to configure V Series Proxy for the monitoring domain. For details, refer to [Configure GigaVUE V Series Proxy](#)

Azure Fabric Launch Configuration Check Permissions Save C

Connections
Centralized Virtual Network
Authentication Type
SSH Public Key
Resource Group
Security Groups
Enable Custom Certificates
Prefer IPv6
Configure a V Series Proxy

Select a Connection
Select a Virtual Network
sshPublicKey
Enter your SSH Public Key
Select resource group...
Select management subnet security group...
☐ Disabled
☐ No
☐ No

UCT-V Controller ①

Add
Image
Size
Number of Instances
IP Address Type
Subnet
Agent CA
Additional Subnets
Tags

V Series Node

Add
Image
Size
Disk Size (GB)
IP Address Type
Subnet
Data Subnets
Tags
Min Number of Instances
Max Number of Instances



To deploy GigaVUE fabric components (GigaVUE V Series Nodes, UCT-V Controller, and GigaVUE V Series Proxies) in GigaVUE-FM, you must accept the terms of the GigaVUE fabric components from the Azure marketplace using the Azure CLI or PowerShell. For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).

For details, refer to the following topics:

- [Configure UCT-V Controller](#)
- [Configure GigaVUE V Series Proxy](#)
- [Configure GigaVUE V Series Node](#)

Configure UCT-V Controller

A UCT-V Controller manages multiple UCT-Vs and orchestrates the flow of mirrored traffic to GigaVUE V Series nodes.

NOTE: A single UCT-V Controller can manage up to 500 UCT-Vs. The recommended minimum instance type is Standard_B4ms for UCT-V Controller.

A UCT-V Controller can only manage UCT-Vs that has the same version.

To configure the UCT-V Controllers,

NOTE: You can configure UCT-V Controller only if your **Traffic Acquisition Method** is **UCT-V**.

In the **Azure Fabric Launch Configuration** page, enter or select the appropriate values for the UCT-V Controller as described in the following table.

Controller Version(s)	Add
	Image 1.8.4 ▼ Size Standard_B1... ▼ Number of Instances 1
Management Subnet	IP Address Type ☒ Private ☐ Public Subnet mgmt ▼
Additional Subnets	Add Subnet
Tags	Add

Fields	Description
Controller Version(s)	Ensure that the UCT-V Controller version you configure is always the same as the UCT-Vs' version number deployed in the VM machines. If multiple versions of UCT-Vs are deployed in the VM machines, then you must configure multiple versions of UCT-V Controllers that matches the version numbers of the UCT-Vs. Note: If the version of UCT-V Controllers doesn't match with UCT-Vs, GigaVUE-FM cannot detect the agents in the instances. To add UCT-V Controllers: - Under Controller Versions, click Add. - From the Image drop-down list, select a UCT-V Controller image that matches with the version number of UCT-Vs installed in the instances. - From the Size drop-down list, select a size for the UCT-V Controller. The default size is Standard_B1s. - In Number of Instances, specify the number of UCT-V Controllers to launch. The minimum number you can specify is 1.
Management Subnet	IP Address Type: Select one of the following IP address types: - Select Private if you want to assign an IP address that is not reachable over Internet. You can use private IP address for communication between the UCT-V Controller instances and GigaVUE-FM instances in the same network. - Select Public if you want the IP address to be assigned from Azure's pool of public IP address. The public IP address gets changed every time the instance is stopped and restarted. On selecting Public IP address type, you must select all the required Public IPs. Subnet: Select a Subnet for UCT-V Controller. The subnet that is used for communication between the UCT-V Controllers and the UCT-Vs, as well as to communicate with GigaVUE-FM. Every fabric component (both controllers and the nodes) needs a way to talk to each other and GigaVUE-FM. So, they should share at least one management plane/subnet. Note: Some instance types are supported in Azure platform. For details, refer to Microsoft Azure documentation to learn on supported instance types.
Agent Tunnel Type	The type of tunnel used for sending the traffic from UCT-Vs to GigaVUE V Series Nodes.

Fields	Description
Agent Tunnel CA	The Certificate Authority (CA) that should be used in the UCT-V Controller for connecting the tunnel.
Additional Subnet(s)	(Optional) If UCT-Vs are on subnets that are not IP routable from the management subnet, you must specify additional subnets so that the UCT-V Controller can communicate with all the UCT-Vs. Select Add to specify additional data subnets, if needed. Also, make sure that you specify a list of security groups for each additional subnet.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your Azure environment. For example, you might have UCT-V Controllers deployed in many regions. To distinguish these UCT-V Controllers based on the regions, you can provide a name that is easy to identify such as us-west-2-uctv-controllers. To add a tag: - Click Add. - In the Key field, enter the key. For example, enter Name. - In the Value field, enter the key value. For example, us-west-2-uctv-controllers.

Configure GigaVUE V Series Proxy

GigaVUE V Series Proxy can manage multiple GigaVUE V Series Nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Proxies to communicate with the GigaVUE V Series nodes.

NOTE: A single GigaVUE V Series Proxy can manage up to 100 GigaVUE V Series nodes. The recommended minimum instance type is Standard_B1s for V Series Proxy.

To configure the GigaVUE V Series Proxy,

1. In the **Azure Fabric Launch Configuration** page, select **Yes to Configure a V Series Proxy**. The GigaVUE V Series Proxy fields appear.
2. Enter or select the appropriate values for the V Series Proxy. For details, refer to the [UCT-V Controller field descriptions](#).

Configure GigaVUE V Series Node

GigaVUE V Series node is a visibility node that aggregates mirrored traffic from multiple UCT-Vs. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaVUE Cloud Suite for Azure using the standard VXLAN tunnels.

To launch a GigaVUE V Series node:

In the **Azure Fabric Launch Configuration** page, enter or select the appropriate values for the GigaVUE V Series Node.

V Series Node

Image	gigavue-gigavue-series-node-3.10-340871
Size	Standard_D4s_v4 8 NICs
Disk Size (GB)	>= 30
IP Address Type	☒ Private ☐ Public
Management Subnet	Subnet mgmt
Data Subnets	Add Subnet
Tool Subnet	☒ Tool Subnet ⓘ
Subnet 1	dataout
Security Groups	sg-1234567890 x
Tags	Add

Fields	Description
Image	From the Image drop-down list, select a GigaVUE V Series Node image.
Size	From the Size down-down list, select a size for the GigaVUE V Series Node. The default size for GigaVUE V Series Node configuration is Standard_D4s_v4 .
Disk Size (GB)	The size of the storage disk. The default disk size is 30GB. Note: When using Application Metadata Exporter, the minimum recommended Disk Size is 80GB.
IP Address Type	Select one of the following IP address types: - Select Private if you want to assign an IP address that is not reachable over Internet. You can use private IP address for communication between the GigaVUE V Series Node instances and GigaVUE-FM instances in the same network. - Select Public if you want the IP address to be assigned from Azure's pool of public IP address. On selecting Public IP address type, you must select the number of Public IPs defined in the Maximum Instance.
Management Subnet	Subnet: Select a management subnet for GigaVUE V Series. The subnet that is used for communication between the UCT-Vs and the GigaVUE V Series Nodes, as well as to communicate with GigaVUE-FM. Every fabric component (both controllers and the nodes) need a way to talk to each other and GigaVUE-FM. So, they should share at least one management

Fields	Description
	plane/subnet.
Data Subnet(s)	The subnet that receives the mirrored VXLAN tunnel traffic from the UCT-Vs. Select a Subnet and the respective Security Groups. Click Add to add additional data subnets. Note: Using the Tool Subnet checkbox you can indicate the subnets to be used by the GigaVUE V Series Node to egress the aggregated/manipulated traffic to the tools.
Tag(s)	(Optional) The key name and value that helps to identify the GigaVUE V Series Node instances in your Azure environment. For example, you might have GigaVUE V Series Nodes deployed in many regions. To distinguish these GigaVUE V Series Nodes based on the regions, you can provide a name that is easy to identify. To add a tag: - Click Add. - In the Key field, enter the key. For example, enter Name. - In the Value field, enter the key value.
Min Instances	The minimum number of GigaVUE V Series Nodes to be launched in the Azure connection. The minimum number of instances that can be entered is 1. Note: Nodes are launched when a monitoring session is deployed if GigaVUE-FM discovers some targets to monitor. The minimum amount is launched at that time. The GigaVUE-FM will delete the nodes if they are idle for over 15 minutes.
Max Instances	The maximum number of GigaVUE V Series Nodes that can be launched in the Azure connection. When the number of instances per V Series node exceeds the max instances specified in this field, increase the number in the Max Instances to Launch. When additional V Series nodes are launched, GigaVUE-FM re-balances the instances assigned to the nodes. This can result in a brief interruption of traffic.

Click **Save** to complete the Azure Fabric Launch Configuration.

For details regarding the Check permissions to validate IAM and security-group settings, refer to [Check Permissions while Configuring GigaVUE Fabric Components using GigaVUE-FM](#).

What to do Next

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

Acquire Traffic using vTAP as Traffic Acquisition Method

This section outlines the workflow to acquire traffic with vTAP and deploy GigaVUE Fabric Components using Azure. Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Obtain GigaVUE-FM Image](#)
- [Acquire Traffic using vTAP as Traffic Acquisition Method](#)
- [Check the Required Permissions and Privileges](#)
- [Create the Azure Credentials](#)
- [Create Monitoring Domain](#)
- [Configure GigaVUE Fabric Components in GigaVUE-FM](#)
- [Acquire Traffic using vTAP as Traffic Acquisition Method](#)

Obtain GigaVUE-FM Image

The image for the GigaVUE Cloud Suite Cloud is available in both the Azure Public Cloud and in the Azure Government portal.

GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud

GigaVUE Cloud Suite Cloud is available in the Azure Marketplace with the Volume Based License options.

GigaVUE Cloud Suite Cloud Suite in Azure Government

Azure Government is an isolated Azure region that contains specific regulatory and compliance requirements of the US government agencies.

To monitor the VMs that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the Azure Government (US) Region, the Azure Government solution provides the same robust features in Azure Government as in the Azure public cloud.

Install GigaVUE-FM on Azure

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

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Check the Required Permissions and Privileges

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. Refer to [Permissions and Privileges \(Azure\)](#) for more details.

Create the Azure Credentials

You can monitor workloads across multiple Azure subscriptions within one monitoring domain. All the deployed GigaVUE fabric components are shared among many Azure subscriptions to reduce the cost. Earlier, each Azure subscription carried a set of GigaVUE fabric components.



- After launching GigaVUE-FM in Azure, the **Managed Identity** authentication credential is automatically added to the Azure Credential page as the default credential.
- You can only add the **Application ID with Client Secret** authentication credentials to the Azure Credential page.

To create Azure credentials,

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Credential**.
3. In the Azure Credential page, select **Add**.

The **Configure Credential** wizard appears.

Name*	Credential Name
Authentication Type	Application ID with Client Secret
Tenant ID*	Tenant ID
Application ID*	Application ID
Application Secret*	Application Secret
Azure Environment	Azure Environment...

Azure

AZURE_US_GOVERNMENT

4. Enter or select the appropriate information for the Azure credential:

a. **Name:** An alias used to identify the Azure credential.

b. **Authentication Type:**

Application ID with Client Secret: Connection with Azure with a service principal. Enter the values for the following fields.

- **Tenant ID:** A unique identifier of the Azure Active Directory instance.
- **Application ID:** A unique identifier of an application in Azure platform.
- **Application Secret:** a password or key to request tokens.

For details on how to create service principal and assign custom roles, refer to [Application ID with client secret](#)

c. **Azure Environment:** Select an Azure environment where your workloads are located. For example, Azure_US_Government.

5. Select **Save**.

You can view the list of available credentials in the Azure Credential page.

Create Monitoring Domain

Before configuring, you must establish a connection between GigaVUE-FM and your Azure environment. A Monitoring Domain in GigaVUE-FM allows you to define and manage this connection. Once established, GigaVUE-FM can deploy and manage UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes within your specified VNets and Resource Groups.

GigaVUE-FM connects to Azure using either an Application ID and Client Secret (Service Principal) or the Managed Service Identity (MSI) authentication method.

To create an Azure monitoring domain in GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**

2. Select **Monitoring Domain**.

The **Monitoring Domain** page appears.

3. In the Monitoring Domain page, select **New**.

The **Azure Monitoring Domain Configuration** wizard appears.

4. Enter or select the appropriate information for the Monitoring Domain:

- **Monitoring Domain:** An alias used to identify the monitoring domain.
- **Traffic Acquisition Method:** Select one of the following Tapping methods:
 - **vTAP:** If you select vTAP as the tapping method, traffic tapping is performed by the Azure platform and sent to the GigaVUE V Series Node. GigaVUE-FM creates the necessary configurations in Azure to enable this.
- **Traffic Acquisition Tunnel MTU:** 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 default value is 1450.

- When using IPv4 tunnels, the maximum MTU value is 1450. Ensure that the UCT-V tunnel MTU is 50 bytes less than the UCT-V destination interface MTU size.
- When using IPv6 tunnels, the maximum MTU value is 1430. Ensure that the UCT-V tunnel MTU is 70 bytes less than the UCT-V destination interface MTU size.
- Enable the **Use FM to Launch Fabric** toggle, to deploy GigaVUE Fabric Components using GigaVUE-FM.
- Enable IPv6 Preference: Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes. This option appears only when **Use FM to Launch Fabric** is disabled and Traffic Acquisition Method is UCT-V.

- **Connections:**

Connections

▼

Name*

Credential*

Credential Name...
▼

Subscription ID*

Subscription ID...
▼

Region*

Region Name...
▼

Resource Groups*

☒ Discovered
 ☐ Regex ⓘ

Resource Groups...
▼

+
-

- A Monitoring Domain can have multiple connections, however only one connection can have **Managed Service Identity** as the **Credential**.
- The connections in a monitoring domain can be a combination of multiple **Application ID with Client Secret** (Service Principal) accounts, or one **Managed Service Identity** and multiple **Application ID with Client Secret** (Service Principal) accounts.
- Each connection can have only one **Subscription ID**.
- **Name:** An alias used to identify the connection.
- **Credential:** Select an Azure credential. For details, refer to [Create Azure Credentials](#).
- **Subscription ID:** A unique alphanumeric string that identifies your Azure subscription.
- **Region:** Azure region for the monitoring domain. For example, West India.
- **Resource Groups:** Select the Resource Groups of the corresponding VMs to monitor.

**Notes:**

- This field is not applicable if you select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
- When you remove and re-add a resource group in IAM, it won't appear in GigaVUE-FM automatically. To refresh the list and make the resource group selectable, reselect the Subscription ID from the drop down. This action triggers the UI to reload the resource groups associated with the selected subscription.

5. Select **Save**.

The **Azure Fabric Launch Configuration** wizard appears.

**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 only view and delete the existing configuration for GigaVUE V Series Node 1. You cannot perform any other actions on the existing configuration for GigaVUE V Series Node 1 as the features are deprecated from GigaVUE-FM.

You can use the **Copy** button to copy the permissions to the clipboard. Also, you can use the **Download** button to download the permission in JSON format.

NOTE: After updating the IAM Policy, it takes around 5 minutes for the changes to reflect on the Check Permissions screen.

Configure GigaVUE Fabric Components in GigaVUE-FM

After configuring the Monitoring Domain, you reach the Azure Fabric Launch Configuration page. In the same page, you can configure all the GigaVUE fabric components.

Enter or select the required information as described in the following table.

Fields	Description
Connections	A connection that you created in the monitoring domain page. For details, refer to Create Monitoring Domain .
Centralized Virtual Network	Alias of the centralized VNet in which the UCT-V Controllers, V Series Proxies, and the GigaVUE V Series nodes are launched.
Authentication Type	Select SSH Public Key as the Authentication Type to connect with the Centralized VNet.
SSH Public Key	The SSH public key for the GigaVUE fabric components.
Resource Group	The Resource Groups created in Azure for communication between the controllers, nodes, and GigaVUE-FM.
Security Groups	The security group created for the GigaVUE fabric components.
Enable Custom Certificates	Enable this option to validate the custom certificate during SSL Communication. GigaVUE-FM validates the Custom certificate with the trust store. If the certificate is not available in Trust Store, communication does not happen, and an handshake error occurs. Note: If the certificate expires after the successful deployment of the fabric components, then the fabric components moves to failed state.

Fields	Description
Certificate	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes, GigaVUE V Series Proxy, and UCT-V Controllers. For details, refer to Secure Communication between GigaVUE Fabric Components .
Prefer IPv6	Enables IPv6 to deploy all the Fabric Controllers, and the tunnel between hypervisor to GigaVUE V Series Nodes using IPv6 address. If the IPv6 address is unavailable, it uses an IPv4 address. Note: This option can be enabled only when deploying a new GigaVUE V Series Node. If you wish to enable this option after deploying the GigaVUE V Series Node, then you must delete the existing GigaVUE V Series Node and deploy it again with this option enabled.

Select **Yes** to configure V Series Proxy for the monitoring domain. For details, refer to [Configure GigaVUE V Series Proxy](#)

Azure Fabric Launch Configuration Check Permissions Save C

Connections
Centralized Virtual Network
Authentication Type
SSH Public Key
Resource Group
Security Groups
Enable Custom Certificates
Prefer IPv6
Configure a V Series Proxy

Select a Connection
Select a Virtual Network
sshPublicKey
Enter your SSH Public Key
Select resource group...
Select management subnet security group...
☐ Disabled
☐ No
☐ No

UCT-V Controller ①

Add
Image
Size
Number of Instances
IP Address Type
Subnet
Agent CA
Additional Subnets
Tags

V Series Node

Add
Image
Size
Disk Size (GB)
IP Address Type
Subnet
Data Subnets
Tags
Min Number of Instances
Max Number of Instances



To deploy GigaVUE fabric components (GigaVUE V Series Nodes, , and GigaVUE V Series Proxies) in GigaVUE-FM, you must accept the terms of the GigaVUE fabric components from the Azure marketplace using the Azure CLI or PowerShell. For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).

For details, refer to the following topics:

- [Configure GigaVUE V Series Proxy](#)
- [Configure GigaVUE V Series Node](#)

Configure GigaVUE V Series Proxy

GigaVUE V Series Proxy can manage multiple GigaVUE V Series Nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Proxies to communicate with the GigaVUE V Series nodes.

NOTE: A single GigaVUE V Series Proxy can manage up to 100 GigaVUE V Series nodes. The recommended minimum instance type is Standard_B1s for V Series Proxy.

To configure the GigaVUE V Series Proxy,

1. In the **Azure Fabric Launch Configuration** page, select **Yes** to **Configure a V Series Proxy**. The GigaVUE V Series Proxy fields appear.
2. Enter or select the appropriate values for the V Series Proxy. For details, refer to the [UCT-V Controller field descriptions](#).

Configure GigaVUE V Series Node

GigaVUE V Series node is a visibility node that aggregates mirrored traffic from multiple UCT-Vs. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaVUE Cloud Suite for Azure using the standard VXLAN tunnels.

To launch a GigaVUE V Series node:

In the **Azure Fabric Launch Configuration** page, enter or select the appropriate values for the GigaVUE V Series Node.

V Series Node

Image

gigamon-gigamon-series-node-3.10-34GB1 | v

Size

Standard_D4ds_v4 | 16 vCPUs

Disk Size (GB)

>= 30

IP Address Type

☒ Private ☐ Public

Management Subnet

Subnet mgmt | v

Data Subnets

Add Subnet

Tool Subnet

☒ Tool Subnet ⓘ

Subnet 1

dataout | v

Security Groups

sg-xxxxxx x | v

Tags

Add

Fields	Description
Image	From the Image drop-down list, select a GigaVUE V Series Node image.
Size	From the Size down-down list, select a size for the GigaVUE V Series Node. The default size for GigaVUE V Series Node configuration is Standard_D4s_v4 .
Disk Size (GB)	The size of the storage disk. The default disk size is 30GB. Note: When using Application Metadata Exporter, the minimum recommended Disk Size is 80GB.
IP Address Type	Select one of the following IP address types: - Select Private if you want to assign an IP address that is not reachable over Internet. You can use private IP address for communication between the GigaVUE V Series Node instances and GigaVUE-FM instances in the same network. - Select Public if you want the IP address to be assigned from Azure's pool of public IP address. On selecting Public IP address type, you must select the number of Public IPs defined in the Maximum Instance.
Management Subnet	Subnet: Select a management subnet for GigaVUE V Series. The subnet that is used for communication between the UCT-Vs and the GigaVUE V Series Nodes, as well as to communicate with GigaVUE-FM. Every fabric component (both controllers and the nodes) need a way to talk to each other and GigaVUE-FM. So, they should share at least one management

Fields	Description
	plane/subnet.
Data Subnet(s)	The subnet that receives the mirrored VXLAN tunnel traffic from the UCT-Vs. Select a Subnet and the respective Security Groups. Click Add to add additional data subnets. Note: Using the Tool Subnet checkbox you can indicate the subnets to be used by the GigaVUE V Series Node to egress the aggregated/manipulated traffic to the tools.
Tag(s)	(Optional) The key name and value that helps to identify the GigaVUE V Series Node instances in your Azure environment. For example, you might have GigaVUE V Series Nodes deployed in many regions. To distinguish these GigaVUE V Series Nodes based on the regions, you can provide a name that is easy to identify. To add a tag: - Click Add. - In the Key field, enter the key. For example, enter Name. - In the Value field, enter the key value.
Min Instances	The minimum number of GigaVUE V Series Nodes to be launched in the Azure connection. The minimum number of instances that can be entered is 1. Note: Nodes are launched when a monitoring session is deployed if GigaVUE-FM discovers some targets to monitor. The minimum amount is launched at that time. The GigaVUE-FM will delete the nodes if they are idle for over 15 minutes.
Max Instances	The maximum number of GigaVUE V Series Nodes that can be launched in the Azure connection. When the number of instances per V Series node exceeds the max instances specified in this field, increase the number in the Max Instances to Launch. When additional V Series nodes are launched, GigaVUE-FM re-balances the instances assigned to the nodes. This can result in a brief interruption of traffic.

Click **Save** to complete the Azure Fabric Launch Configuration.

For details regarding the Check permissions to validate IAM and security-group settings, refer to [Check Permissions while Configuring GigaVUE Fabric Components using GigaVUE-FM](#).

What to do Next

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

Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method

This section outlines the workflow to acquire traffic with Customer Orchestrated Source and deploy GigaVUE Fabric Components using GigaVUE-FM. In this case you can use tunnels as a source where the traffic is directly tunneled to V Series nodes without deploying UCT-V or UCT-V controllers.

Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Obtain GigaVUE-FM Image](#)
- [Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method](#)
- [Check the Required Permissions and Privileges](#)
- [Create Monitoring Domain](#)
- [Configure GigaVUE Fabric Components in GigaVUE-FM](#)
- [Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method](#)

Obtain GigaVUE-FM Image

The image for the GigaVUE Cloud Suite Cloud is available in both the Azure Public Cloud and in the Azure Government portal.

GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud

GigaVUE Cloud Suite Cloud is available in the Azure Marketplace with the Volume Based License options.

GigaVUE Cloud Suite Cloud Suite in Azure Government

Azure Government is an isolated Azure region that contains specific regulatory and compliance requirements of the US government agencies.

To monitor the VMs that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the Azure Government (US) Region, the Azure Government solution provides the same robust features in Azure Government as in the Azure public cloud.

Install GigaVUE-FM on Azure

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

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Check the Required Permissions and Privileges

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. Refer to [Permissions and Privileges \(Azure\)](#) for more details.

Create Monitoring Domain

Before configuring, you must establish a connection between GigaVUE-FM and your Azure environment. A Monitoring Domain in GigaVUE-FM allows you to define and manage this connection. Once established, GigaVUE-FM can deploy and manage UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes within your specified VNets and Resource Groups.

GigaVUE-FM connects to Azure using either an Application ID and Client Secret (Service Principal) or the Managed Service Identity (MSI) authentication method.

To create an Azure monitoring domain in GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**
2. Select **Monitoring Domain**.

The **Monitoring Domain** page appears.

3. In the Monitoring Domain page, select **New**.

The **Azure Monitoring Domain Configuration** wizard appears.

4. Enter or select the appropriate information for the Monitoring Domain:

- **Monitoring Domain:** An alias used to identify the monitoring domain.
- **Traffic Acquisition Method:** Select one of the following Tapping methods:
 - **Customer Orchestrated Source:** If you select Customer Orchestrated Source as the tapping method, you can select the tunnel as a source where the traffic is directly tunneled to GigaVUE V Series nodes without deploying UCT-Vs or UCT-V Controllers.

NOTE: Select the **Traffic Acquisition Method** as **Customer Orchestrated Source** if you wish to use Application Metadata Exporter (AMX) application.

- **Traffic Acquisition Tunnel MTU:** 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 default value is 1450.

- When using IPv4 tunnels, the maximum MTU value is 1450. Ensure that the UCT-V tunnel MTU is 50 bytes less than the UCT-V destination interface MTU size.
- When using IPv6 tunnels, the maximum MTU value is 1430. Ensure that the UCT-V tunnel MTU is 70 bytes less than the UCT-V destination interface MTU size.
- Enable the **Use FM to Launch Fabric** toggle, to deploy GigaVUE Fabric Components using GigaVUE-FM.
- Enable IPv6 Preference: Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes. This option appears only when **Use FM to Launch Fabric** is disabled and Traffic Acquisition Method is UCT-V.

- **Connections:**

Connections

▼

Name*

Credential*

Subscription ID*

Region*

Resource Groups*

Enter a connection name

Credential Name... | ▼

Subscription ID... | ▼

Region Name... | ▼

☒ Discovered
 ☐ Regex ⓘ

Resource Groups... | ▼

+
-

- A Monitoring Domain can have multiple connections, however only one connection can have **Managed Service Identity** as the **Credential**.
- The connections in a monitoring domain can be a combination of multiple **Application ID with Client Secret** (Service Principal) accounts, or one **Managed Service Identity** and multiple **Application ID with Client Secret** (Service Principal) accounts.
- Each connection can have only one **Subscription ID**.
- **Name:** An alias used to identify the connection.
- **Credential:** Select an Azure credential. For details, refer to [Create Azure Credentials](#).
- **Subscription ID:** A unique alphanumeric string that identifies your Azure subscription.
- **Region:** Azure region for the monitoring domain. For example, West India.
- **Resource Groups:** Select the Resource Groups of the corresponding VMs to monitor.

**Notes:**

- This field is not applicable if you select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
- When you remove and re-add a resource group in IAM, it won't appear in GigaVUE-FM automatically. To refresh the list and make the resource group selectable, reselect the Subscription ID from the drop down. This action triggers the UI to reload the resource groups associated with the selected subscription.

5. Select **Save**.

The **Azure Fabric Launch Configuration** wizard appears.

**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 only view and delete the existing configuration for GigaVUE V Series Node 1. You cannot perform any other actions on the existing configuration for GigaVUE V Series Node 1 as the features are deprecated from GigaVUE-FM.

You can use the **Copy** button to copy the permissions to the clipboard. Also, you can use the **Download** button to download the permission in JSON format.

NOTE: After updating the IAM Policy, it takes around 5 minutes for the changes to reflect on the Check Permissions screen.

Configure GigaVUE Fabric Components in GigaVUE-FM

After configuring the Monitoring Domain, you reach the Azure Fabric Launch Configuration page. In the same page, you can configure all the GigaVUE fabric components.

Enter or select the required information as described in the following table.

Fields	Description
Connections	A connection that you created in the monitoring domain page. For details, refer to Create Monitoring Domain .
Centralized Virtual Network	Alias of the centralized VNet in which the UCT-V Controllers, V Series Proxies, and the GigaVUE V Series nodes are launched.
Authentication Type	Select SSH Public Key as the Authentication Type to connect with the Centralized VNet.
SSH Public Key	The SSH public key for the GigaVUE fabric components.
Resource Group	The Resource Groups created in Azure for communication between the controllers, nodes, and GigaVUE-FM.
Security Groups	The security group created for the GigaVUE fabric components.
Enable Custom Certificates	Enable this option to validate the custom certificate during SSL Communication. GigaVUE-FM validates the Custom certificate with the trust store. If the certificate is not available in Trust Store, communication does not happen, and an handshake error occurs. Note: If the certificate expires after the successful deployment of the fabric components, then the fabric components moves to failed state.

Fields	Description
Certificate	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes, GigaVUE V Series Proxy, and UCT-V Controllers. For details, refer to Secure Communication between GigaVUE Fabric Components .
Prefer IPv6	Enables IPv6 to deploy all the Fabric Controllers, and the tunnel between hypervisor to GigaVUE V Series Nodes using IPv6 address. If the IPv6 address is unavailable, it uses an IPv4 address. Note: This option can be enabled only when deploying a new GigaVUE V Series Node. If you wish to enable this option after deploying the GigaVUE V Series Node, then you must delete the existing GigaVUE V Series Node and deploy it again with this option enabled.

Select **Yes** to configure V Series Proxy for the monitoring domain. For details, refer to [Configure GigaVUE V Series Proxy](#)

Azure Fabric Launch Configuration Check Permissions Save C

Connections
Centralized Virtual Network
Authentication Type
SSH Public Key
Resource Group
Security Groups
Enable Custom Certificates
Prefer IPv6
Configure a V Series Proxy

Select a Connection
Select a Virtual Network
sshPublicKey
Enter your SSH Public Key
Select resource group...
Select management subnet security group...
☐ Disabled
☐ No
☐ No

UCT-V Controller ①

Add
Image
Size
Number of Instances
IP Address Type
Subnet
Agent CA
Additional Subnets
Tags

V Series Node

Add
Image
Size
Disk Size (GB)
IP Address Type
Subnet
Data Subnets
Tags
Min Number of Instances
Max Number of Instances



To deploy GigaVUE fabric components (GigaVUE V Series Nodes, , and GigaVUE V Series Proxies) in GigaVUE-FM, you must accept the terms of the GigaVUE fabric components from the Azure marketplace using the Azure CLI or PowerShell. For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).

