Gigamon®

GigaVUE Cloud Suite Deployment Guide - OpenStack

GigaVUE Cloud Suite

Product Version: 6.14.xx

Document Version: 2.0

(See Change Notes for document updates.)

Copyright © 2026 Gigamon Inc. All rights reserved.

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

Trademark Attributions

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

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

Change Notes

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

Product Version	Document Version	Date Updated	Change Notes
6.14.01/6.14.02	2.0	09/07/2026	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 o Security Group for OpenStack
6.14.00	1.0	07/10/2026	The original release of this document with 6.14.00 GA.

Contents

GigaVUE Cloud Suite Deployment Guide - OpenStack **1**
Change Notes 3
Contents 4
Get Started **12**
GigaVUE Cloud Suite for OpenStack 13
Overview of GigaVUE Cloud Suite for OpenStack 13
GigaVUE Fabric Components and Deployment Options 14
GigaVUE-FM Orchestrated Fabric Deployment 14
GigaVUE-FM 15
UCT-V 15
UCT-V Controller 17
GigaVUE V Series Node 17
GigaVUE V Series Proxy 17
Third Party Orchestrated Fabric Deployment 18
Components for Third Party Orchestration 19
Key Concepts 20
Secure Communication between GigaVUE Fabric Components 20
GigaVUE-FM acts as the PKI 22
Bring Your Own CA 22
Secure Communication in FMHA Mode 22
Supported Platforms 22
Supported Components 22
Rules and Notes 23
Precryption™ 23
How Gigamon Precryption Technology Works 24
Why Gigamon Precryption 24
Key Features 25
Key Benefits 25
How Gigamon Precryption Technology Works 25
Supported Platforms 27
Prerequisites 28
Secure Tunnels 29
Prefiltering 30
Deployment Overview 31
Deployment Planning 32
Identify the Deployment Method 32

GigaVUE-FM Orchestration 33
Third Party Orchestration 33
Define the Deployment Location 33
Identify the Traffic Acquisition Method 33
Identify the Deployment Model 33
Identify the Networks and Subnets Involved 34
Identify the Required Tokens and Credentials 34
Create the Keypairs 34
Identify the UCT-V Deployment Details 34
Deployment Prerequisites 35
Licensing GigaVUE Cloud Suite 35
Default Trial Licenses 35
Volume Based License (VBL) 37
Install GigaVUE-FM on OpenStack 44
Supported Hypervisor for OpenStack 44
Supported Compute Requirements for OpenStack 45
Network Requirements 46
Virtual Network Interface Cards (vNICs) 47
Security Group for OpenStack 47
Reduced Control-Plane Ports 47
GigaVUE-FM 49
UCT-V Controller 49
UCT-V 50
GigaVUE V Series Node 50
GigaVUE V Series Proxy (Optional) 50
UCT-V OVS Controller 51
UCT-V OVS Module 51
GigaVUE FM 51
UCT-V Controller 52
UCT-V 53
GigaVUE V Series Node 54
GigaVUE V Series Proxy (Optional) 55
UCT-V OVS Controller 56
UCT-V OVS Module 56
Legacy Control-Plane Ports 57
GigaVUE FM 57
UCT-V Controller 60
UCT-V 61
GigaVUE V Series Node 62
Giga VUE V Series Proxy(Optional) 63
UCT-C Controller - deployed in Kubernetes worker mode 64
UCT-V Controller 65
GigaVUE V Series Node 65
GigaVUE V Series Proxy(Optional) 65

Key Pairs ... 66
Prerequisites for OVS Mirroring ... 66
OVS-Mirror Requirements ... 67
OpenStack Cloud Environment Requirements ... 67
GigaVUE-FM Version Compatibility ... 69
Default Login Credentials for GigaVUE Fabric Components ... 69
Configure Tokens ... 69
Prerequisite ... 70
Rules and Notes ... 70
Create Token ... 70
Revoke Tokens ... 71
Export Token ... 72
Install GigaVUE-FM on OpenStack ... 73
Upload Fabric Images ... 73
Install GigaVUE-FM on OpenStack ... 76
Initial GigaVUE-FM Configuration ... 78
78
Prerequisite ... 80
Upgrade UCT-V Controller ... 81
To change UCT-V Controller version between different major versions ... 81
To change UCT-V Controller version with in the same major version ... 82
Upgrade GigaVUE V Series Nodes and GigaVUE V Series Proxy ... 83
Pre-Configuration Checklist for OpenStack ... 85
Deploy ... 86
Deploy GigaVUE Cloud Suite for OpenStack using GigaVUE-FM Orchestration ... 86
Acquire Traffic using UCT-V as Traffic Acquisition Method ... 86
Install GigaVUE-FM on OpenStack ... 87
Install UCT-Vs ... 87
Install UCT-V from Ubuntu/Debian Package ... 91
Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS ... 93
Create Monitoring Domain ... 102
Configure GigaVUE Fabric Components in GigaVUE-FM ... 104
Configure UCT-V Controller ... 105
Configure GigaVUE V Series Proxy ... 107
Configure GigaVUE V Series Node ... 107
Configure and Deploy Monitoring Session ... 110
Uninstall UCT-V ... 110
Upgrade or Reinstall UCT-V ... 111
Acquire Traffic using OVS Mirroring as Traffic Acquisition Method ... 114
Install GigaVUE-FM on OpenStack ... 114
Install UCT-V OVS Mirroring Module ... 115

Create Monitoring Domain ... 118
Configure GigaVUE Fabric Components in GigaVUE-FM ... 120
Configure UCT-V Controller ... 121
Configure GigaVUE V Series Proxy ... 123
Configure GigaVUE V Series Node ... 123
Configure and Deploy Monitoring Session ... 126
Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method ... 126
Install GigaVUE-FM on OpenStack ... 126
Create Monitoring Domain ... 127
Configure GigaVUE Fabric Components in GigaVUE-FM ... 129
Configure UCT-V Controller ... 130
Configure GigaVUE V Series Proxy ... 132
Configure GigaVUE V Series Node ... 132
Configure and Deploy Monitoring Session ... 135
Deploy GigaVUE Cloud Suite for OpenStack using Third Party Orchestration ... 135
Generic Mode ... 136
Integrated Mode ... 136
Deploy GigaVUE Fabric Components using OpenStack - Generic Mode ... 136
Acquire Traffic using UCT-V - Without creating Monitoring Domain ... 137
Prerequisites ... 139
Linux Network Firewall Requirements ... 139
Linux UCT-V Installation Scenarios ... 140
Install Linux UCT-Vs ... 140
Install Linux UCT-Vs using Installation Script ... 140
Install Linux UCT-Vs using Manual Configuration ... 143
Post Deployment Check: ... 147
Rules and Notes ... 148
Windows UCT-V Installation Scenarios ... 149
Windows Network Firewall Requirements ... 149
Install Windows UCT-Vs ... 149
Install Windows UCT-Vs using Installation Script ... 149
Acquire Traffic using Customer Orchestrated Source - Without creating Monitoring Domain ... 162
Acquire Traffic using UCT-V - By creating Monitoring Domain ... 173
Prerequisites ... 176
Linux Network Firewall Requirements ... 176
Linux UCT-V Installation Scenarios ... 177
Install Linux UCT-Vs ... 177
Install Linux UCT-Vs using Installation Script ... 177
Install Linux UCT-Vs using Manual Configuration ... 179
Post Deployment Check: ... 183

Rules and Notes ... 184
Windows UCT-V Installation Scenarios ... 185
Windows Network Firewall Requirements ... 185
Install Windows UCT-Vs ... 185
Install Windows UCT-Vs using Installation Script ... 185
Create a Custom User ... 202
Create a Custom Role and Assign the User ... 202
Create a Custom User Group and Assign the Custom Role and the User ... 203
Linux UCT-V Installation ... 204
Windows UCT-V Installation ... 204
Recommended Instance Type ... 207
Configure UCT-V Controller in OpenStack ... 208
Register UCT-V Controller during Instance Launch ... 208
Register UCT-V Controller after Instance Launch ... 209
Assign Static IP address for UCT-V Controller ... 209
Configure UCT-V in OpenStack ... 211
Configure GigaVUE V Series Nodes and V Series Proxy in OpenStack ... 212
Register V Series Nodes or V Series Proxy during Instance Launch ... 212
Register V Series Node or V Series Proxy after Instance Launch ... 213
Assign Static IP address for GigaVUE V Series ... 214
What to do Next ... 216
Deployment GigaVUE Fabric Components using OpenStack - Integrated Mode ... 216
Acquire Traffic using UCT-V as Traffic Acquisition Method ... 217
Prerequisites ... 218
Linux Network Firewall Requirements ... 218
Linux UCT-V Installation Scenarios ... 219
Install Linux UCT-Vs ... 219
Install Linux UCT-Vs using Installation Script ... 219
Install Linux UCT-Vs using Manual Configuration ... 221
Post Deployment Check: ... 225
Rules and Notes ... 226
Windows UCT-V Installation Scenarios ... 227
Windows Network Firewall Requirements ... 227
Install Windows UCT-Vs ... 227
Install Windows UCT-Vs using Installation Script ... 227
Deploy GigaVUE Cloud Suite using Terraform ... 242
License ... 243
Supportability ... 243
Supported Deployment Models ... 243
Rules and notes ... 244
Prerequisites ... 244
Get Started ... 245

Install Terraform ... 245
Deployment Using Terraform ... 246
Deploy Third Party Orchestration by using Terraform ... 246
Best Practices ... 255
Verification and Troubleshooting ... 257
Verify the deployment ... 257
Troubleshoot the deployment ... 257
Configure ... 258
Configure UCT-V Features ... 258
Configure Precryption in UCT-V ... 258
Create Precryption Template for UCT-V ... 260
Rules and Notes: ... 260
Create Precryption Template for Filtering based on Applications ... 260
Create Precryption Template for Filtering based on L3-L4 details ... 261
Create Prefiltering Policy Template ... 265
Configure Secure Tunnel (OpenStack) ... 266
Precrypted Traffic ... 266
Mirrored Traffic ... 267
Prerequisites ... 267
Notes ... 267
Configure Secure Tunnel from UCT-V to GigaVUE V Series Node ... 267
Configure Secure Tunnel between GigaVUE V Series Nodes ... 269
Viewing Status of Secure Tunnel ... 272
Configure Monitoring Session ... 273
Create a Monitoring Session (OpenStack) ... 273
Monitoring Session Page (OpenStack) ... 274
Configure Monitoring Session Options (OpenStack) ... 275
Create Ingress and Egress Tunnels (OpenStack) ... 280
Create Raw Endpoint (OpenStack) ... 289
Create a New Map (OpenStack) ... 290
Example- Create a New Map using Inclusion and Exclusion Maps ... 293
Map Library ... 295
Add Applications to Monitoring Session (OpenStack) ... 296
Application Filtering Intelligence ... 301
Application Metadata Exporter ... 306
Hardware (AMI) ... 308
Hardware (Control Plane Metadata) ... 308
Private Cloud (VMware) ... 309
Access GigaVUE-FM ... 328
Create a Custom User ... 329
Create a Custom Role and Assign the User ... 329
Create a Custom User Group and Assign the Custom Role and the User ... 330

Register V Series Nodes or V Series Proxy during Instance Launch 331
Register V Series Node or V Series Proxy after Instance Launch 332
Assign Static IP address for GigaVUE V Series 333
Scenario 1 - Export Application Metadata in JSON format over HTTPS/Kafka 336
Scenario 2 - Export GigaVUE Enriched Metadata for Mobile Networks and Export over HTTPS/Kafka 337
Scenario 3 - Export GigaVUE Enriched Metadata for Cloud Workloads and Export over HTTPS/Kafka 337
Scenario 4 - Transform NetFlow/IPFIX records to JSON format and export over HTTPS/Kafka 338
Set up the Third Party Tool for AMX Integration 339
Configure Exporter to Export Attributes to Third Party tools 339
Configure Cloud Tool Exporter 339
Deploy the Monitoring Session 339
Cloud Tool Exporter - Field References 340
Configure Kafka Exporter 340
Deploy the Monitoring Session 341
Kafka Exporter - Field References 341
Extend Your AMI and AMX Deployment 342
Verify Your Deployment 343
Application Metadata Intelligence 363
Extend Your AMI and AMX Deployment 371
Application Visualization 373
De-duplication 378
What's Next 381
GigaSMART NetFlow Generation 382
GRE-In-UDP Tunnel Decapsulation 385
Header Stripping 387
What's Next 391
Load Balancing 392
Masking 398
What's Next 400
Slicing 401
What's Next 402
SSL Decrypt 404
Prerequisite 409
5G-Service Based Interface Application 413
Adding CSV file for IP Mapping 433
Adding CSV file for IP Mapping 441
5G-Cloud Application 444
Configuration of GVHTTP2 Application for Casa vTAP 446
Downloading Logs for GVHTTP2 application 448
Configuration of 5G-Cloud Application for Casa vTAP 448