For details, refer to the following topics:

- [Configure GigaVUE V Series Proxy](#)
- [Configure GigaVUE V Series Node](#)

Configure GigaVUE V Series Proxy

GigaVUE V Series Proxy can manage multiple GigaVUE V Series Nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Proxies to communicate with the GigaVUE V Series nodes.

NOTE: A single GigaVUE V Series Proxy can manage up to 100 GigaVUE V Series nodes. The recommended minimum instance type is Standard_B1s for V Series Proxy.

To configure the GigaVUE V Series Proxy,

1. In the **Azure Fabric Launch Configuration** page, select **Yes** to **Configure a V Series Proxy**. The GigaVUE V Series Proxy fields appear.
2. Enter or select the appropriate values for the V Series Proxy. For details, refer to the [UCT-V Controller field descriptions](#).

Configure GigaVUE V Series Node

GigaVUE V Series node is a visibility node that aggregates mirrored traffic from multiple UCT-Vs. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to GigaVUE Cloud Suite for Azure using the standard VXLAN tunnels.

To launch a GigaVUE V Series node:

In the **Azure Fabric Launch Configuration** page, enter or select the appropriate values for the GigaVUE V Series Node.

V Series Node

Image

gigamon-gigamon-series-node-3.10-340871

Size

Standard_D4ds_v4 | 16 vCPUs

Disk Size (GB)

>= 30

IP Address Type

☒ Private ☐ Public

Management Subnet

Subnetmgmt

Data Subnets

Add Subnet

Tool Subnet

☒ Tool Subnet ⓘ

Subnet 1

dataout

Security Groups

sg-12345678 x

Tags

Add

Fields	Description
Image	From the Image drop-down list, select a GigaVUE V Series Node image.
Size	From the Size down-down list, select a size for the GigaVUE V Series Node. The default size for GigaVUE V Series Node configuration is Standard_D4s_v4 .
Disk Size (GB)	The size of the storage disk. The default disk size is 30GB. Note: When using Application Metadata Exporter, the minimum recommended Disk Size is 80GB.
IP Address Type	Select one of the following IP address types: - Select Private if you want to assign an IP address that is not reachable over Internet. You can use private IP address for communication between the GigaVUE V Series Node instances and GigaVUE-FM instances in the same network. - Select Public if you want the IP address to be assigned from Azure's pool of public IP address. On selecting Public IP address type, you must select the number of Public IPs defined in the Maximum Instance.
Management Subnet	Subnet: Select a management subnet for GigaVUE V Series. The subnet that is used for communication between the UCT-Vs and the GigaVUE V Series Nodes, as well as to communicate with GigaVUE-FM. Every fabric component (both controllers and the nodes) need a way to talk to each other and GigaVUE-FM. So, they should share at least one management

Fields	Description
	plane/subnet.
Data Subnet(s)	The subnet that receives the mirrored VXLAN tunnel traffic from the UCT-Vs. Select a Subnet and the respective Security Groups. Click Add to add additional data subnets. Note: Using the Tool Subnet checkbox you can indicate the subnets to be used by the GigaVUE V Series Node to egress the aggregated/manipulated traffic to the tools.
Tag(s)	(Optional) The key name and value that helps to identify the GigaVUE V Series Node instances in your Azure environment. For example, you might have GigaVUE V Series Nodes deployed in many regions. To distinguish these GigaVUE V Series Nodes based on the regions, you can provide a name that is easy to identify. To add a tag: - Click Add. - In the Key field, enter the key. For example, enter Name. - In the Value field, enter the key value.
Min Instances	The minimum number of GigaVUE V Series Nodes to be launched in the Azure connection. The minimum number of instances that can be entered is 1. Note: Nodes are launched when a monitoring session is deployed if GigaVUE-FM discovers some targets to monitor. The minimum amount is launched at that time. The GigaVUE-FM will delete the nodes if they are idle for over 15 minutes.
Max Instances	The maximum number of GigaVUE V Series Nodes that can be launched in the Azure connection. When the number of instances per V Series node exceeds the max instances specified in this field, increase the number in the Max Instances to Launch. When additional V Series nodes are launched, GigaVUE-FM re-balances the instances assigned to the nodes. This can result in a brief interruption of traffic.

Click **Save** to complete the Azure Fabric Launch Configuration.

For details regarding the Check permissions to validate IAM and security-group settings, refer to [Check Permissions while Configuring GigaVUE Fabric Components using GigaVUE-FM](#).

What to do Next

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


Acquire Traffic using Inline V Series Solution

NOTE: Click the following prompt to generate a quick deployment checklist for this topic.

Create a comprehensive checklist based on the deployment steps outlined in the Deploy Inline V Series Solution in Azure topic.

This section outlines the work flow for acquiring traffic using Inline V Series Node and deploying GigaVUE Fabric Components using Third Party Orchestration. It provides instructions to configure traffic acquisition, processing, and forwarding to your desired destination.

Note: The work flow described in this section is based on the topology shown below. The approach for acquiring traffic depends on your specific requirements. Additionally, you can decide whether to route traffic through public load balancer based on your topology.



Refer to the following topics to deploy Inline V Series in Azure:

- [Create a Resource Group](#)
- [Create Virtual Networks](#)
- [Create a Virtual Network Peering](#)
- [Create a Network Security Group](#)
- [Create a Load Balancer](#)
- [Install GigaVUE-FM on Azure](#)
- [Enable System Assigned Managed Identity](#)
- [Create GigaMon Custom Role](#)

- [Assign Role to Resource Group](#)
- [Assign Role to GigaVUE-FM Instance](#)
- [Create a GigaVUE-FM Token](#)
- [Modify Virtual Machine Scale Set \(VMSS\) Cloud Initialization Template](#)
- [Create a Virtual Machine Scale Set for Inline GigaVUE V Series Node \(Tier 1\)](#)
- [Assign VMSS to GWLB Backend Pools](#)
- [Create Monitoring Domain](#)
- [Deploy GigaVUE V Series Nodes for Inline V Series Solution](#)
- (Optional) [Create a Virtual Machine Scale Set for Out-of-Band GigaVUE V Series Node \(Tier 2\)](#)
- [Configure Monitoring session](#)

Create a Resource Group

The resource group is a container that holds all the resources for a solution. Select an existing resource group or create a new resource group. For navigation steps and detailed instructions, refer to [Create a resource group](#) topic in the Azure Documentation.

NOTE: We recommend creating a dedicated Resource Group for GigaVUE Visibility Fabric components such as GigaVUE-FM, V Series Nodes, Gateway Load Balancer, and others.

Create Virtual Networks

Azure Virtual Network (VNet) is the fundamental building block for your private network in Azure. You can choose to:

- Select an existing VNet: When you choose an existing VNet, Azure automatically populates the Subnet and Public IP fields.
- Create a new VNet: To create a new VNet, follow the steps in the [Create a virtual network](#) topic in the Azure Documentation

NOTE: We recommend setting up two separate virtual networks to support the GigaVUE Visibility Fabric. The first VNet (Viz_VNet) will host the GigaVUE visibility components such as GigaVUE-FM and V Series Nodes. The second VNet (App_VNet) will contain the virtual machines responsible for generating traffic.

Create a Virtual Network Peering

When workload virtual machines are distributed across multiple virtual networks, you must enable Virtual Network Peering between each workload VNet and the VNet hosting the GigaVUE V Series Node. This peering setup allows seamless communication between VNets within Azure.

To set up peering between the two VNets you created, refer to [Manage a virtual network peering](#) topic in Azure documentation.

Create a Network Security Group

Network Security Groups (NSGs) filter inbound and outbound traffic to subnets and network interfaces using security rules. You should configure the Network Security Group to allow GigaVUE-FM to communicate with the rest of the components. Select an existing network security group or create a new network security group. For navigation steps and detailed instructions, refer to [Create a network security group](#) topic in the Azure Documentation.

Create a Load Balancer

Gigamon deploys its solution as a Network Virtual Appliance positioned behind a Gateway Load Balancer. To ensure traffic flows through the V Series, it must be routed to the GWLB. You can achieve this routing either via a Standard Load Balancer or directly from a virtual machine's public interface. Microsoft recommends using a Standard Load Balancer for optimal integration.

Refer to the following sections for information on creating a Gateway Load Balancer and Standard Load Balancer:

Create a Gateway Load Balancer

Set up a Gateway Load Balancer (GWLB), which will later forward traffic to the Tier 1 V Series Node. Once deployed, the Tier 1 V Series Node will mirror packets from the workload virtual machines, enabling traffic visibility.

The following table lists the specific options you must select when creating a Gateway Load Balancer for an inline V Series deployment. For navigation steps and detailed instructions, refer to [Create a Gateway Load Balancer](#) topic in Azure documentation.

Parameters	Description	Mandatory field
Basics		
Subscription	Select your subscription	Yes
Resource Group	Select the Inline V Series Resource Group that you created	Yes
Name	Enter a valid name	Yes
Region	Select the region	Yes
SKU	Select Gateway	Yes
Type	Select Internal	Yes
Tier	Select Regional (default)	Yes
FrontEnd IP Configuration - Add a frontend IP configuration		
Name	Enter a valid name	Yes
IP Version	Select based on the requirement	Yes

Parameters	Description	Mandatory field
Virtual Network	Select your virtual network	Yes
Subnet and IP Assignment	Select your subnet and choose Dynamic for assignment	Yes
Backend Pools - Add a backend pool		
Name	Enter a valid name	Yes
Virtual Network	Default	Yes
Backend Pool Configuration	Select NIC	Yes
Gateway load balancer configuration		
Type	Choose Internal and External .	Yes
Internal and External Ports	Use default values Note: If you change the port values here, update the same ports in the Custom data and cloud-init field when creating the Virtual Machine Scale Set	Yes
NOTE: Add backend pool without IP Configuration (you will attach NICs later in this step Assign VMSS to GWLB Backend Pools).		
Inbound Rules - Add a load balancing rule		Yes
Name	Enter a valid name	Yes
IP Version	Select IPv4 or IPv6	Yes
Frontend IP Address	Select an existing Frontend IP from the drop-down list	Yes
Backend Pool	Select an existing Backend pool from the drop-down list	Yes
Session Persistence	Select None	Yes
Health Probe	Select Create New and enter the following details: - Protocol - Select HTTP as the protocol - Port - Enter 8888 as the port - Path: /health - Interval - Enter 5 seconds as the approximate amount of time, in seconds	Yes
Idle timeout (minutes)	Leave default or adjust as required	Yes

(Optional) Create a Standard Load Balancer

NOTE: Routing traffic through a public load balancer is optional and applies only to the topology shown at the beginning of this document. Alternatively, you can route traffic to the Gateway Load Balancer from a VM that has an interface with a public IP.

The following table lists the specific options you must select when creating a Standard Load Balancer for an inline V Series deployment. For navigation steps and detailed instructions, refer to [Create a Public Gateway Load Balancer](#) topic in Azure documentation.

For details regarding traffic flow from Gateway Load Balancer to Standard Load Balancer, refer to [Gateway Load Balancer](#) topic in Azure documentation.

Parameters	Description	Mandatory field
Basics		
Subscription	Select your subscription	Yes
Resource Group	Select the Inline V Series Resource Group that you created	Yes
Name	Enter a valid name	Yes
Region	Select the region	Yes
SKU	Select Standard	Yes
Type	Select Public (validated type)	Yes
Tier	Select Regional	Yes
FrontEnd IP Configuration - Add a frontend IP configuration		
Name	Enter a valid name	Yes
IP Version	Select IPv4	Yes
IP Type	Select IP Address as the IP type	Yes
Public IP address	Select the public IP address from the drop-down list. If required, you can create a new IP address	Yes
Gateway Load Balancer	Select the Gateway Load Balancer you created in the previous step to associate it with the frontend IP configuration	Yes
Backend Pools - Add a backend pool		
Name	Enter a valid name	Yes
Virtual Network	Select the same virtual network as the Standard Load Balancer—typically the App VNet—to ensure communication with the workloads.	Yes
Backend Pool Configuration	Select IP Address	Yes
IP Address	Specify the private IP address of the source/customer VM	
Inbound Rules - Add a load balancing rule		Yes
Name	Enter a valid name	Yes
IP Version	Select IPv4	Yes
Frontend IP Address	Select an existing Frontend IP from the drop-down list	Yes
Backend Pool	Select an existing Backend pool from the drop-down list	Yes

Parameters	Description	Mandatory field
Protocol	Select TCP as the protocol	Yes
Port	Enter 80 as the port	Yes
Backend Port	You can configure the backend port to match the frontend port. Enter a value based on your traffic requirements	Yes
Health Probe	Select Create new and create a new Health Probe with TCP Protocol, Port 80, and 5-second attempt interval	Yes
Session Persistence	Select None	Yes
Idle timeout (minutes)	Enter or select 4	Yes
Outbound Rules -		
Name	Enter a valid name	Yes
IP Version	Select IPv4	Yes
Frontend IP Address	Select an existing Frontend IP from the drop-down list.	Yes
Protocol	All	Yes
Idle timeout (minutes)	Enter or select 4	Yes
TCP Reset	Enabled	Yes
Backend Pool	Select an existing Backend pool from the drop-down list.	Yes
Port Allocation	Use the default number of outbound ports	Yes

Install GigaVUE-FM on Azure

To install GigaVUE-FM using the Azure Marketplace:

1. Go to Azure Marketplace and search for Gigamon. The latest version of Gigamon GigaVUE Cloud Suite for Azure appears. Click **Get it Now**.
2. In the pop-up window, select the **GigaVUE-FM (Fabric Manager) vX.XX - BYOL** option and click Continue.
3. Select the "**Want to deploy programmatically? Get started**" link.
4. Review the terms of service and the subscription name, and then select Enable. Click **Save**.
5. Verify the selected **Subscription** and **Plan**, then click **Create**.

6. Configure the GigaVUE-FM VM details. Most fields are pre-populated, but some require manual input. Enter the details as mentioned in [Table 1: GigaVUE-FM Installation Steps](#). For detailed instructions, refer to [Create a Linux virtual machine in the Azure](#) in Azure Documentation.

Table 1: GigaVUE-FM Installation Steps

Field	Description
Basics	
Subscription	Select your subscription.
Resource Group	Select the Inline V Series Resource Group that you created.
System-assigned managed identity	Use a system-assigned managed identity when a resource needs to authenticate to other services, and you want the identity to be created and deleted with the resource. Note: If you update any role it would take more than an hour to reflect in GigaVUE-FM, however, if you use APP registration it would take between 5-10 minutes to update in GigaVUE-FM.
Virtual machine name	Enter a name for the VM.
Region	Select a region for Azure VM.
Availability Zone	Choose your availability zone
Security Type	To enable UEFI secure boot, select Trusted launch virtual machines from the drop-down list. Click Configure security features and ensure that the Enable secure boot check box is enabled.
Image	Select the latest GigaVUE-FM images. NOTE: You cannot select multiple images for a VM.
Size	Select the recommended instance type: GigaVUE-FM - Standard_D4s_v3
Authentication Type	We support only SSH public key authentication type - SSH public key - Enter the administrator username for the VM. - Enter the SSH public key pair name. - Password - Enter the administrator username for the VM. - Enter the administrator password. NOTE: The username "gigamon" is reserved for internal usage. Do not create a user with the name "gigamon"
Disks	

Field	Description
Disk Size	The required disk size for GigaVUE-FM is 2 x 40GB .
Networking	
Virtual Network	Select the virtual network that you created.
Configure network security group	Select the network security group that you created.

NOTE: Verify the summary before proceeding to create. It will take several minutes for the VM to initialize. After the initialization is completed, you can verify the VM through the Web interface.

After the deployment, navigate to the VM overview page, copy the **Public IP address**, and paste it in a new web browser tab.

If GigaVUE-FM is deployed in Azure, use **admin123A!!** as the password for the **admin** user to login to GigaVUE-FM. You must change the default password after logging in to GigaVUE-FM.

Enable System Assigned Managed Identity

Managed Identity (MSI) is a feature of Azure Active Directory. When you enable MSI on an Azure service, Azure automatically creates an identity for the service VM in the Azure AD tenant used by your Azure subscription.

To enable MSI on the VM running in GigaVUE-FM using the Azure portal, refer to [Configure managed identities using the Azure portal](#) in the Azure documentation.

Create Gigamon Custom Role

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. This is used for GigaVUE-FM to integrate with Azure APIs and to automate the fabric deployment and management.

IMPORTANT: "Microsoft.Authorization/roleAssignments/read" permission is required for validating the required permissions. Ensure to include "Microsoft.Authorization/roleAssignments/read" permission in your IAM policy.

The 'built-in' roles provided by Microsoft are open to all resources. Refer to [Create or update Azure custom roles](#) topic in the Azure documentation to update the policy with the relevant IAM service.

After completing the configuration in the Basics, Permissions, and Assignable Scopes sections, copy the permissions listed below and paste them into the JSON code to create a custom role.

NOTE: Ensure to assign the roles for the permissions listed below at the subscription level.

```
{
  "Name": "CustomRoleForInline",
  "description": "Minimum requirements for FM in inline tapping",
  "assignableScopes": [
    "/subscriptions/<Subscription ID>"
  ],
  "permissions": [
    {
      "actions": [
        "Microsoft.Resources/subscriptions/read",
        "Microsoft.Resources/subscriptions/resourceGroups/read",
        "Microsoft.Network/virtualNetworks/read",
        "Microsoft.Network/loadBalancers/read",
        "Microsoft.Network/loadBalancers/backendAddressPools/read",
        "Microsoft.Network/loadBalancers/backendAddressPools/backendPoolAddresses/read",
        "Microsoft.Compute/virtualMachineScaleSets/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/networkInterfaces/read",
        "Microsoft.Compute/virtualMachineScaleSets/virtualMachines/networkInterfaces/ipConfigurations/read",
        "Microsoft.Compute/virtualMachines/read"
      ],
      "notActions": [],
      "dataActions": [],
      "notDataActions": []
    }
  ]
}
```

NOTE: Ensure you maintain the exact indentation and order shown in the code snippet when copying and pasting it. Any changes in formatting or order may cause errors or prevent the code from working correctly. We recommend copying the indentation from the online documentation for complete accuracy.

Assign Role to Resource Group

You can specify where the custom role is available for assignment such as a management group, subscription, or resource group and assign the created role to the resource group you created. In Azure, locate your newly created Role by navigating to “Management Groups” and selecting your subscription Id.

For further navigation and instructions to assign a role to your resource group, refer to the steps listed in **Update a custom role** and **Assignable scopes** sections of the [Create or update Azure custom roles](#) topic in the Azure documentation.

Assign Role to GigaVUE-FM Instance

You can assign the custom role you have created to GigaVUE-FM instance in Azure. For navigation path and detailed instructions, follow the steps listed in **Open the Add role assignment page**, **Select the appropriate role**, and **Select who needs access** sections of the [Assign Azure roles using the Azure portal](#) topic in the Azure documentation.

Create a GigaVUE-FM Token

GigaVUE-FM allows you to generate a token only if you are an authenticated user and based on your privileges in accessing the GigaVUE-FM. You can create multiple tokens if required.

To create a token in GigaVUE-FM, follow these steps:

1. Go to , select **Authentication > GigaVUE-FM User Management**. The **User Management** page appears.
2. In the **User Management** page, select **Tokens**.

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

3. Select **New Token**.
4. Enter a name for the new token in the **Name** field.
5. Enter the days until the token is valid in the **Expiry** field.
6. Select the user group for which you are privileged to access GigaVUE-FM from the **User Group** drop-down list.
7. Select **OK** to generate a new token. The generated token appears on the **Tokens** page.
8. Select the token to copy and use it to authenticate the GigaVUE-FM REST APIs.
 - a. Go to **Actions > Copy Token**.
 - b. Paste the copied token in the required fields to complete authentication.

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

Modify Virtual Machine Scale Set (VMSS) Cloud Initialization Template

The script given below will be essential for creating the VMSS, which uses this data to deploy the V Series node. This information is also required when configuring the Monitoring Domain in GigaVUE-FM.

Ensure you save these values in a secure location for reference in the upcoming steps. You can modify only the following parameters in the template:

- groupName
- subGroupName
- remoteIP
- Token

Custom Template:

NOTE: Ensure you maintain the exact indentation and order shown in the code snippet when copying and pasting it. Any changes in formatting or order may cause errors or prevent the code from working correctly. We recommend copying the indentation from the online documentation for complete accuracy.

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      remoteIP: <IP address of the GigaVUE-FM>
      remotePort: 443
      token: <token>

- path: /etc/vseries-inline.conf
  owner: root:root
  permissions: '0644'
  content: ""
```

Create a Virtual Machine Scale Set for Inline GigaVUE V Series Node (Tier 1)

V Series Nodes are created as part of a Virtual Machine Scale Set (VMSS) to allow scaling based on demand. This setup enables you to scale out or scale in the number of V Series Nodes as needed for each tier. Therefore, each tier will have a dedicated VMSS. The inline V Series solution does not support standalone V Series Nodes that are not part of a scale set.

The following table lists the specific options you must select when creating a Virtual Machine Scale Set for an inline V Series deployment. For navigation steps and detailed instructions, refer to [Create a Virtual Machine Scale Set](#) topic in Azure documentation.

Parameters	Description	Mandatory field
Subscription	Select your subscription	Yes
Resource Group	Select the Inline V Series Resource Group	Yes
Virtual Machine	Enter a valid name	

Parameters	Description	Mandatory field
Scale Set Name		
Region	Select a desired region	
Availability Zones	Choose if you want to use zones for high availability.	No
Orchestration		
Orchestration Mode	Select Uniform as the orchestration mode.	Yes
Security Type	Select Standard mode.	Yes
Scaling		
Scaling Mode	Choose Autoscaling .	Yes
Scaling Configuration	Click Configure to edit the scaling conditions.	Yes
Default Condition	Enter the Initial Instance Count as 0. NOTE: Once the Monitoring Domain and connection is configured, edit this value to the number of GigaVUE V Series Node that you need to deploy in this Monitoring Domain.	Yes
Condition	Choose a metric-based scaling condition (For example, CPU usage, network traffic).	Yes
Metric Source	Select the metric (For example, Average CPU Percentage).	Yes
Scale out	Set conditions like greater than 70% for scaling up.	Yes
Scale in	Set conditions like less than 20%.	
Cooldown Period	Set a cooldown period to prevent rapid scaling.	Yes
Instance Details		
Instance Type	Choose Standard_D4S_v4 as the VM size.	
Image	Select the GigaVUE V Series Node image.	
Authentication Type	Choose SSH public key.	
Username	Enter a user name. Do not use admin or gigamon.	
Networking		
Virtual Network	Select the required VNet.	Yes
Subnet Selection	Choose the appropriate subnet for Inline V Series Node.	Yes
NIC Configuration	GigaVUE V Series Node requires two NICs—one for management and one for mirrored data traffic. To configure the Data NIC, add a second network interface, select the appropriate subnet and network security group (NSG), and enable Accelerated Networking.	Yes
Management		

Parameters	Description	Mandatory field
Upgrade Mode	Choose Automatic .	Yes
Advanced		
Custom data and cloud init	Enter the Custom Template modified in step Modify Virtual Machine Scale Set (VMSS) Cloud Initialization Template as text in the following format and deploy the instance. The GigaVUE V Series Nodes uses this user data to generate config files (/etc/gigamon-cloud.conf and /etc/vseries-inline.conf) and register with GigaVUE-FM using Third Party Orchestration. NOTE: Ensure you maintain the exact indentation and order shown in the code snippet when copying and pasting it. Any changes in formatting or order may cause errors or prevent the code from working correctly. We recommend copying the indentation from the online documentation for complete accuracy. <pre>#cloud-config write_files: - path: /etc/gigamon-cloud.conf owner: root:root permissions: '0644' content: Registration: groupName: <Monitoring Domain Name> subGroupName: <Connection Name> remoteIP: <IP address of the GigaVUE-FM> remotePort: 443 token: <token> - path: /etc/vseries-inline.conf owner: root:root permissions: '0644' content: ""</pre> Custom Data with Internal and External Ports If you have modified the internal and external port values in the Gateway Load Balancer, use the following custom data: <pre>#cloud-config write_files: - path: /etc/gigamon-cloud.conf owner: root:root permissions: '0644' content: </pre>	Yes

Parameters	Description	Mandatory field
	<pre> Registration: groupName: <Monitoring Domain Name> subGroupName: <Connection Name> remoteIP: <IP address of the GigaVUE-FM> remotePort: 443 token: <token> - path: /etc/vseries-inline.conf owner: root:root permissions: '0644' content: tunnel: vxlan external_port : <Enter the port value> external_vni : <Enter the port value> internal_port : <Enter the port value> internal_vni : <Enter the port value> </pre>	

Assign VMSS to GWLB Backend Pools

Identify the Gateway Load Balancer you created in [Create a Load Balancer](#) step. To attach the Data NIC of the VMSS to the Gateway Load Balancer:

1. In the Azure portal, navigate to the Gateway Load Balancer you created earlier.
2. From the main page, go to **Settings** and select **Backend Pool**.
3. Locate the Virtual Network (VNet) where your VMSS is deployed.
4. Click **+ Add** under **IP Configuration**.
5. In the pop-up window, search for the Data NIC associated with your VMSS.
6. Select the NIC and click **Add**.
7. Click Save to apply the changes.

Create Monitoring Domain

Azure Load Balancer launches and manages GigaVUE V Series Node that is registered with GigaVUE-FM.

To deploy GigaVUE V Series Node with Gateway Load Balancing in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Monitoring Domain**.
3. On the Monitoring Domain page, select **New**.
4. On the **Monitoring Domain Configuration** page, select **Inline** as the Traffic Acquisition method.

5. Enter the **Monitoring Domain** Name and the **Connection** Name as mentioned in the user data provided during the template launch in Azure. Refer to Advanced section in [Create a Virtual Machine Scale Set for Inline GigaVUE V Series Node \(Tier 1\)](#).
6. (Optional) Turn on the **Use FM to launch Proxy** toggle to launch the GigaVUE V Series Proxy using GigaVUE-FM.

NOTE: You can use GigaVUE V Series proxy if GigaVUE-FM cannot directly reach the GigaVUE V Series Nodes (management interface) directly over the network. GigaVUE V Series Proxy is an optional component.

- a. From the **Image** drop-down list, select the required image.
 - b. From the **Size** drop-down list, select the instance size.
 - c. For **Number of Instances**, specify the required number of instances.
 - d. For **Management Subnet**:
 - a. Select the **IP Address Type** as Private or Public.
 - b. From the **Subnet** drop-down list, select the management subnet.
 - c. Select **Add Subnet** under **Additional Subnets** to add additional subnets.
 - e. Select **Add** under **Tags** to assign tags for resource identification.
7. Select **Save**.

Deploy GigaVUE V Series Nodes for Inline V Series Solution

When the Monitoring Domain is created successfully you will be navigated to the **Azure Fabric Launch Configuration** page in GigaVUE-FM automatically.

1. From the **Connections** drop-down list, select the required connection that you have configured.



Note:

After configuring the Monitoring Domain and establishing the connection, update the **Default Condition** count in the following steps based on the number of GigaVUE V Series Nodes you want to deploy within the Monitoring Domain:

- [Create a Virtual Machine Scale Set for Inline GigaVUE V Series Node \(Tier 1\)](#)
- [\(Optional\) Create a Virtual Machine Scale Set for Out-of-Band GigaVUE V Series Node \(Tier 2\)](#)

Adjust the instance count in the scale set to match the required number of nodes for your deployment.

2. Select the required resource group from the **Resource Group** drop-down list.

3. From the **Gateway Load Balancer** drop-down list, select the Load Balancer configured in Azure.
4. Under **Node Groups**, you can configure multiple node groups based on the deployment use case.
 - o **Inline Node Group:** This node group is used for the Inline V Series Node that is used for traffic acquisition.
 - a. In the **Inline Node Group Name** field, enter a name for the node group.
 - b. From the **Inline Auto Scaling Group** drop-down list, select the auto scaling group where you deploy the Inline V Series Node.
 - o **(Optional) Node Group:** You can configure this section if you wish to process the traffic using GigaVUE V Series Node. You can add or delete node groups using the + and - buttons.
 - a. In the **Node Group Name** field, enter a name for the node group.
 - b. From the **Auto Scaling Group** drop-down list, select the VMSS created in Azure.
5. Select **Save**.

NOTE: You can configure a maximum of eight Node groups.

(Optional) Create a Virtual Machine Scale Set for Out-of-Band GigaVUE V Series Node (Tier 2)

The Tier 2 V Series Node (Out-of-Band) processes mirrored traffic using GigaSMART operations to enrich and optimize data before forwarding it to the tool set. If only filtering is required, the Tier 1 V Series Node can handle it and send the traffic directly to the tool—eliminating the need for a Tier 2 node.

The following table lists the specific options you must select when creating a Virtual Machine Scale Set for Out-of-Band for an inline V Series deployment. For navigation steps and detailed instructions, refer to [Create a Virtual Machine Scale Set](#) topic in Azure documentation.

Parameters	Description	Mandatory field
Availability Zones	Choose if you want to use zones for high availability.	No
Orchestration		
Orchestration Mode	Select Uniform as the orchestration mode.	Yes
Security Type	Select Standard mode.	Yes
Scaling		

Parameters	Description	Mandatory field
Scaling Mode	Choose Autoscaling .	Yes
Scaling Configuration	Click Configure to edit the scaling conditions.	Yes
Default Condition	Enter the Initial Instance Count as 0. NOTE: Once the Monitoring Domain and connection is configured, edit this value to the number of GigaVUE V Series Node that you need to deploy in this Monitoring Domain.	Yes
Condition	Choose a metric-based scaling condition (For example, CPU usage, network traffic).	Yes
Metric Source	Select the metric (For example, Average CPU Percentage).	Yes
Scale out	Set conditions like greater than 70% for scaling up.	Yes
Scale in	Set conditions like less than 20%.	
Cooldown Period	Set a cooldown period to prevent rapid scaling.	Yes
Instance Details		
Instance Type	Choose Standard_D4S_v4 as the VM size.	Yes
Image	Select the GigaVUE V Series Node image.	Yes
Authentication Type	Choose SSH public key.	Yes
Username	Enter a user name. Do not use admin or gigamon.	Yes
Networking		
Virtual Network	Select the required VNet.	Yes
Subnet Selection	Choose the appropriate subnet for V Series Node.	Yes
NIC Configuration	GigaVUE V Series Node requires two NICs—one for management and one for mirrored data traffic. To configure the Data NIC, add a second network interface, select the appropriate subnet and network security group (NSG), and enable Accelerated Networking.	Yes
Management		

Parameters	Description	Mandatory field
Upgrade Mode	Choose Automatic .	
Advanced		
Custom data and cloud init	Enter the Custom data as text in the following format and deploy the instance. The GigaVUE V Series Nodes uses this user data to generate config files (/etc/gigamon-cloud.conf) and register with GigaVUE-FM using Third Party Orchestration. NOTE: Ensure you maintain the exact indentation and order shown in the code snippet when copying and pasting it. Any changes in formatting or order may cause errors or prevent the code from working correctly. We recommend copying the indentation from the online documentation for complete accuracy. <pre>#cloud-config write_files: - path: /etc/gigamon-cloud.conf owner: root:root permissions: '0644' content: Registration: groupName: <Monitoring Domain Name> subGroupName: <Connection Name> remoteIP: <IP address of the GigaVUE-FM> remotePort: 443 token: <token></pre>	Yes

Configure Monitoring session

When the **Traffic Acquisition Method** is **Inline**, the **UCT-I** application is available on the canvas by default. You can configure up to three tiers in a Monitoring Session and define multiple Sub Policies. Each Sub Policy can have its own ingress and egress tunnels and traffic processing applications.



Notes:

- You can configure a maximum of three tiers in a Monitoring Session.
- Tier 1 supports only Maps. Inline traffic is disabled and reserved for future use.
- You can configure a maximum of 8 Sub Policies in a Monitoring Session.
- Each Sub Policy can have its own Ingress Tunnels, Egress Tunnels, and Applications.
- Traffic from an out-of-band endpoint can either:
 - Pass through a Map and send to a tool using an Egress Tunnel.



- (Optional) Send to the GigaVUE V Series Node of the next tier for further processing.

To configure the Monitoring Session for Inline V Series Solution:

Tier 1 Monitoring Session:

1. Perform one of the following options:
 - Create a new Monitoring Session. Refer to [Create a Monitoring Session \(Azure\)](#) for details.
 - On an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab.


The GigaVUE-FM Monitoring Session canvas page appears.

When the **Traffic Acquisition Method** is **Inline**, the **UCT-I** application is available on the canvas by default.

2. Drag and drop the following items to the canvas as required for Tier 1 or Sub Policy 1:
 - Maps from the **new map** section. Refer to [Create a New Map \(Azure\)](#) for details.
 - Egress tunnels from the **new tunnel** section. When configuring Egress Tunnel, configure the **Remote Tunnel IP** if you intend to send the traffic directly from Tier 1 to the tool. Refer to [Create Ingress and Egress Tunnels \(Azure\)](#) for details.

NOTE: If sending traffic to Tier 2, Remote IP is optional. GigaVUE-FM will automatically add the remote IPs internally.

3. Now create a connection between the three tiles by dragging a line from the Inline-Source tile labeled “Out-Band” to the newly created Map and from Map to Egress tunnel.



Deploy Monitoring Session (Tier 1)


1. From the Actions drop-down list, select **Deploy**. The Deploy Monitoring Session pop-up appears.
2. Enter the following details:
 - In the **Policy Name** field, verify the auto-generated policy name or enter a custom name.
 - From the **Node Group** drop-down list, select the appropriate node group associated with this policy.
 - Under **Interface Mapping**, configure the interfaces:
 - i. From the **Ingress - <Tunnel>** drop-down list, select the input interface.
 - ii. From the **Egress - <Tunnel>** drop-down list, select the output interface.
3. Select **Deploy** the Monitoring Session.

To view the GigaVUE V Series Node associated with each Sub Policy, navigate to the **V SERIES NODES** tab and select a policy from the **Select a Sub policy** drop-down menu.

Tier 2 Monitoring Session (Optional):

You can send the filtered traffic to a Tier 2 V Series node, where GigaVUE-FM enriches and optimizes the data further.

1. In the same Monitoring Session canvas, drag and drop the following items to the canvas as required for Tier 2 or Sub Policy 2:
 - Ingress tunnel (as a source) from the **New** section. Refer to [Create Ingress and Egress Tunnels \(Azure\)](#) for details.
 - Maps from the **New Map** section. Refer to [Create a New Map \(Azure\)](#) for details.
 - GigaSMART apps from the Applications section. Refer to [Add Applications to Monitoring Session \(Azure\)](#).
 - Egress tunnels from the **new tunnel** section. Enter the **Remote Tunnel IP** address.
2. Create a link from the Ingress Tunnel to the Map or Application, and then connect it to the Egress Tunnel.
3. Create a direct link between the Egress Tunnel of Tier 1 and the Ingress Tunnel of Tier 2. The Blue Dot serves as an identifier to differentiate between tiers.
4. Repeat the above steps to configure a third tier, if required.



Deploy Monitoring Session Tier 1 to Tier 2

1. From the Actions drop-down list, select **Deploy**.

The Deploy Monitoring Session pop-up appears.

2. For each Policy (Tier) configured in the Monitoring Session, enter the following details:
 - In the **Policy Name** field, verify the auto-generated policy name or enter a custom name.
 - From the **Node Group** drop-down list, select the appropriate node group associated with this policy.
 - Under **Interface Mapping**, configure the interfaces:
 - i. From the **Ingress - <Tunnel>** drop-down list, select the input interface.
 - ii. From the **Egress - <Tunnel>** drop-down list, select the output interface.
3. Select **Deploy** the Monitoring Session.

To view the GigaVUE V Series Node associated with each Sub Policy, navigate to the **V SERIES NODES** tab and select a policy from the **Select a Sub policy** drop-down menu.

What to do Next:

To view Monitoring Session Statistics and Dashboards for Inline V Series Solution, refer to:

- [View Monitoring Session Statistics \(Azure\)](#)
- [Analytics for Inline V Series Solution \(Azure\)](#)

Deploy GigaVUE Cloud Suite for Azure using Third Party Orchestration

The GigaVUE Cloud Suite fabric components can be deployed in generic and integrated modes.

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.

Integrated Mode

In integrated mode, using GigaVUE-FM, you create a monitoring domain in your respective GigaVUE Cloud Suite. Then, use your own orchestration system just deploy GigaVUE V Series nodes.

GigaVUE-FM displays the deployed Nodes under the monitoring domain you created in your respective GigaVUE Cloud Suite. In this case, ensure that the monitoring domain and the connection name given in GigaVUE-FM matches the groupName and subGroupName in the user data provided in your orchestration system.

You can also create a monitoring domain and connection under Third party Orchestration in GigaVUE-FM. Then, use the monitoring domain name and connection name as the groupName and subGroupName in the registration data used in your respective cloud platform.

Refer to the following topics on more detailed information on how to deploy your fabric components in the Azure:

- [Deploy GigaVUE Fabric Components using Azure - Generic Mode](#)
- [Deployment GigaVUE Fabric Components using Azure -Integrated Mode](#)

Deploy GigaVUE Fabric Components using Azure - Generic Mode

This topic provides details on how you can configure GigaVUE Cloud Suite for Third Party Orchestration in multiple ways to gain visibility into both physical and virtual traffic. The configuration options depend on two factors:

- The traffic acquisition method
- The deployment method for fabric components

GigaVUE-FM supports four deployment options based on these variables.

- [Acquire Traffic using UCT-V - Without creating Monitoring Domain](#)
- [Acquire Traffic using Customer Orchestrated Source - Without creating Monitoring Domain](#)
- [Acquire Traffic using UCT-V - By Creating Monitoring Domain](#)
- [Acquire Traffic using Customer Orchestrated Source - By Creating Monitoring Domain](#)

Deployment GigaVUE Fabric Components using Azure -Integrated Mode

This section provides detailed information on the multiple ways in which you can configure GigaVUE Cloud Suite for Azure to provide visibility for physical and virtual traffic. Using three different options, you can configure GigaVUE Cloud Suite for Azure

based on the traffic acquisition method and the method in which you want to deploy fabric components. For details, refer to the [Prerequisites for GigaVUE Cloud Suite for Azure](#).

You can use your own Azure Orchestrator to deploy the GigaVUE fabric components instead of using GigaVUE-FM to deploy your fabric components.

The third-party orchestration feature allows you to deploy GigaVUE fabric components using your own Azure orchestration system. These fabric components register themselves with GigaVUE-FM with user information. Once the nodes are registered with GigaVUE-FM, you can configure monitoring sessions and related services in GigaVUE-FM.

You can deploy the fabric components using one of the following options:

- Manually using a configuration file.
- Use the Azure portal to launch the instances and deploy the fabric components using Custom data. Using your Custom data, the fabric components register themselves with the GigaVUE-FM. Based on the group name and the subgroup name details provided in the Custom data, GigaVUE-FM groups these fabric components under their respective monitoring domain and connection name.

The heartbeat messages sent from the respective nodes help GigaVUE-FM determine the health status of the registered nodes.

For more detailed information and the work flow, refer to the following topics:

- [Deployment GigaVUE Fabric Components using Azure -Integrated Mode](#)
 - [Deployment GigaVUE Fabric Components using Azure -Integrated Mode](#)
 - [Deployment GigaVUE Fabric Components using Azure -Integrated Mode](#)

Acquire Traffic using UCT-V as Traffic Acquisition Method

This section outlines the workflow to acquire traffic with UCT-V and deploy GigaVUE Fabric Components using Azure. Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Obtain GigaVUE-FM Image](#)
- [GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud](#)
- [Check the Required Permissions and Privileges](#)
- [Install UCT-Vs](#)
- [Create the Azure Credentials](#)
- [Create Monitoring Domain](#)
- [Configure GigaVUE Fabric Components in Azure](#)
- [GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud](#)

Obtain GigaVUE-FM Image

The image for the GigaVUE Cloud Suite Cloud is available in both the Azure Public Cloud and in the Azure Government portal.

GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud

GigaVUE Cloud Suite Cloud is available in the Azure Marketplace with the Volume Based License options.

GigaVUE Cloud Suite Cloud Suite in Azure Government

Azure Government is an isolated Azure region that contains specific regulatory and compliance requirements of the US government agencies.

To monitor the VMs that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the Azure Government (US) Region, the Azure Government solution provides the same robust features in Azure Government as in the Azure public cloud.

Install GigaVUE-FM on Azure

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

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Check the Required Permissions and Privileges

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. Refer to [Permissions and Privileges \(Azure\)](#) for more details.

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 [Gigamon Technical Support](#).
- b. Copy this package to your instance and Install the package with root privileges. For example,

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

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

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

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

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

sudo uctv-wizard

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

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

Options	Use Command	Description
pre-check	sudo uctv-wizard pre-check	Checks the status of the required packages and firewall requirements. - If any package is missing, it displays an appropriate message with the missing package details. - If installation includes all the packages, it displays a success message indicating that UCT-V is ready for configuration.
pkg-install	sudo uctv-wizard pkg-install	Displays the missing package and version details. To proceed with the installation, you can choose between the following: If you wish to skip the prompts and proceed with the system update, enter your option as y . The console interface installs the missing packages and restarts the UCT-V service. Enter N if you wish to install it manually. For details, refer to Install Linux UCT-Vs using Manual Configuration .
configure	sudo uctv-wizard configure	First, it checks for any existing configured file in the tmp directory (file named gigamon-cloud.conf in the C:\Users\<username>\AppData\Local (Windows) or /etc/uctv/uctv.conf (Linux) location). If available, UCT-V uses that configuration. If unavailable, UCT-V automatically adds the interface configuration in uctv.conf file, excluding the loopback (lo) interface, with all permissions enabled (source ingress, source egress, and destination). You can add the required policy for the available port if a firewall is installed. If you wish to skip the prompts to add the

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

**Notes:**

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

Linux UCT-V Installation Scenarios

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

Install Linux UCT-Vs using Manual Configuration

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

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

5. Save the file.

6. Restart the UCT-V service.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

```
$ sudo service uctv restart
```

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

```
$ sudo service uctv status
```



Notes:

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

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

Post Deployment Check:

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

1. Enter the command:

```
sudo uctvl uctv-show
```

2. Manually execute the following command:

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

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 [Gigamon 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 Install Linux UCT-Vs using Manual Configuration.
adapter-restore	uctv-wizard adapter-restore	Using this command, you can restore the backup copy of the network adapter buffer size configuration saved in the in the uctv-wizard adapter-setup step. NOTE: You need to manually restart the network adapters for changes to take effect immediately. You can choose between the following: • If you wish to skip the prompts for restoring the buffer size of the compatible network adapters, enter the option as y.

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

**Notes:**

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

```
uctv-wizard help
```

Windows UCT-V Installation Scenarios

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

Install Windows UCT-Vs using Manual Configuration

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

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

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



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

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

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

For IPv4:

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

For IPv6:

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

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

For IPv4:

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

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

For IPv6:

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

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

4. Save the file.

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

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

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 the Azure Credentials

You can monitor workloads across multiple Azure subscriptions within one monitoring domain. All the deployed GigaVUE fabric components are shared among many Azure subscriptions to reduce the cost. Earlier, each Azure subscription carried a set of GigaVUE fabric components.



- After launching GigaVUE-FM in Azure, the **Managed Identity** authentication credential is automatically added to the Azure Credential page as the default credential.
- You can only add the **Application ID with Client Secret** authentication credentials to the Azure Credential page.

To create Azure credentials,

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Credential**.
3. In the Azure Credential page, select **Add**.

The **Configure Credential** wizard appears.

Field	Value
Name*	Credential Name
Authentication Type	Application ID with Client Secret
Tenant ID*	Tenant ID
Application ID*	Application ID
Application Secret*	Application Secret
Azure Environment	Azure Environment... Azure AZURE_US_GOVERNMENT

4. Enter or select the appropriate information for the Azure credential:

a. **Name:** An alias used to identify the Azure credential.

b. **Authentication Type:**

Application ID with Client Secret: Connection with Azure with a service principal. Enter the values for the following fields.

- **Tenant ID:** A unique identifier of the Azure Active Directory instance.
- **Application ID:** A unique identifier of an application in Azure platform.
- **Application Secret:** a password or key to request tokens.

For details on how to create service principal and assign custom roles, refer to [Application ID with client secret](#)

c. **Azure Environment:** Select an Azure environment where your workloads are located. For example, Azure_US_Government.

5. Select **Save**.

You can view the list of available credentials in the Azure Credential page.

Create Monitoring Domain

Before configuring, you must establish a connection between GigaVUE-FM and your Azure environment. A Monitoring Domain in GigaVUE-FM allows you to define and manage this connection. Once established, GigaVUE-FM can deploy and manage UCT-V

Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes within your specified VNets and Resource Groups.

GigaVUE-FM connects to Azure using either an Application ID and Client Secret (Service Principal) or the Managed Service Identity (MSI) authentication method.

To create an Azure monitoring domain in GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**
2. Select **Monitoring Domain**.

The **Monitoring Domain** page appears.

3. In the Monitoring Domain page, select **New**.

The **Azure Monitoring Domain Configuration** wizard appears.

4. Enter or select the appropriate information for the Monitoring Domain:

- **Monitoring Domain:** An alias used to identify the monitoring domain.
- **Traffic Acquisition Method:** Select one of the following Tapping methods:
 - **UCT-V:** If you select UCT-V as the tapping method, the traffic is acquired from the UCT-Vs installed on your standard VMs in the Resource Group or in the Scale Sets. Then, the acquired traffic is forwarded to the GigaVUE V Series nodes. You must configure the UCT-V Controller to monitor the UCT-Vs.
- **Traffic Acquisition Tunnel MTU:** 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 default value is 1450.

- When using IPv4 tunnels, the maximum MTU value is 1450. Ensure that the UCT-V tunnel MTU is 50 bytes less than the UCT-V destination interface MTU size.
- When using IPv6 tunnels, the maximum MTU value is 1430. Ensure that the UCT-V tunnel MTU is 70 bytes less than the UCT-V destination interface MTU size.
- Disable the **Use FM to Launch Fabric** toggle, to deploy GigaVUE Fabric Components using Third Party Orchestration.
- Enable IPv6 Preference: Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes. This option appears only when **Use FM to Launch Fabric** is disabled and Traffic Acquisition Method is UCT-V.

- **Connections:**

Connections

Name*

Enter a connection name

Credential*

Credential Name... | v

Subscription ID*

Subscription ID... | v

Region*

Region Name... | v

Resource Groups*

☒ Discovered
 ☐ Regex *i*

Resource Groups... | v

+

-

- A Monitoring Domain can have multiple connections, however only one connection can have **Managed Service Identity** as the **Credential**.
- The connections in a monitoring domain can be a combination of multiple **Application ID with Client Secret** (Service Principal) accounts, or one **Managed Service Identity** and multiple **Application ID with Client Secret** (Service Principal) accounts.
- Each connection can have only one **Subscription ID**.
- **Name:** An alias used to identify the connection.
- **Credential:** Select an Azure credential. For details, refer to [Create Azure Credentials](#).
- **Subscription ID:** A unique alphanumeric string that identifies your Azure subscription.
- **Region:** Azure region for the monitoring domain. For example, West India.
- **Resource Groups:** Select the Resource Groups of the corresponding VMs to monitor.

**Notes:**

- This field is not applicable if you select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
- When you remove and re-add a resource group in IAM, it won't appear in GigaVUE-FM automatically. To refresh the list and make the resource group selectable, reselect the Subscription ID from the drop down. This action triggers the UI to reload the resource groups associated with the selected subscription.

5. Select **Save**.

The **Azure Fabric Launch Configuration** wizard appears.



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 only view and delete the existing configuration for GigaVUE V Series Node 1. You cannot perform any other actions on the existing configuration for GigaVUE V Series Node 1 as the features are deprecated from GigaVUE-FM.

You can use the **Copy** button to copy the permissions to the clipboard. Also, you can use the **Download** button to download the permission in JSON format.

NOTE: After updating the IAM Policy, it takes around 5 minutes for the changes to reflect on the Check Permissions screen.

Configure GigaVUE Fabric Components in Azure

This section provides information on how to register GigaVUE fabric components using Azure Portal or a configuration file.

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

- [Prerequisites](#)
- [Configure UCT-V Controller in Azure](#)
- [Configure UCT-V Controller in Azure](#)
- [Configure GigaVUE V Series Node and GigaVUE V Series Proxy in Azure](#)

Prerequisites

Ensure that GigaVUE V Series Node has at least have two Networks Interfaces (NIC) attached to it:

- A management NIC
- A data NIC with Accelerated Networking enabled.

When creating a virtual machine for GigaVUE V Series Node using CLI, you can attach Management NIC and Data NIC at the time of the virtual machine creation. However, if you are using Azure GUI to create the virtual machine for GigaVUE V Series Node, then you can attach the data NIC only after creating the virtual machine.

For details, refer to the following topics:

- [Create GigaVUE V Series Node with Management and Data NIC Attached using CLI](#)
- [Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI](#)



NOTE:

- Accelerated Networking must be enabled in the Data NIC only when deploying GigaVUE V Series Nodes using Third Party Orchestration.
- Accelerated Networking is not required for Management NIC.

Create GigaVUE V Series Node with Management and Data NIC Attached using CLI

1. Create the management NIC.

```
az network nic create -g <resource group> --vnet-name <VNet Name> --subnet  
<Subnet name> -n <Mangement NIC Name>
```

2. Create data NIC with Accelerated Networking enabled.

```
az network nic create -g <resource group> --vnet-name <VNet> --subnet <Subnet> -  
n <Data NIC> --accelerated-networking true
```

3. Create GigaVUE V Series Node virtual machine using the above NICs.

```
az vm create --resource-group <Resource group> --size <Standard_D4s_  
v4/Standard_D8S_V4> --name <GigaVUE V Series Node> --admin-username gigamon -  
-generate-ssh-keys --image gigamon-inc:gigamon-gigavue-cloud-suite-  
v2:vseries-node-v6.14.02:6.14.02 --plan-name vseries-node-v6.14.02 --plan-  
product gigamon-gigavue-cloud-suite-v2 --plan-publisher gigamon-inc --nics  
<Management NIC and Data NIC>
```

NOTE: You can use the following command to view all the images from Gigamon.

```
az vm image list --all --publisher gigamon-inc
```

Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI

Enable Management NIC when creating the GigaVUE V Series Node virtual machine.

For details, refer to the [Create virtual machine](#) topic in Azure Documentation.

Perform the following steps to attach the data NIC:

1. Select the GigaVUE V Series Node virtual machine from the Resources Page.
2. Stop the Virtual Machine using the **Stop** button.
3. Navigate to **Setting > Networking** from the left navigation pane.
The **Networking** page appears.
4. In the **Networking** page, select **Attach network interface**.

5. Select an existing network interface for Data NIC and select **OK**.
To enable accelerated networking, refer to [Manage Accelerated Networking through the portal](#).
6. Start the Virtual Machine.

Keep in mind the following when deploying the fabric components using third party orchestration in integrated mode:

- Create tokens in the **User Management** page in GigaVUE-FM. Refer to [Token-based Authentication](#) for more detailed information.
- When GigaVUE-FM is 6.10.00 or above and the Fabric Components are on (n-1) or (n-2) versions, you must create a **Username** and **Password** instead of using tokens in the registration data. For more details, refer to the Configure Role-Based Access for Third-Party Orchestration section in the 6.9 Documentation.
- When configuring UCT-V Controller, select **UCT-V** as the Traffic Acquisition Method.
- When you select Customer Orchestrated Source as your Traffic Acquisition Method, UCT-V and UCT-V Controller registration are not applicable.
- When you deploy the fabric components using third party orchestration, you cannot delete the monitoring domain without unregistering the GigaVUE V Series Nodes or UCT-V Controllers.
- Deployment of UCT-V Controller, GigaVUE V Series Node, and GigaVUE V Series Proxy through a third-party orchestrator is supported only on Linux platform.
- Deployment of UCT-V through a third-party orchestrator is supported on Linux and Windows platforms. Refer to [Linux UCT-V Installation](#) and [Windows UCT-V Installation](#) for detailed information.
- When creating virtual machine for deploying the fabric components in Azure, **SSH public key** must only be used as the **Authentication type** in Azure.

Configure UCT-V Controller in Azure

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

You can register UCT-V Controller in Azure Portal using one of the following methods:

- [Register UCT-V Controller during Virtual Machine Launch](#)
- [Register UCT-V Controller after Virtual Machine Launch](#)

Register UCT-V Controller during Virtual Machine Launch

In your Azure portal, to launch the UCT-V Controller init virtual machine and register UCT-V Controller using custom data, follow these steps:

1. In the Virtual machines page of the Azure Portal, select **Create > Virtual machine**. The **Create a Virtual Machine** Page appears. For details, refer to the [Create virtual machine](#) topic in Azure Documentation.

2. On the **Advanced** tab, enter the Custom Data as text in the following format and deploy the virtual machine

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      token: <Token>
      remoteIP: <IP address of the GigaVUE-FM>
      sourceIP: <IP address of UCT-V Controller> (Optional Field)
      remotePort: 443
```

- Enter the monitoring domain name and the connection name of the monitoring domain created earlier as the groupName and the subGroupName in the Custom Data.
- The UCT-V Controller uses this custom data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM

The UCT-V Controller deployed in your Azure portal appears on the Monitoring Domain page of GigaVUE-FM.

☐ Monitoring Domain	Connection	Fabric	Management IP	Fabric Version	Status
☐ MD1					
	pubintapi-vpc				✓ Connected
		G-vTapController	34.219.250.146	1.7-304	✓ Ok
		Gigamon-VSeriesProxy-1	34.211.211.49	2.1.0	✓ Ok
		Gigamon-VSeriesNode-1	172.30.24.188	2.2.0	✓ Ok

Register UCT-V Controller after Virtual Machine Launch

To register UCT-V Controller after launching a Virtual Machine using a configuration file,

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>
  remoteIP: <IP address of the GigaVUE-FM>
  sourceIP: <IP address of UCT-V Controller> (Optional Field)
  remotePort: 443
```

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

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

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

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

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

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    <interface>:                # Replace with your actual interface name (e.g., eth0)
      dhcp4: no
      dhcp6: no
      addresses:
        - <IPv4/24>              # e.g., 192.168.1.10/24
        - <IPv6/64>              # e.g., 2001:db8:abcd:0012::1/64
      nameservers:
        addresses:
          - <DNS_IPV4>            # e.g., 8.8.8.8
          - <DNS_IPV6>            # e.g., 2001:4860:4860::8888
      routes:
        - to: 0.0.0.0/0
          via: <IPv4_GW>          # e.g., 192.168.1.1
        - to: ::/0
          via: <IPv6_GW>          # e.g., 2001:db8:abcd:0012::fffe
```

Example netplan config:

```
network:
  version: 2
  renderer: NetworkManager
  ethernets:
    ens3:
      addresses:
```

```
- 10.114.53.24/21
dhcp4: no
dhcp6: no
accept-ra: false
routes:
- to: 10.114.48.1/32
  scope: link
- to: default
  via: 10.114.48.1
```

4. Save the file.
5. Apply the configuration.

```
$ sudo netplan apply
```

Configure UCT-V in Azure



Notes:

- You can deploy UCT-Vs using a third-party orchestrator on Linux and Windows platforms. For details, refer to [Linux UCT-V Installation](#) and [Windows UCT-V Installation](#).
- You can configure more than one UCT-V Controller for a UCT-V. If one UCT-V Controller goes down, the UCT-V registration happens through another active UCT-V Controller.

To register UCT-V after launching a Virtual Machine using a configuration file, follow these steps:

1. Install the UCT-V on the Linux or Windows platform. For detailed instructions, refer to [Linux UCT-V Installation](#) and [Windows UCT-V Installation](#).
2. Log in to the UCT-V. For details, refer to [Default Login Credentials](#).

3. Create a local configuration file and enter the following custom data.



- **/etc/gigamon-cloud.conf** is the local configuration file on the Linux platform.
- **C:\ProgramData\uctv\gigamon-cloud.conf** is the local configuration file on the Windows platform.

Registration:

```
groupName: <Monitoring Domain Name>
subGroupName: <Connection Name>
token: <Token>
remoteIP: <IP address of the UCT-V Controller 1>,
          <IP address of the UCT-V Controller 2>
sourceIP: <IP address of UCT-V> (Optional Field)
```



If you are using multiple interface in UCT-V and UCT-V Controller is not connected to the primary interface, then add the following to the above registration data:

```
localInterface:<Interface to which UCT-V Controller is connected>
```

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, then GigaVUE-FM unregisters UCT-V and it is removed from GigaVUE-FM.

Configure GigaVUE V Series Node and GigaVUE V Series Proxy in Azure

It is not mandatory to register GigaVUE V Series Nodes via GigaVUE V Series



- You can register your nodes using GigaVUE V Series Proxy when,
 - A large number of nodes are connected to GigaVUE-FM or
 - The user does not wish to reveal the IP addresses of the nodes.
- In this case, GigaVUE-FM communicates with GigaVUE V Series Proxy to manage the GigaVUE V Series Nodes.



- When deploying GigaVUE V Series Node using GigaVUE V Series Proxy, deploy the GigaVUE V Series Proxy first and provide the IP address of the proxy as the Remote IP of the GigaVUE V Series Node.

Register GigaVUE V Series Node and GigaVUE V Series Proxy in Azure Portal using one of the following methods:

- [Register Linux UCT-V](#)
- [Register GigaVUE V Series Node and GigaVUE V Series Proxy after Virtual Machine Launch](#)

Register GigaVUE V Series Node and GigaVUE V Series Proxy during Virtual Machine Launch

To register GigaVUE V Series Node and GigaVUE V Series Proxy using the custom data in Azure Portal, follow these steps:

1. In the Virtual machines page of the Azure Portal, select **Create** then **Virtual machine**.
2. On the **Create a Virtual Machine** Page, configure the settings.
For detailed information, refer to [Create virtual machine](#) topic in Azure Documentation.
3. On the **Advanced** tab, enter the Custom Data as text in the following format and deploy the virtual machine.