Downloading Logs for 5G-Cloud application ... 451
Configuration of GVHTTP2 Application for Oracle SCP/Nokia SCP ... 454
Download Logs - GVHTTP2 application ... 456
Configuration of 5G-Cloud Application for Oracle SCP/Nokia SCP ... 457
Adding 5G-Cloud Application in Oracle SCP/Nokia SCP ... 458
Download Logs - 5G-Cloud application ... 463
Adding CSV file for IP Mapping ... 463
Add 5G-Cloud Application in Ericsson SCP ... 471
Download Logs - 5G-Cloud application ... 472
Add CSV file for IP Mapping ... 472
5G Cloud Ericsson SCP Configuration - Field References ... 474
How Nokia HEP3 IMS Solution works ... 477
Configure Nokia IMS HEP3 Solution (Non-SBI) ... 478
Reference - Configuration Table ... 480
Add CSV file for IP Mapping ... 485
Recommended VM Specifications ... 487
Reference - Service Map Field Definitions ... 488
FHA Dashboards for 5G-Cloud Applications ... 489
How Nokia HEP3 CMM Solution works ... 490
Nokia CMM Solution ... 490
Configure Nokia CMM HEP3 Solution (SBI and Non-SBI) ... 491
Download Logs - 5G-Cloud / GvHTTP2 Application ... 499
Add CSV file for IP Mapping ... 499
Appendix - Casa vTap Statistics ... 500
Appendix-5G Cloud Nokia HEP3 Support (SBI) ... 506
Interface Mapping (OpenStack) ... 507
Deploy Monitoring Session ... 508
Monitor and Troubleshoot ... 511
Monitor Cloud Health ... 511
View Health Status ... 511
View Health Status of an Application ... 511
View Operational Health Status of an Application ... 512
View Health Status for Individual GigaVUE V Series Nodes ... 513
View Application Health Status for Individual V Series Nodes ... 514
View Monitoring Session Statistics ... 514
Visualize the Network Topology (OpenStack) ... 516
Administer GigaVUE Cloud Suite for OpenStack ... 518
Manage Monitoring Domain ... 518
Monitoring Domain ... 518
Connections Domain ... 519
Fabric ... 520
UCT-Vs ... 521
Configure Certificate Settings ... 521

Configure the OpenStack Settings 522
Shutdown or Restart of OVS traffic 524
Manual shutdown or restart of OVS traffic 524
Automatic shutdown or restart of OVS traffic 524
Role Based Access Control 526
About Audit Logs 527
About Events 528
High Availability for Cloud Infrastructure 530
About GigaVUE-FM High Availability 531
Supportability Information 533
Recommendations, Rules, Notes, and Limitations 534
Licensing Information 536
Hostname Setups 537
Configure GigaVUE-FM HA 537
Configure Secure Communication between Fabric Components in FMHA 546
GigaVUE-FM HA Landing Page 546
Remove Standby GigaVUE-FM Instance 546
Disassemble GigaVUE-FM High Availability Group 547
GigaVUE-FM High Availability States 548
Fail over Mechanism 550
GigaVUE-FM High Availability Scenarios 550
Troubleshoot GigaVUE-FM High Availability Issues 551
Analytics for Virtual Resources 552
Virtual Inventory Statistics and Cloud Applications Dashboard 552
Debuggability and Troubleshooting 556
Sysdumps 556
Sysdumps—Rules and Notes 557
Generate a Sysdump File 557
Troubleshooting 559
OpenStack Connection Failed 559
Handshake Alert: unrecognized_name 559
GigaVUE V SeriesNode or UCT-V Controller is Unreachable 559
FAQs - Secure Communication between GigaVUE Fabric Components (OpenStack) 560

Get Started

GigaVUE Cloud Suite for OpenStack

This guide describes how to install, configure and deploy the GigaVUE Cloud solution on OpenStack. Use this document for instructions on configuring the GigaVUE Cloud components and setting up the traffic monitoring sessions for OpenStack.

Refer to the following sections for details:

- Overview of GigaVUE Cloud Suite for OpenStack
- GigaVUE Fabric Components and Deployment Options
- Key Concepts
- Deployment Prerequisites
- Install GigaVUE-FM on OpenStack
- Deploy GigaVUE Cloud Suite for OpenStack using GigaVUE-FM Orchestration
- Deploy GigaVUE Cloud Suite for OpenStack using Third Party Orchestration
- Configure UCT-V Features
- Configure Monitoring Session
- Monitor Cloud Health
- Administer GigaVUE Cloud Suite for OpenStack
- Analytics for Virtual Resources
- Debuggability and Troubleshooting
- FAQs - Secure Communication between GigaVUE Fabric Components (OpenStack)

Overview of GigaVUE Cloud Suite for OpenStack

GigaVUE-FM fabric manager is a web-based fabric management interface that provides a single-pane-of-glass visibility and management of both the physical and virtual traffic. GigaVUE-FM is a key component of the GigaVUE Cloud Suite for OpenStack.

The OpenStack software is designed for multi-tenancy (multiple projects), where a common set of physical compute and network resources are used to create project domains that provide isolation and security. Characteristics of a typical OpenStack deployment include the following:

- Projects are unaware of the physical hosts on which their instances are running.
- A project can have several virtual networks and may span across multiple hosts.

In a multi-project OpenStack cloud, where project isolation is critical, the Gigamon solution extends visibility for the project's workloads without impacting others by doing the following:

- Support project-wide monitoring domains—a project may monitor any of its instances.
- Honor project isolation boundaries—no traffic leakage from one project to any other project during monitoring.
- Monitor traffic without needing cloud administration privileges. There is no requirement to create port mirror sessions and so on.
- Monitor traffic activity of one project without adversely affecting other projects.

Refer Deploying Gigamon CloudSuite on OpenStack to scale-inand Open vSwitch with Hardware offload and scale-out monitoring tools for more detailed information.

GigaVUE Fabric Components and Deployment Options

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

GigaVUE Cloud Suite for OpenStack 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 OpenStack environment.

GigaVUE Cloud Suite for OpenStack 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 for Virtual Machines (**UCT-V**), formerly known **as G-vTAP Agent**, is a standalone service that is installed on a virtual machine (VM). UCT-V mirrors selected traffic from the VM and sends it to the GigaVUE V Series for processing, analysis, and distribution to monitoring tools. UCT-V is available as Debian (.deb), Red Hat Package Manager (.rpm), ZIP, and MSI packages.

Next Generation UCT-V, formerly known as Next Generation G-vTAP Agent, is a lightweight traffic acquisition solution that enhances UCT-V mirroring performance. It includes a prefiltering capability that allows traffic to be filtered at the tap level before it is sent to the GigaVUE V Series, reducing both network traffic and processing load on the GigaVUE V Series. Next Generation UCT-V is enabled automatically on Windows systems and on Linux systems running Kernel version 5.4 or later.

Prefiltering helps reduce bandwidth consumption and monitoring costs by filtering traffic on the UCT-V before it is forwarded to the GigaVUE V Series Node. Using GigaVUE-FM, you can create a prefiltering policy template and apply it to a monitoring session. Prefiltering is supported only on Next Generation UCT-V and is available on Windows and Linux systems running Kernel version 4.18 or later.

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 Controller6.14.02 can only manage UCT-Vs 6.14.02. If you have the previous version of UCT-V still deployed in the Virtual Network, you must configure both UCT-V Controller6.14.02 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. 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. • 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 OpenStack using Third Party Orchestration.

Key Concepts

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

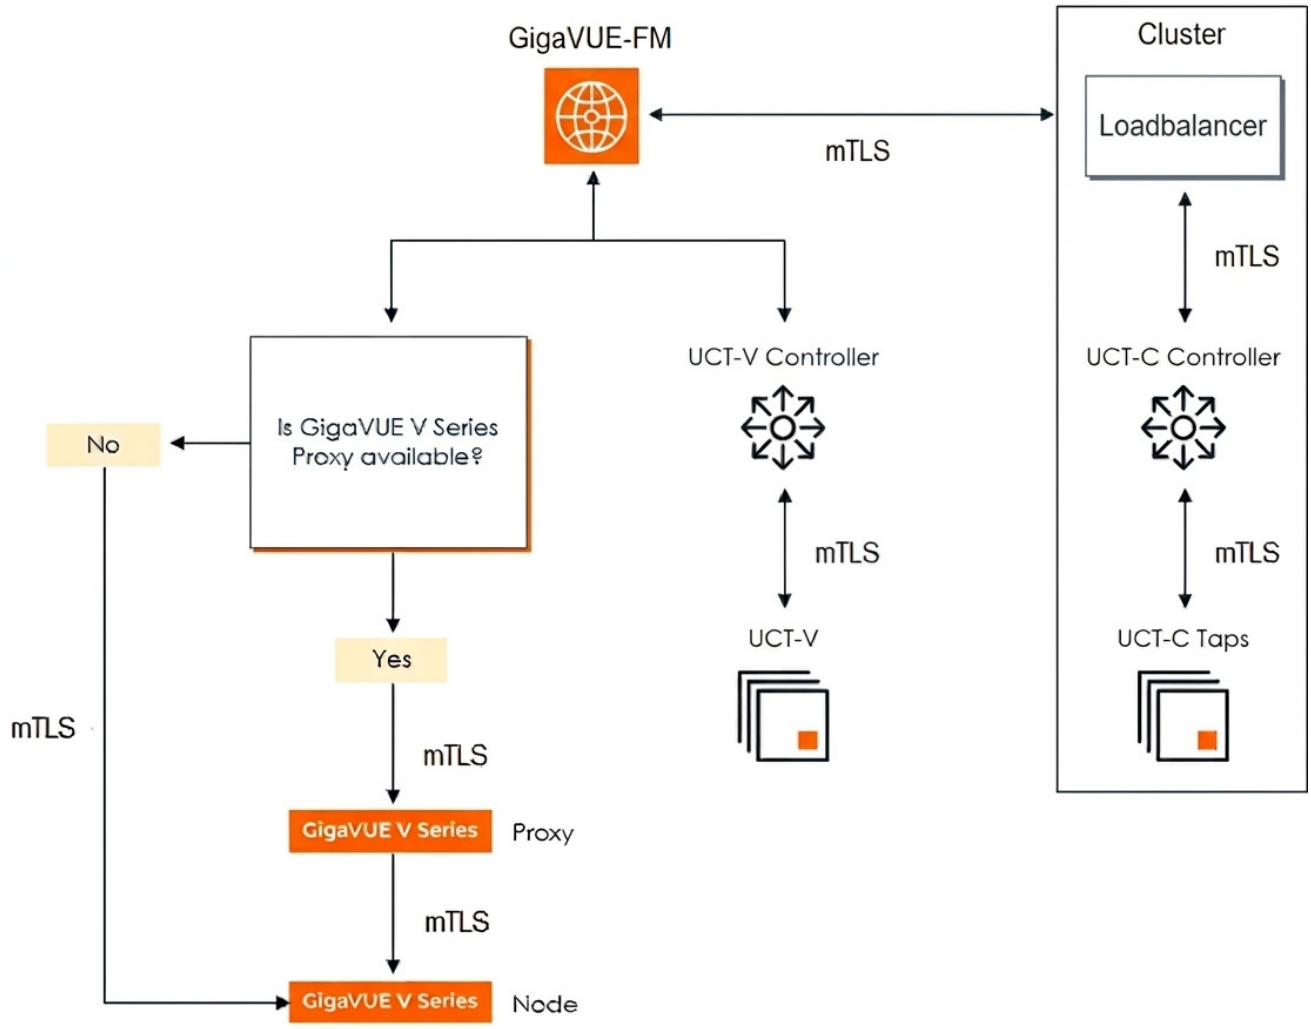
For details, refer to the following section:

- Secure Communication between GigaVUE Fabric Components
- Precryption™
- Secure Tunnels
- Prefiltering
- Customer Orchestrated Source - Use Case

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 theUCT-V Controller, establishing an mTLS connection with UCT-V.

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

GigaVUE-FM acts as the PKI

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

Bring Your Own CA

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

For more details on how to integrate your PKI infrastructure with GigaVUE-FM, refer to Integrate Private CA

Secure Communication in FMHA Mode

In FMHA (Fabric Manager High Availability) mode:

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

Supported Platforms

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

Supported Components

- GigaVUE V Series Node
- GigaVUE V Series Proxy

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

Rules and Notes

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

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

Precryption™

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

Gigamon Precryption™ technology[1] provides you clear-text visibility into encrypted network traffic without the need for traditional decryption. It works across virtual, cloud, and container environments, helping you get the full security stack without added complexity

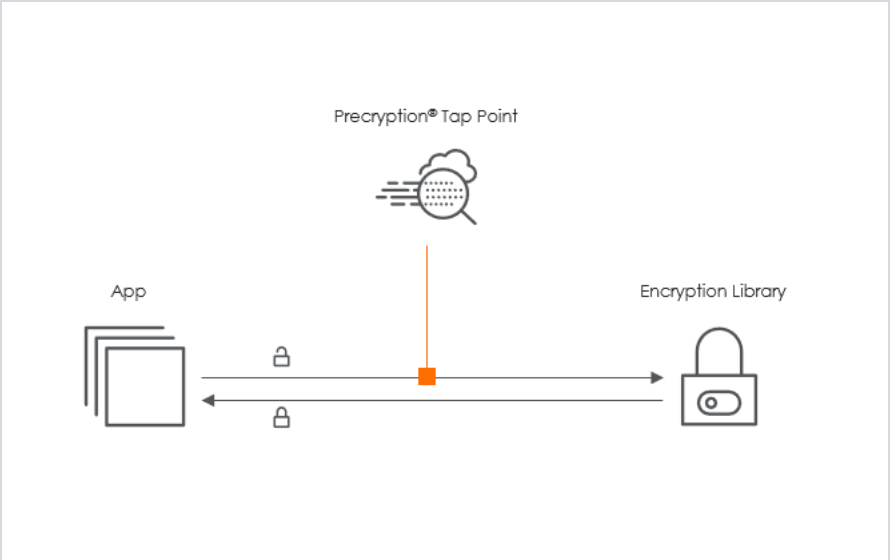
In this section:

- How Gigamon Precryption Technology Works
- Why Gigamon Precryption
- Key Features
- Key Benefits
- Precryption Technology on Single Node
- Precryption Technology on Multi-Node
- Supported Platforms
- Prerequisites

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

How Gigamon Precryption Technology Works

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



Key Highlights

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

Why Gigamon Precryption

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

Precryption helps you:

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

Key Features

The following are the key features of this technology:

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

Key Benefits

The following are the key benefits of this technology:

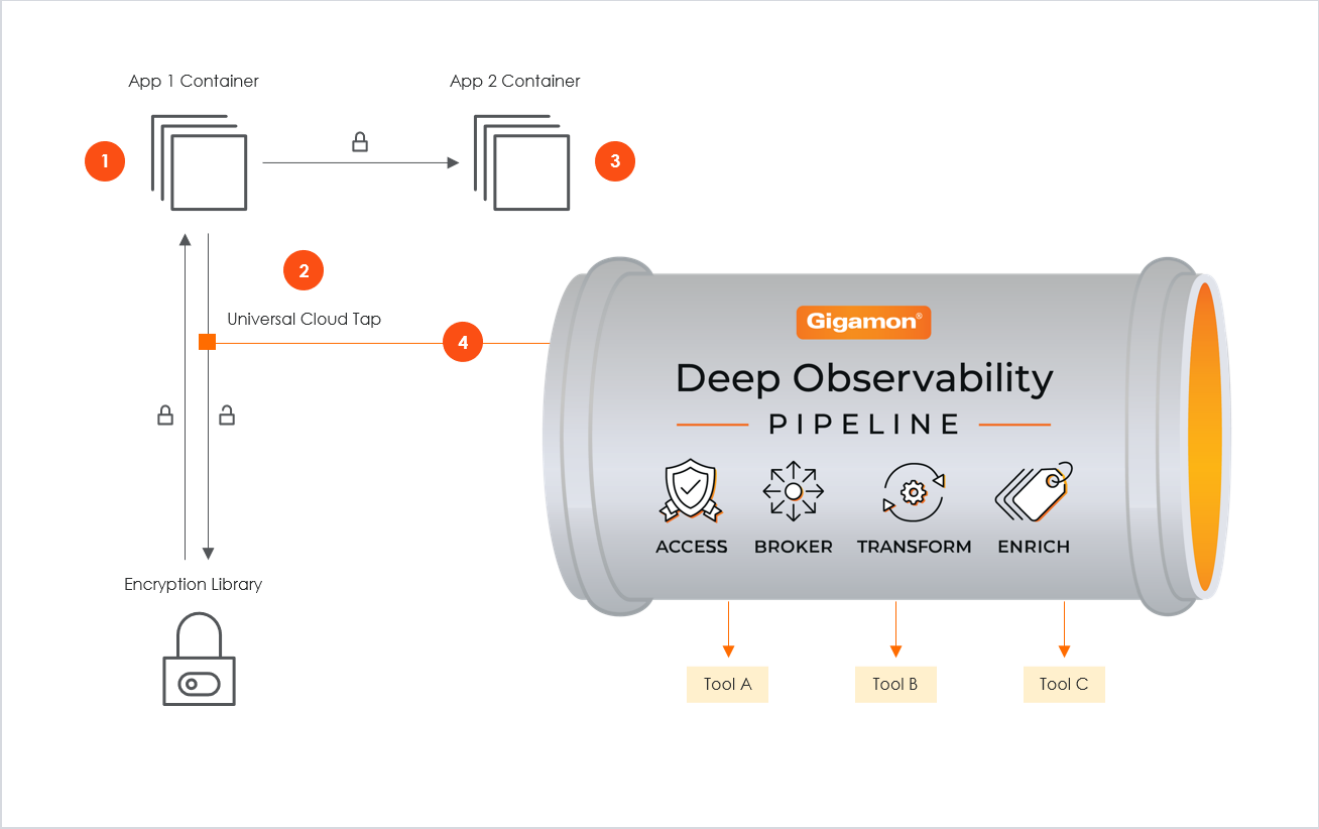
- Eliminates blind spots for encrypted East-West (lateral) and North-South communications, including traffic that may not cross firewalls.
- Monitors application communications with an independent approach that enhances development team velocity.
- Extends security tools' visibility to all communications,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

How Gigamon Precryption Technology Works

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

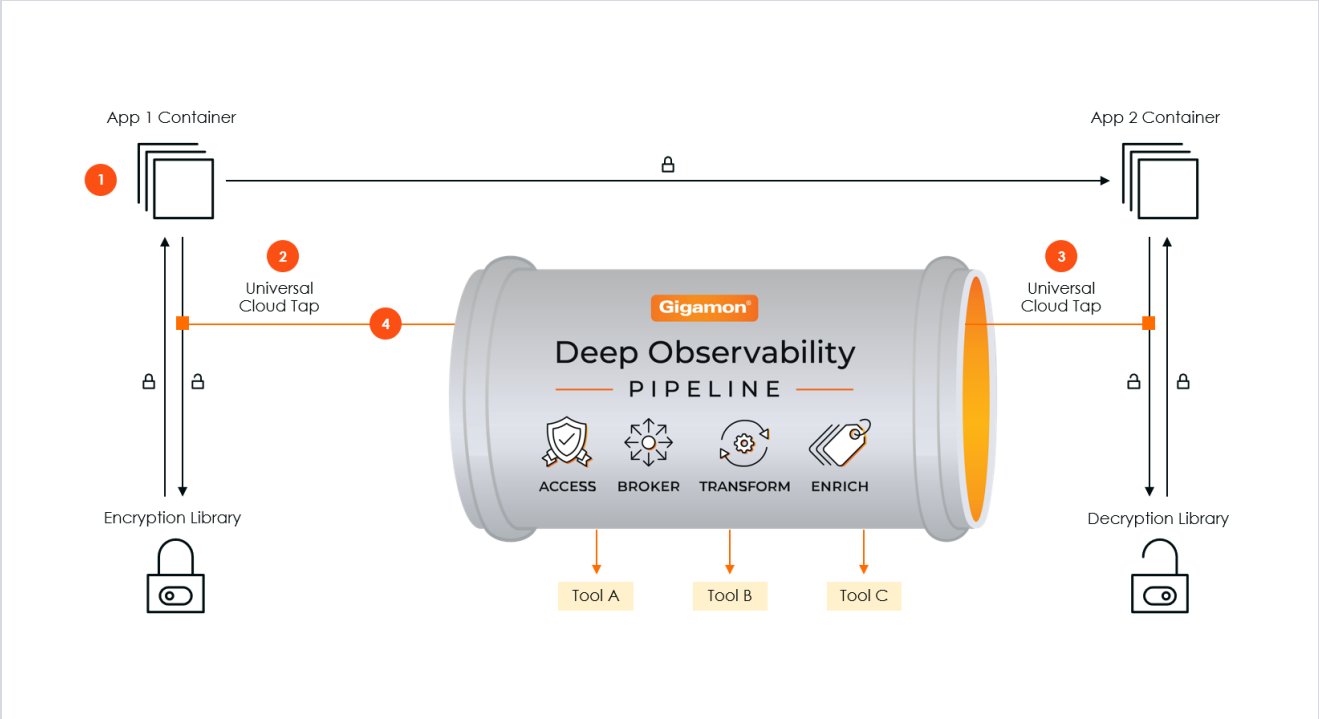
- Precryption Technology on Single Node
- Precryption Technology on Multi-Node

Precryption Technology on Single Node



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

Precryption Technology on Multi-Node



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

Supported Platforms

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

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

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

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

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

Prerequisites

Points to Note

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

License Prerequisite

- Precryption™ requires a SecureVUE Plus license.

Supported Kernel Version

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

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

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

For more details, refer to Gigamon TV.

Note

- See the Configure Precryption in UCT-V section for details on how to enable Precryption™ in VM environments.
- See the Configure Precryption in UCT-C section for details on how to enable Precryption™ in container 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.

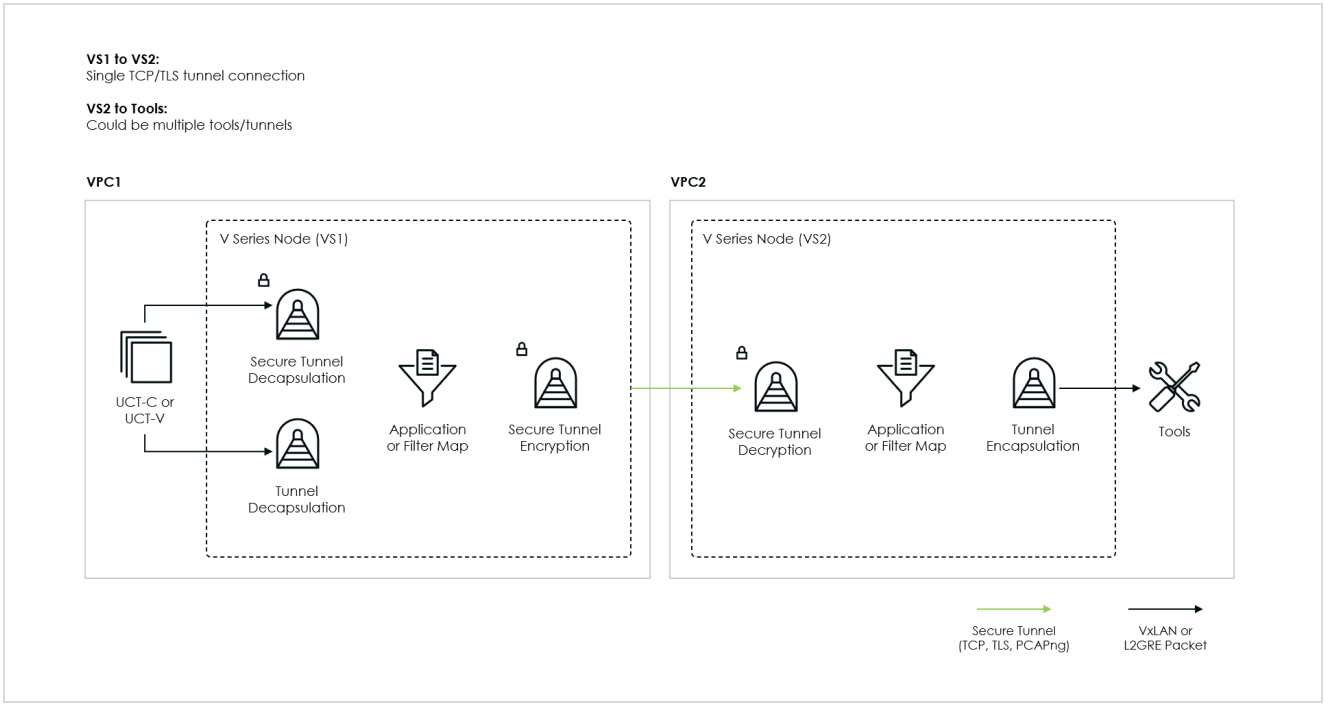
When sending traffic between two V Series Nodes, the source node captures and encapsulates the packets in PCAPng format. It then sends them to the destination V Series Node that decapsulates and processes the traffic based on your configuration.