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      token: <Token>
      remoteIP: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
      remotePort: 443
```

- Enter the monitoring domain name and the connection name of the monitoring domain created earlier as the groupName and the subGroupName in the Custom Data.
- The GigaVUE V Series Node and GigaVUE V Series Proxy uses this custom data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

Register GigaVUE V Series Node and GigaVUE V Series Proxy after Virtual Machine Launch

To register GigaVUE V Series Proxy after launching the virtual machine using a configuration file, follow these steps:

1. Log in to the GigaVUE V Series Node or Proxy. For details on UCT-V Controller default login credentials, refer to [Default Login Credentials](#).
2. Create a local configuration file (`/etc/gigamon-cloud.conf`) and enter the following custom data.

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



- You can register your GigaVUE V Series Node directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series with GigaVUE-FM.
- If you wish to register GigaVUE V Series Node 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 Node using GigaVUE V Series Proxy then, enter the `remotePort` value as 8891 and `remoteIP` as <IP address of the Proxy>.
- You need to configure User and Password in the **User Management** page. For details, refer to [Configure Role-Based Access for Third Party Orchestration](#).
- Enter the Username and Password created in the **Add Users** Section.

3. Restart the GigaVUE V Series Proxy service.
 - GigaVUE V Series Node:
`$ sudo service vseries-node restart`
 - GigaVUE V Series Proxy:
`$ sudo service vps restart`

The deployed GigaVUE V Series Node or Proxy registers with the GigaVUE-FM. After successful registration, the GigaVUE V Series Node or 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 Node or Proxy. If that fails, then GigaVUE-FM unregisters the GigaVUE V Series Node or Proxy and removes it from GigaVUE-FM.

If you are using Azure GUI to create the virtual machine for GigaVUE V Series Node, then you need to attach data NIC to GigaVUE V Series Node after creating the virtual machine.

For details, refer to [Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI](#).

What to do Next

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

Acquire Traffic using vTAP as Traffic Acquisition Method

This section outlines the workflow to acquire traffic with vTAP and deploy GigaVUE Fabric Components using Azure. Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- [Obtain GigaVUE-FM Image](#)
- [GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud](#)
- [Check the Required Permissions and Privileges](#)
- [GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud](#)
- [Create the Azure Credentials](#)
- [Create Monitoring Domain](#)
- [Configure GigaVUE Fabric Components in Azure](#)
- [Configure and Deploy Monitoring Session](#)

Obtain GigaVUE-FM Image

The image for the GigaVUE Cloud Suite Cloud is available in both the Azure Public Cloud and in the Azure Government portal.

GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud

GigaVUE Cloud Suite Cloud is available in the Azure Marketplace with the Volume Based License options.

GigaVUE Cloud Suite Cloud Suite in Azure Government

Azure Government is an isolated Azure region that contains specific regulatory and compliance requirements of the US government agencies.

To monitor the VMs that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the Azure Government (US) Region, the Azure Government solution provides the same robust features in Azure Government as in the Azure public cloud.

Install GigaVUE-FM on Azure

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

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Check the Required Permissions and Privileges

When you first connect GigaVUE-FM to Azure, you need the appropriate authentication for Azure to verify your identity and check if you have permission to access the resources that you are requesting. Refer to [Permissions and Privileges \(Azure\)](#) for more details.

Create the Azure Credentials

You can monitor workloads across multiple Azure subscriptions within one monitoring domain. All the deployed GigaVUE fabric components are shared among many Azure subscriptions to reduce the cost. Earlier, each Azure subscription carried a set of GigaVUE fabric components.



- After launching GigaVUE-FM in Azure, the **Managed Identity** authentication credential is automatically added to the Azure Credential page as the default credential.
- You can only add the **Application ID with Client Secret** authentication credentials to the Azure Credential page.

To create Azure credentials,

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Credential**.
3. In the Azure Credential page, select **Add**.

The **Configure Credential** wizard appears.

Field	Value
Name*	Credential Name
Authentication Type	Application ID with Client Secret
Tenant ID*	Tenant ID
Application ID*	Application ID
Application Secret*	Application Secret
Azure Environment	Azure Environment... Azure AZURE_US_GOVERNMENT

4. Enter or select the appropriate information for the Azure credential:

a. **Name:** An alias used to identify the Azure credential.

b. **Authentication Type:**

Application ID with Client Secret: Connection with Azure with a service principal. Enter the values for the following fields.

- **Tenant ID:** A unique identifier of the Azure Active Directory instance.
- **Application ID:** A unique identifier of an application in Azure platform.
- **Application Secret:** a password or key to request tokens.

For details on how to create service principal and assign custom roles, refer to [Application ID with client secret](#)

c. **Azure Environment:** Select an Azure environment where your workloads are located. For example, Azure_US_Government.

5. Select **Save**.

You can view the list of available credentials in the Azure Credential page.

Create Monitoring Domain

Before configuring, you must establish a connection between GigaVUE-FM and your Azure environment. A Monitoring Domain in GigaVUE-FM allows you to define and manage this connection. Once established, GigaVUE-FM can deploy and manage UCT-V

Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes within your specified VNets and Resource Groups.

GigaVUE-FM connects to Azure using either an Application ID and Client Secret (Service Principal) or the Managed Service Identity (MSI) authentication method.

To create an Azure monitoring domain in GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**
2. Select **Monitoring Domain**.

The **Monitoring Domain** page appears.

3. In the Monitoring Domain page, select **New**.

The **Azure Monitoring Domain Configuration** wizard appears.

4. Enter or select the appropriate information for the Monitoring Domain:

- **Monitoring Domain:** An alias used to identify the monitoring domain.
- **Traffic Acquisition Method:** Select one of the following Tapping methods:
 - **vTAP:** If you select vTAP as the tapping method, traffic tapping is performed by the Azure platform and sent to the GigaVUE V Series Node. GigaVUE-FM creates the necessary configurations in Azure to enable this.
- **Traffic Acquisition Tunnel MTU:** 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 default value is 1450.

- When using IPv4 tunnels, the maximum MTU value is 1450. Ensure that the UCT-V tunnel MTU is 50 bytes less than the UCT-V destination interface MTU size.
- When using IPv6 tunnels, the maximum MTU value is 1430. Ensure that the UCT-V tunnel MTU is 70 bytes less than the UCT-V destination interface MTU size.
- Disable the **Use FM to Launch Fabric** toggle, to deploy GigaVUE Fabric Components using Third Party Orchestration.
- Enable IPv6 Preference: Enable this option to create IPv6 tunnels between UCT-V and the GigaVUE V Series Nodes. This option appears only when **Use FM to Launch Fabric** is disabled and Traffic Acquisition Method is UCT-V.

- **Connections:**

Connections

▼

Name*

Credential*

Subscription ID*

Region*

Resource Groups*

Enter a connection name

Credential Name... | ▼

Subscription ID... | ▼

Region Name... | ▼

☒ Discovered
 ☐ Regex ⓘ

Resource Groups... | ▼

+ -

- A Monitoring Domain can have multiple connections, however only one connection can have **Managed Service Identity** as the **Credential**.
- The connections in a monitoring domain can be a combination of multiple **Application ID with Client Secret** (Service Principal) accounts, or one **Managed Service Identity** and multiple **Application ID with Client Secret** (Service Principal) accounts.
- Each connection can have only one **Subscription ID**.
- **Name:** An alias used to identify the connection.
- **Credential:** Select an Azure credential. For details, refer to [Create Azure Credentials](#).
- **Subscription ID:** A unique alphanumeric string that identifies your Azure subscription.
- **Region:** Azure region for the monitoring domain. For example, West India.
- **Resource Groups:** Select the Resource Groups of the corresponding VMs to monitor.

☰

Notes:

- This field is not applicable if you select **Customer Orchestrated Source** as the **Traffic Acquisition Method**.
- When you remove and re-add a resource group in IAM, it won't appear in GigaVUE-FM automatically. To refresh the list and make the resource group selectable, reselect the Subscription ID from the drop down. This action triggers the UI to reload the resource groups associated with the selected subscription.

5. Select **Save**.

The **Azure Fabric Launch Configuration** wizard appears.



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 only view and delete the existing configuration for GigaVUE V Series Node 1. You cannot perform any other actions on the existing configuration for GigaVUE V Series Node 1 as the features are deprecated from GigaVUE-FM.

You can use the **Copy** button to copy the permissions to the clipboard. Also, you can use the **Download** button to download the permission in JSON format.

NOTE: After updating the IAM Policy, it takes around 5 minutes for the changes to reflect on the Check Permissions screen.

Configure GigaVUE Fabric Components in Azure

This section provides information on how to register GigaVUE fabric components using Azure Portal or a configuration file.

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

- [Prerequisites](#)
- [GigaVUE Cloud Suite Cloud Suite in Azure Public Cloud](#)
- [Configure GigaVUE V Series Node and GigaVUE V Series Proxy in Azure](#)

Prerequisites

Ensure that GigaVUE V Series Node has at least have two Networks Interfaces (NIC) attached to it:

- A management NIC
- A data NIC with Accelerated Networking enabled.

When creating a virtual machine for GigaVUE V Series Node using CLI, you can attach Management NIC and Data NIC at the time of the virtual machine creation. However, if you are using Azure GUI to create the virtual machine for GigaVUE V Series Node, then you can attach the data NIC only after creating the virtual machine.

For details, refer to the following topics:

- [Create GigaVUE V Series Node with Management and Data NIC Attached using CLI](#)
- [Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI](#)



NOTE:

- Accelerated Networking must be enabled in the Data NIC only when deploying GigaVUE V Series Nodes using Third Party Orchestration.
- Accelerated Networking is not required for Management NIC.

Create GigaVUE V Series Node with Management and Data NIC Attached using CLI

1. Create the management NIC.

```
az network nic create -g <resource group> --vnet-name <VNet Name> --subnet  
<Subnet name> -n <Mangement NIC Name>
```

2. Create data NIC with Accelerated Networking enabled.

```
az network nic create -g <resource group> --vnet-name <VNet> --subnet <Subnet> -  
n <Data NIC> --accelerated-networking true
```

3. Create GigaVUE V Series Node virtual machine using the above NICs.

```
az vm create --resource-group <Resource group> --size <Standard_D4s_  
v4/Standard_D8S_V4> --name <GigaVUE V Series Node> --admin-username gigamon -  
-generate-ssh-keys --image gigamon-inc:gigamon-gigavue-cloud-suite-  
v2:vseries-node-v6.14.02:6.14.02 --plan-name vseries-node-v6.14.02 --plan-  
product gigamon-gigavue-cloud-suite-v2 --plan-publisher gigamon-inc --nics  
<Management NIC and Data NIC>
```

NOTE: You can use the following command to view all the images from Gigamon.

```
az vm image list --all --publisher gigamon-inc
```

Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI

Enable Management NIC when creating the GigaVUE V Series Node virtual machine.

For details, refer to the [Create virtual machine](#) topic in Azure Documentation.

Perform the following steps to attach the data NIC:

1. Select the GigaVUE V Series Node virtual machine from the Resources Page.
2. Stop the Virtual Machine using the **Stop** button.
3. Navigate to **Setting > Networking** from the left navigation pane.
The **Networking** page appears.
4. In the **Networking** page, select **Attach network interface**.

5. Select an existing network interface for Data NIC and select **OK**.
To enable accelerated networking, refer to [Manage Accelerated Networking through the portal](#).
6. Start the Virtual Machine.

Keep in mind the following when deploying the fabric components using third party orchestration in integrated mode:

- Create tokens in the **User Management** page in GigaVUE-FM. Refer to [Token-based Authentication](#) for more detailed information.
- When GigaVUE-FM is 6.10.00 or above and the Fabric Components are on (n-1) or (n-2) versions, you must create a **Username** and **Password** instead of using tokens in the registration data. For more details, refer to the Configure Role-Based Access for Third-Party Orchestration section in the 6.9 Documentation.
- When configuring UCT-V Controller, select **UCT-V** as the Traffic Acquisition Method.
- When you select Customer Orchestrated Source as your Traffic Acquisition Method, UCT-V and UCT-V Controller registration are not applicable.
- When you deploy the fabric components using third party orchestration, you cannot delete the monitoring domain without unregistering the GigaVUE V Series Nodes or UCT-V Controllers.
- Deployment of UCT-V Controller, GigaVUE V Series Node, and GigaVUE V Series Proxy through a third-party orchestrator is supported only on Linux platform.
- Deployment of UCT-V through a third-party orchestrator is supported on Linux and Windows platforms. Refer to [Linux UCT-V Installation](#) and [Windows UCT-V Installation](#) for detailed information.
- When creating virtual machine for deploying the fabric components in Azure, **SSH public key** must only be used as the **Authentication type** in Azure.

Configure GigaVUE V Series Node and GigaVUE V Series Proxy in Azure

It is not mandatory to register GigaVUE V Series Nodes via GigaVUE V Series



- You can register your nodes using GigaVUE V Series Proxy when,
 - A large number of nodes are connected to GigaVUE-FM or
 - The user does not wish to reveal the IP addresses of the nodes.
- In this case, GigaVUE-FM communicates with GigaVUE V Series Proxy to manage the GigaVUE V Series Nodes.
- When deploying GigaVUE V Series Node using GigaVUE V Series Proxy, deploy the GigaVUE V Series Proxy first and provide the IP address of the proxy as the Remote IP of the GigaVUE V Series Node.

Register GigaVUE V Series Node and GigaVUE V Series Proxy in Azure Portal using one of the following methods:

- [Register GigaVUE V Series Node and GigaVUE V Series Proxy during Virtual Machine Launch](#)
- [Register GigaVUE V Series Node and GigaVUE V Series Proxy after Virtual Machine Launch](#)

[Register GigaVUE V Series Node and GigaVUE V Series Proxy during Virtual Machine Launch](#)

To register GigaVUE V Series Node and GigaVUE V Series Proxy using the custom data in Azure Portal, follow these steps:

1. In the Virtual machines page of the Azure Portal, select **Create** then **Virtual machine**.
2. On the **Create a Virtual Machine** Page, configure the settings. For detailed information, refer to [Create virtual machine](#) topic in Azure Documentation.
3. On the **Advanced** tab, enter the Custom Data as text in the following format and deploy the virtual machine.

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      token: <Token>
      remoteIP: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
      remotePort: 443
```

- Enter the monitoring domain name and the connection name of the monitoring domain created earlier as the groupName and the subGroupName in the Custom Data.
- The GigaVUE V Series Node and GigaVUE V Series Proxy uses this custom data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

[Register GigaVUE V Series Node and GigaVUE V Series Proxy after Virtual Machine Launch](#)

To register GigaVUE V Series Proxy after launching the virtual machine using a configuration file, follow these steps:

1. Log in to the GigaVUE V Series Node or Proxy. For details on UCT-V Controller default login credentials, refer to [Default Login Credentials](#).
2. Create a local configuration file (`/etc/gigamon-cloud.conf`) and enter the following custom data.

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



- You can register your GigaVUE V Series Node directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series with GigaVUE-FM.
- If you wish to register GigaVUE V Series Node 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 Node using GigaVUE V Series Proxy then, enter the `remotePort` value as 8891 and `remoteIP` as <IP address of the Proxy>.
- You need to configure User and Password in the **User Management** page. For details, refer to [Configure Role-Based Access for Third Party Orchestration](#).
- Enter the Username and Password created in the **Add Users** Section.

3. Restart the GigaVUE V Series Proxy service.

- GigaVUE V Series Node:
`$ sudo service vseries-node restart`
- GigaVUE V Series Proxy:
`$ sudo service vps restart`

The deployed GigaVUE V Series Node or Proxy registers with the GigaVUE-FM. After successful registration, the GigaVUE V Series Node or 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 Node or Proxy. If that fails, then GigaVUE-FM unregisters the GigaVUE V Series Node or Proxy and removes it from GigaVUE-FM.

If you are using Azure GUI to create the virtual machine for GigaVUE V Series Node, then you need to attach data NIC to GigaVUE V Series Node after creating the virtual machine.

For details, refer to [Create GigaVUE V Series Node with Management and Data NIC Attached using Azure GUI](#).

Configure and Deploy Monitoring Session

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

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

What to do Next

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

Deploy GigaVUE Cloud Suite for VMware ESXi using Terraform

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

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

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

License

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

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

Supportability

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

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

Supported Deployment Models

GigaVUE-FM Orchestration

- VMware ESXi-Traffic acquisition using Platform Mirroring.

Third-Party Orchestration (Generic Mode)

Supported platforms

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

Supported traffic acquisition methods

- UCT-V
- Customer-Orchestrated Source

Supported GigaSMART Applications

The Gigamon Terraform provider supports the following GigaSMART applications:

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

Rules and notes

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

Prerequisites

Before starting a Terraform deployment, ensure the following:

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

Get Started

Install Terraform

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

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

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

```
terraform version
```

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

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

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

Create a working directory

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

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

Deployment Using Terraform

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

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 Monitoring Session

This chapter describes how to setup ingress and egress tunnels, maps, and applications in a Monitoring Session to receive and send traffic to the GigaVUE Cloud Suite V Series node. It also describes how to filter, manipulate, and send the traffic from the V Series node to monitoring tools.

Refer to the following sections for details:

- [Create a Monitoring Session \(Azure\)](#)
- [Configure Monitoring Session for Inline V Series](#)
- [Create Ingress and Egress Tunnels \(Azure\)](#)
- [Create Raw Endpoint \(Azure\)](#)
- [Create a New Map \(Azure\)](#)
- [Add Applications to Monitoring Session \(Azure\)](#)
- [Interface Mapping \(Azure\)](#)
- [Deploy Monitoring Session \(Azure\)](#)
- [View Monitoring Session Statistics \(Azure\)](#)
- [Visualize the Network Topology \(Azure\)](#)

Create a Monitoring Session (Azure)

You must create a Monitoring Domain before creating a Monitoring Session. Refer to [Create Monitoring Domain](#).

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 (Azure)

The following table outlines the functional tabs available on the Azure Monitoring Session page, each designed to support specific aspects of network monitoring and 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 (Azure) .
Sources	Displays the sources and target details monitored by the Monitoring Session. You can view and filter the connection details of the Monitoring Session. You can view the deployment status, number of targets, and targets source health. In the Selection Status section, you can view the VM status. The status indicates whether the VM is supported, not supported, selected, or not selected. When you hover over the status, a tooltip displays the reason for that status. NOTE: In the case of OVS Mirroring, the Sources tab also displays the Hypervisor details along with the Instances.
Traffic Acquisition	You can enable or disable Prefiltering, Precryption, and Secure Tunnel here. You can also create a Prefiltering and Precryption templates and apply them to the Monitoring Session. Refer to Enable Prefiltering. Note: Traffic Acquisition is only applicable for Monitoring Domain created with UCT-V as Acquisition method.
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 templates, enable user defined applications, and enable or disable distributed De-duplication. Refer to Enable Prefiltering .
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 Node name, Health status (Configuration health + Traffic health), Host VPC, Management IP and Deployment Failure Message (if applicable). You can also change the interfaces mapped to an individual GigaVUE V Series Node. Refer to Interface Mapping (Azure) .
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 (Azure) .

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.

Button	Description
Deploy	Deploys the selected Monitoring Session.
Undeploy	Undeploys the selected Monitoring Session.

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

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

Configure Monitoring Session Options (Azure)

In the Monitoring Session page, you can perform the following actions in the **TRAFFIC ACQUISITION** and **TRAFFIC PROCESSING** tabs:

- Enable Prefiltering
- Enable Precryption
- Apply Threshold Template
- Enable User-defined applications
- Enable Distributed De-duplication

TRAFFIC ACQUISITION

To navigate to **TRAFFIC ACQUISITION** tab,

1. Go to **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. Select the required Monitoring Session from the list view on the left pane and select the **TRAFFIC ACQUISITION** tab.

You can perform the following actions in the **TRAFFIC ACQUISITION** page:

- [Enable Prefiltering](#)
- [Enable Precryption](#)

Enable Prefiltering

To enable Prefiltering:

1. In the **TRAFFIC ACQUISITION** page, go to **Mirroring > Edit Mirroring**.
2. Enable the **Mirroring** toggle button.
3. Enable **Secure Tunnel** option if you wish to use Secure Tunnels. Refer to the *Configure Secure Tunnel* section in the respective GigaVUE Cloud Suite Deployment Guide.
4. Select an existing Prefiltering template from the **Template** drop-down menu, or create a new template using **Add Rule** option and apply it. For details, refer to .
5. Select the **Save as Template** to save the newly created template.
6. Select **Save** to apply the template to the Monitoring Session.

Enable Precryption

Consideration before you enable Precryption:

- To avoid packet fragmentation, change the option precryption-path-mtu in 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.
- If you wish to use IPv6 tunnels, ensure that the versions of GigaVUE-FM and the fabric components are 6.6.00 or above.

NOTE: We recommend to enable the secure tunnel feature whenever the Precryption feature is enabled. Secure tunnel helps to securely transfer the cloud captured packets or Precryption data to a GigaVUE V Series Node. For more information, refer to *Secure Tunnels* in the respective GigaVUE Cloud Suite Deployment Guide.

To enable Precryption:

1. In the **TRAFFIC ACQUISITION** page, select **Precryption** tab and click **Edit Precryption**.
2. Enable the **Precryption** toggle button. Refer to Precryption™ topic in the respective cloud guides for details.

3. Apply Precryption to a few selective components based on the traffic:

NOTE: If you wish to use Selective Precryption, ensure that the versions of GigaVUE-FM and the fabric components are 6.8.00 or above.

Applications:

- a. Select the **APPLICATIONS** tab.
The **Pass All Applications** is enabled by default. If you wish to use selective Precryption, disable this option.
- b. Select any one of the following options from **Actions**:
 - i. Include: Select to include the traffic from the selected applications for Precryption.
 - ii. Exclude: Select to exclude the traffic from the selected applications for Precryption.
- c. Select **Add**. The **Add Application** widget opens.
- d. Select **csv** as the **Type**, if you wish to add the applications using a .csv file.
- e. Select **Choose File** and upload the file.
- f. Select **Manual** as the **Type**, if you wish to add the applications manually.
- g. Enter the **Application Name** and select + icon to add more applications.
- h. Select **Save**.

L3-L4

You can select an existing Precryption template from the **Template** drop-down list, or you can create a new template and apply it. For details, refer to [Create Precryption Template for UCT-V](#).

4. Enable the **Secure Tunnel** option if you wish to use Secure Tunnels. Refer to the *Configure Secure Tunnel* section in the respective GigaVUE Cloud Suite Deployment Guide.

Validate Precryption connection

To validate the Precryption connection, follow the steps:

- To confirm it is active, navigate to the Monitoring Session **Overview** tab and check the Traffic Acquisition Options.
- Select **Precryption**, to view the rules configured.

Limitations

During Precryption, UCT-V generates a TCP message with the payload being captured in clear text. Capturing the L3/L4 details of this TCP packet by probing the SSL

connect/accept APIs. The default gateway's MAC address is the destination MAC address for the TCP packet when SSL data is received on a specific interface. If the gateway is incorrectly configured, the destination MAC address is all Zeros.

TRAFFIC PROCESSING

To navigate to **TRAFFIC PROCESSING** tab:

1. Go to **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
2. Select the required Monitoring Session from the list view on the left side of the screen and click **TRAFFIC PROCESSING** tab.

You can perform the following actions in the **TRAFFIC PROCESSING** page:

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

Apply Threshold Template

To apply threshold:

1. In the **TRAFFIC PROCESSING** page, select **Thresholds** under **Options** menu.
2. You can select an existing threshold template from the **Select Template** drop-down list, or you can create a new template using **New Threshold Template** option and apply it.
For more details on Threshold Template, refer to the [Traffic Health Monitoring](#) section.
3. Select **Save** to save the newly created template.
4. Select **Apply** to apply the template to the Monitoring Session.

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.

You can also view the related details of the applied thresholds, such as Traffic Element, Metric, Type, Trigger Values, and Time Interval in the **Threshold** window. Select **Clear Thresholds** to clear the applied thresholds across the selected Monitoring Session.

Enable User Defined Applications

To enable user defined application:

1. In the **TRAFFIC PROCESSING** page, click **User Defined Applications** under **Options** menu.
2. Enable the **User-defined Applications** toggle button.
3. Add from the existing applications or create new User-Defined Application from the **Actions** drop-down. Refer to [Supported Protocols and Attributes](#).

Enable Distributed De-duplication

In the TRAFFIC PROCESSING page, click **Distributed De-duplication** under **Options** menu. 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. Refer to [Distributed De-duplication](#).



Notes:

- Distributed De-duplication is only supported on V Series version 6.5.00 and later.
- From version 6.9.00, Traffic Distribution option is renamed to Distributed De-duplication.

Tool Exclusion

Tool Exclusion helps prevent traffic loops by ensuring monitoring tools are not mistakenly selected as traffic targets during Automatic Target Selection (ATS). This feature is available only when the traffic acquisition method is VPC Traffic Mirroring.


You can exclude tool instances using either of the following methods:

1. Using AWS Tag Key

During deployment, apply the AWS tag key **GigamonExclude:Value** (Any Value) to any instance that acts as a monitoring tool. This tag ensures the system automatically excludes these instances from ATS.

2. Using the Tool Exclusion Feature in UI

During deployment, if the same instance IP is configured as both source (ingress) and tool (egress), the system prompts you to manually identify and exclude tools. Also, you can use the **Tool Exclusion** option to include or exclude tools and targets manually.



Create Ingress and Egress Tunnels (Azure)

Traffic from the GigaVUE V Series Node is distributed to tunnel endpoints in a Monitoring Session. A tunnel endpoint can be created using a standard L2GRE, VXLAN, UDPGRE, UDP, 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 2: Details - Add Tunnel Specifications

Field	Description
Alias	The name of the tunnel endpoint.
Description	The description of the tunnel endpoint.
Admin State Note: This option appears only after the Monitoring session deployment.	Use this option to send or stop the traffic from GigaVUE-FM to the egress tunnel endpoint. Admin State is enabled by default. You can use this option to stop sending traffic to unreachable or down tools. Each egress tunnel configured on the GigaVUE V SeriesNode has an administrative state that enables GigaVUE-FM to halt the tunnel's traffic flow. GigaVUE-FM only disable the tunnels when it receives a notification via REST API indicating that a tool or group of tools is down. Note: This option is not supported for TLS-PCAPNG tunnels.
Type	The type of the tunnel. Select from the options below to create a tunnel. ERSPAN, L2GRE, VXLAN, TLS-PCAPNG, UDP, or UDPGRE.

VXLAN

Traffic Direction

The direction of the traffic flowing through the GigaVUE V Series Node.

Note: In the scenario where secure tunnels need to be established between a GigaVUE V Series Node and a GigaVUE HC Series, you can utilize the **Configure Physical Tunnel** option provided in the GigaVUE V Series Secure Tunnel page. This allows you to configure secure tunnels on your physical device conveniently. For details, refer to [Secure Tunnels](#).

In	Choose In (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.	
	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 Raw Endpoint (Azure)

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. To deploy the Monitoring Session after adding the Raw Endpoint:
 - a. Select **Deploy** from the **Actions** drop-down list on the **TRAFFIC PROCESSING** page. The **Deploy Monitoring Session** dialog box appears.
 - b. Select the V Series Nodes for which you wish 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.

Create a New Map (Azure)

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 (Azure)

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

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

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

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

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

GigaSMA RT Operatio n	GigaVU E Cloud Suite for AWS	GigaVU E Cloud Suite for Azure	GigaVUE Cloud Suite for OpenSta ck	GigaVU E Cloud Suite for VMwar e (VMwa re vCente r)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestrati on	GigaVU E Cloud Suite for Nutani x
the retrospectiv e GigaVUE Cloud Suite Deployment Guide.							
UDPGRE Tunnel Decapsulati on Refer to <i>Create Ingress and Egress Tunnels</i> section on the retrospectiv e GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✗
GigaSMART NetFlow Generation	✓	✓	✓	✓	✓	✓	✓
GRE-In-UDP Tunnel Decapsulati on	✗	✗	✓	✓	✗	✓	✗
GENEVE Decapsulati on	✓	✗	✗	✗	✓	✗	✗
Header Stripping	✓	✓	✓	✓	✓	✓	✓
Load Balancing	✓	✓	✓	✓	✓	✓	✓


GigaSMA RT Operatio n	GigaVUE Cloud Suite for AWS	GigaVUE Cloud Suite for Azure	GigaVUE Cloud Suite for OpenStack	GigaVUE Cloud Suite for VMware (VMware vCenter)	GigaVUE Cloud Suite for VMware (NSX-T)	GigaVUE Cloud Suite for Third Party Orchestrati on	GigaVUE Cloud Suite for Nutanix
Masking	✓	✓	✓	✓	✓	✓	✓
SSL Decrypt	✓	✓	✓	✓	✗	✓	✓
Secure Tunnels Encapsulati on	✗	✓	✗	✓	✓	✓	✓
Secure Tunnels Decapsulati on	✗	✓	✗	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓	✓	✓
5G-Cloud Application	✗	✗	✓	✓	✗	✓	✗
5G-Service Based Interface Application	✗	✗	✓	✓	✓	✓	✗

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.



The screenshot shows the 'Masking' application configuration window. The window has a title bar with 'X' and 'Application' buttons, and a 'Save' button. The main area is divided into two columns. The left column contains a list of application types: 'Application', 'Alias', 'Protocol', 'Offset', 'Pattern', and 'Length'. The right column contains the configuration details for the selected 'Masking' application. The 'Alias' field is set to 'masking'. The 'Protocol' field is a dropdown menu with 'none' selected. The 'Offset', 'Pattern', and 'Length' fields are empty text input boxes.

Field	Value
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.



Application	Slicing
Alias	slicing
Protocol	none
Offset	64
Enhanced Name	none

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

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

4. Click **Save**.

What's Next

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

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

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

De-duplication

De-duplication application targets, identifies, and eliminates duplicate packets, blocking unnecessary duplication and sending optimized flows to your security and network monitoring tools. De-duplication lets you detect and choose the duplicate packets to count or drop in a network analysis environment.

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output. They can also appear when packets are gathered from multiple collection points along a path. The de-duplication application lets you eliminate these packets, only forwarding a packet once and thus reducing the processing load on your tools.

Refer to the following topics:

- [Feature Overview](#)
- [Configure De-duplication Application](#)
- [Distributed De-duplication](#)

Feature Overview

There are two actions that can be specified for handling the duplicate packets detected:

- drop, which drops the duplicate packets
- count, which counts the duplicate packets, but does not drop them

A time interval can be configured within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination.

For example, if two of the same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates.


For IPv4 and IPv6 packets, to determine if a packet is considered to be a duplicate, parts of the IP headers (Layer 3 and Layer 4), as well as part of the payload are compared.

For non-IP packets, a packet is considered to be a duplicate if it is identical.

Configure De-duplication Application

To add a de-duplication application:

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



The screenshot shows the 'Application' configuration window for the 'Dedup' application. The window has a title bar with an 'X' button, the text 'Application', and a 'Save' button. On the left is a sidebar with a grid background and a 'dedup' icon (a blue hexagon with a red gear and a red dot). The main area contains the following configuration fields:

Application	Dedup ⓘ
Alias	dedup
Action	☐ Count ☒ Drop
IP Tclass	☒ Include ☐ Ignore
IP TOS	☒ Include ☐ Ignore
TCP Sequence	☒ Include ☐ Ignore
VLAN	☐ Include ☒ Ignore
Timer	50000

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

Parameter	Description
Alias	Enter a name for the application.
Action	Specifies whether duplicate packets are to be counted or dropped as follows: - o Count– The de-duplication application counts the duplicate packets, but does not drop them. - o Drop– The de-duplication application drops the duplicate packets. The default is drop.
IP Tclass IP TOS TCP Sequence VLAN	These options are useful when applying de-duplication operations to packets in a NAT environment. Different NAT implementations can change certain packet header fields (for example, the TCP sequence number). If you want to be able to detect duplicates without requiring that these fields match (ToS field, TCP sequence number, VLAN ID), you can disable the corresponding option. - o IP Tclass – Ignore or include IPv6 traffic class. Use for IPv6. The default is include. - o IP TOS – Ignore or include the IP ToS bits when detecting duplicates. Use for IPv4. The default is include. - o TCP Sequence – Ignore or include the TCP Sequence number when detecting duplicates. The default is include. - o VLAN – Ignore or include the VLAN ID when detecting duplicates. The default is ignore. Include means the field will be included when the application compares packets. Ignore means the field will be ignored when the application compares packets.
Timer <Value: 10-500000 µs>	Configures the time interval within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination. The default is 50,000µs. For example, if two same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates. NOTE: Retransmissions are not counted as duplicates.

4. Click **Save**.

Distributed De-duplication

In distributed de-duplication, when you set up a Monitoring Session with a de-duplication app, the traffic is first sent to a component which distributes the traffic based on a consistent mechanism to ensure that packets from a particular traffic go to the same de-duplication instance. The packets are shared across the GigaVUE V Series Nodes. The distributed de-duplication more efficient that the existing de-duplication.

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

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

Important:

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

Limitations

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

What's Next

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

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

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

Application Intelligence

Application Intelligence provides a comprehensive solution that:

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

Application Intelligence provides the following capabilities for virtual nodes:

- [Configure Application Visualization for Virtual Environment](#)
- [Configure Application Filtering Intelligence for Virtual Environment](#)
- [Configure Application Metadata Intelligence for Virtual Environment](#)
- [Application Metadata Exporter](#)

Points to Note for Application Intelligence

Point to note when configuring Application Intelligence:

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

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

Cloud Platform	Instance Size	Maximum Number of Flows
VMware	Large (8 vCPU and 16GB RAM)	200k
Azure	Large (Standard_D8s_V4)	500k
	Medium (Standard_D4s_v4)	100k
Nutanix	Large (8 vCPU and 16GB RAM)	200k

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

8. Small Form Factor for VMware ESXi is not supported when using applications like Application Visualization, Application Metadata, Application Filtering. Refer to *Configure GigaVUE V Series Nodes for VMware ESXi* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more detailed information on how to deploy GigaVUE V Series Node, where you select the Form Factor.

Application Visualization

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

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

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

Configure Application Visualization for Virtual Environment

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

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

view appears.

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

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

4. Click **Save**.

Configure Filtering and Metadata Export for Selected Applications in Application Visualization

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

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

Filter Traffic for Selected Applications

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

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

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

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

Export Metadata for Selected Applications

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

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

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

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

View Application Statistics for Application Visualization

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

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

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

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

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

User Defined Application

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

Refer to the following topic for more detailed information:

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

Supported Protocols and Attributes

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

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

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

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

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

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

			ia ted Service (DS) Field in IP header		String	
	resolv_ name	DNS Name	Server's DNS name		REGEXP	gigamon.com
ipv6	address	Server IP Addres s	IP address of the server		IPV6 as REGEXP String	2001:0:9d38:6ab8:307b:16a 4:9c66:5f4 2001:0:9d38::9c66:5f4/64
	dscp	DSCP Value	DSCP from Different ia ted Service (DS) Field in IP header		UINT8 as REGEXP String	43

Mindata

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

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

Supported RegExp Syntax

Pattern	Description
.	Matches any symbol
*	Searches for 0 or more occurrences of the symbol or character set that

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

Limitations

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

Create Rules for User Defined Application

To create a new application:

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

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

NOTE: The lowest value has the highest priority.

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

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

- i. Click **Save**.

To add the created applications to the Monitoring Session:

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

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

Application Filtering Intelligence

Application Filtering Intelligence allows filtering of traffic based on the application (such as YouTube, Netflix, Sophos, or Facebook) or application family (such as antivirus, web, erp, or

instant-messaging) or application tags. Enables traffic filtering by layer 7 applications, which means you can filter out high-volume, low-risk traffic from reaching the tools and distribute high-risk network traffic of interest to the right tool at the right time.

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

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

Configure Application Filtering Intelligence for Virtual Environment

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

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

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

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

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

- 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 Supported Protocols and Attributes topic for more detailed information on what is user defined applications and how to configure it. - In this Add Application dialog box, select User Defined Applications from the Application Family list.

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

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

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

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

View Application Statistics for Application Filtering

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

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

Application Metadata Intelligence

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

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

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

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

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

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

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

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

Configure Application Metadata Intelligence for Virtual Environment

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



Notes:

- Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
- Flow Direction/Behaviour under Application Metadata Settings must be configured as **Bidirectional** for exporting application Metadata. Each exporter can be configured to export application metadata for up to 256 applications, and up to 64 attributes can be exported for each application. GigaVUE-FM allows selecting metadata for applications



that support exporting metadata. Applications which do not support exporting metadata are not listed in the GigaVUE-FM. By default, application name is exported for all such applications.

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

view appears.

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

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

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

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

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

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

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

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

7. Click **Save**.

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

1. 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.
2. 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.

3. Enter the **MTU** value.
4. **L4 Destination IP Address** - Enter the IP address of the tool port. If using Application Metadata Exporter (AMX), enter the data subnet IP address of the V Series Node where AMX application is deployed. Refer to [Application Metadata Exporter](#) for details.
5. **Source L4 Port** - The port from which the connection will be established to the target. For example, if A is the source and B is the destination, this port value belongs to A.

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

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

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

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

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

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

Extend Your AMI and AMX Deployment

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

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

View Application Statistics for Application Metadata

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

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

Application Metadata Exporter

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

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

You can use AMX to:

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

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

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

Application Metadata Exporter Deployment Options

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

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


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

On-Premises

Hardware (AMI)


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

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




Hardware (Control Plane Metadata)

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




Private Cloud (VMware)

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



Public Cloud

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




Application Metadata Exporter - Export Scenarios

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

- Export Application Metadata in JSON format over HTTPS/Kafka
- Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka
- Export of 3G/4G/5G Control Plane Metadata by AMX
- Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka
- Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka
- Identity Enrichment using Zscaler Private Access (ZPA)

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.

Identity Enrichment using Zscaler Private Access (ZPA)


Zscaler Private Access (ZPA) is a cloud-delivered Zero Trust Network Access (ZTNA) platform that brokers users' access to private applications without putting users on a flat VPN network.

With this feature, GEM enriches application metadata with user details such as username, location, network access policy, etc., AMX ingests Zscaler User Activity Logs through the ZPA Log Streaming Service (LSS), correlates them with application metadata from AMI, and

exports the enriched records to tools (for example, SIEM or Observability tools). To reduce tool-side ingestion costs, enrichment is exported at the start of a flow, not for the entire session.

This enrichment can be combined with GEM for Cloud Workloads for VMware, enabling a unified identity enrichment strategy across both ZPA-brokered access and virtualized cloud environments.

The following image shows the ZPA deployment architecture.



The workflow of the ZPA enriched data export is given as follows:

1. A user launches a private application for example OKTA, Jira, which the Zscaler Client Connector identifies as ZPA-protected traffic.
2. The client establishes a tunnel to the ZPA cloud, where the user is authenticated through the IdP and access policies are evaluated.
3. If access is permitted, ZPA selects an appropriate Application Connector with connectivity to the private application, and the Application Connector proxies the traffic to the application.
4. Gigamon physical or virtual taps mirror the traffic to Application Metadata Intelligence, which generates application metadata in CEF format.
5. For the same sessions, ZPA produces LSS User Activity Logs and streams them over TCP to Application Metadata Exporter on the configured listening port.

6. AMX correlates the application metadata with the ZPA User Activity logs using the following 5-tuples, enriches with the selected attributes, and exports records to tools.
 - Connector IP
 - Connector Port
 - Server IP
 - Server Port
 - IP Protocol

Rules and Notes

- With any AMI entitlement (AMI SKU, ZTA bundle, or SecureVUE PLUS bundle), you are also entitled to deploy AMX.
- The GigaVUE V Series Node deployed must be entirely dedicated to the AMX application, it cannot have other applications in it.
- The monitoring session can only have Raw End Point (REP), it cannot have other applications, maps, or tunnels when using the AMX application. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more detailed information on how to add a REP to the monitoring session and how to configure it.
- When using this application for production usage, it is recommended to use large size Virtual Machines.
- If you reload the GigaVUE V Series Node after configuring the AMX application, then the Ingestor in the AMX application fails.
- When using GigaVUE Enriched Metadata for Cloud Workloads, if duplicate IP address is assigned to workloads, the metadata received from the most recently queried workload will be enriched. Hence, it is not recommended to use the feature in environments that are prone to have duplicate IP addresses.
- When using GigaVUE Enriched Metadata for Cloud Workloads, endpoints and the GigaVUE V Series Node running AMI and AMX should be co-located in the same cloud platform.

Configure Metadata Export Filter for AMX

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

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

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

Rules and Notes

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

To configure the export filter:

1. Create or edit the `/etc/amx_settings.yaml` file on the V Series Node where AMX is deployed.
2. Define the `set_export_filters` key with a valid ClickHouse sub-query as its value. For example, `set_export_filters: "mdata['app_name'] LIKE 'http%'"`.
3. To apply the new filter settings, undeploy the monitoring session and then deploy it again.

Deploy Application Metadata Exporter in Azure using FM Orchestration

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

Predeployment Checklist

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

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

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

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

5. Select the appropriate deployment mode for the GigaVUE V Series node based on your requirements:
 - Single physical uplink for Management, Data and Tool connectivity.
 - Single physical uplink for Management and Tool connectivity, and another uplink for Data connectivity.
 - (Recommended) Separate uplinks (three) for Management, Data, and Tool connectivity.
6. If static IP addressing is used, obtain the following details from your server team:
 - IP addresses for the management, data, and tool interfaces.
 - IP address of the DNS server and the gateway connected to the tool interface.
 - Subnet Mask and MTU.
7. Verify whether remote web console access to the V Series is available. Console access is required for troubleshooting. If remote access is not available, request the server team to enable SSH access.

8. Review the deployment guide for the tool specific configuration needed to export records from AMX. Refer to [AMX Tool Integration](#).
9. Review the connectivity requirements for GigaVUE-FM, AMI, AMX, and the associated tools in your environment. Ensure that the required ports are open and any external connectivity dependencies are addressed.
10. Ensure that the standalone GigaVUE-FM deployment meets the recommended requirements shown in the following table. For medium and large scale deployments, refer to [Recommended Requirements for GigaVUE-FM Deployments](#). This step is optional. A standalone GigaVUE-FM instance is required only when an existing GigaVUE-FM is not available for managing AMX.

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

Refer to the following topics to deploy AMX using GigaVUE-FM Orchestration:

- [Deploy Application Metadata Exporter in Azure using FM Orchestration](#)
- [Upload GigaVUE V Series Node Image into GigaVUE-FM](#)
- [\(Optional\) Configure Metadata Export Filter for AMX](#)
- [\(Optional\) Integrate Private CA](#)
- [Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source](#)
- [Deploy Application Metadata Exporter in Azure using FM Orchestration](#)

- [Create a Monitoring Session to Configure Application Metadata Exporter](#)
- [Configure Ingestor in AMX to Process Traffic](#)
- [Integrate AMX with Third Party Tool](#)
- [Configure Exporter to Export Attributes to Third Party tools](#)
- [Verify Your Deployment](#)

Install GigaVUE-FM on Azure

Refer to [Install GigaVUE-FM on Azure](#) for steps to install GigaVUE-FM on Azure and the steps to start GigaVUE-FM instance and configure it.

Upload GigaVUE V Series Node Image into GigaVUE-FM

This step is optional. You can upload the GigaVUE V Series Node image to GigaVUE-FM even while deploying the GigaVUE V Series Node.

To upload the V Series image into GigaVUE-FM,

1. Go to **Inventory > VIRTUAL > Azure**, and select **Settings > OVA Files**.
2. In the **OVA Files** page, select **Browse**, and then select the image file. For example, *gigamon-gigavue-vseries-node-x.x.x-0-xxxxxx.ova*. You can upload three OVA files to GigaVUE-FM for VMware vCenter.
3. Select **Upload to Server**. You have uploaded the OVA image file to the GigaVUE-FM server.

(Optional) Configure Metadata Export Filter for AMX

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

(Optional) Integrate Private CA

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

Rules and Notes

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

- Include all intermediate CAs in a single file in the correct order.
- Place the last intermediate CA in the chain at the top.
- Place the preceding CAs in descending order.

To integrate private CA

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

Generate CSR

To create an intermediate CA certificate:

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

The CSR is downloaded successfully.

Upload CA Certificate

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

To upload the signed CA certificate to GigaVUE-FM:

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

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

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

Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source

To create a Monitoring Domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > Azure**, and select **Monitoring Domain**. The **Create Monitoring Domain Configuration** page appears.
2. On the **Monitoring Domain** page, select **New**. The **Azure Configuration page** appears.
3. Enter the following fields:
 - a. **Monitoring Domain** - Enter a name for the monitoring domain.
 - b. **Connection** - Enter a name for the connection.
 - c. **Virtual Center** - Enter the IP address or FQDN of the vCenter
 - d. **Username** - Enter the username of the vCenter user.
 - e. **Password** - Enter the password to connect to the vCenter.
4. Select the Traffic Acquisition Method as Customer Orchestrated Source.
5. Click **Save**. Once you save the changes, the **Virtual Fabric Launch Configuration** page appears. The configuration details are explained in the following section.

Deploy V Series Node

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

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

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

- b. Select **Actions > Deploy Fabric**.

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

VMware vCenter ESXi Fabric Launch Monitoring Domain Settings

VMware Fabric Launch Configuration Deploy Cancel

Datacenter* DataCenter

Cluster* Cluster-2_Scaled

V Series Node Image* gigamon-gigavue-vseries-node-6.13.0... Add

Form Factor* Large, 8vCPU, 16GB RAM, 80GB Disk

Enable Custom Certificates ☐

Select All Select None

10.115.81.126

Common Configuration

10.115.81.126

VSeries.10.115.81.126-1

☒ Datastores ☐ Datastore Clusters

Datastore* datastore1 (2)

Name Server Use Comma to Separate

V Series Node Name* VSeries.10.115.81.126-1

Management

Interface Name Network adapter 1

Network* VM Network

IP Type Static

IP Address* 10.115.85.38

Gateway IP* 10.115.80.1

CIDR Prefix* 21

MTU Enter MTU Value

Data interfaces

Interface Name Network adapter 2

Network* DPortGroup_tunnel

IP Type Static

IP Address* 172.21.1.220

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

Field	Description																				
Datacenter	Select the vCenter datacenter with the ESXi hosts to be provisioned with GigaVUE V Series Nodes.																				
Cluster	Select the cluster where you want to deploy the GigaVUE V Series Nodes.																				
V Series Node Image	Select the OVA file uploaded in the Upload GigaVUE V Series Node Image into GigaVUE-FM , from the dropdown list.																				
Form Factor	Select Large, 8vCPU, 16GB RAM, 80GB Disk option.																				
Enable Custom Certificates	Enable this option to validate the custom certificate during SSL Communication. GigaVUE-FM validates the Custom certificate with the trust store. If the certificate is unavailable in Trust Store, communication does not happen, and a handshake error occurs. NOTE: If the certificate expires after the successful deployment of the fabric components, then the fabric components move to failed state.																				
Certificate NOTE: This field appears only when Enable Custom Certificates option is enabled.	Select the custom certificate from the drop-down menu. You can also upload the custom certificate for GigaVUE V Series Nodes. For more information, refer to Secure Communication between GigaVUE Fabric Components.																				
Hosts	Select the ESXi hosts for GigaVUE V Series Node deployment. The Common Configuration dropdown wizard and the specific host dropdown wizard appears. Expand the individual host and configure the GigaVUE V Series Node. <table border="1"> <tr> <td colspan="2">ESXi Host</td></tr> <tr> <td>Datastore</td><td>Network datastore shared among all ESXi hosts.</td></tr> <tr> <td>Name Server</td><td>The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.</td></tr> <tr> <td>V Series Node Name</td><td>Enter a name for the GigaVUE V SeriesNode.</td></tr> <tr> <td colspan="2">Management</td></tr> <tr> <td>Network</td><td>Management network for GigaVUE V Series Nodes.</td></tr> <tr> <td>IP Type</td><td>Select the management network IP type as Static or DHCP.</td></tr> <tr> <td>IP Address</td><td>Enter the IP address of the management port.</td></tr> <tr> <td>Gateway IP</td><td>Gateway IP address of the management network.</td></tr> <tr> <td colspan="2"> NOTE: This field appears only when the Management IP type is Static. </td></tr> </table>	ESXi Host		Datastore	Network datastore shared among all ESXi hosts.	Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.	V Series Node Name	Enter a name for the GigaVUE V SeriesNode.	Management		Network	Management network for GigaVUE V Series Nodes.	IP Type	Select the management network IP type as Static or DHCP.	IP Address	Enter the IP address of the management port.	Gateway IP	Gateway IP address of the management network.	NOTE: This field appears only when the Management IP type is Static.	
ESXi Host																					
Datastore	Network datastore shared among all ESXi hosts.																				
Name Server	The server that stores the mapping between the domain names and the IP addresses. The maximum number of name servers that can be entered is three. Enter valid IPv4 addresses, separated by a comma.																				
V Series Node Name	Enter a name for the GigaVUE V SeriesNode.																				
Management																					
Network	Management network for GigaVUE V Series Nodes.																				
IP Type	Select the management network IP type as Static or DHCP.																				
IP Address	Enter the IP address of the management port.																				
Gateway IP	Gateway IP address of the management network.																				
NOTE: This field appears only when the Management IP type is Static.																					

Field	Description	
	CIDR Prefix NOTE: This field appears only when you select the IP Type as Static.	Enter the CIDR prefix of the IP address.
	MTU (Optional)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
	Data Interfaces - When using Customer Orchestrated Source as the Traffic Acquisition Method, you must configure two data interfaces, one for data port and other for tool port.	
	Interface 1	
	Network	Tunnel Network for the GigaVUE V Series Nodes.
	IP Type	Select the tunnel network IP address type as Static or DHCP.
	IP Address	Enter the IP address of the data port.
	Gateway IP	Gateway IP address of the tunnel network.
	CIDR Prefix NOTE: This field appears only when you select the IP Type as Static.	Enter the CIDR prefix of the IP address.
	MTU (Optional)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
	Use IPv6 NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.	Enable this option to use IPv6 address.
	Interface 2	
	Network	Tunnel Network for the GigaVUE V Series Nodes.
	IP Type IP Address Gateway IP CIDR Prefix	Select the tunnel network IP address type as Static or DHCP. Enter the IP address of the tool port. Gateway IP address of the tool network. Enter the CIDR prefix of the IP address.

Field	Description
	NOTE: This field appears only when you select the IP Type as Static.
	MTU (Optional) The Maximum Transmission Unit (MTU) is the maximum size of each packet that can be transferred as a single entity in a network connection. Enter value between 1280 to 9000.
	Use IPv6 Enable this option to use IPv6 address.
	NOTE: This field appears only when Customer Orchestrated Source as the Traffic Acquisition Method.
	Virtual Disk Format Select the Virtual Disk Format from the dropdown list.
	Deployment Folder Enter the folder name in vCenter, under which the GigaVUE V SeriesNodes must be deployed.
	Password Enter the password you wish to use for the GigaVUE V Series Node.

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

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

Create a Monitoring Session to Configure Application Metadata Exporter

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

1. **Access the Monitoring Session Page:**
 - a. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**.
 - b. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
2. **Create an ingress REP:**

- a. In the canvas, from the **New** expand menu, drag and drop New Raw Endpoint to the graphical workspace.
- b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
- c. Enter the Alias and Description details for the Ingress Raw End Point and click **Save**.

3. Add AMX Application:

- a. In the canvas, from the **Applications** expand menu, drag and drop AMX to the graphical workspace.
- b. Configure the ingestor settings to process traffic. Refer to [Configure Ingestor in AMX to Process Traffic](#) for details.
- c. Configure the required exporter settings to export attributes to Third Party tools. Refer to [Configure Exporter to Export Attributes to Third Party tools](#) for details.
- d. Create a link between the ingress REP and Application Metadata Exporter application.

4. Create an egress REP:

- a. From the **New** expand menu, drag and drop New Raw Endpoint to the graphical workspace.
- b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
- c. Enter the Alias and Description details for the Egress Raw End Point and click **Save**.
- d. Create a link between the Application Metadata Exporter application and egress REP



Notes:

- Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
- When deploying AMX using the Third-Party Orchestration method, Monitoring Session must also be configured in the same work flow.

Configure Ingestor in AMX to Process Traffic

The following scenarios provide more information about the different ways to configure Ingestor AMX:

- [Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka](#)
- [Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka](#)
- [Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka](#)
- [Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka](#)

Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka

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

To configure ingestor for AMI:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.

Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka

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

To export GigaVUE Enriched Metadata for Mobile Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.

4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which application metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.
9. Select the Enable option to enable the enrichment.
10. Select **Mobility** from the **Type** drop-down menu.
11. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Mobile Networks](#) for more details.

Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka

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

To export GigaVUE Enriched Metadata for Cloud Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which Application Metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.
9. Select the Enable option to enable the enrichment.
10. Select **Workload** from the **Type** drop down menu.

11. Select the platform as **VMware VCenter** from the drop down list.
12. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Cloud Workloads](#) for more details.

Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka

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

To configure ingestor for NetFlow/IPFIX:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select the input type as **NetFlow/IPFIX**. Only the elements specified in the mapping file will be included in the JSON output. Editing or customizing the mapping file is not supported. Refer to [Supported Element Types for NetFlow Integrator Element Mapping](#).
3. Enter the port number on which NetFlow/IPFIX must be ingested. The default port value is 2055.
4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor. A maximum of 12 NetFlow/IPFIX type ingestors can be configured per deployment.

NOTE: If you select NetFlow/IPFIX as the ingestor type, Metadata Enrichment is disabled by default. Metadata Enrichment is not supported for NetFlow/IPFIX inputs.

Integrate AMX with Third Party Tool

This section explains how to set up the third-party tool and configure the exporter settings required to route application metadata attributes for analysis and visualization.

Set up the Third Party Tool for AMX Integration

You must set up the third party tools to ingest and visualize application metadata.

Refer to [AMX Tool Integration](#) for configuration guidance and the verification steps required to confirm that data export and visualization are functioning correctly.

Analyze Metadata with Gigamon Insights: Use [Gigamon Insights](#) to query your AMI and AMX metadata using plain-language prompts and get immediate, AI-powered answers — no dashboards or manual queries needed.

Configure Exporter to Export Attributes to Third Party tools

AMX supports the following types of exporters:

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

Configure Cloud Tool Exporter

To configure Cloud Tool Export settings in AMX:

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

Deploy the Monitoring Session

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

Cloud Tool Exporter - Field References

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable Secure Endpoint option to hide the API keys from the URL. Example: • Normal Endpoint: http://examples.com/mykey/key123 • Masked Endpoint: http://examples.com/mykey/****
Headers	Enter the custom header value of the cloud tool instance.

Fields	Description
	Example: Add Content-Type: application/json in a separate header field. Ensure that the Secure Keys option is enabled.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to transform NetFlow/IPFIX records to JSON format.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value. Note: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Zip	Enable the box to compress the output file. Note: Enable this field when using New Relic as the cloud tool.
Interval (sec)	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Cloud Tool. Both IPv4 and IPv6 are supported. This field is optional.

Configure Kafka Exporter

To configure Kafka Export settings in AMX:

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

Deploy the Monitoring Session

1. Click **Deploy** from the **Actions** menu.

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

Kafka Exporter - Field References

Fields	Description
Alias	Enter the alias name for the Kafka Export.
Topic	The topic name to push JSON streams to, which is generally given to users part of the Kafka administration.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to export NetFlow/IPFIX ingestor types.
Brokers	The URL that contains the Kafka cluster endpoints. Click  to add another broker and click  to remove an existing broker.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds. The default time interval is 30 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.

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

Verify Your Deployment

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

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

Deploy Application Metadata Exporter in Azure using Third Party Orchestration

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

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.
 - o **8 vCPUs, 16 GB memory, and 80 GB disk space** for exporting Application Metadata in JSON format.
 - o **40 vCPUs, 64 GB RAM, 128 GB disk space** for exporting NetFlow/IPFIX records in JSON format.

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

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

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

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

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

- [\(Optional\) Install GigaVUE-FM on Azure](#)
- [\(Optional\) 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 on Azure

You can launch GigaVUE-FM from the Azure VM dashboard or Azure Marketplace.

Install GigaVUE-FM Using Azure VM Dashboard

To install,

1. Go to **Azure VM Dashboard > Virtual Machines** and select **Create** to create an Azure Virtual Machine.

For details, refer to [Create a Linux virtual machine in the Azure](#) in Azure Documentation.

2. Enter the details as mentioned in [Table 1: GigaVUE-FM Installation Steps](#).

Install GigaVUE-FM Using Azure Marketplace

You can install GigaVUE-FM using the Azure Marketplace.

To install,

1. Go to Azure Marketplace and search for Gigamon.

The latest version of Gigamon GigaVUE Cloud Suite for Azure appears.

2. Open the latest version of GigaVUE-FM.

3. Review and accept the terms for Gigamon GigaVUE Cloud Suite for Azure.

For details, refer to [Enable Subscription for GigaVUE Cloud Suite for Azure](#).

4. For details, refer to [Create a Linux virtual machine in the Azure](#) in Azure Documentation.

5. Enter the details as mentioned in [Table 1: GigaVUE-FM Installation Steps](#).

The following table describes the important fields.

Table 1: GigaVUE-FM Installation Steps

Field	Description
Basics	
Subscription	Select your subscription.
Resource Group	Select an existing resource group or create a new resource group. For more information, refer to Create a resource group topic in the Azure Documentation.
System-assigned managed identity	Use a system-assigned managed identity when a resource needs to authenticate to other services, and you want the identity to be created and deleted with the resource. Note: If you update any role it would take more than an hour to reflect in GigaVUE-FM, however, if you use APP registration it would take between 5-10 minutes to update in GigaVUE-FM.
Virtual machine name	Enter a name for the VM.
Region	Select a region for Azure VM.
Availability Zone	Choose your availability zone
Security Type	To enable UEFI secure boot, select Trusted launch virtual machines from the drop-down list. Click Configure security features and ensure that the Enable secure boot check box is enabled.
Image	Select the latest GigaVUE-FM images. Note: You cannot select multiple images for a VM. Refer to Configure GigaVUE Fabric Components in Azure for more details on configuring GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller in Azure.
Size	The recommended instance types are as follows: • GigaVUE-FM - Standard_D4s_v3 • UCT-V Controller - Standard_B4ms • V Series Node - Standard_D4s_v4 • V Series Proxy - Standard_Bms
Authentication Type	We support only SSH public key authentication type

Field	Description
	- SSH public key - Enter the administrator username for the VM. - Enter the SSH public key pair name. - Password - Enter the administrator username for the VM. - Enter the administrator password.
Disks	
Disk Size	The required disk size for GigaVUE-FM is 2 x 40GB .
Networking	
Virtual Network	Select an existing VNet or create a new VNet. For more information, refer to Create a virtual network topic in the Azure Documentation. On selecting an existing VNet, the Subnet and the Public IP values are auto-populated.
Configure network security group	Select an existing network security group or create a new network security group. For more information, refer to Network Security Groups. Configure the Network Security Group to allow GigaVUE-FM to communicate with the rest of the components.

NOTE: Verify the summary before proceeding to create. It will take several minutes for the VM to initialize. After the initialization is completed, you can verify the VM through the Web interface.

After the deployment, navigate to the VM overview page, copy the **Public IP address**, and paste it in a new web browser tab.

If GigaVUE-FM is deployed in Azure, use **admin123A!!** as the password for the **admin** user to login to GigaVUE-FM. You must change the default password after logging in to GigaVUE-FM.