How Does it Work!

- The source V Series forwards packets to the destination V series for further analysis.
- The destinatin 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 (OpenStack).

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 and the policy template can be applied 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.

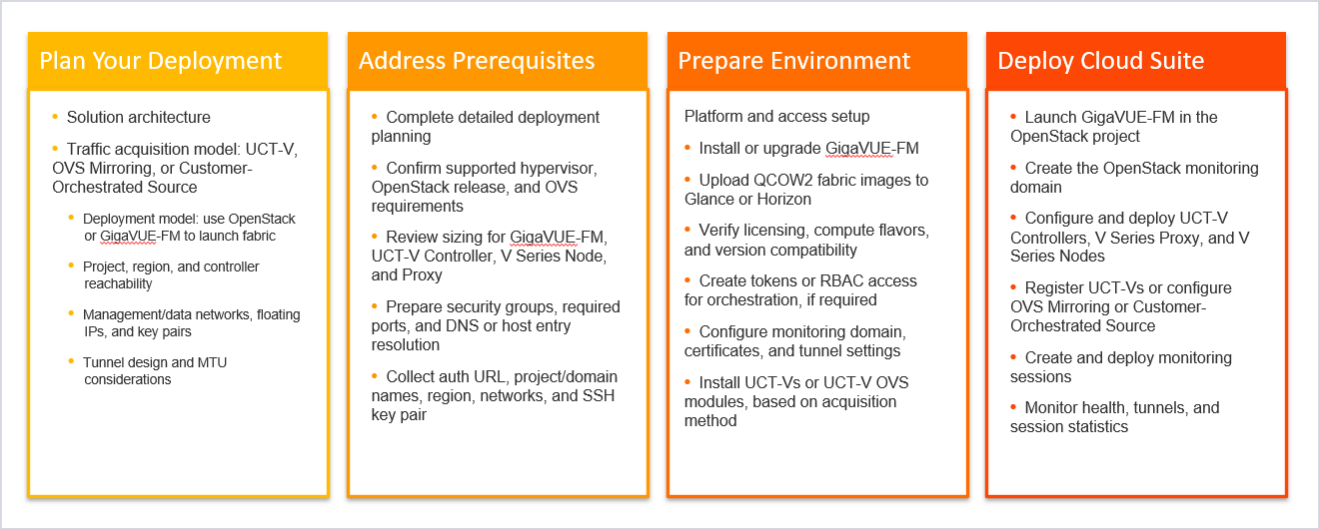
Some of the points that must be remembered for prefiltering in Next Generation UCT-Vs are:

- 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 more information on configuring a prefilter, refer to Create Prefiltering Policy Template

Deployment Overview

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



Deployment Planning

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

Refer to the following sections for more detailed information:

- Identify the Deployment Method
- Define the Deployment Location
- Identify the Traffic Acquisition Method
- Identify the Deployment Model
- Identify the Networks and Subnets Involved
- Identify the Required Tokens and Credentials
- Create the Keypairs
- Identify the UCT-V Deployment Details

Identify the Deployment Method

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

- GigaVUE-FM Orchestration
- Third Party Orchestration

In either scenario, you need to first launch GigaVUE-FM and apply the required license. You can deploy GigaVUE-FM as an OpenStack instance from the qcow2 image, or use an existing GigaVUE-FM instance to manage the deployment.

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.

Third Party Orchestration

In this method, you use your own OpenStack orchestration system or an external process to deploy the GigaVUE fabric components. This is typically a combination of Ansible and Terraform products, OpenStack Heat templates, or similar products. All of the GigaVUE fabric component images are accessible in OpenStack to support this method.

Define the Deployment Location

You can deploy GigaVUE-FM inside or outside of OpenStack, such as an existing GigaVUE-FM instance on-premises. Deploying a new instance of GigaVUE-FM in OpenStack eliminates upgrade requirements for an existing on-premises GigaVUE-FM instance and keeps management traffic local to the OpenStack environment. As an alternative, you can use an external GigaVUE-FM instance to manage the OpenStack deployment.

Identify the Traffic Acquisition Method

GigaVUE Cloud Suite for OpenStack supports traffic acquisition using the Universal Cloud Tap (UCT-V) and customer orchestrated sources. When creating a Monitoring Domain in GigaVUE-FM, choose the Traffic Acquisition Method that matches your environment:

- **UCT-V** – Select this method to deploy UCT-V Controller and UCT-V for agent-based traffic acquisition from VMs.
- **Customer Orchestrated Source** – Select this method when traffic is delivered to the GigaVUE V Series Node through tunnels from your own sources.

You can only choose one traffic acquisition method per Monitoring Domain. To use different traffic acquisition methods, create multiple Monitoring Domains.

Identify the Deployment Model

GigaVUE Cloud Suite for OpenStack provides flexibility in the deployment location of fabric components. Plan the deployment around these supported components:

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

For UCT-V traffic acquisition, deploy at least one UCT-V Controller to act as a communication proxy between GigaVUE-FM and the UCT-Vs. The UCT-V Controller connects to each UCT-V and maintains a single connection to GigaVUE-FM. You can configure more than one UCT-V Controller for a UCT-V so that registration can continue through another active controller if one controller goes down.

Identify the Networks and Subnets Involved

Based on the chosen traffic acquisition method, identify the networks and subnets involved in the deployment. A GigaVUE V Series Node in OpenStack must have at least two attached NICs: one management NIC and one data NIC.

If you use UCT-V, the UCT-V is registered through the UCT-V Controller and communicates through port 8891. The Traffic Acquisition Tunnel MTU is set to 1500 by default for OpenStack third-party orchestration deployments. If needed, you can edit the Monitoring Domain later to change the MTU value.

Identify the Required Tokens and Credentials

Create the required tokens in the GigaVUE-FM User Management page before registering the OpenStack fabric components.

Create the Keypairs

GigaVUE fabric components support administrative SSH access via SSH key pairs. When deploying these components, provide the name of an OpenStack stored SSH key to provision the public key on each component. You can use the associated private key for administrative access to these components.

Identify the UCT-V Deployment Details

The UCT-V is provided as a Debian (.deb) and RedHat Package Manager (.rpm) package file. The UCT-V deployment configuration requires knowledge of the OS network interface naming conventions, such as eth0 or ens5. Before deploying UCT-V:

- Identify the VMs where UCT-V deployment is required.
- Familiarize yourself with OS network interface naming conventions.
- Identify the deployment method.

Deployment Prerequisites

This section describes the requirements and prerequisites for configuring the GigaVUE Cloud Suite for OpenStack. Refer to the following section for details.

- Supported Hypervisor for OpenStack
- Supported Compute Requirements for OpenStack
- Network Requirements
- Licensing GigaVUE Cloud Suite
- Virtual Network Interface Cards (vNICs)
- Security Group for OpenStack
- Key Pairs
- Prerequisites for OVS Mirroring
- GigaVUE-FM Version Compatibility
- Default Login Credentials for GigaVUE Fabric Components
- Configure Tokens
- Configure Role-Based Access for Third Party Orchestration

Licensing GigaVUE Cloud Suite

You can license the GigaVUE Cloud Suite using one of the following method:

- Purchase GigaVUE Cloud Suite using CPPO
- Volume Based License (VBL)

For purchasing licenses with the Volume-Based License (VBL) option, contact our Sales. Refer to Contact Sales. For instructions on how to generate and apply license refer to the *GigaVUE Administration Guide* and the GigaVUE Licensing Guide.

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 SKUs [1]. 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	√	√	√

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

Feature	CoreVUE	NetVUE	SecureVUE Plus
Deduplication	-	√	√
NetFlow Generation	-	√	√
Adaptive Packet Filtering (APF)	-	-	√
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

Install GigaVUE-FM on OpenStack

You can install and upgrade the GigaVUE-FM fabric manager on cloud or on-premises. You can also upgrade GigaVUE-FM deployed in OpenStack environment.

- **Cloud**—To install GigaVUE-FM inside your OpenStack environment, you can simply launch the GigaVUE-FM instance in your Project.

  > **NOTE:** You cannot upgrade your 5.7.00 or lower versions of the GigaVUE-FM instance deployed in OpenStack environment to GigaVUE-FM 5.8.00 or higher versions. You must perform a fresh installation of GigaVUE-FM 5.8.00 or higher versions.

- **On-premises**—To install and upgrade GigaVUE-FM in your enterprise data center, refer to *GigaVUE-FM Installation and Upgrade Guide* available in the Gigamon Documentation Library.

Refer to the following sections for more details:

- Upload Fabric Images
- Install GigaVUE-FM on OpenStack
- Initial GigaVUE-FM Configuration

Supported Hypervisor for OpenStack

The following table lists the hypervisor with the supported versions for UCT-V.

Hypervisor	Supported Versions
KVM	**UCT-V**—Pike through Stein releases **OVS Mirroring**—Rocky and above, RHOSP 16.2 and 17.1, Kolla-ansible

The OpenStack and KVM configurations listed in this table are the officially validated and supported configurations for GigaVUE Cloud Suite. Validation applies only to the specific OpenStack/KVM combination listed and does not extend to, or guarantee compatibility with, third-party OpenStack distributions, management platforms, or orchestration layers, including Platform9, Canonical OpenStack, Huawei FusionSphere OpenStack, and Mirantis.

Configurations that differ from the listed combinations have not been officially validated by Gigamon and may be supported on a best-effort basis. Compatibility and full feature support are not guaranteed for non-validated configurations.

Supported Compute Requirements for OpenStack

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

Compute Instances	vCPU	Memory	Disk Space	Description
UCT-V	2 vCPU	4GB	N/A	Available as rpm or Debian package. Instances can have a single vNIC or dual vNICs configured for monitoring the traffic.
UCT-V Controller	2 vCPU	4GB	10GB	Based on the number of agents being monitored, multiple controllers will be required to scale out horizontally.
GigaVUE V Series Node	4 vCPU	4GB	10GB	NIC 1: Monitored Network IP; Can be used as Tunnel IP NIC 2: Tunnel IP (optional) NIC 3: Management IP
GigaVUE V Series Proxy	1 vCPU	4GB	10GB	Based on the number of GigaVUE V Series nodes being monitored, multiple controllers will be required to scale out horizontally.
GigaVUE-FM	4 vCPU	16GB	80GB	GigaVUE-FM must be able to access the controller instance for relaying the commands. Use a flavor with a root disk of minimum 40GB and an ephemeral disk of minimum 41GB.

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

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

Network Requirements

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

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

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

Virtual Network Interface Cards (vNICs)

OpenStack Cloud Instances with UCT-V can be configured with one or more vNICs.

- **Single vNIC**—If there is only one interface configured on the instance with the UCT-V, the UCT-V sends the mirrored traffic out using the same interface.
- **Multiple vNICs**—If there are two or more interfaces configured on the instance with the UCT-V, the UCT-V monitors any number of interfaces. It provides an option to send the mirrored traffic out using any one of the interfaces or using a separate, non-monitored interface. When multiple interfaces are added to the controller, floating IP is used to make the first interface as management interface.

NOTE: vNICs are only applicable if the UCT-V is installed on the instances being monitored. It is not applicable for OVS Mirroring or OVS Mirroring +DPDK.

Security Group for OpenStack

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

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
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	Allows GigaVUE-FM to

				Node IP	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.
Outbound	TCP	443	9900	OVS Controller	Controller management communication.

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, or OVS components) 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.

The following table list the Network Firewall or Security Group requirements when using OVS Mirroring.

Direction	Protocol	Port	CIDR	Purpose
UCT-V OVS Controller				
Inbound	TCP	9900	GigaVUE-FM IP	Allows GigaVUE-FM to communicate with UCT-V OVS Controllers
UCT-V OVS Module				
Inbound	TCP	9901	UCT-V OVS Controller IP	Allows UCT-V OVS Controllers to communicate with UCT-V OVS Modules

NOTE: The Security Group Rules table lists only the ingress rules. Make sure the egress ports are open for communication. Along with the ports listed in the Security Group Rules table, make sure the suitable ports required to communicate with Service Endpoints such as Identity, Compute, and Cloud Metadata are also open.

Key Pairs

A key pair consists of a public key and a private key. You must create a key pair and select the name of this key pair when you launch the UCT-V Controllers, GigaVUE V Series nodes, and GigaVUE V Series Controllers from GigaVUE-FM. Then, you must provide the private key to connect to these instances. For information about creating a key pair, refer to OpenStack documentation.