(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

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

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

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



To deploy and register:

1. Log into the vSphere or VMware ESXi web interface.
2. Right-click the ESXi Host, Cluster or data center (if using vSphere) on which you want to deploy the GigaVUE V Series Node and then select **Create/Register VM**. The **New Virtual Machine** wizard appears.
3. In the **Select Creation Type** page, select Deploy a virtual machine from an OVF and OVA file option. Downloaded the AMX supported OVA from the [VUE Gigamon Community portal](#). Extract the OVA to obtain the OVF and VMDK files, then select **vseries-node-file-ovf15.ovf** (AMX-supported OVF that brings up three interfaces) for the AMX deployment.
4. Click **Next**. The **Select OVF and VMDK files** page appears.
5. In the **Select and VMDK files** page, enter a unique name for the virtual machine and upload the .ovf and .vmdk files from your local machine.
6. Click **Next** and the **Select storage** page appears.
7. Select a datastore where the virtual machine's files will be stored.
8. Click **Next**. The **Deployment Options** page appears.
9. In the **Deployment Options** page, select the management port, data port, and tool port as referenced in the prerequisites.
 - a. Select the **Deployment Type** from the list:
 - Do Not Use DHCP – Select this option if you want to use static IP addresses for the management, data, and tool ports.

- Management, Data and Tool Port DHCP – Select this option if you want to use dynamic IP addresses for the management, data, and tool ports.
 - Management Port DHCP – Select this option if you want to use dynamic IP address only for the management port.
 - Tool Port DHCP - Select this option if you want to use dynamic IP address only for the tool port.
 - Data Port DHCP – Select this option if you want to use dynamic IP address only for the data port.
- b. Select Thin/Thick in the **Disk Provisioning** field.
 - c. Unselect the **Power on automatically** check box. This is selected by default. It is recommended to disable the Power on automatically option and review all configuration before switching on the virtual machine.
10. Click **Next**. The **Additional Settings** page appears.
 11. Do the following configuration in the **Additional Settings** page.
 - a. In the **System** section, enter the hostname of the V Series node instance in the **Hostname** field and create a new admin password for the V Series node instance in the **Administrative Login Password** field.

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

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

- c. Enter the required value in the **Management Port MTU size in bytes** field. The default value is 1500B.
- d. Enter the DNS server address in the **Nameserver** field to resolve the domain name of the tool destination URL.
- e. In the **Optional Parameters** section, enter the Monitoring Domain name in the **GroupName** field and connection name in the **SubGroupName** field. The Monitoring Domain and connection name corresponds to the domain name and connection that you created during the [Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source](#) configuration.

- f. Enter the token created in GigaVUE-FM in the **JWT Token used for registration** field. Refer to [Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM](#).
 - g. Enter the GigaVUE-FM IP address and remote port in the **RemoteIP** and **RemotePort** fields respectively. In case of any error, you should re-deploy the V Series.
 - h. In the **Custom node properties** field enter `app_mode=linux_apps` (mandatory).
12. Click **Next** and the **Ready to complete** page appears.
 13. Review all the entered information and modify the configuration if required.
 14. Click **Finish**. When the operation completes, you have successfully deployed a GigaVUE V Series Node.

Verify GigaVUE V Series Node Registration

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

To check the logs:

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

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

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

Create a Monitoring Session to Configure Application Metadata Exporter

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

1. **Access the Monitoring Session Page:**
 - a. In GigaVUE-FM, go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**.
 - b. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.

2. Create an ingress REP:

- a. In the canvas, from the **New** expand menu, drag and drop New Raw Endpoint to the graphical workspace.
- b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
- c. Enter the Alias and Description details for the Ingress Raw End Point and click **Save**.

3. Add AMX Application:

- a. In the canvas, from the **Applications** expand menu, drag and drop AMX to the graphical workspace.
- b. Configure the ingestor settings to process traffic. Refer to [Configure Ingestor in AMX to Process Traffic](#) for details.
- c. Configure the required exporter settings to export attributes to Third Party tools. Refer to [Configure Exporter to Export Attributes to Third Party tools](#) for details.
- d. Create a link between the ingress REP and Application Metadata Exporter application.

4. Create an egress REP:

- a. From the **New** expand menu, drag and drop New Raw Endpoint to the graphical workspace.
- b. On the new raw endpoint icon, select the  menu button and select **Details**. The **Raw** quick view page appears.
- c. Enter the Alias and Description details for the Egress Raw End Point and click **Save**.
- d. Create a link between the Application Metadata Exporter application and egress REP



Notes:

- Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
- When deploying AMX using the Third-Party Orchestration method, Monitoring Session must also be configured in the same work flow.

Configure Ingestor in AMX to Process Traffic

The following scenarios provide more information about the different ways to configure Ingestor AMX:

- [Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka](#)
- [Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka](#)
- [Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka](#)
- [Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka](#)
- [Scenario 5 - Configure Identity Enrichment Metadata using Zscaler Private Access \(ZPA\)](#)

Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka

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

To configure ingestor for AMI:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.

Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka

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

To export GigaVUE Enriched Metadata for Mobile Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the **Type** field, select **AMI** as the input type.

3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which application metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.
9. Select the Enable option to enable the enrichment.
10. Select **Mobility** from the **Type** drop-down menu.
11. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Mobile Networks](#) for more details.

Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka

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

To export GigaVUE Enriched Metadata for Cloud Workloads:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select **AMI** as the input type.
3. Enter the port number on which Application Metadata must be ingested. The default port value is 514.
4. Click  to add another ingestor.
5. In the Type field, select **Mobility Control Plane** as the input type.
6. Enter the port number on which Application Metadata must be ingested.
7. Click  again to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
8. In the **Metadata Enrichment** section, enter a unique name for the enrichment.

9. Select the Enable option to enable the enrichment.
10. Select **Workload** from the **Type** drop down menu.
11. Select the platform as **VMware VCenter** from the drop down list.
12. Select the required attributes from the list. You can use **Select All** to choose all available attributes. Refer to the [Attributes for GigaVUE Enriched Metadata for Cloud Workloads](#) for more details.

Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka

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

To configure ingestor for NetFlow/IPFIX:

1. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
2. In the Type field, select the input type as **NetFlow/IPFIX**. Only the elements specified in the mapping file will be included in the JSON output. Editing or customizing the mapping file is not supported. Refer to [Supported Element Types for NetFlow Integrator Element Mapping](#).
3. Enter the port number on which NetFlow/IPFIX must be ingested. The default port value is 2055.
4. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor. A maximum of 12 NetFlow/IPFIX type ingestors can be configured per deployment.

NOTE: If you select NetFlow/IPFIX as the ingestor type, Metadata Enrichment is disabled by default. Metadata Enrichment is not supported for NetFlow/IPFIX inputs.

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

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

Prerequisites

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

Zscaler Infrastructure

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

- **Acquire Zscaler Tenant Access:** Ensure that you have admin access to the Zscaler Private Access admin portal (ZPA tenant). Refer to [Zscaler Tenant Access](#) for more information.
- **Integrate Identity Provider (IdP):** An IdP (such as Okta, Microsoft Azure AD, or any supported IdP) must be configured and integrated with ZPA. ZPA uses the IdP to authenticate users before granting access to private applications. Refer to [About External Identity Providers](#) for more information.
- **Configure App Connector Group :** At least one App Connector group must be created and have active App Connectors registered to it. You can use an existing App Connector group or create a new one before proceeding. For instructions, refer to the Zscaler documentation. Refer to [Configuring App Connector Groups](#) for more information.
- **Install App Connector:** The Zscaler App Connector package must be downloaded and installed on a Linux-based VM (Ubuntu, plain Linux, or similar) in the domain where your private applications reside. Refer to the Zscaler App Connector deployment guide for platform-specific instructions (Linux, Docker, GCP, Kubernetes, and so on). Refer to [App Connectors](#) for more information.
- **Define Private Applications:** Application segments must be created for all private applications you want to monitor (for example, internal web apps, Jira, databases). Each application segment must include the application IP address and port range, and must be linked to a server group and an App Connector group. Refer to [Configuring Defined Application Segments](#) for more information.

Gigamon Infrastructure

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

Integrate ZPA with Gigamon AMX

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

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

Parameters	Description
Log Receiver	
Name	Enter a name for the log receiver.
Domain or IP Address	Enter the domain or IP address to the IP address of the VSN node on which AMX is deployed (not the ESXi host IP address).
TCP Port	Enter the same port that is configured as the ZPA LSS ingestor port on AMX.
App Connector Groups	Select an app connector group that was created.
Log Stream	
Log Type	Select User Activity as the log type.
Log Template	Select the Custom template.
Log Stream Content	Add the new user activity log ("ConnectionSetupTime": %j {ConnectionSetupTime} }\n) in the content.

Configure Ingestor for AMX

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

To export identity enrichment metadata using ZPA :

1. Configure an ingestor for AMX
 - a. On the AMX quick view in the **Ingestor** section, enter an alias name for the ingestor.
 - b. In the **Type** field, select **AMI** as the input type.
 - c. Enter the port number on which application metadata must be ingested. The default port value is 514.
 - d. Click  to add another ingestor to add multiple inputs to the AMX application and click  to remove an existing ingestor.
2. Configure identity metadata enrichment using ZPA
 - a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
 - b. Select the **Enable** option to enable the enrichment.
 - c. In the **Type** field, select **Identity** as the input type.
 - d. Select **Zscaler Private Access** from the **Platform** dropdown list.

- e. Select the attributes for identity enrichment from the **Attributes** tab.
 - f. In the **Source Information** tab, enter the port number on which AMX receives the ZPA LSS logs.
 - g. Click  to add another ZPA source information to add multiple inputs to the AMX application and click  to remove an existing source information.
3. (Optional) Configure GigaVUE Enriched Metadata workload for ESXi
 - a. On the AMX quick view in the **Metadata Enrichment** section, enter an alias for the enrichment.
 - b. Select the Enable option to enable the enrichment.
 - c. In the **Type** field, select **Workload** as the input type.
 - d. Select **VMware vCenter** from the **Platform** dropdown list.

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

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

Integrate AMX with Third Party Tool

This section explains how to set up the third-party tool and configure the exporter settings required to route application metadata attributes for analysis and visualization.

Set up the Third Party Tool for AMX Integration

You must set up the third party tools to ingest and visualize application metadata.

Refer to [AMX Tool Integration](#) for configuration guidance and the verification steps required to confirm that data export and visualization are functioning correctly.

Analyze Metadata with Gigamon Insights: Use [Gigamon Insights](#) to query your AMI and AMX metadata using plain-language prompts and get immediate, AI-powered answers — no dashboards or manual queries needed.

Configure Exporter to Export Attributes to Third Party tools

AMX supports the following types of exporters:

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

Configure Cloud Tool Exporter

To configure Cloud Tool Export settings in AMX:

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

Deploy the Monitoring Session

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

Cloud Tool Exporter - Field References

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable Secure Endpoint option to hide the API keys from the URL. Example: • Normal Endpoint: <code>http://examples.com/mykey/key123</code> • Masked Endpoint: <code>http://examples.com/mykey/****</code>
Headers	Enter the custom header value of the cloud tool instance.

Fields	Description
	Example: Add Content-Type: application/json in a separate header field. Ensure that the Secure Keys option is enabled.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to transform NetFlow/IPFIX records to JSON format.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value. Note: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Zip	Enable the box to compress the output file. Note: Enable this field when using New Relic as the cloud tool.
Interval (sec)	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Cloud Tool. Both IPv4 and IPv6 are supported. This field is optional.

Configure Kafka Exporter

To configure Kafka Export settings in AMX:

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

Deploy the Monitoring Session

1. Click **Deploy** from the **Actions** menu.

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

Kafka Exporter - Field References

Fields	Description
Alias	Enter the alias name for the Kafka Export.
Topic	The topic name to push JSON streams to, which is generally given to users part of the Kafka administration.
Type	Select any one of the following: AMI - Select this option to export AMI or identity enrichment metadata Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. NetFlow/IPFIX - Select this option to export NetFlow/IPFIX ingestor types.
Brokers	The URL that contains the Kafka cluster endpoints. Click  to add another broker and click  to remove an existing broker.
Labels	Click Add . Enter the following details: - o Enter the Key. - o Enter the Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds. The default time interval is 30 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.

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

Verify Your Deployment

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

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

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.

Parameters	Description
File Upload Errors	Displays the number of times the file was not uploaded to the cloud tools or Kafka consumers due to errors.
Re Exported Entries (Only for Kafka Tool Exporter)	Displays the number of records re-exported from the errors directory after the Kafka tool is restored. Errored files are re-exported automatically at a rate of 10 files per 5 seconds.
Exported Entries Rolling 15 Sec	Displays the rolling average of entries exported per second over the past 15 seconds.
Exported Entries Rolling 30 Sec	Displays the rolling average of entries exported per second over the past 30 seconds.
Exported Entries Rolling 60 Sec	Displays the rolling average of entries exported per second over the past 60 seconds.
Exported MBytes	Total number of megabytes exported.
Exported MBytes Rolling 15 Sec	Rolling average of megabytes exported over the past 15 seconds.
Exported MBytes Rolling 30 Sec	Rolling average of megabytes exported over the past 30 seconds.
Exported MBytes Rolling 60 Sec	Rolling average of megabytes exported over the past 60 seconds.

Ingestor - Statistics to monitor ingestion


Parameters	Description
Name	Displays the name of the ingestor created for this application.
Packets Received	Displays the packets received from the AMI to AMX ingestor.
Packets Dropped	Displays the packets dropped when the traffic passes from AMI to AMX ingestor.
Octets Received	Displays the octets received from the AMI to AMX ingestor.
Octets Dropped	Displays the octets dropped when the traffic passes from AMI to AMX ingestor.
Records Added	Displays the number of workload information pulled from the workload environment.
Records Dropped	Displays the number of the workload details that were dropped from the workload environment.
Request Success	Displays if the workload information was fetched from the cloud environment. The default interval at which the AMX fetches the workload information is 60

Parameters	Description
	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
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_	User Plane Uplink Tunnel Endpoint Identifier

up_link_teid	
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:

- [Azure](#)
- [Azure](#)
- [VMware vCenter](#)
- [Azure Kubernetes Service \(AKS\)](#)
- [Zscaler Private Access \(ZPA\) Attributes](#)

For VPC/Organization Unit centric use cases			For aws service centric use cases			For data volumes use cases			
account_id	vpc/subnet_id	region/az_id	inst_type/id	dns	aws_service	app_name	up_bytes	down_bytes	...
244919184467	vpc-079c6ec	r-d7b3a	t3a.xl/i-dfd23	ec2.us-q	eks	icloud	3455	4566778	...
For account centric use cases			For service centric use cases		For app centric use cases				

- [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.

Azure

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	✓	✓

id			
azure_private_dns_name	Private DNS name assigned to the VM instance	✓	✓
azure_private_ip	Private IP address assigned to the VM instance	✓	✓
azure_public_dns_name	Public DNS name assigned to the VM instance	✓	✓
azure_public_ip	Public IP address assigned to the VM instance	✓	✓
azure_instance_state	Current state of the VM instance	✓	✓
azure_subnet_id	Subnet ID of the VM instance	✓	✓
azure_vnet_id	ID of the VNET in which the VM instance is running	✓	✓
azure_network_if_id	ID of the network interface	✓	✓
azure_network_if_attach_id	ID of the network interface attachment	✓	✓
azure_tags	Tags assigned to the VM instance in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
azure_flat_tags	Tags assigned to the VM instance extracted to top level keys and values	✓	✓
azure_owner_id	Owner of the VM instance	✓	✓
azure_location	Resource Location	✓	✓
azure_scaleset_name	Scale set name	✓	✓
azure_scaleset_mode	Scale set orchestration mode	✓	✓
workload_platform	Platform Name - Azure	✓	✓

VMware vCenter

Attribute	Description	Source	Destination
vmware_vm_name	Name of the guest VM	✓	✓
vmware_vm_network	Network name of the guest VM	✓	✓
vmware_vm_status	Current status of the guest VM	✓	✓
vmware_vm_tags	Tags assigned to the guest VM in kv format with kv delimiter ':' and pair delimiter ','	✓	✓

vmware_vm_flat_tags	Tags assigned to the guest VM extracted to top level keys and values	✓	✓
vmware_vm_os_family_name	OS type of the guest VM (eg: Linux)	✓	✓
vmware_vm_os_distro_name	OS distribution name of the guest VM	✓	✓
vmware_vm_os_distro_version	OS distribution version of the guest VM	✓	✓
vmware_vm_os_distro_pretty_name	OS distribution pretty name of the guest VM	✓	✓
vmware_vm_dns_host_name	Host DNS name of the guest VM	✓	✓
vmware_vm_dns_host_domain_name	Host DNS domain name of the guest VM	✓	✓
vmware_host_name	Host name of the guest VM	✓	✓
vmware_host_ip	Host IP address of the guest VM	✓	✓
vmware_datacenter_name	Datacenter name of the guest VM	✓	✓
vmware_cluster_name	Cluster name of the guest VM	✓	✓
vmware_vcenter_name	vCenter name of the guest VM	✓	✓
workload_platform	Platform Name - VMware	✓	✓

Azure Kubernetes Service (AKS)

Attribute	Description	Source	Destination
k8s_cluster_name	Name identifying the Kubernetes cluster	✓	✗
k8s_cluster_version	Version of Kubernetes running on the cluster	✓	✗
k8s_cluster_status	Operational state of the Kubernetes cluster	✓	✗
k8s_node_name	Identifier for the node within the cluster	✓	✗
k8s_node_internal_ip	Internal IP address assigned to the node	✓	✗
k8s_node_external_ip	External IP address of the node	✓	✗
k8s_node_os_version	Operating system version installed on the node	✓	✗
k8s_node_os_image	Name of the OS image used by the node	✓	✗
k8s_node_kernel_version	Kernel version of the node	✓	✗
k8s_node_cri_version	Version of the container runtime	✓	✗

	interface used on the node		
k8s_node_status	Health and availability status of the node	✓	✗
k8s_node_labels	Node labels parsed into key-value pairs for metadata enrichment	✓	✗
k8s_pod_name	Name assigned to the Kubernetes pod	✓	✗
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✓	✗
k8s_pod_namespace	Namespace in which the pod is deployed	✓	✗
k8s_pod_status	Operational state of the pod	✓	✗
k8s_pod_ip	IP address assigned to the pod	✓	✗
k8s_service_name	Name of the Kubernetes service	✓	✗
k8s_service_type	Type of service	✓	✗
k8s_container_name	Name of the container running inside the pod	✓	✗
k8s_container_image	Image used to create the container	✓	✗
k8s_container_port	Container port details	✓	✗
k8s_cluster_name	Name identifying the Kubernetes cluster	✗	✓
k8s_cluster_version	Version of Kubernetes running on the cluster	✗	✓
k8s_cluster_status	Operational state of the Kubernetes cluster	✗	✓
k8s_node_name	Identifier for the node within the cluster	✗	✓
k8s_node_internal_ip	Internal IP address assigned to the node	✗	✓
k8s_node_external_ip	External IP address of the node	✗	✓
k8s_node_os_version	Operating system version installed on the node	✗	✓
k8s_node_os_image	Name of the OS image used by the node	✗	✓
k8s_node_kernel_version	Kernel version of the node's operating system	✗	✓
k8s_node_cri_version	Version of the container runtime interface used on the node	✗	✓
k8s_node_status	Health and availability status of the node	✗	✓
k8s_node_labels	Node labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_name	Name assigned to the Kubernetes	✗	✓

	pod		
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_namespace	Namespace in which the pod is deployed	✗	✓
k8s_pod_status	Operational state of the pod	✗	✓
k8s_pod_ip	IP address assigned to the pod	✗	✓
k8s_service_name	Name of the Kubernetes service	✗	✓
k8s_service_type	Type of service	✗	✓
k8s_container_name	Name of the container running inside the pod	✗	✓
k8s_container_image	Docker image used to create the container	✗	✓
k8s_container_port	Container port details	✗	✓

Zscaler Private Access (ZPA) Attributes

Attribute	Description
z_iss_username	Username of Zscaler Private Access user.
z_iss_client_private_ip	Private IP address of the Zscaler Private Access client
z_iss_client_public_ip	Public IP address of the Zscaler Private Access client.
z_iss_customer	Name of Zscaler Private Access customer.
z_iss_connection_id	ID of the Zscaler Private Access connection.
z_iss_connection_status	Status of the Zscaler Private Access connection.
z_iss_connection_status_code	Status code of the Zscaler Private Access connection.
z_iss_connector	Name of the Zscaler Private Access application connector.
z_iss_hostname	Name of the host device as given by the Zscaler client connector.
z_iss_client_country_code	Country code of the Zscaler Private Access client location.
z_iss_policy	Name of the Zscaler Private Access policy.
z_iss_app_group	Group name of the Zscaler Private Access application.
z_iss_policy_processing_time	Processing time of the Zscaler Private Access policy.
z_iss_idp	Identity provider used by Zscaler Private Access.
z_iss_client_city	City of the Zscaler Private Access client location.
z_iss_platform	Client platform of Zscaler Private Access.
z_iss_connection_setup_time	Connection setup time for Zscaler Private Access.

Supported Element Types for NetFlow Integrator Element Mapping

This section outlines the element types (v5, v9, and v10) supported by the AMX application for element mapping from the ingest NetFlow/IPFIX flow records to JSON on the following platforms. These elements are predefined and cannot be configured:

- VMware ESXi
- VMware NSX-T

Elements	Description	Supported for NetFlow v5	Supported for NetFlow v9	Supported for NetFlow v10/IPFIX
type	Type of flow message	✓	✓	✓
sampler_address	Address of the device that generated the packet	✓	✓	✓
sampling_rate	Sampling rate of the flow	✓	✓	✓
export_time_ns	Export time in Nano seconds	✓	✓	✓
time_received_ns	Timestamp in nanoseconds of when the message was received	✓	✓	✓
sequence_num	Sequence number of the flow packet	✓	✓	✓
source_id	An identifier that uniquely identifies the flow source. For IPFIX it refers to the observation domain ID.	✗	✓	✓
etype	Ethernet type (0x86dd for IPv6...)	✓	✓	✓
src_addr	Source address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_SRC_ADDR(8), IPV6_SRC_ADDR(27) • NetFlow v10/IPFIX - sourceIPv4Address(8), sourceIPv6Address(27)	✓	✓	✓
dst_addr	Destination address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_DST_ADDR (12), IPV6_DST_ADDR (28) • NetFlow v10/IPFIX -	✓	✓	✓

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

	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)	x	✓	✓
in_if	Input interface - NetFlow v5 – input - NetFlow v9 - INPUT_SNMP (10) - NetFlow v10/IPFIX - ingressInterface (10)	✓	✓	✓
out_if	Output interface - NetFlow v5 – output - NetFlow v9 - OUTPUT_SNMP (14) - NetFlow v10/IPFIX - egressInterface (14)	✓	✓	✓
time_flow_start_ns	Time the flow started in nanoseconds. Refer to Important Notes . - NetFlow v5 – System uptime and first - NetFlow v9 - System uptime and FIRST_SWITCHED (22) - NetFlow v10 - flowStartXXX (150, 152, 154, 156)	✓	✓	✓
time_flow_end_ns	Time the flow ended in nanoseconds. Refer to Important Notes . - NetFlow v5 – System uptime and last - NetFlow v9 - System uptime and LAST_SWITCHED (21) - NetFlow v10/IPFIX - flowEndXXX (151, 153, 155, 157)	✓	✓	✓
tcp_flags	TCP flags - NetFlow v5 – tcp_flags - NetFlow v9 - TCP_FLAGS (6) - NetFlow v10/IPFIX - tcpControlBits (6)	✓	✓	✓
ip_length_min	Minimum IP packet length - NetFlow v9 - MIN_PKT_LENGTH (25) - NetFlow v10/IPFIX - minimumIpTotalLength (25)	x	✓	✓
ip_length_max	Maximum IP packet length - NetFlow v9 - MAX_PKT_LENGTH (26) - NetFlow v10/IPFIX - maximumIpTotalLength (26)	x	✓	✓
next_hop	Nexthop IP address - NetFlow v5 – nexthop - NetFlow v9 - IPV4_NEXT_HOP (15), IPV6_NEXT_HOP (62) - NetFlow v10/IPFIX - ipNextHopIPv4Address	✓	✓	✓

	(15), ipNextHopIPv6Address (62)			
ipv4_next_hop	Nexthop IPv4 address - NetFlow v9 - IPV4_NEXT_HOP (15) - NetFlow v10/IPFIX - ipNextHopIPv4Address (15)	x	✓	✓
ipv6_next_hop	Nexthop IPv6 address - NetFlow v9 - IPV6_NEXT_HOP (62) - NetFlow v10/IPFIX - ipNextHopIPv6Address (62)	x	✓	✓
flow_id	An identifier of a flow that is unique within an Observation Domain - NetFlow v9 - flowId (148) - NetFlow v10/IPFIX - flowId (148)	x	✓	✓
firewall_event	Indicates a firewall event. Allowed values are listed in the firewall Event registry. firewallEvent(233)	x	✓	✓
icmp_type	Type of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpTypeXXX (176, 178), icmpTypeCodeXXX (32, 139)	x	✓	✓
icmp_code	Code of the ICMP message - NetFlow v9 - ICMP_TYPE (32) - NetFlow v10/IPFIX - icmpCodeXXX (177, 179) icmpTypeCodeXXX (32, 139)	x	✓	✓
flow_end_reason	The reason for flow termination. Values are listed in the flowEndReason registry.	x	✓	✓
application_id	Specifies an Application ID - NetFlow v9 - APPLICATION TAG (95) - NetFlow v10/IPFIX - applicationId (95)	x	✓	✓
application_description	Specifies the description of an application - NetFlow v9 - APPLICATION DESCRIPTION (94) - NetFlow v10/IPFIX - applicationDescription (94)	x	✓	✓
application_name	Specifies the name of an application - NetFlow v9 - APPLICATION NAME (96) - NetFlow v10/IPFIX - applicationName (96)	x	✓	✓
tcp_window_size	The window field in the TCP header tcpWindowsSize (186)	x	✓	✓
viptela_vpn_id	Private enterprise number for vIPtela Inc. (41916)	x	x	✓

Migrate Application Intelligence Session to Monitoring Session

Starting from Software version 6.5.00, you must configure the Application Intelligence solution from Monitoring Session Page. After upgrading to 6.5.00, you cannot create a new Application Intelligence Session or edit an existing Application Intelligence Session for a virtual environment from the **Application Intelligence** page.

The following actions are available only when using the existing Application Intelligence Session:

- View Details
- Delete
- Forced Delete

We recommend to migrate the existing sessions to Monitoring Session for full functionality. GigaVUE-FM seamlessly migrates all your virtual Application Intelligence sessions and their connections. If migration fails, all sessions return to their original states.



Points to Note:

- You must have write access for the **Traffic Control Management** Resource in GigaVUE-FM to perform this migration. For details, refer to Create Roles section In GigaVUE Administration Guide
- The migration does not proceed:
 - If any of the existing Application Intelligence Session is in PENDING or SUSPENDED. Resolve the issue and start the migration process.
 - If any of the existing Application Intelligence Session is in FAILED state due to incorrect configuration. Resolve the issue and start the migration process.
 - If an existing Monitoring Session has the same name as the Application Intelligence Session. Change the existing Monitoring Session name to continue with the migration process.
- You cannot continue the session if any of the existing Application Intelligence Session has Application Filtering configured with Advanced Rules as Drop Rule and No Rule Match Pass All in the 5th rule set. In the Monitoring Session, the fifth Rule Set supports either Pass All or Advanced Rules as Drop. Delete this session and start the migration.
- When migrating the Application Intelligence Session, in rare scenarios, the migration process might fail after the pre-validation. In such cases, all the Application Intelligence Session roll back to the Application Intelligence page. Contact Technical Support for assistance.

Migrate your existing Application Intelligence Session to Monitoring Session Page

Follow these steps:

1. In the left navigation pane, select **Traffic > Solutions > Application Intelligence**. You cannot create a new Application Intelligence Session from this page. When you have an existing virtual Application Intelligence Session in the above page, the **Migrate Virtual Application Intelligence** dialog box appears.
2. Review the message and select **Migrate**. The **Confirm Migration** dialog box appears with the list of Application Intelligence Session that you need to migrate.
3. Review the list and select **Migrate**. GigaVUE-FM verifies the requirements and then migrates the Application Intelligence Sessions to the Monitoring Session Page.
4. Select **Go to Monitoring Session Page**.

You can view that all the virtual Application Intelligence Sessions in the Application Intelligence page are migrated to the Monitoring Session Page.

Post Migration Notes for Application Intelligence

After migrating Application Intelligence session to Monitoring Session page, consider the following:

1. **Secure Tunnels in the Options page** If you wish to enable Secure tunnels after migrating the Session, follow these steps:
 - a. Go to **Traffic > Virtual > Orchestrated Flows > Select your cloud platform**.
 - b. Select a Monitoring Session from the Monitoring Sessions list view on the left pane and select the **TRAFFIC ACQUISITION** tab.
 - c. Enable **Secure tunnels**. For information about how to enable secure tunnel for a Monitoring Session, refer to the *Configure Monitoring Session Options* topic in the respective GigaVUE Cloud Suite Deployment Guide.
 - d. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
 - e. From the **Monitoring Sessions** page, select the Monitoring Session for which you enabled Secure Tunnels.
 - f. Select **Actions > Undeploy**. The Monitoring Session is undeployed.
 - g. Select the Monitoring Session for which you enabled Secure Tunnels and edit the Monitoring Session.
 - h. Add the Application Intelligence applications.
 - i. Modify the Number of Flows as per the below table.

Cloud Platform	Instance Size	Maximum Number of Flows
Azure	Large (Standard_D8s_v4)	500k
	Medium (Standard_D4s_v4)	100k

- 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.
- j. Select **Deploy**. For details on how to deploy, refer to [Application Intelligence](#) Application Intelligence section in the GigaVUE V Series Applications Guide.

2. Temporary Loss of Statistics with Version Mismatch

When GigaVUE-FM version is 6.5.00, and the GigaVUE V Series Node version is below 6.5.00, after migrating the Application Intelligence Session to the Monitoring Session and redeploying the monitoring session, a momentary loss in the statistical data of the Application Visualization application appears while redeploying the monitoring session.