Prerequisites for OVS Mirroring

This section is only applicable if you wish to use OVS Mirroring as your traffic acquisition method. The following items are required to deploy a UCT-V OVS module:

- An existing OpenStack cloud environment should be available with admin project and login credentials to create a monitoring domain.
- A user with OVS access is required to enable OVS-Mirror. The user can be an admin or can be a user with a custom role that has the permissions and the ability to list projects.
- A working GigaVUE-FM with latest build.

OVS-Mirror Requirements

Operating System	OVS Versions	OpenStack Version	Operating System Version
Ubuntu	2.9.8, 2.13.8	4.0.2, 5.5.1	20.04
Red Hat OpenStack platform	2.15.5	16.2, 17.1	Red Hat Enterprise Linux 8.4, Red Hat Enterprise Linux 9.2 (Plow)

OpenStack Cloud Environment Requirements

- ML2 mechanism driver: Open vSwitch.
- You must have the following role privileges as shown in the table for the respective files to enable OVS mirroring:

File	Command
/etc/nova/policy.json	"os_compute_api:os-hypervisors": "role:gigamon", "os_compute_api:servers:detail:get_all_tenants": "role:gigamon", "os_compute_api:servers:index:get_all_tenants": "role:gigamon", "os_compute_api:servers:allow_all_filters":"role:gigamon", "os_compute_api:os-extended-server-attributes":"role:gigamon"
/etc/keystone/policy.json	"identity:list_projects": "role:admin or role:gigamon", "identity:list_user_projects": "role:admin or role:gigamon or rule:owner", "identity:list_users": "role:admin or role:gigamon"
/etc/neutron/policy.json	"context_is_advsvc": "role:advsvc or role:gigamon", "get_subnet": "rule:admin_or_owner or rule:shared or role:gigamon", "get_network": "rule:admin_or_owner or rule:shared or rule:external or rule:context_is_advsvc", "update_floatingip": "rule:admin_or_owner or role:gigamon", "get_floatingip": "rule:admin_or_owner or role:gigamon", "get_security_groups": "rule:admin_or_owner or role:gigamon", "get_security_group": "rule:admin_or_owner or or role:gigamon", "get_port": "rule:context_is_advsvc or rule:admin_owner_or_network_owner", "get_port:binding:vif_details":"rule:admin_only or rule:context_is_gigamon"

- Here are the APIs and commands required for OVS mirroring

OpenStack CLI command	Supported API/Action	Description
openstack hypervisor list	GET /os-hypervisors	Should list all hypervisors in the domain.
openstack server list --all -- host	GET /servers	Should list all the servers on a specified host

OpenStack CLI command	Supported API/Action	Description
<hostname>		
openstack server list-all	GET /servers	Should list servers of all projects in the domain.
openstack project list	GET /v3/projects	Should list all projects in the domain.
openstack project list - user <user with custom role>	GET /v3/projects	Should list all projects that a specified user (user specified in GigaVUE-FM config) is associated with
openstack user show <userName>	GET /v3/users	Should list all users by username
openstack subnet list	GET /subnets	Should list all subnets for all projects in the domain.
openstack network list	GET /network	Should list all networks for all projects in the domain.
openstack floating ip list	GET /floatingips	Should list all floating ips for all projects in the domain.
openstack floating ip set-port <portid> <floating ip>	PUT /floatingips/{floatingip_ID}	Used to attach floating ip to fabric nodes.
openstack security group list	GET /security-groups	Should list security groups for all projects in the domain
openstack security group show <security group id>	GET /security-groups/{security_group_id}	Should list details of specified security group
openstack port list	GET /ports	Should list ports for all projects in the domain
openstack port show <portID>	GET /ports/{portID}	Should list port details including bridge name.
openstack server create	POST /servers	Launch fabric nodes
openstack server <action> <serverName>	POST /servers/{server_id}/action	stop/start/reboot fabric nodes
openstack server delete <serverName>	DELETE /servers/{serverID}	Delete fabric nodes
openstack server set	PUT /servers/{serverID}/metadata	Update visibility node metadata
openstack flavor list	GET /flavors	Get list of flavors
openstack availability zone list	GET /os-availability-zone	Get list of availability zones
openstack keypair list	GET /os-keypairs	Get list of keypairs

If the OpenStack CLI command **`openstack hypervisor list`** does not return a reachable IP for the hypervisors that are being monitored, you must manually enter a reachable

IP for each hypervisor in OpenStack CLI using project properties. For each hypervisor you will need to add a key value pair property in the following format:

- key: value
- key: must be in the form gigamon-hv-<hypervisorID>
- value: reachable IP for hypervisor

For example: `openstack project set --property gigamon-hv-1=1.2.3.4 project-name`

GigaVUE-FM Version Compatibility

GigaVUE-FM version 6.14.02 supports the latest version (6.14.02) 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.

Default Login Credentials for GigaVUE Fabric Components

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 and GigaVUE V Series proxy	You can login to the GigaVUE V Series Node and GigaVUE V Series proxy by using ssh. The default username and password is: Username: gigamon Password: Gigamon123!
UCT-V Controllers	You can login to the UCT-V Controller by using ssh. The default username and password is: Username: gigamon Password: Gigamon123!

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.

Install GigaVUE-FM on OpenStack

You can install and upgrade the GigaVUE-FM fabric manager on cloud or on-premises. You can also upgrade GigaVUE-FM deployed in OpenStack environment.

- **Cloud**—To install GigaVUE-FM inside your OpenStack environment, you can simply launch the GigaVUE-FM instance in your Project.

  > **NOTE:** You cannot upgrade your 5.7.00 or lower versions of the GigaVUE-FM instance deployed in OpenStack environment to GigaVUE-FM 5.8.00 or higher versions. You must perform a fresh installation of GigaVUE-FM 5.8.00 or higher versions.

- **On-premises**—To install and upgrade GigaVUE-FM in your enterprise data center, refer to *GigaVUE-FM Installation and Upgrade Guide* available in the Gigamon Documentation Library.

Refer to the following sections for more details:

- Upload Fabric Images
- Install GigaVUE-FM on OpenStack
- Initial GigaVUE-FM Configuration

Upload Fabric Images

First, you must fetch the images from Gigamon Customer Portal using FTP, SCP, or other desired methods and copy it to your cloud controller. After fetching the images, you must source the credentials file and then upload the qcow2 images to Glance.

For example, you can source the credentials file with admin credentials using the following command:

$ source admin_openrc.sh

To upload the qcow2 images to Glance, use the following commands:

GigaVUE-FM - **glance image-create --disk-format qcow2 --visibility public --container-format bare --progress --name fm-6.14.xx-uefiSec-boot --file gfm-6.14.xx.qcow2**

GigaVUE-FM with UEFI Boot - **glance image-create --disk-format qcow2 --visibility public --container-format bare --property hw_firmware_type=uefi --progress --name fm-6.14.xx-uefiSec-boot --file gfm-6.14.xx.qcow2**

GigaVUE-FM with UEFI secure Boot - **glance image-create --disk-format qcow2 --visibility public --container-format bare --property hw_firmware_type=uefi --property os_secure_boot=required --progress --name fm-6.14.xx-uefiSec-boot --file gfm-6.14.xx.qcow2**

> **NOTE:** UEFI secure boot is supported on RHSOP 17.01 version and above. UEFI secure boot is not supported on Wallaby.

Refer to the below table for the OpenStack flavors and the versions that support UEFI and UEFI secure boot:

OpenStack Flavor	Version	UEFI	UEFI Secure boot
kolla Ansible	Wallaby	Yes	No
Rhosp	17.01	Yes	Yes
Rhosp	16.2	No	No

UCT-V OVS Controller - **glance image-create --disk-format qcow2 --visibility public --container-format bare --progress --name gigamon-gigavue-uctv-ovs-cntlr-6.14.xx --file gigamon-gigavue-uctv-ovs-cntlr-6.14.xx.qcow2**

UCT-V Controller - **glance image-create --disk-format qcow2 --visibility public --container-format bare --progress --name gigamon-gigavue-uctv-cntlr-6.14.xx --file gigamon-gigavue-uctv-cntlr-6.14.xx.qcow2**

GigaVUE V Series Node - **glance image-create --disk-format qcow2 --visibility public --container-format bare --progress --name gigamon-gigavue-vseries-node-6.14.xx --file gigamon-gigavue-vseries-node-6.14.xx.qcow2**

While uploading images to OpenStack, the names of the image files should be of the following format:

- gigamon-gigavue-vseries-node-6.14.xx
- gigamon-gigavue-vseries-proxy-6.14.xx
- gigamon-gigavue-uctv-cntlr-6.14.xx
- gigamon-gigavue-uctv-ovs-cntlr-6.14.xx

> **NOTE:** Always use '-' after the build number when providing designation details. For example, "**gigamon-gigavue-uctv-ovs-cntlr-buildNumber-**". Failure to follow this pattern will result in controller's deployment failure.

Upload Images using OpenStack Horizon Dashboard:

To upload the qcow2 using OpenStack Horizon Dashboard:

1. Log into Horizon.
2. From the Horizon GUI, select the appropriate project, and select **Compute > Images**.
3. Click **Create Image**. In the dialog box, provide the following details:
   a. Name - Enter a descriptive name for the image

b. Image Source - Select File and browse to the location of the .qcow2 image file

c. Format - Select QCOW2

d. Visibility - Set as required (e.g., Public or Private)

4. Click **Create Image**. The image will be uploaded and listed in the **Images** page.

Refer to Install GigaVUE-FM on OpenStack to launch GigaVUE-FM instance using Horizon Dashboard.

Install GigaVUE-FM on OpenStack

To launch the GigaVUE-FM instance inside the cloud:

1. Log into Horizon.
2. From the Horizon GUI, select the appropriate project, and select **Compute > Images**. The list of existing images is displayed.
3. Select the GigaVUE-FM image and click **Launch**. The Launch Instance dialog box is displayed.
4. In the **Details** tab, enter the following information and Click **Next**.

Parameter	Attribute
Instance Name	Initial hostname for the instance
Availability Zone	Availability zone where the image will be deployed.
Count	Number of instances to be launched

5. In the **Source** tab, verify that the selected GigaVUE-FM image is displayed under **Allocated** section and click **Next**.
6. In the **Flavor** tab, select a flavor complying the Supported Compute Requirements for OpenStack and then move the flavor from the **Available** section to the **Allocated** section. The selected GigaVUE-FM flavor is displayed under Allocated and click **Next**.
7. In the **Networks** tab, select the specific network for the GigaVUE-FM instance from the **Available** section and then move the Network to the **Allocated** section. The selected network is displayed under Allocated and Click **Next**.
8. In the **Network Ports** tab, click **Next** again.
9. In the **Security Groups** tab, select the appropriate security group for the GigaVUE-FM instance from the **Available** section and then move the Security Group to the **Allocated** section. For information about the security groups, refer to Security Group for OpenStack . The selected security group is displayed under Allocated. Click **Next**.
10. In the **Key Pair** tab, select the existing key pair from the **Available** section and then move the Key Pair to the **Allocated** section. or create a new key pair. For information about the key pairs, refer to Key Pairs. The selected key pair is displayed under Allocated. Click **Next**.
11. Click **Launch Instance**. The GigaVUE-FM instance takes few minutes to fully initialize.
12. From the Horizon GUI, navigate to **Compute > Instances**. You can view the launched instance displayed in the **Instances** page. During the initial boot-up sequence, click **Associate Floating IP**. The **Manage Floating IP Associations** dialog box appears.

13. In the Manage Floating IP Associations dialog box, enter the following information and click **Associate**.

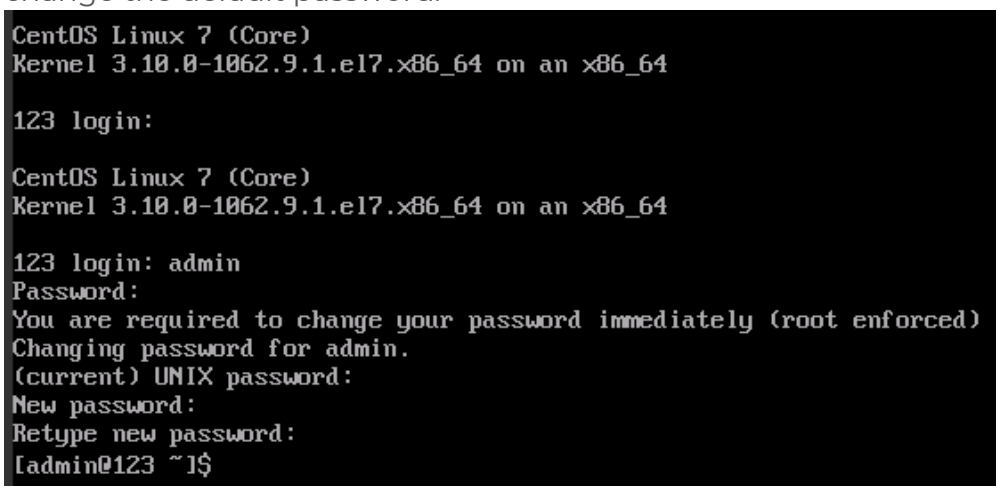
Parameter	Attribute
IP Address	Floating IP address of the instance
Port to be associated	Port for the GigaVUE-FM instance

The Floating IP is then displayed in the **IP Address** column of the corresponding Instance.

Initial GigaVUE-FM Configuration

After you have deployed a new GigaVUE-FM instance, you need to perform an initial configuration before you can start using GigaVUE-FM. This is a one-time activity that must be performed for each GigaVUE-FM instance deployed.

1. From the Horizon GUI, navigate to **Compute > Instances**.
2. In the Instances page, click the GigaVUE-FM instance name. The GigaVUE-FM instance **Overview** tab is displayed by default.
3. Click the **Console** tab and the **Instance Console** appears.
4. Log in as admin with password as admin123A!! and then the console prompts you to change the default password.



> **NOTE:** You can also choose to perform the IP Networking and NTP configurations by running the **fmctl set ip** command after you power on the GigaVUE-FM instance

5. To access GigaVUE-FM GUI, enter **wget -q -O - http://169.254.169.254/latest/meta-data/instance-id** command in the Instance Console and retrieve the instance ID in the format of **i-000000##** which is the default password for the admin user. If GigaVUE-FM is deployed inside OpenStack, use the **Instance ID** as the password for the admin user to login to GigaVUE-FM, however if GigaVUE-FM is deployed outside OpenStack, use admin123A!! as the default admin password.

Upgrade GigaVUE Fabric Components in GigaVUE-FM for OpenStack

This chapter describes how to upgradeGigaVUE V Series Proxy and GigaVUE V Series Nodes. For more detailed information about UCT-V Controller, GigaVUE V Series Proxy and Node Version refer GigaVUE-FM Version Compatibility Matrix.

> **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. Refer to Install UCT-Vfor more details.
- Create Tokens for deploying the Fabric Components using Third Party Orchestration. Refer to Configure Tokens for more details.
- Open the required ports in the cloud platform. Refer to Network Firewall Requirement for Cloud Suite] for more details.
- When using FMHA configuration, follow the steps given provided in the Configure Secure Communication between Fabric Components in FMHA section.

Refer to the following topic for more information:

- Prerequisite
- Upgrade UCT-V Controller
- Upgrade GigaVUE V Series Nodes and GigaVUE V Series Proxy

Prerequisite

Before you upgrade the GigaVUE V Series Proxy and GigaVUE V Series nodes, you must upgrade GigaVUE-FM to software version 5.13. For better performance, Gigamon recommends you to upgrade to the latest version.

Upgrade UCT-V Controller

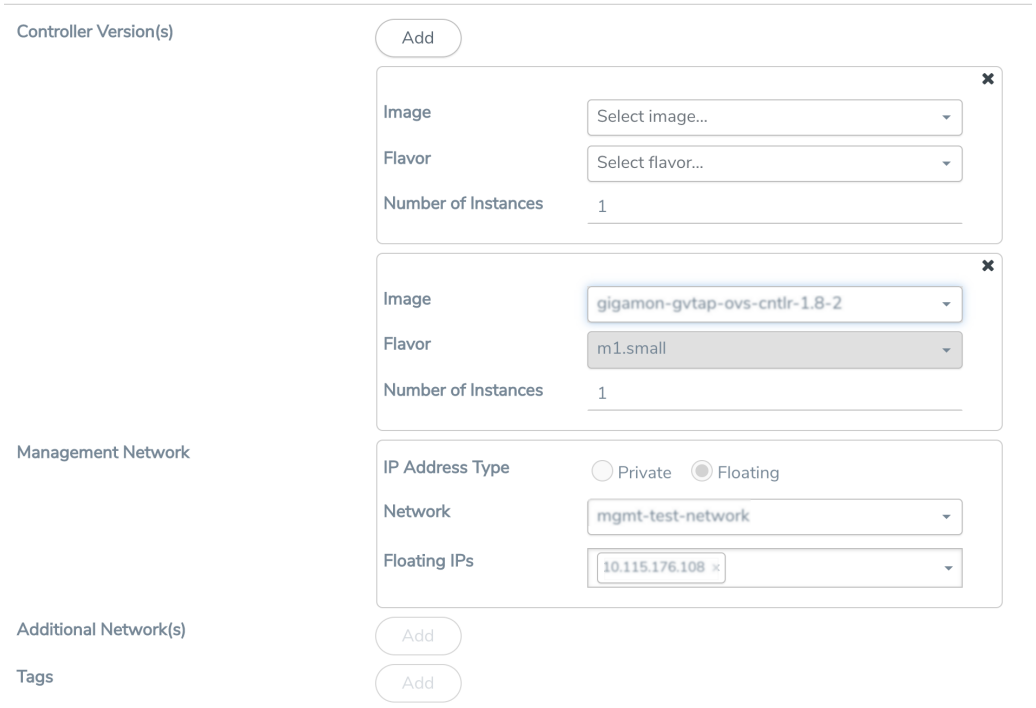
NOTE: UCT-V Controllers cannot be upgraded. Only a new version that is compatible with the UCT-V's version can be added or removed in the **OpenStack Fabric Launch Configuration** page.

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.

a. Under **Controller Versions**, click **Add**.
b. 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.
c. From the **Flavor** drop-down list, select a size for the UCT-V Controller.
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 Networks details, provided at the time of UCT-V Controller configuration.

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

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

To 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 **Network** 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 Nodes and GigaVUE V Series Proxy

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

There are multiple ways to upgrade the GigaVUE V Series Proxy and nodes. You can:

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

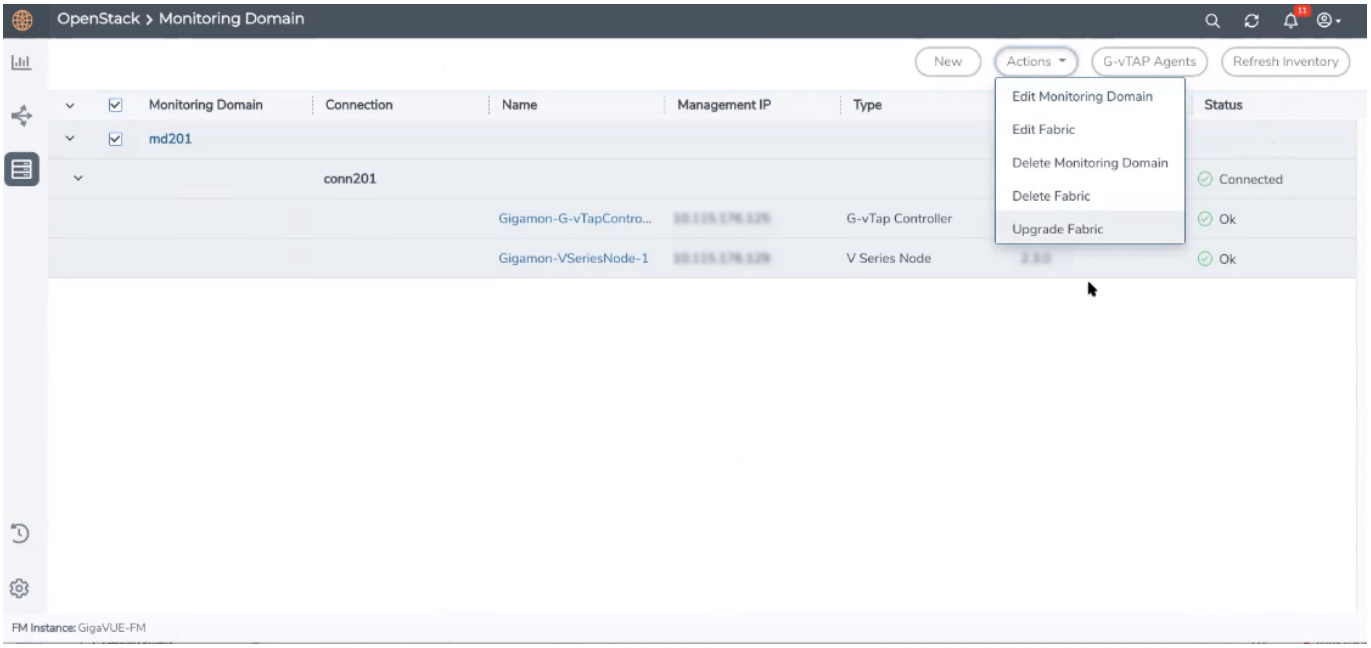
  For example, if you have 1 GigaVUE V Series Proxy and 10 GigaVUE V Series nodes in your project, you can upgrade all of them at once. First, the new version of GigaVUE V Series Proxy is launched. Next, the new version of GigaVUE V Series Nodes are launched. Then, the old version of V Series Proxy and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

  NOTES:
  - When the new version of nodes and controllers are launched, the old version still exists in the project until they are deleted. Make sure the flavor determined during the configuration can accommodate the total number of new and old fabric nodes present in the project. If the flavor cannot support so many Virtual Machines, you can choose to upgrade in multiple batches.
  - If there is an error while upgrading the complete set of controllers and nodes present in the project, the new version of the fabric is immediately deleted and the old version of the fabric is retained as before.
  - Prior to upgrading the GigaVUE V Series Proxy and Nodes, you must ensure that the required number of floating IP addresses are available in the respective subnets. Otherwise, the upgrade will fail.
- Launch and replace the nodes and controllers in multiple batches.

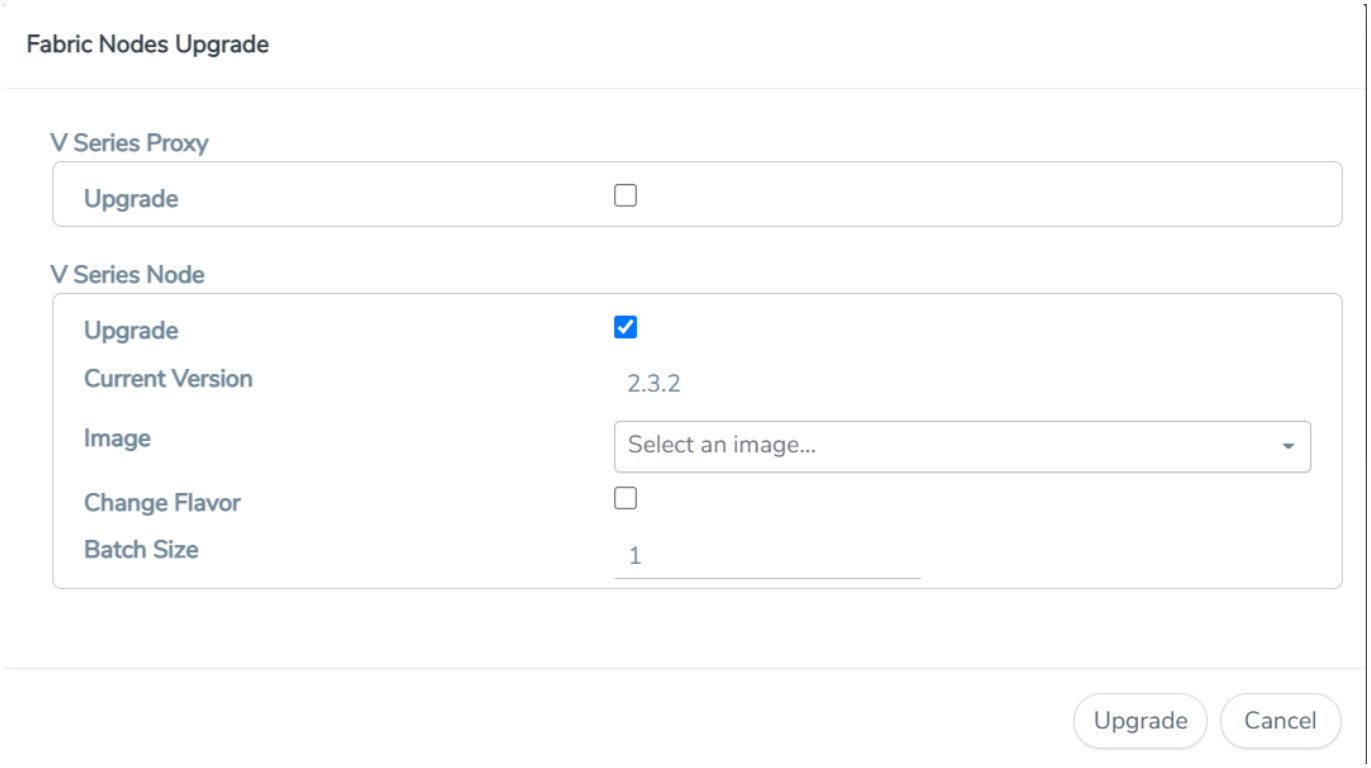
  For example, if there are 18 GigaVUE V Series Nodes to be upgraded, you can specify how many you want to upgrade per batch.