3. **Configuration Changes Post-Migration**

After migrating the Application Intelligence Session to monitoring session, if you wish to make any configuration changes, then make sure that the GigaVUE V Series Node version is greater than or equal to 6.3.00.

Header Stripping

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

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

List of Protocols that are supported for stripping:

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



Note:

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

- VMware vCenter
- VMware NSX-T
- Nutanix

Limitation

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

Configure Header Stripping Application

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

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

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

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

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

4. Click **Save**.

Create Custom Port Template

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

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

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

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

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

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

What's Next

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

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

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

GigaSMART NetFlow Generation

NetFlow Generation is a simple and effective way to increase visibility into traffic flows and usage patterns across systems. The flow-generated data can be used to build relationships and usage patterns between nodes on the network.

Refer to the following topic for step-by-step instructions on how to configure NetFlow for Virtual Environment:

- [Configure Application Metadata Intelligence for Virtual Environment](#)- For SecureVUE Plus Base Bundle
- [Create NetFlow Session for Virtual Environment](#) - ForNetVUE Base Bundle

Create NetFlow Session for Virtual Environment

NOTE: This configuration is applicable only when using NetVUE Base Bundle.

To create an NetFlow session, follow these steps:

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

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

Parameter	Description														
Name	Enter a name for the application.														
Description	Enter the description.														
Application Metadata Settings															
Flow Direction	Enable or Disable Bi-Directional Flow behavior. Bi-Directional is enabled by default. Disable this option for Uni-Directional Flow behavior.														
Timeout	Specify the traffic flow inactivity timeout, in seconds. The session will be removed due to inactivity when no packets match.														
Data Link	If you want to include the VLAN ID along with the 5-tuple to identify the traffic flow, select the Data Link and enable the VLAN option.														
Observation ID	Enter a value to identify the source from where the metadata is collected. The range is from 0 to 255. The calculated value of Observation Domain Id in Hexadecimal is 00 01 02 05 , and in Decimal is 66053 .														
Advanced Settings															
Number of Flows	The number of flows supported by the application. Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms. <table border="1"> <thead> <tr> <th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr> </thead> <tbody> <tr> <td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> <tr> <td rowspan="2">Azure</td><td>Large (Standard_D8s_V4)</td><td>500k</td></tr> <tr> <td>Medium (Standard_D4s_v4)</td><td>100k</td></tr> <tr> <td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr> </tbody> </table> - Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k.. - If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.	Cloud Platform	Instance Size	Maximum Number of Flows	VMware	Large (8 vCPU and 16GB RAM)	200k	Azure	Large (Standard_D8s_V4)	500k	Medium (Standard_D4s_v4)	100k	Nutanix	Large (8 vCPU and 16GB RAM)	200k
Cloud Platform	Instance Size	Maximum Number of Flows													
VMware	Large (8 vCPU and 16GB RAM)	200k													
Azure	Large (Standard_D8s_V4)	500k													
	Medium (Standard_D4s_v4)	100k													
Nutanix	Large (8 vCPU and 16GB RAM)	200k													

NOTE: When using NetVUE Base Bundle, Multi-Collect, Fast Mode, and Aggregate round-trip time fields are disabled.

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

Parameter	Description
Exporter Name	Enter a name for the Exporter.
Actions	Using this option, you can perform the following functions: • Add Exporter - Use to add a new Exporter to this Application Metadata Intelligence Application • Apply Template - Use to select the tool template. Refer to Create Custom Tool Templates for more details on what are tool templates and to create custom tool templates. • Save as New Template - Use to save the current configuration as a new custom tool template. • Delete this Exporter - Use to delete the Exporter.
APPLICATION ID	Enable to export the data with Application Id.
Format	Select NetFlow
NetFlow: Select this option to use NetFlow	
Record / Template type	• Segregated - The application-specific attributes and the generic attributes will be exported as individual records to the tool. • Cohesive- The application-specific attributes and the generic attributes will be combined as a single record and exported to the tool.
Active Timeout	Enter the active flow timeout value in seconds.
Inactive Timeout	Enter the inactive flow timeout in seconds.
Version	Select the NetFlow version.
Template Refresh Interval	Enter the time interval at which the template must be refreshed in seconds
APPLICATION & ATTRIBUTES: Select the applications and their attributes for traffic filtering by layer seven applications. You can select a maximum of 64 attributes for each application. (Not applicable when using Netflow V5, V9, Netflow IPFIX(V10), or CEF when the flow direction is Uni-Directional in the above Template drop-down menu.)	
Add Application	Click on the Add Application button. The Add Application dialog box opens. Select a Type. The available options are: • Application Family: Each application is mapped only mapped to one Application Family ▪ Select an Application Family and the Applications that needs to be filtered from the traffic. ▪ Attributes for the selected application is displayed in the Attribute column. You can select the required attributes. • Application Tag: Each application can be mapped to one or more Application Tags. ▪ Select an Application Tag and the Applications that needs to be filtered from the traffic. ▪ Attributes for the selected application is displayed in the

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

5. Click **Save**.


Load Balancing

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

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

To add a load balancing application:

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



X Application Save

Application Load Balancing

Alias lb

Hash Fields ipOnly

Field Location outer

Load balancing groups

+ - Application Endpoint ID 2 Weight 1-100

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

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

4. Click **Save**.

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

Enhanced Load Balancing

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

To enable to enhanced load balancing application:

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

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

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

× Load Balancing: lb
 Save

DETAILS

THRESHOLDS

HEALTH STATUS

STATISTICS

Application Load Balancing

Alias* lb

☐ Stateless
 ☒ Enhanced Load balancing

ELB profiles* Select... ▼

Load balancing groups

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

⏪

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

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

3. Click **Save**.

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

What's Next

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

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

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

SSL Decrypt

License: For information on licensing, refer to the) section of respective GigaVUE Cloud Suite Deployment Guide.

SSL Decrypt application delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. Secure Socket Layer (SSL) is a cryptographic protocol that adds security to TCP/IP communications such as Web browsing and email. The protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them.

SSL encryption secures traffic between a client and a server, such as a Web server. SSL decryption uses keys to decode the traffic between the client and server.

SSL and Transport Layer Security (TLS) protocols consist of a set of messages exchanged between a client and server to set up and tear down the SSL connection between them. To set up the connection, the client and server use the Public Key Infrastructure (PKI) to exchange the bulk encryption keys needed for data transfer.

IMPORTANT: To use SSL Decrypt application in GigaVUE-FM 6.3.00, install new GigaVUE-FM 6.3.00 image. Refer to *GigaVUE-FM Installation and Upgrade Guide* for step-by-step instructions on how to install GigaVUE-FM. SSL Decrypt application does not work if you upgrade from any previous GigaVUE-FM version to GigaVUE-FM 6.3.00.

Keep in mind the following when using SSL Decrypt application:

1. On updating the keys, service, or key maps which are already used in a monitoring session, the monitoring session is dynamically updated, and you need not re-deploy the monitoring session. You can also see if the updated keys, services, or key maps were successfully updated to the monitoring session and the respective GigaVUE V Series Nodes on the **All Events** page. Refer to [Overview of Events](#) section in the *GigaVUE Administration Guide* for detailed information on Events.
2. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the key will be removed from the Key Map. If that key is the only available entry in the Key Map, then it will not be removed.
3. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session that is not deployed, then the key will be removed from the Key Map and if that key is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
4. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the service will be removed from the Key Map. If that service is the only available entry in the Key Map, then it will not be removed.

5. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is not deployed, then the service will be removed from the Key Map and if that service is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
6. In VMware NSX-T platform, the throughput of SSL Decrypt application is improved to 480 Mbps.

Refer to the following topics for more detailed information:

- [Supported Protocols, Algorithms, and Ciphers for SSL Decrypt](#)
- [Configure SSL Decrypt](#)

Supported Protocols, Algorithms, and Ciphers for SSL Decrypt

The supported protocols are as follows:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2

The supported authentication (Au) is as follows:

- RSA

The supported key exchange (Kx) is as follows:

- RSA

The supported encryption algorithms (Enc) are as follows:

- NULL
- RC4
- DES
- 3DES
- AES (including GCM mode)
- CAMELLIA
- SEED
- IDEA

The supported compression algorithm is as follows:

- NULL

The supported digest algorithms are as follows:

- MD5
- SHA1
- SHA2

The supported key sizes are 128, 256, 512, 1024, 2048, and 4096.

The supported TLS extensions are as follows:

- Extended Master Secret, RFC 7627
- Encrypt-then-MAC, RFC 7366

The following table lists the supported ciphers:

Table 1: Supported Ciphers forSSL 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_SHA	RSA	RSA	CAMELLIA_128_CBC	128	SHA
TLS_RSA_WITH_CAMELLIA_256_CBC_SHA	RSA	RSA	CAMELLIA_256_CBC	256	SHA
TLS_RSA_WITH_SEED_CBC_SHA	RSA	RSA	SEED_CBC	128	SHA
TLS_RSA_WITH_NULL_SHA256	RSA	RSA	NULL	0	SHA256
TLS_RSA_WITH_AES_128_CBC_SHA256	RSA	RSA	AES_128_CBC	128	SHA256
TLS_RSA_WITH_AES_256_CBC_SHA256	RSA	RSA	AES_256_CBC	256	SHA256
TLS_RSA_WITH_AES_128_GCM_SHA256	RSA	RSA	AES_128_GCM	128	SHA256
TLS_RSA_WITH_AES_256_GCM_SHA384	RSA	RSA	AES_256_GCM	256	SHA384

All algorithms used for SSL Decrypt application are FIPS 140-3 compliant.

All key URLs must point to an RSA private key stored in the PEM or PKCS12 format, as follows:

- <http://keyserver.domain.com/path/keyfile.pem>
- <https://keyserver.domain.com/path/keyfile.pem>
- <ftp://keyserver.domain.com/path/keyfile.pem>
- <tftp://keyserver.domain.com/path/keyfile.pem>
- [scp://username\[:password\]@keyserver.domain.com/path/keyfile.pem](scp://username[:password]@keyserver.domain.com/path/keyfile.pem)

The supported applications are as follows:

- HTTPS
- FTPS
- SMTP, IMAP, and POP3 with StartTLS

Configure SSL Decrypt

To configure SSL Decrypt Application using GigaVUE-FM follow the steps given below:

- [Upload SSL Key Pairs](#)
- [Create SSL Service](#)
- [Key Mapping](#)
- [SSL Key Pair Store](#)
- [Add SSL Decrypt to Monitoring Session](#)

Upload SSL Key Pairs

To upload an SSL private key, do the following:

1. Go to **Inventory > Resources > Security > SSL Keys**.
2. Click **Add**. The **Create SSL Key Pair** page appears.
3. Enter the following details:

Field	Description
Key Alias	Enter a name for the key.
Comment	Enter a description
Key Type	Select the either of the key type:
PEM	
PassPhrase (optional)	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.
Private key	Enter the Private Key using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
Certificate	Enter the Certificate using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
PKCS12	
PassPhrase	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.

4. Click **Save**.

NOTE: SSL Decrypt application does not support HSM.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Key pair, select the key from the list in the SSL Key Pairs page and click the Edit button.
Delete	To delete a SSL Key pair, select the key from the list in the SSL Key pairs page and click the Delete button.
Delete all	Use this button to delete all the SSL Key pairs in the SSL Key pairs page.
View certificate	To view the certificate associated with the particular SSL Key Pair, select the key from the list in the SSL Key Pairs page and click the View Certificate button.

Create SSL Service

After uploading a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number.

Prerequisite

Before creating a service, upload a private key as described in [Upload SSL Key Pairs](#)

To create a SSL service, do the following:

1. Go to **Inventory > Resources > Security > SSL Service**. The SSL Services page appears.
2. Click **Add**. The **Create SSL Service** page appears.
3. On the Create SSL Service page, enter the following details:

Field	Description
Alias	Enter a name for the SSL Service.
Default Service	Enable this to use default service.
Server IP Address	Enter the IP address of the server in which the service runs.
Server Port	Enter the port number of the server.

4. Click **Save**.

The following actions can also be performed from the SSL Service Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the Edit button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the Delete button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

Key Mapping

After adding the SSL Service, now you map the private key with the service using Key Mapping.

To map a key with the service, follow the steps given below,

1. Go to **Inventory > Resources > Security > SSL Key Mapping**. The SSL Key Pair Mapping page appears.
2. Click **Add**.

3. Enter the Key Mapping Alias.
4. Select the SSL Service and Key Alias from the drop-down.
5. Click **Save**.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the Edit button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the Delete button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

SSL Key Pair Store

SSL Key Pair Store is a repository, that allows you to save all the key under a single location. You can create multiple key stores and in each key store you can store multiple keys.

1. Go to **Inventory > Resources > Security > SSL Key Store**. The SSL Key Pair Store page appears.
2. Click **Add**.
3. Enter the **Key Store Alias** and **Comment**.
4. Click **Save**.

The following actions can also be performed from the SSL Key Pair Store Page:

Field	Description
Edit	To edit a SSL Key pair Store, select the Key Store from the list in the SSL Key Pair Store page and click the Edit button.
Delete	To delete a SSL 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**.



XApplicationSave

Application	SSL Decrypt
Alias	ssl-decrypt
Enable	☒
Key Map	▼
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**.

Interface Mapping (Azure)

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.

Deploy Monitoring Session (Azure)

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.

Monitor and Troubleshoot

Monitor Cloud Health

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

- [Configuration Health Monitoring](#)
- [Traffic Health Monitoring](#)
- [View Health Status](#)
- [View Monitoring Session Statistics \(Azure\)](#)
- [Visualize the Network Topology \(Azure\)](#)

View Health Status

You can view the health status of the Monitoring Session on the Monitoring Session details page. The health status of the Monitoring Session is healthy only if both the configuration health and traffic health are healthy.

View Health Status of an Application

To view the health status of an application across an entire Monitoring Session,

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
2. Select a Monitoring Session and navigate to the **TRAFFIC PROCESSING** tab.
3. Select the application for which you wish to see the health status and select **Details**. The quick view page appears.
4. Select the **HEALTH STATUS** tab.

This displays the application's **Configuration Health**, **Traffic Health**, and the **Operational Health**, along with the thresholds applied to each.

NOTE: The secure tunnel status is refreshed every 5 minutes, and the GigaVUE-FM does not display UCT-V secure tunnel status that is older than 7 minutes. If the secure tunnel in the UCT-V is removed, it takes up to 7 minutes to reset the status on the GigaVUE-FM.

View Operational Health Status of an Application

When you configure the Application Metadata Exporter to use **Kubernetes** as the workload platform, the V Series Node transmits failure and error events to GigaVUE-FM, which processes them and updates the node's health status on the Monitoring Session page.

When interacting with Kubernetes workloads, the system may encounter errors while retrieving resources such as pods, services, nodes, or endpoints. Refer to [Errors](#) for additional error details.

Operational events for Exporter:

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

Format: <Server Type>_<Message>

Server Types: CLOUD EXPORT, KAFKA

Message	Description
UPLOAD_MAX_TRIES_EXCEED	Upload retries exceeded the maximum limit Example: CLOUDEXPORTER_UPLOAD_MAX_TRIES_EXCEED
REACHABILITY_FROM_AMX_TO_TOOLS	AMX failed to reach the tool (Cloud Exporter server or Kafka server) Example: CLOUDEXPORTER_REACHABILITY_FROM_AMX_TO_TOOLS
NO_IP_ADDRESS	No IP address was configured on the interface Example: CLOUDEXPORTER_NO_IP_ADDRESS
EXPORTER_UPLOAD_ERROR	Upload to the exporter failed Example: CLOUDEXPORTER_EXPORTER_UPLOAD_ERROR
EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED	The exporter health-monitor probe or upload to the tool timed out. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED
EXPORTER_ERR_EVENT_CONTEXT_CANCELED	The exporter health-monitor probe or upload was canceled before completion, typically due to shutdown or reconfiguration. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_CANCELED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_CANCELED

Message	Description
EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL	The hostname of the configured Kafka broker or Cloud Exporter endpoint could not be resolved due to a DNS failure. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL KAFKA_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL
EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE	The exporter probe failed at the network layer, for example, due to no route to the host or an unreachable network. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE KAFKA_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE
EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR	The exporter probe failed at the transport layer, such as connection refusal, TCP reset, or TLS handshake failure. Example: CLOUDEXPORTER_EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR KAFKA_EXPORTER_ERR_EVENT_TRANSPORT_LAYER_ERR

Operational events for Enrichment:

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

Format: <Operation Type>_<Message>

Operation Types: GETSERVICES, GETPODS, GETNODES, GETENDPOINTSlices, 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 (Azure)

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.

View FQDN resolution details

Use the **FQDN tab** under the **V SERIES NODES tab** to view the FQDNs configured for supported egress tunnel destinations in the monitoring session. This tab displays the configured FQDN, the current resolution status, and the IP address to which the FQDN currently resolves.

- Use this tab to verify that the system resolved the configured FQDN and to identify DNS resolution issues.
- If the system resolves the FQDN successfully, the status shows OK and the Resolve IP field displays the mapped IP address.

- If the system cannot resolve the FQDN, the status shows Failed and the Resolve IP field displays 0.0.0.0.

Following is the list of V Series Nodes where MS1 is deployed.

V SERIES NODE	IP ADDRESS	VPC	HEALTH STATUS	DEPLOYMENT STATUS
Gigamon-VSeriesNode-	54,189.50.107	devtest_FQA-de7c	Healthy	Success

1 v series nodes total

V Series Node: Gigamon-VS... Details **FQDN**

FQDN	STATUS	RESOLVE IP	NOTE
sslqa1.com	Resolved	22.0.3.110	—
sslqa2.com	Resolved	2600:1f14:388a:9c01:b477:c431...	—
sslqa3.com	Resolved	22.0.3.250	—

Visualize the Network Topology (Azure)

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 Azure

You can perform the following administrative tasks:

- [Manage Monitoring Domain](#)
- [Configure Certificate Settings](#)

- [Check Permissions while Configuring GigaVUE Fabric Components using GigaVUE-FM](#)
- [Set Up Email Notifications](#)
- [Configure Proxy Server](#)
- [Configure Azure Settings](#)
- [Role Based Access Control](#)
- [About Events](#)
- [About Audit Logs](#)
- [High Availability for Cloud Infrastructure](#)

Manage Monitoring Domain

You can view the details of the monitoring domain created in the list view. The list view details is based on:

- [Monitoring Domain](#)
- [Connections Domain](#)
- [Connections Domain](#)
- [UCT-Vs](#)

You can also filter the monitoring domain based on a specified criterion. In the monitoring domain page there are two filter options as follows:

- Right filter: Select the **Filter** button on the right to filter the Monitoring Domain based on a specific criterion.
- Left filter: Select  to filter the based on the Monitoring Domain and Connections. You can click **+** to create a new monitoring domain. This filter once applied also works even when the tabs are swapped.

To edit or delete a specific monitoring domain, select the monitoring domain, click the ellipses "...".

When you click a monitoring domain, you can view details of it in a split view of the window. In the split view window, you can view the details such as **Configuration**, **Launch Configuration** and **V Series configuration**.

Monitoring Domain

The list view shows the following information in the monitoring domain page:

- Monitoring Domain
- Connections
- Tunnel MTU
- Acquisition Method
- Centralized connection

- Management Network

NOTE: Click the  to select the columns that should appear in the list view.

Use the following buttons to manage your Monitoring Domain:

Button	Description
New	Use to create new connection
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.
Actions	You can select a monitoring domain and then perform the following options: • Edit Monitoring Domain- Select a monitoring domain and then click Edit Monitoring domain to update the configuration. • Delete Monitoring Domain - You can select a monitoring domain or multiple monitoring domains to delete them. • Deploy Fabric - -You can select a monitoring domain to deploy a fabric, you cannot choose multiple monitoring domains at the same time to deploy fabrics. This option is only enabled when there is No FABRIC (launch configuration) for that specific monitoring domain and GigaVUE-FM orchestration is enabled.. You must create a fabric in the monitoring domain, if the option is disabled • Upgrade Fabric-You can select a monitoring domain or multiple monitoring domains to upgrade the fabric. You can upgrade the V Series nodes using this option. • Delete Fabric- You can delete all the fabrics associated with the monitoring domain of the selected Fabric. • Edit SSL Configuration - You can use this option to add Certificate Authority and the SSL Keys. • View Permission Status Report - The View Permission Status Report monitors, audits, and reviews the current status of permissions assigned to users or roles, ensuring proper access control and compliance with security policies.
Filter	Filters the monitoring domain based on the list view options that are configured: • Tunnel MTU • Acquisition Method • Load Balancer • Centralised Connection • Management Subnet You can view the filters applied on the top of the monitoring domain page as a button. You can remove the filters by closing the button.

Connections Domain

To view the connection related details for a monitoring domain, click the **Connections** tab.

The list view shows the following details:

- Connections
- Monitoring Domain
- Status
- Fabric Nodes
- User Name
- Region

Fabric

To view the fabric related details for a monitoring domain, click the **Fabric** tab.

The list view shows the following details:

- Connections
- Monitoring Domain
- Fabric Nodes
- Type
- Management IP
- Version
- Status - Click to view the upgrade status for a monitoring domain.
- Security groups

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

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

You can use the **Actions** buttons in this page to perform the following actions in the Monitoring domain page:

Buttons	Description
Edit Fabric	Use to edit a GigaVUE V Series Nodes.
Upgrade Fabric	Use to upgrade GigaVUE V Series Nodes. Refer to Upgrade GigaVUE V Series Node in GigaVUE-FM for ESXi for more detailed information on how to upgrade.

Buttons	Description
Delete Fabric	Use to delete a GigaVUE V Series Node.
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

UCT-Vs

To view all the UCT-Vs associated with the available Monitoring Domains click the **UCT-Vs** tab.

The list view shows the following details:

- Monitoring Domain
- IP address
- Registration time
- Last heartbeat time
- Agent mode
- Status

For details on **Settings**, refer to [Configure Azure Settings](#).

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.

Configure Proxy Server

GigaVUE-FM cannot reach Azure API endpoints if the virtual network (VNet) hosting GigaVUE-FM does not have internet access. To enable connectivity, you must configure a proxy server.

To create a proxy server,

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Proxy Server Configuration**.
3. In the **Proxy Server Configuration** page, select **Add**.

The **Configure Proxy Server** page appears.

Configure Proxy Server Save Cancel

Alias	Alias
Host	IP Address
Port	0 - 65535
Username	Username
Password	Password

☐ NTLM

4. Select or enter the appropriate information for the following fields.
 - a. **Alias:** The name of the proxy server.
 - b. **Host:** The host name or the IP address of the proxy server.
 - c. **Port:** The port number that the proxy server uses for Internet connection.
 - d. **Username:** (Optional) The username of the proxy server.
 - e. **Password:** The password of the proxy server.
 - f. **NTLM:** (Optional) The type of the proxy server used to connect to the VNet.
 - g. **Domain:** The domain name of the client accessing the proxy server.
 - h. **Workstation:** (Optional) The name of the workstation or the computer accessing the proxy server.
5. Select **Save**.

GigaVUE-FM adds the new proxy server configuration to the Proxy Server Configuration page. The proxy server is also listed in the Azure Connection page.


NOTE: If you change any of the fields in the Proxy Server Configuration page after the initial connection is established between the GigaVUE-FM and Azure, then you must also edit the connection and select the proxy server again and save (in the Azure Connection Page). Otherwise, GigaVUE-FM does not use the new saved configuration that was saved, and may lose connectivity to the Azure platform.

Configure Azure Settings

This section provides information on how to customize Azure-related settings in GigaVUE-FM to control refresh intervals, connection limits, tunnel ranges, and UCT-V behavior.

Access Advanced Azure Settings

1. Go to **Inventory > VIRTUAL > Azure**.
2. Select **Settings > Advanced Settings**.



Advanced Settings		Edit
Refresh interval for instance target selection inventory (secs)	120	
Refresh interval for fabric deployment inventory (secs)	900	
Number of UCT-Vs per V Series Node	100	
Refresh interval for UCT-V inventory (secs)	900	
Traffic distribution tunnel range start	8000	
Traffic distribution tunnel range end	8512	
Traffic distribution tunnel MTU	9001	
Permission status purge interval in days	30	
Use UCT-V conf file ①	Enabled	
Un-Registration Timeout For FabricNode (secs)	150	
Reboot threshold limit for UCT-V Controller down ①	2	
Monitoring Session redeployment wait time during FMHA switchover, if the MS deployment in-progress.	300	

The page displays the configuration listed below.

Settings	Description
Refresh interval for VM target selection inventory(secs)	Specifies the frequency for updating the state of Virtual Machines target selection in Azure.
Refresh interval for fabric deployment inventory (secs)	Specifies the frequency for updating the state of fabric deployment information such as subnets, security groups, images, and VNets.
Number of UCT-Vs per GigaVUE V	Specifies the maximum number of instances that you can assign to the GigaVUE V Series node.

Settings	Description
Series Node	
Refresh interval for UCT-V inventory (secs)	Specifies the frequency for discovering the UCT-Vs available in the VNet. 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.
Traffic distribution tunnel MTU	Specifies the MTU value for the traffic distribution tunnel.
Permissions status purge interval in days	Specifies the number of days GigaVUE-FM retains permissions report before auto purging.
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, create a new Monitoring Domain and configure the Monitoring Session. - When changing the UCT-V conf file option from disabled to enabled, perform 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
Un-Registration Timeout For FabricNode (secs)	Specify the unregistration wait time between 150 to 900 seconds to control how long GigaVUE-FM waits before removing an unhealthy node; the default is 180 seconds.
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 the down state. GigaVUE-FM retries every 60 seconds.

Role Based Access Control

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

which the access rights of a user depends on the following:

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

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

Resource Category	Cloud Configuration Task
Physical Device Infrastructure Management: This includes the following cloud infrastructure resources: • Cloud Connections • Cloud Proxy Server • Cloud Fabric Deployment • Cloud Configurations • Sys Dump • Syslog • Cloud licenses • Cloud Inventory	• Configure GigaVUE Cloud Components • Create Monitoring Domain and Launch Visibility Fabric • Configure Proxy Server
Traffic Control Management: This includes the following traffic control resources: • Monitoring session • Threshold Template • Stats • Map library • Tunnel library • Tools library • Inclusion/exclusion Maps	• Create, Clone, and Deploy Monitoring Session • Create and Apply Threshold Template • Add Applications to Monitoring Session • Create Maps • View Statistics • Create Tunnel End Points
Third Party Orchestration: This includes the following resource: • Cloud Orchestration	Deploy the fabric components using Third Party Orchestration. For more details, refer to Deploy GigaVUE Cloud Suite for Azure using Third Party Orchestration .

NOTE: Cloud APIs are also RBAC enabled.

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

About Events

The Events page displays all the events occurring in the virtual fabric component, VM Domain, and VM manager. An event is an incident that occurs at a specific point in time. Examples of events include:

- Cloud provider License Expiry
- UCT-V Inventory Update Completed
- Cloud provider Connection Status Changed

An Alarm is a response to one or more related events. If an event is considered of high severity, then GigaVUE-FM raises an alarm. An example of alarm is your cloud provider license expiry.

The alarms and events broadly fall into the following categories: Critical, Major, Minor, or info.

Access Event

Navigate to **Dashboard > SYSTEM > Events**. The Event page appears.

Source	Time	Event Type	Severity	Affected Entity T...	Affected Entity	Alias	Device IP	Host Name	Scope	Description	Tags
FM	2022-08-10 0...	Licenses Expir...	Info	Floating License					FM	4 Floating	
FM	2022-08-09 0...	Licenses Expir...	Info	Floating License					FM	4 Floating	
FM	2022-08-08 0...	Licenses Expir...	Info	Floating License					FM	4 Floating	
FM	2022-08-07 0...	Licenses Expir...	Info	Floating License					FM	4 Floating	
FM	2022-08-06 0...	Licenses Expir...	Info	Floating License					FM	4 Floating	
FM	2022-08-05 1...	FM Applicatio...	Info	fm application ...				fmha1	fmService	CMS service f...	
FM	2022-08-04 1...	FM Applicatio...	Info	fm application ...				fmha1	fmService	CMS service f...	
FM	2022-08-04 1...	Alarm Delete ...	Critical	VSeries Node	vc-obc-pod2.u...				Alarm	Node Down. P...	

The following table describes the parameters recording for each alarm or event. You can also use filters to narrow down the results.

Controls/ Parameters	Description
Source	The source from where the events are generated. The criteria are: • FM - indicates the event that the GigaVUE-FM fabric manager flagged. • VMM - indicates the event that the Virtual Machine Manager flagged. • FM Health - indicates the event that the health status change of GigaVUE-FM flagged.
Duration	The timestamp when the event occurred or the duration of the event. IMPORTANT: Timestamps or the duration appear in the time zone of the client browser's computer and not the time zone of the node reporting the event. The timestamp is based on the correctly configured clock on the GigaVUE-FM server

Controls/ Parameters	Description
	and converted from UTC to the client computer's configured time zone.
Scope	The category to which the events belong. Events can belong to the following categories: Domain, Node, Card, Port, Stack, Cluster, Chassis, GigaVUE-FM, GigaVUE-VM, and so on. For example, if an event generates a notification for port utilization low threshold, the scope is displayed as Physical Node.
Alarm Type	The type of events that generates the alarms. The types of alarms are Abnormal Fan Operation, Card Unhealthy, Circuit Tunnel Unhealthy, CPU Over Loaded, Device Upgrade Failed.
Event Severity	The severity is one of Critical, Major, Minor, Warning, or Info. Info is informational messages. For example, when a power status change notification is displayed, the message is Info.
Event Status	The status of the event. The status is either Acknowledged or Unacknowledged.
Event Type	The type of event that generated the events. The types of events are CPU utilization high, cluster updated, device discovery failed, fan tray changed, netflow generation statistics, and so on.
Affected Entity Type	The resource type associated with the event. For example, when a low disk space notification is generated, 'Chassis' is displayed as the affected entity type.
Cluster ID	Enter the Cluster ID.
Affected Entity	The resource ID of the affected entity type. For example, when low disk space notification is generated, the IP address of the node with the low disk space is displayed as the affected entity.
Device IP	The IP address of the device.
Host Name	The host name of the device.
Alias	Event Alias
Monitoring Domain	The name of the Monitoring Domain.
Connection	The name of the Connection.
Show Non-taggable Entities	Enable to display the events for entities that you cannot tag. For example, Policies, GigaVUE-FM instance, and other such entities.
Tags	Select the Key and the Value from the drop-down list.