To upgrade the GigaVUE V Series Proxy and GigaVUE V Series Nodes:

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



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



4. To upgrade the GigaVUE V Series Nodes/Proxy, select the **Upgrade** checkbox.
5. From the **Image** drop-down list, select the latest version of the GigaVUE V SeriesProxy/Nodes.

6. Select the **Change Flavor** checkbox to change the flavor of the nodes/proxy, only if required.
7. To upgrade the GigaVUE V Series Nodes/Proxy, specify the batch size in the **Batch Size** box.

   For example, if there are 7 GigaVUE V Series Nodes, you can specify 7 as the batch size and upgrade all of them at once. 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. Click **Upgrade**.

   The upgrade process takes a while depending on the number of GigaVUE V SeriesProxy and Nodes upgrading in your OpenStack 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. In the V Series Proxy page, click the link under Progress to view the upgrade status.

   The monitoring session is deployed automatically.

Pre-Configuration Checklist for OpenStack

The following table provides information that you would need while launching the visibility components using GigaVUE-FM. Obtaining this information will ensure a successful and efficient deployment of the GigaVUE Cloud Suite for OpenStack.

You can log in to GigaVUE-FM and use the CLI command: **ip host <controller-hostname> <ip-address of the controller>**. (For example: `ip host os-controller1 192.168.2.3`.) Then, add the connection to the OpenStack tenant.

In order for GigaVUE-FM to make a connection to an OpenStack tenant, GigaVUE-FM must be able to resolve the hostname of the OpenStack controller, even if using an IP address in the Identity URL. For example, if GigaVUE-FM is configured to use DNS, and that controller hostname is in the DNS, this will work, and no further configuration will be needed. If not, then you must add a host entry to GigaVUE-FM.

NOTE: If you are not using DNS, you must manually enter the host entry in /etc/hosts on GigaVUE-FM for the OpenStack Controller. On using DNS you can directly enter the host entry in GigaVUE-FM.

	Required Information
☐	Authentication URL
☐	Project Name

	Required Information
☐	Floating IP
☐	Region name for the Project
☐	Domain
☐	SSH Key Pair
☐	Networks
☐	Security groups

Deploy

Deploy GigaVUE Cloud Suite for OpenStack using GigaVUE-FM Orchestration

This section describes various options to configure GigaVUE Cloud Suite for OpenStack to provide visibility for both physical and virtual traffic. You can configure GigaVUE Cloud Suite for OpenStack in the following ways, based on the traffic acquisition method and your approach to deploying fabric components. For pre-requisite details, refer to the Deployment Prerequisites

Refer to the following deployment options for more detailed information:

- Acquire Traffic using UCT-V as Traffic Acquisition Method
- Acquire Traffic using OVS Mirroring as Traffic Acquisition Method
- Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method

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

- Install GigaVUE-FM on OpenStack
- Install UCT-Vs
- Create Monitoring Domain

- Configure GigaVUE Fabric Components in GigaVUE-FM
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

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. For details, refer to Network Firewall Requirement.

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

After you deploy the UCT-V VM, you can also use the following commands to add the network firewall rule:

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

Linux UCT-V Installation Scenarios

1. **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.
2. **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.
3. **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

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**6.14.02** 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-6.14.02-amd64.deb
    $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
    ```

2. **To install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS:**
    a. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
    $ sudo rpm -i gigamon-gigavue-uctv-6.14.02-x86_64.rpm
    ```

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

```
sudo uctv-wizard
```

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

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

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

Options	Use Command	Description
		If you wish to skip the prompts to add the required firewall policy, enter your option as **y**. The console interface adds the firewall rules automatically. Enter **N** if you wish to configure manually. For details, refer to Install Linux UCT-Vs using Manual Configuration section.
uninstall	sudo uctv-wizard uninstall	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
  ```

Install Linux UCT-Vs using Manual Configuration

- Install UCT-V from Ubuntu/Debian Package
- Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

1. Download the UCT-V**6.14.02** Debian (.deb) 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-6.14.02-amd64.deb
   $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
   ```

4. Modify the file **/etc/uctv/uctv.conf** to configure and register the source and destination interfaces.
   The following examples registers eth0 as the mirror source for both ingress and egress traffic and eth1 as the destination for this traffic:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

1. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
$ sudo rpm -i gigamon-gigavue-uctv-6.14.02-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

After installing the UCT-V package, 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:

- **GigaVUE-FM Orchestration**: Perform the following steps:
  a. Log in to the UCT-V.
  b. Create a local configuration file and enter the following user data. **/etc/gigamon-cloud.conf** is the local configuration file in the Linux platform.

```
Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        remoteIP: <IP address of the UCT-V Controller>
        remotePort: 443
        token: <Token>
```

  c. Restart the UCT-V service.
  Linux platform:

```
$ sudo service uctv restart
```

  d. Verify UCT-V is running:

```
sudo service uctv status
```

For more details on how to create tokens, refer to .

Verify that the service status is Active (running). After the UCT-V VM successfully connects to GigaVUE-FM, it appears in **Inventory** > **Resources** > **UCT-V** in the GigaVUE-FM.

- **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 toGigaVUE-FM using Third Party Orchestration with the following two modes:
    1. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
    2. 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

Rules and Notes

- VXLAN is the only tunnel type supported for Windows UCT-V.
- Loopback interface is not supported on 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 workflows.

Windows UCT-V Installation Scenarios

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

Windows Network Firewall Requirements

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

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

Install Windows UCT-Vs

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

- Install Windows UCT-Vs using Installation Script
- Install Linux UCT-Vs using Manual Configuration