To filter the alarms and events,

1. Select **Filter**.

The Filter quick view is displayed.

2. Select the filtering criteria, and then select **Apply Filter**.

The result appears on the Events page.

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 map info	Map	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 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	• Azure • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix
Third-party Orchestration	Generic Mode • Azure • Azure • OpenStack • VMware ESXi • VMware NSX-T • Nutanix • GCP • UCT-C • GCB Integrated Mode • Azure • 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.



- 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.

- 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](#).
- 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

- On the first instance, click the ellipses icon and select **Edit Node**.
- Enter the following details: (These fields are optional)
 - Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:

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

High Availability

Create New There must be atleast three active eligible nodes that are reachable to each other. Cancel Submit

Group Name *

Secure Communication ☒ Pre-Shared Key ☐ Certificate Based

Add GigaVUE-FM Instances Add New Node

IP ADDRESS / DNS NAME	TUNNEL IP ADDRESS / D...	EXTERNAL IP ADDRESS / ...	ROLE	USERNAME	PASSWORD	ENTITY ID	IDP METADATA URL
fcab:0:1:2200:96ae:3c8...	—	—	Active Eligible	admin	*****	—	—

Add New Node

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

- When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.

**Notes:**

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

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

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

High Availability

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

No supporting nodes in cluster ⓘ

Group Name: fmha

No of Nodes: 3

Encryption Authentication Method: Pre-Shared Key ⓘ

Tunnel Status: All tunnels are healthy ⓘ

Change ⓘ

+ 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 2: Additional Fields

Field	Description
Status and Reachability	Status of the GigaVUE-FM instances. Can be: • Active • Standby • Supporting Refer to Active-Eligible instances for details. Reachability: • - indicates that the instances are reachable. • - indicates that the instances are unreachable.
Host Name	Host name of the GigaVUE-FM instance.
Application Status	Status of application services that are responsible for handling the user requests: • Up: Indicates service is ready to handle the user requests • Down: Indicates service is not ready to handle the user requests with one or more underlying services being down. • In Progress: A state that appears between up and down indicating that the services are initializing.
Software Version	Software version of the GigaVUE-FM instance.
System Uptime	Time since the instance is up and active.
Third Party Authentication URL	Third party authentication URL required for SSO configuration.
Upgrade Status	Upgrade status of GigaVUE-FM. This is displayed only when the upgrade is triggered. Otherwise, this field is empty.

You can also remove the instances from the HA group and reduce the size of the HA group as required. However, while removing the instances ensure to retain the quorum¹ number of instances.

¹Quorum number of instances is the minimum number of instances that must be active-eligible in a HA group. In a HA group with five instances, the quorum number is 3 (based on the formula $n/2+1$).

NOTE: To scale components such as Fabric Health Analytics, topology visualization and to improve the performance of GigaVUE-FM by fully utilizing the resource available, it is recommended to configure three active eligible instances and additional supporting instances.

The following are the behavioral changes observed in the HA group with active, standby and supporting instances:

Task	Scenario	Behavior
Upgrade Operation	Few instances not being up (either active-eligible or supporting instances not up)	Upgrade will not be initiated.
	Few supporting instances are removed from the HA group after formation of the group	Upgrade will be initiated, and the supporting instances present in HA group will be upgraded first followed by the active-eligible instances.
	All instances are up and running	Upgrade will be initiated in an active instance. Upgrade will be done in a rolling fashion starting with supporting instances and ending with the active-eligible instances.
System Configurations	All instances are active-eligible	All GigaVUE-FM's GUI will be accessible, and user can configure the system through the same
	HA with active-eligible and few Supporting instances	For supporting instances, GUI will not be available (as application service is shut down). An option is provided in the GUI to turn on the application service and its dependent service(s) and the status can be seen in HA page. You can then configure the system related configuration from the logged in GigaVUE-FM instance once the application status of the supporting instance is up.
Background processes		

-FM Service Distribution	All GigaVUE-FMs are active-eligible instances.	Services will be equally distributed/redistributed among the instances based on the availability of the instances.
	HA group with active-eligible instances and supporting instances.	Services will be distributed/re-distributed only to active eligible instances
-Load Balancer	All GigaVUE-FMs are active-eligible instances.	All GigaVUE-FM API requests will be distributed to all instances in the HA group.
	With supporting instances.	All GigaVUE-FM API requests will be distributed only to active-eligible instances.

Connect GigaVUE-FM High Availability to a Monitoring Domain using GigaVUE-FM Orchestration

When fabric nodes need to communicate with GigaVUE-FM using public IP addresses, you should enable the **Use Public IP** option during Monitoring Domain creation. When **Use Public IP** option is enabled, GigaVUE-FM pushes its public IP address to Fabric Nodes. Fabric Nodes use this IP to communicate back to FM. FM does not initiate communication using its private IP. This setting is essential in cloud or hybrid environments where private IPs are not reachable, such as across cloud platforms. It ensures that GigaVUE-FM can send notifications and accept registrations using an externally accessible IP. Refer to [Supportability Information](#) for the list supported platforms.

Connect GigaVUE-FM High Availability to a Monitoring Domain using Third Party Orchestration

Fabric Nodes can be configured with either the IP address (remoteIP) or the Fully Qualified Domain Name (FQDN) (remoteAddress) of GigaVUE-FM to establish communication. The node sends registration requests to GigaVUE-FM until it successfully registers. In High Availability mode, if a node fails to register with one GigaVUE-FM instance, it automatically attempts registration with the next available GigaVUE-FM in its list until successful. Refer to [Supportability Information](#) for the list supported platforms.

Use the following data to manually configure **/etc/gigamon-cloud.conf** with registration configurations to register the fabric components with GigaVUE-FM.

```
Registration:
  groupName: <Monitoring Domain Name>
  subGroupName: <Connection Name/Alias>
  token: <Token>
  remoteAddress: <remoteIPs/FQDN of the GigaVUE-FM>
```



```

remotePort: 443
sourceIP: <IP address of the fabric node>
nameServer: <DNS server's IP address>
usePublicIP: <true/false>
preferV6: <true/false><If set to true IPv6 will be preferred over IPv4, used only for DNS
resolution>

```

**Notes:**

- Configure either remoteIP or remoteAddress—not both. Using both prevents the node from registering with GigaVUE-FM.
- Remote IP will continue to work as expected. However, if you plan to use FQDN, update the configuration to use the Remote Address field instead.
- When configuring a remote address, you can specify multiple IP addresses. However, if you use a FQDN, ensure that you configure only one address
- FQDN configuration requires a valid nameserver.
- Ensure that the FQDN configured for Third Party Orchestration resolves to all IP addresses in the High Availability cluster.

Integrate GigaVUE-FM High Availability with UCT-C

UCT-C supports integration with GigaVUE-FM in high availability mode. You can configure the controller with either a Fully Qualified Domain Name (FQDN) or a comma-separated list of GigaVUE-FM IP addresses in the Helm configuration. The controller iterates through the provided FQDN or IP list and attempts to register with the master GigaVUE-FM node. If the GigaVUE-FM IP changes, you must update the configuration again. This enhancement enables UCT-C to maintain connectivity and recover from GigaVUE-FM failover events in HA deployments.

Below are the communication parameters:

- **fm_fqdn** – Specifies the GigaVUE-FM HA FQDN value used by the UCT-C Controller for communication with the GigaVUE-FM HA group. When this parameter is set, the controller uses the IP addresses resolved from the FQDN.
- **fmha_ip_list** – Specifies a comma-separated list of GigaVUE-FM HA IP addresses used by the UCT-C Controller for communication with the GigaVUE-FM HA group. If fm_fqdn is not set, the controller uses the IP addresses provided in fmha_ip_list by default.

Integrate GigaVUE-FM High Availability with Gigamon Containerized Broker

GigaVUE-FM High Availability group can be configured in an Gigamon Containerized Broker through a DNS Server in which the GigaVUE-FM instances are launched and deployed in the Gigamon Containerized Broker environment.

For GCB to make use of high availability feature of GigaVUE-FM, you must configure the FQDN (Fully Qualified Domain Name) of the GigaVUE-FM.

In a standalone GigaVUE-FM, the GCB uses the FQDN name of the GigaVUE-FM (if configured). It uses the legacy method of using the configured IP address when:

- The FQDN is not configured or
- The GCB fails to resolve to an IP address.

For example:

GCB Controller YAML file:

```
env:
  - name: GCB_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 log in to any instance in the HA group, the Dashboard page appears. On specific pages, you can use the instance dropdown list in the page header to view information for a particular instance.

Remove Standby GigaVUE-FM Instance

To remove or replace a standby GigaVUE-FM instance from the GigaVUE-FM HA group follow the below listed steps:

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 1 Remove Standby GigaVUE-FM Instance](#).



The screenshot shows the 'High Availability' dashboard. At the top, a green checkmark indicates the group is 'Healthy'. Below this, a table lists the nodes in the group. The first node is in a 'Standby' state. A context menu is open for this node, showing options: 'Edit', 'Reboot', and 'Remove from group'. A red arrow points to the 'Remove from group' option.

STATUS	IP ADDRESS / D...	TUNNEL IP ADD...	ROLE	EXTERNAL IP A...	ENTITY ID	HOST NAME	APPLICATION STA...	SOFTWARE VERSI...	SYSTEM UPTIME	IDP METADATA ...	
Standby		—	Active Eligible		FM3	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...
Active		—	Active Eligible		FM1	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...
Standby		—	Active Eligible		FM2	fm-6.12.00-549...	Up		3 days, 3 hours, ...		...

Figure 1 Remove Standby GigaVUE-FM Instance

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 3: 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.

State	Number of GigaVUE-FM Instances	Description
		NOTE: You must recover the standby GigaVUE-FM instance within 3 days (72 hours). Failure to do so will cause the instance to move out of the High Availability group.
Incomplete	One GigaVUE-FM instance is up and running	Only one GigaVUE-FM instance is in active state. The other two GigaVUE-FM instances have either not joined the HA group or have left the HA group.
Standalone	One GigaVUE-FM instance is up and running	HA is not configured on the GigaVUE-FM instance.
Suspended	Two or more GigaVUE-FM instances are up and running	HA is configured, but an active GigaVUE-FM instance is yet to be elected. NOTE: The GigaVUE-FM high availability group can move to a Suspended state from the At Risk state when the HA group has only one active-eligible and reachable GigaVUE-FM instance. Events will not be captured for the status of the GigaVUE-FM instance that went down (which causes the HA group to move into a suspended state) as there is no active/leader instance in the High Availability group.



In case of FMHA going to suspended mode from At-Risk mode, the reachability status of node which went down (due to which HA went into suspended mode) won't be audited as writes are not accepted if there is no active/leader instance in the DB cluster

Figure 2 *High Availability States*

Fail over Mechanism

The active GigaVUE-FM instance in the high availability group may fail at times resulting in one of the standby instances to take over and become the active instance. This process is called failover.

The following table provides the reasons for failover:

Reason for Failover	Description
Reloading the active GigaVUE-FM instance	An active GigaVUE-FM instance is reloaded (using the Reboot option) to bring back the HA group to healthy state again.
Planned downtime of the active GigaVUE-FM instance	An active GigaVUE-FM instance is brought down due to various reasons, for example to upgrade to a newer software version.
Monitoring session deployment during HA leader switchover	The new leader automatically redeploys Monitoring Sessions that failed during the switchover.
Node registration during HA leader switchover	Registration requests may temporarily fail or be delayed. Fabric nodes automatically retry registration. The new leader's IP address is configured as the RMQ notification target on all fabric nodes.
Certificate handling or fabric 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.

Scenario	Changes in GUI
	NOTE: Refresh the page after a minute to view the new active GigaVUE-FM instance. However, the page will get updated automatically after 5 minutes.
What happens to an active GigaVUE-FM instance when it fails (either by itself or if failover is triggered manually)?	- One of the standby GigaVUE-FM instances changes to the active state. The GigaVUE-FM instance that was initially in the active state changes to the standby state. It takes a few seconds for this transition. - The GigaVUE-FM GUI of the new active instance will have all the menus and dashboards.
What happens to the GigaVUE-FM instances that were previously in standby state?	- One of the standby GigaVUE-FM instances changes to the active state - The other standby GigaVUE-FM instance remains in the standby state.
What happens to the embedded devices when a new Active instance takes over?	There are no changes to the devices except that they are being managed by the new active GigaVUE-FM instance.
How do you trigger a failover?	Click on the 'Reboot' option on the current active GigaVUE-FM instance to trigger a failover.

Troubleshoot GigaVUE-FM High Availability Issues

Use the following table to troubleshoot issues that you might encounter while working with the HA feature.

Problem	Solution
Unable to add a license to GigaVUE-FM HA group after a failover Reason: Using the MAC Address in the About page to generate the license	Always use the Challenge MAC Address in the Licenses page of the active GigaVUE-FM instance to generate licenses. Add the licenses to the HA group.
Unable to join the GigaVUE-FM HA group after changing the password	Always logout of GigaVUE-FM after changing the password and login again with the new password before joining the HA group.

Problem	Solution
Reason: Not logging out of GigaVUE-FM after changing the password and before joining the HA group	
(missing or bad snippet)(missing or bad snippet)	

Analytics for Virtual Resources

Analytics in GigaVUE-FM is a standalone service that provides data visualization capabilities. Using Analytics¹, you can create visual elements such as charts that are embedded as visualizations. The visualizations are grouped together in dashboards. You can also create search objects using Analytics. Dashboards, Visualizations and Search Objects are called Analytics objects. For details, refer to the [Analytics](#) section in *GigaVUE Fabric Management Guide*.

Rules and Notes:

- You cannot edit or delete these default dashboards. However, you can clone the dashboards and visualizations.

For details, refer to the [Clone Dashboard](#) section in GigaVUE-FM Installation and Upgrade Guide

- Use the Time Filter option to select the required time interval for which you need to view the visualization.

Virtual Inventory Statistics and Cloud Applications Dashboard

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly. For details on dashboard and visualization about the Discover page and Reports page. refer to the [Analytics](#) section.

To access the dashboards,

- Go to  -> **Analytics -> Dashboards.**
- Select the required dashboard to view the visualizations.

The following table lists the various virtual dashboards:

¹Analytics uses the OpenSearch front-end application to visualize and analyze the data in the OpenSearch database of GigaVUE-FM.

Dashboard	Displays	Visualizations	Displays
Inventory Status (Virtual)	Statistical details of the virtual inventory based on the platform and the health status. You can view the following metric details at the top of the dashboard: - Number of Monitoring Sessions - Number of V Series Nodes - Number of Connections - Number of GCB Nodes You can filter the visualizations based on the following control filters: - Platform - Health Status	<i>V Series Node Status by Platform</i>	Number of healthy and unhealthy GigaVUE V Series Nodes for each of the supported cloud platforms.
		<i>Monitoring Session Status by Platform</i>	Number of healthy and unhealthy monitoring sessions for each of the supported cloud platforms
		<i>Connection Status by Platform</i>	Number of healthy and unhealthy connections for each of the supported cloud platforms
		<i>GCB Node Status by Platform</i>	Number of healthy and unhealthy GCB nodes for each of the supported cloud platforms
V Series Node Statistics	Displays the Statistics of the V Series node such as the CPU usage, trend of the receiving and transmitting packets of the V Series node. You can filter the visualizations based on the following control filters: - Platform - Connection - V Series Node	<i>V Series Node Maximum CPU Usage Trend</i>	Line chart that displays maximum CPU usage trend of the GigaVUE V Series Node in 5 minutes interval, for the past one hour. NOTE: The maximum CPU Usage trend refers to the CPU usage for service cores only. Small form factor V Series nodes do not have service cores, therefore the CPU usage is reported as 0.
		<i>V Series Node with Most CPU Usage For Past 5 minutes</i>	Line chart that displays Maximum CPU usage of the GigaVUE V Series Node for the past 5 minutes. NOTE: You cannot use the time based filter options to filter and visualize the data.
		<i>V Series Node Rx</i>	Receiving trend of the

Dashboard	Displays	Visualizations	Displays
		<i>Trend</i>	GigaVUE V Series Node in 5 minutes interval, for the past one hour.
		<i>V Series Network Interfaces with Most Rx for Past 5 mins</i>	Total packets received by each of the GigaVUE V Series Node network interface for the past 5 minutes. NOTE: You cannot use the time based filter options to filter and visualize the data.
		<i>V Series Node Tunnel Rx Packets/Errors</i>	Displays the reception of packet at the Tunnel RX. This is the input to GigaVUE V Series Node, Grouping by tunnel identifier comprising {monDomain, conn, GigaVUE V Series Node, tunnelName}, before aggregation.
		<i>V Series Node Tunnel Tx Packets/Errors</i>	TX is for output tunnels from GigaVUE V Series Node. GigaVUE V Series Node Tunnel Tx Packets/Errors
Dedup	Displays visualizations related to Dedup application. You can filter the visualizations based on the following control filters: - Platform - Connection - V Series Node	<i>Dedup Packets Detected/Dedup Packets Overload</i>	Statistics of the total de-duplicated packets received (ipV4Dup, ipV6Dup and nonIPDup) against the de-duplication application overload.
		<i>Dedup Packets Detected/Dedup Packets Overload Percentage</i>	Percentage of the de-duplicated packets received against the de-duplication application overload.
		<i>Total Traffic In/Out Dedup</i>	Total incoming traffic against total outgoing traffic
Tunnel (Virtual)	Displays visualizations related to the tunneled traffic in both bytes as well as	<i>Tunnel Bytes</i>	Displays received tunnel traffic vs transmitted

Dashboard	Displays	Visualizations	Displays
	the number of packets. You can select the following control filters, based on which the visualizations will get updated: • Monitoring session: Select the required monitoring session. The cloud platform, monitoring domain and connection within the monitoring domain that the V Series node uses are shown in square brackets, comma-separated, after the name, to distinguish the whole path to it. • V Series node: Management IP of the V Series node. Choose the required V Series node from the drop-down. • Tunnel: Select any of the tunnels shown in the Tunnel drop-down. The direction for each tunnel is shown with the prefix in or out. The following statistics are displayed for the tunnel: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets		tunnel traffic, in bytes. • For input tunnel, transmitted traffic is displayed as zero. • For output tunnel, received traffic is displayed as zero.
		Tunnel Packets	Displays packet-level statistics for input and output tunnels that are part of a monitoring session.
App (Virtual)	Displays Byte and packet level statistics for the applications for the chosen monitoring session on the selected V Series node. You can select the following control filters, based on which the visualizations will get updated: • Monitoring session • V Series node • Application: Select the required application. By default, the visualizations displayed includes all the applications. By default, the following statistics are displayed:	App Bytes	Displays received traffic vs transmitted traffic, in Bytes.

Dashboard	Displays	Visualizations	Displays
	- Received Bytes - Transmitted Bytes - Received Packets - Transmitted Packets - Errored Packets - Dropped Packets		
		<i>App Packets</i>	Displays received traffic vs transmitted traffic, as the number of packets.
End Point (Virtual)	Displays Byte and packet level statistics for the un-tunneled traffic deployed on the V Series nodes. The following statistics that are shown for Endpoint (Virtual): - Received Bytes - Transmitted Bytes - Received Packets - Transmitted Packets - Received Errored Packets - Received Dropped Packets - Transmitted Errored Packets - Transmitted Dropped Packets The endpoint drop-down shows <V Series Node Management IP address : Network Interface> for each endpoint. You can select the following control filters that help to update the visualizations: Monitoring session V Series node Endpoint: Management IP of the V Series node followed by the Network Interface (NIC)	<i>Endpoint Bytes</i>	Displays received traffic vs transmitted traffic, in Bytes.
		<i>Endpoint Packets</i>	Displays received traffic vs transmitted traffic, as the number of packets.

NOTE: The Tunnel (Virtual), App (Virtual) and Endpoint (Virtual) dashboards do not show data from the previous releases if the *Monitoring Session [Platform : Domain : Connection]* dashboard filter is applied. This filter relies on the new attributes in the OpenSearch database available only from software version 5.14.00 and beyond.

Analytics for Inline V Series Solution (Azure)

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly.

Analytics support is available for the following cloud platforms:

- AWS
- Azure

To access the dashboards:

1. From the left navigation pane, go to  -> **Analytics -> Dashboards.**
2. Navigate to **System Dashboards -> Inline.**
3. From the **Load Balancer** drop-down list, select the Gateway load Balancer configured in AWS.
4. From the **Monitoring Session** drop-down list, select the Monitoring Session in which Inline V Series solution is configured.
5. From the **Node Name** drop-down list, select the GigaVUE V Series Node.

The following tables lists the various visualizations for Inline V Series solution:

Table 4: Overall 5G Apps Dashboard

Dashboard	Description	Visualizations	Details
Inline Source (Packets)	Displays the overall visualization details of Inline V Series Solution	LoadBalancer to Inline Source Average Packets	Displays the Inline traffic received from the Load balancer to the Inline V Series Node interface in packets.
		Inline Source to Load Balancer Average Packets	Displays the Inline traffic sent back from the Inline V Series Node interface to the Load balancer in packets.
		LoadBalancer to Inline Source App Average Packets	Displays the Inline traffic received from the Inline V Series Node interface to the IVTAP application in packets.
		Inline Source to LoadBalancer App Average Packets	Displays the Inline traffic sent back from the IVTAP application to the Inline V Series Node interface in packets.
		Average IVTAP App Total Packets Drop	Displays the IVTAP application total packet drops while processing the Inline traffic received from Inline V Series Node interface.
		Average Inline Source IVTAP Errors	Displays the IVTAP application errors while processing the Inline traffic received from Inline V Series Node interface
		Average Out-of-band Ingress Tunnel rx	Displays the Out-of-Band traffic (Mirrored traffic) received from

Dashboard	Description	Visualizations	Details
		Packets	Inline V Series Node interface.
		Average Tool Tunnel tx Packets	Displays the bytes transmitted to the tool from GigaVUE V Series Node of the last tier.
		Average Out-of-band Ingress Tunnel Packets Drop	Displays the Out-of-Band traffic packet drops while receiving traffic (Mirrored traffic) from Inline V Series Node interface.
		Average Out-of-band Ingress Tunnel Errors	Displays the Out-of-Band errors while receiving traffic (Mirrored traffic) from Inline V Series Node interface.
Inline Source (Bytes)	Displays the overall visualization details of Inline V Series Solution	Load Balancer to Inline Source Average Bytes	Displays the Inline traffic received from the Load balancer to the Inline V Series Node interface in bytes.
		Inline Source to Load Balancer Average Bytes	Displays the Inline traffic sent back from the Inline V Series Node interface to the Load balancer in bytes.
		LoadBalancer to Inline Source App Average Bytes	Displays the Inline traffic received from the Inline V Series Node interface to the IVTAP application in bytes.
		Inline Source to LoadBalancer App Average Bytes	Displays the Inline traffic sent back from the IVTAP application to the Inline V Series Node interface in bytes.
		Average Out-of-band Ingress Tunnel rx Bytes	Displays the Out-of-Band traffic (Mirrored traffic) received from Inline V Series Node interface in bytes.
		Average Tool Tunnel tx Bytes	Displays the bytes transmitted to the tool from GigaVUE V Series Node of the last tier in bytes.
Heart Beat Analytics		Average LoadBalancer To Inline Source Heart Beat Packets	Displays the Health Check request packets (Heart beat packets) received by Inline V Series Node from Load balancer
		Average Inline Source To LoadBalancer Heart Beat Packets	Displays the Health Check response packets (Heart beat packets) sent by Inline V Series Node to Load balancer.

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**.

The screenshot shows the Azure Fabric Manager interface. At the top, there are tabs for Monitoring Domains, Connections, Fabric, UCT-V, and UCT-V Upgrade. Below the tabs, there are filters for Monitoring Domains (All) and Connections (All). A table lists V Series Nodes with columns: FABRIC NODES, MONITORING DOMAIN, CONNECTIONS, TYPE, and MANAGEMENT IP. The first node, vamsi-inline-vmss_0, is selected. An 'Actions' dropdown menu is open, showing options: Edit Fabric, Delete Fabric, Upgrade Fabric, and Generate Sysdump. The 'Generate Sysdump' option is highlighted.

FABRIC NODES	MONITORING DOMAIN	CONNECTIONS	TYPE	MANAGEMENT IP
☒ vamsi-inline-vmss_0	MD	CN	V Series Node	10.0.0.10
☐ vamsi-inline-vmss_1	MD	CN	V Series Node	10.0.0.12
☐ vamsi-oob-vmss_0	MD	CN	V Series Node	10.0.0.4

Other Actions

- To download a sysdump file, select the file in the lower pane, and then click **Actions > Download**.
- To delete a sysdump file,
 1. Select the file in the lower pane.
 2. Select the desired sysdump file.
 3. Select **Actions > Delete**.
- To bulk delete, select all the sysdump files, and then select **Actions > Delete All**.

FAQs - Secure Communication between GigaVUE Fabric Components (Azure)

This section addresses frequently asked questions about Secure Communication between GigaVUE Fabric Components and GigaVUE-FM. Refer to Secure Communication between GigaVUE Fabric Components section for more details.

1. Is there a change in the upgrade process for GigaVUE-FM and GigaVUE V Series Node?

No. The upgrade process remains unchanged across all supported upgrade paths. You can upgrade your nodes without any additional steps. The upgrade results in the automatic deployment of the appropriate certificates based on the node versions

GigaVUE-FM	GigaVUE V Series Nodes	Custom Certificates Selected (Y/N)	Actual Node Certificate
6.10	6.10	Y	GigaVUE-FM PKI Signed Certificate
6.10	6.9 or earlier	Y	Custom Certificate
6.10	6.9 or earlier	N	Self-Signed Certificate

2. What is the new authentication type used between GigaVUE-FM and the GigaVUE Fabric Components? Is backward compatibility supported?

Backward compatibility is supported, ensuring that fabric components running on version 6.9 or earlier remain compatible with GigaVUE-FM 6.10. The following authentication types are supported across different versions:

GigaVUE-FM	GigaVUE Fabric Components	Authentication
6.10	6.10	Tokens + mTLS Authentication (Secure Communication)
6.10	6.9 or earlier	User Name and Password

3. What are the new ports that must be added to the security groups?

The following table lists the port numbers that must be opened for the respective fabric components:

Component	Port
GigaVUE-FM	9600
GigaVUE V Series Node	80, 8892
GigaVUE V Series Proxy	8300, 80, 8892
UCT-V Controller	8300, 80
UCT-V	8301, 8892, 9902 For more details, refer to Prerequisites for GigaVUE Cloud Suite for Azure .

4. Is the registration process different for deploying the fabric components using Third-Party Orchestration?

Yes. Beginning with version 6.10, you must use tokens in the gigamon-cloud.conf file instead of the username and password. To generate the token in GigaVUE-FM, go to **Settings > Authentication > User Management > Token**. For more details, refer to [Configure Tokens](#).

Example Registration Data for UCT-V:

```
#cloud-config
write_files:
- path: /etc/gigamon-cloud.conf
  owner: root:root
  permissions: '0644'
  content: |
    Registration:
      groupName: <Monitoring Domain Name>
      subGroupName: <Connection Name>
      token: <Token>
      remoteIP: <IP address of the UCT-V Controller 1, <IP address of the UCT-V Controller
2>
      sourceIP: <IP address of UCT-V> (Optional Field)
```

5. Are there any changes to the UCT-V manual installation and upgrade process?

Starting from version 6.10, you must add tokens during manual installation and upgrades.

- Create a configuration file named gigamon-cloud.conf with the token and place it in the /tmp directory during UCT-V installation
- After installing UCT-V, you can add the configuration file in the /etc directory.

Important! Without this token, UCT-V cannot register with GigaVUE-FM.

6. Can I use my PKI infrastructure to issue certificates for the Fabric Components?

Direct integration of your PKI with GigaVUE-FM is not supported. However, you can provide your Intermediate Certificate Authority (CA) to sign the node certificate.

7. **What happens to the existing custom certificates introduced in the 6.3 release?**

The custom certificate feature is not supported for the fabric components with version 6.10 or higher, even if a custom certificate is selected in the Monitoring Domain. However, this feature remains available for older versions.

- When upgrading from version 6.9 or earlier with custom certificates upgrades to version 6.10, the system automatically generates and deploys certificates signed by GigaVUE-FM.
- If deploying version 6.9 or earlier components from a 6.10 GigaVUE-FM, custom certificates are still applied.

8. **How to issue certificates after upgrading the fabric components to 6.10?**

When the upgrade process begins, GigaVUE-FM transmits the certificate specifications to the new fabric components using the launch script. The fabric components utilize these specifications to generate their own certificates.

9. **Is secure communication supported in FMHA deployment?**

Yes, it is supported. However, you must follow a few manual steps before upgrading the fabric components to 6.10. For details, refer to [Configure Secure Communication between Fabric Components in FMHA](#).

NOTE: This step is essential if you are using cloud deployments in FMHA mode and would like to deploy or upgrade the fabric components to version 6.10 or later.