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

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

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

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

Install Windows UCT-Vs using Installation Script

1. Download the Windows UCT-V **`6.14.02`** 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
```

Install Windows UCT-Vs using Manual Configuration

1. Download the Windows UCT-V**6.14.02** 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:

- **GigaVUE-FM Orchestration**: Refer to the following steps:

  a. Log in to the UCT-V.

  b. Create a local configuration file and enter the following user data. **C:\ProgramData\uctv\gigamon-cloud.conf** is the local configuration file in Windows platform.

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

  c. Restart the UCT-V service.

  Windows platform: Restart from the Task Manager Service

  For more details on how to create tokens, refer to .

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

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

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

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

Create Monitoring Domain

To create a monitoring domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > OpenStack**. The Monitoring Domain page appears.
2. On the Monitoring Domain page, click **New**. The **Monitoring Domain Configuration** page appears.

3. Enter or select the appropriate information to configure Monitoring Domain for OpenStack. Refer to the following table for field-level details.

NOTE: For the URL, User Domain Name, Project Domain Name, and Region field values, refer to the RC file downloaded from your OpenStack dashboard.

Field	Description
Monitoring Domain	A name for the monitoring domain. **Note**: You can only view and delete the existing configuration for V Series node 1. You cannot create and perform any other actions on the existing configuration for GigaVUE V Series node 1 as the features are deprecated from GigaVUE-FM fabric manager.
Alias	An alias used to identify the monitoring domain.
URL	The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable. Refer to the OpenStack User Manual for more information on retrieving the authentication URL from the OpenStack. **Note**: To ensure the validity of OpenStack certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. Refer to the Trust Store section in GigaVUE Administration Guide for more detailed information.
User Domain Name	The domain name of your OpenStack authentication domain. **Note**: • If you are using a separate domain for AUTH, enter that domain name as User Domain Name. • If you are not using a separate domain, you can use the same domain for User and Project Domain Name.
Project Domain Name	The domain name of your OpenStack project.
Project Name	The name of the project used for OpenStack authentication.
Region	The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version. **keystone endpoint-list** or **openstack endpoint list** or looking at the RC file in OpenStack to view your credentials.
Username	The username used to connect to your OpenStack cloud. **Note**: If you are using OVS mirroring, you must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to OVS Mirroring Prerequisites for more information.
Password	The password of your OpenStack cloud.
Traffic	Select the type of component used to capture traffic for monitoring:

Field	Description
Acquisition Method	• **UCT-V:** If you select UCT-V as the tapping method, the traffic is acquired from the UCT-Vs installed on the VMs. You must configure the UCT-V Controller to monitor the UCT-Vs.
Projects to Monitor (Only for OVS Mirroring traffic acquisition method)	This field only appears for OVS Mirroring traffic acquisition method. • Click the **Get Project List** to view the list of projects. **Note**: The **Get Project List** button will only work if all the OpenStack credentials have been provided. Refer to OVS Mirroring Prerequisites. • Select projects that you want to monitor from the list. • You can click **Select None** to clear existing selections or **Select All** to add all available projects to the connection configuration.
Traffic Acquisition Tunnel MTU (Maximum Transmission Unit)	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 Cloud Suite V Series node. • For GRE, the default value is 1450. • For VXLAN, the default value is 1400. However, the UCT-V tunnel MTU should be 50 bytes less than the default MTU size.

4. Click **Save**. The **OpenStack Fabric Launch Configuration** page appears. Refer to Configure GigaVUE Cloud Suite for OpenStack Components for detailed information.

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.
- If GigaVUE-FM fails to connect to OpenStack, an error message is displayed specifying the cause of failure. The connection status is also displayed in Audit Logs, refer to About Audit Logs for more information.

Configure GigaVUE Fabric Components in GigaVUE-FM

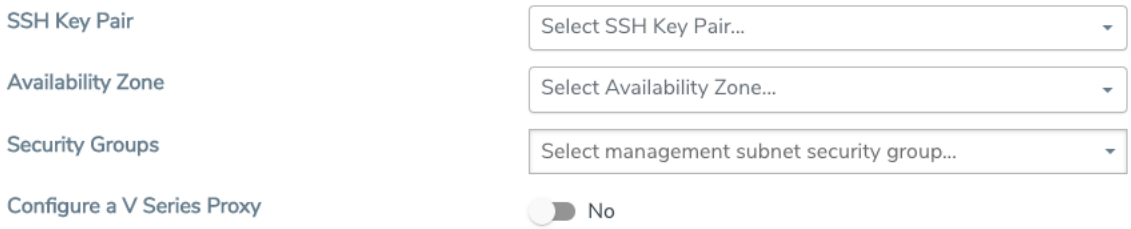
After configuring the Monitoring Domain, you will be navigated to the OpenStack Fabric Launch Configuration page. In the same **OpenStack Fabric Launch Configuration** page, you can configure the following fabric components:

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

In the **OpenStack Fabric Launch Configuration** page, enter or select the required information as described in the following table.

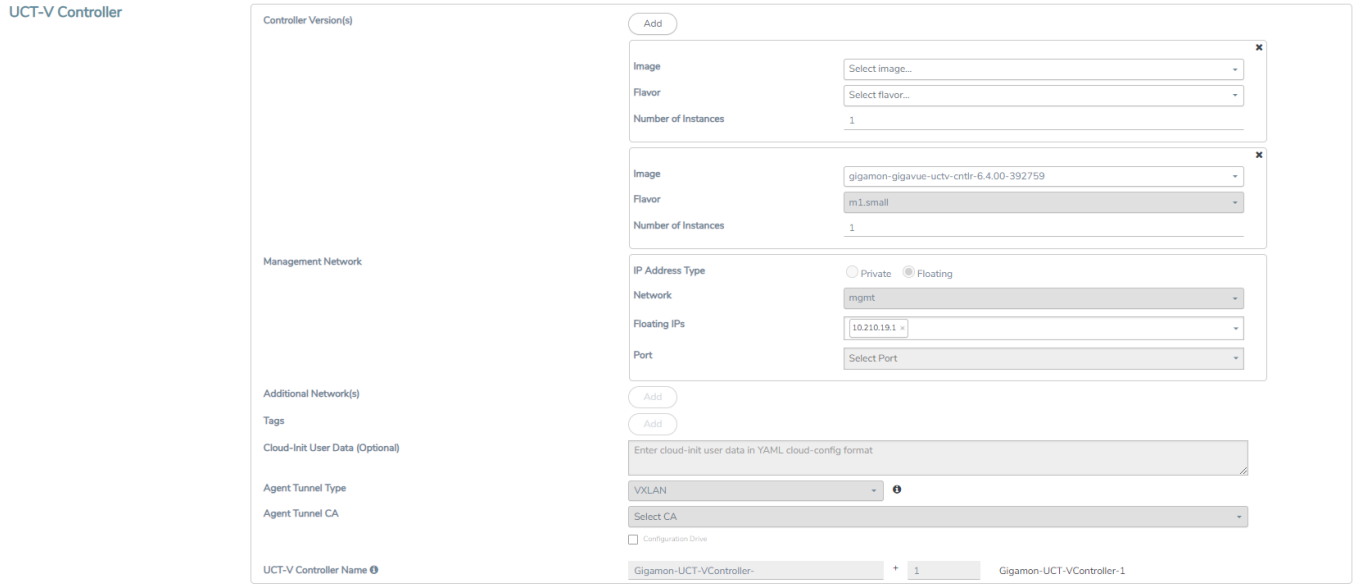
Fields	Description
SSH Key Pair	The SSH key pair for the UCT-V Controller. For more information about SSH key pair, refer to Key Pairs.
Availability Zone	The distinct locations (zones) of the OpenStack region.
Security Groups	The security group created for the UCT-V Controller. For more information, refer to Security Group for OpenStack .
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.
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.
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 more detailed information, refer to .

Select **Yes** to configure a GigaVUE V Series Proxy.



Configure UCT-V Controller

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



- Only if UCT-Vs are used for capturing traffic, then the UCT-V Controllers must be configured in the OpenStack cloud.
- A UCT-V Controller can only manage UCT-Vs that have the same version.

Enter or select the required information in the UCT-V Controller section as described in the following table.

Fields	Description
Controller Version(s)	The UCT-V Controller version that you configure must always have the same version number as the UCT-Vs deployed in the instances. For more detailed information refer GigaVUE-FM Version Compatibility Matrix. **Note**: If there is a version mismatch between the UCT-V Controllers and UCT-Vs, GigaVUE-FM cannot detect the agents in the instances. To add UCT-V Controllers: **a.** Under **Controller Versions**, click **Add**. **b.** 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. **c.** From the **Flavor** drop-down list, select a size for the UCT-V Controller. **d.** In **Number of Instances**, specify the number of UCT-V Controllers to launch. The minimum number you can specify is 1.
Management Network	This segment defines the management network that GigaVUE-FM uses to communicate with UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes. **Network** - Select the management network ID. **Ports** - Select a port, you can choose a port related to the selected management network ID. **IP Address Type**

Fields	Description
	The type of IP address GigaVUE-FM needs to communicate with UCT-V Controllers: o **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy reside inside the same project. o **Floating**—A floating IP is needed only if GigaVUE-FM is not in the same project in the cloud or is outside the cloud. GigaVUE-FM needs a floating IP to communicate with the controllers from an external network.
Additional Network(s)	(Optional) If there are UCT-Vs on networks that are not IP routable from the management network, additional networks or subnets must be specified so that the UCT-V Controller can communicate with all the UCT-Vs. Click **Add** to specify additional networks (subnets), if needed. Also, make sure that you specify a list of security groups for each additional network. **Ports**: Select a port associated with the network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. There is a specific UCT-V Controller Version for OVS Mirroring and OVS Mirroring + DPDK. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name. c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Agent Tunnel Type	The type of tunnel used for sending the traffic from UCT-Vs to GigaVUE V Series nodes. The options are GRE, VXLAN, and Secure tunnels (TLS-PCAPNG).
Agent Tunnel CA	The Certificate Authority (CA) that should be used in the UCT-V Controller for connecting the tunnel.

Configure GigaVUE V Series Proxy

The fields in the GigaVUE V Series Proxy configuration section are the same as those on the UCT-V Configuration page. Refer to Configure UCT-V Controller for the field descriptions.

Configure GigaVUE V Series Node

Creating a GigaVUE V Series node profile automatically launches the V Series node. Enter or select the required information in the GigaVUE V Series Node section as described in the following table.

Prerequisites

Enable **`Host pass through`** by editing the *nova.conf* file and changing the *cpu_mode = host-passthrough*

Parameter	Description
Image	Select the GigaVUE V Series node image file.
Flavor	Select the form of the GigaVUE V Series node.
Management Network	For the GigaVUE V Series Node, the Management Network is what is used by the GigaVUE V Series Proxy to communicate with the GigaVUE V Series Nodes. Select the management network ID. **Ports**— Select a port, you can choose a port related to the selected management network ID. **Note**: When both IPv4 and IPv6 addresses are available, IPv6 address is preferred, however if IPv6 address is not reachable then IPv4 address is used.
Data Network	Click **Add** to add additional networks. This is the network that the GigVUE V Series node uses to communicate with the monitoring tools. Multiple networks are supported. • **Tool Subnet**—Select a tool subnet, this is the default subnet that the GigaVUE-FM use to egress traffic to your tools. This subnet must have proper connectivity to your endpoint. • **IP Address Type** ◦ **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy, or the GigaVUE V Series node 2 reside inside the same project. ◦ **Floating**—A floating IP address specified here will be where V Series node 2x.x can be directly managed by GigaVUE-FM or can optionally managed by controllers. • **Network 1**—Select a network type. • **Ports** —Select a port associated with the network. **Note**: • For OVS Mirroring or OVS Mirroring + DPDK deployments, must select **Floating** in the Data Network section and then specify the IPs in the **Floating IPs** field. You can have multiple Floating IPs. • A network provider that is able to receive the monitored traffic may also be used here for OVS Mirroring and OVS Mirroring + DPDK. In this case, you would not need to provide a floating IP; but could select "private" and choose the provider network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name.

Parameter	Description
	c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Min Instances	The minimum number of GigaVUE V Series nodes to be launched in OpenStack. The minimum number can be 1. • When you deploy an OVS Mirroring or OVS Mirroring + DPDK monitoring session, the V Series nodes will automatically be deployed based on the # of hypervisors being monitored. • When you deploy a UCT-V based monitoring session, the V Series nodes will automatically be deployed based on the # of VMs being monitored and the instance per V Series node ratio defined in the OpenStack Setttings page. **Note**: 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 OpenStack.
V Series Node Name	(Optional) Enter the name of the V Series Node. The V Series Node name must meet the following criteria: o The entire name can be a minimum of 1 to a maximum of 128 characters. o The suffix must only be a numeral and it should range between 0 to 999999999. o When deploying multiple V Series Nodes, the suffix of the consecutive V Series Node name is updated successively. E.g., 000, 001, 002, 003, etc..

Click **Save** to save the OpenStack Fabric Launch Configuration.

To view the fabric launch configuration specification of a visibility node, click on a visibility node or proxy, and a quick view of the Fabric Launch Configuration appears on the Monitoring Domain page.

Setting MTU for OVS Mirroring

When using OVS (Open vSwitch) port mirroring, it's important to configure the MTU (Maximum Transmission Unit) of VM interfaces correctly to avoid packet fragmentation issues in the data path.

- If the original MTU is 1450, reduce it to 1376:
- Calculation: 1450 - (50 bytes VXLAN header + 24 bytes extra headroom)
- If the original MTU is 8950, reduce it to 8816 using the same logic.

For Example: In an OpenStack deployment with a default MTU of 1450, packets are fragmented at the VM based on this MTU. When mirrored via OVS with VXLAN tunneling (which adds ~50 bytes), the total packet size exceeds the path MTU, causing further fragmentation. This can lead to out-of-order IP frame delivery at the VseriesNode.

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

What to do Next

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

Uninstall UCT-V

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

Uninstallation Method

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

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

Upgrade or Reinstall UCT-V

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

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

Refer to the below sections for detailed information and instructions.

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

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

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

Refer to the sections below for details and instructions:

1. Upload the UCT-V Images
2. Upgrade the UCT-V

Rules and Notes:

- Upgrade is allowed only to versions 6.9.00 or later. Ensure that the UCT-V Controller version is compatible with the version to which you are upgrading.
- Do not trigger system upgrades or other upgrades at the same time as the UCT-V upgrade.
- You should have Infrastructure Management permission to upgrade the UCT-Vs.
- Currently, you can upgrade the UCT-Vs to n+2 versions and any number of patch releases through GigaVUE-FM.
- Before you proceed with the upgrade, ensure that the UCT-Vs are in a healthy state.
- Make sure that a UCT-V is performing only one active job at a time. If the selected UCT-V is part of another job, you cannot trigger the immediate job using the same UCT-V.
- You must upload a compatible image type to upgrade the UCT-V; otherwise, the UCT-V is rejected for the upgrade job.
- Upgrade through GigaVUE-FM is not applicable to OVS Modules. For OVS tapping, you should upgrade the UCT-Vs manually.

Upload the UCT-V Images

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

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

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

Delete the file

You can delete one or multiple images.

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

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

Upgrade the UCT-V

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

1. In the **UCT-V Upgrade** drop-down menu, select **Dashboard** to view the UCT-V upgrade landing page.
   In the Dashboard page, you can view the upgrade status of individual UCT-Vs and the stages of the upgrade process (Fetch, Install, Verify). The page also displays the overall progress of the upgrade.
2. Select the required UCT-Vs, and from the **Actions** drop-down menu, select **Upgrade**.
   The **UCT-V Upgrade task** page appears.
3. Enter the task name.

4. In the **Image Version** drop-down menu, select the required version you want to upgrade to from the list of available image versions. You can choose to upgrade immediately or schedule a time for the upgrade to happen.
5. Select the required option in the **Time Selection** field. If you prefer to schedule the upgrade, enter the choice of your date and time in the respective fields.
   Do not schedule the upgrade for a time in the past.
6. Select **Create**.
   The image upgrade task is now created.

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

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

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

You can view the different stages of the upgrade process in UCT-V Upgrade Dashboard page. Each stage is marked with if it is successful and in case of failure. If the upgrade is successful, GigaVUE-FM updates the upgrade status as **Success** for the selected UCT-V.

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

Upgrade UCT-V Manually

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

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

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

NOTE: When the openssl version on the UCT-V is upgraded, Monitoring Session needs to be redeployed

Acquire Traffic using OVS Mirroring as Traffic Acquisition Method

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

- Install GigaVUE-FM on OpenStack
- Install UCT-V OVS Mirroring Module
- Create Monitoring Domain
- Configure GigaVUE Fabric Components in GigaVUE-FM
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Install UCT-V OVS Mirroring Module

This is applicable only if you are using UCT-V OVS module as the source of acquiring traffic. You must have sudo/root access to edit the UCT-V OVS module configuration file. Before installing the UCT-V OVS modules, you must have launched the GigaVUE-FM instance. UCT-V OVS module supports a maximum of 255 source interfaces per OpenStack node.

NOTE: After rebooting your workload VM, you must redeploy the respective Monitoring Sessions to restore the mirror traffic on the respective workload VM interfaces.

You can install theUCT-V OVS module either from Debian or RPM packages as follows:

- Install the UCT-V OVS Module from Ubuntu/Debian Package
- Install the UCT-V OVS Module from RPM package

Install the UCT-V OVS Module from Ubuntu/Debian Package

To install from a Debian package:

1. Download the latest version of UCT-V OVS Module Debian (.deb) package from the Gigamon Customer Portal.
2. Copy this package to OpenStack compute nodes. Install the package with root privileges, for example:

```
$ ls gigamon-gigavue-uctv-ovs-agent_6.14.02_amd64.deb
$ sudo dpkg -i gigamon-gigavue-uctv-ovs-agent_6.14.02_amd64.deb
```

3. Once the UCT-V OVS module package is installed, modify the file **/etc/uctv/uctv.conf** to configure and grant permission to monitor ingress and egress traffic and to transmit the mirrored packets.

   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.

```
br-int mirror-dst

# Changes for OVS Mirroring
# This Value will be used as local Ip in OVS Mirror Config
tunnel-src 172.20.20.11
# This Value will be used as Next Hop for Tunneled Packets
tunnel-gw 172.20.20.1
This Value will be used as local Ipv6 in OVS Mirror Config
tunnel-src-v6 2001::161
This Value will be used as Next Hop ipv6 addr for Tunneled Packets
tunnel-gw-v6 2001::1
# OVS Agent Mode, Values: auto|standard|dpdk|hw-offload
ovs-agent-mode auto
# VLAN Tag value (valid: 0-4094)
ovs-vlan-tag 2020
# Egress Interface for OVS Mirrored Traffic
ovs-egress-if vlan2020
```

4. After modifying the UCT-V OVS config file, start the UCT-V service.

```
$ sudo service uctv start
```

5. The UCT-V OVS module status will be displayed as running. Check the status using the following command:

```
$ sudo service uctv status
UCT-V is running
```

Install the UCT-V OVS Module from RPM package

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

1. Download the UCT-V OVS module RPM (.rpm) package from the Gigamon Customer Portal.
2. Copy this package to OpenStack compute nodes. Install the package with root privileges, for example:

```
$ ls gigamon-gigavue-uctv-ovs-agent_6.14.02_x86_64.rpm
$ sudo rpm -ivh gigamon-gigavue-uctv-ovs-agent_6.14.02_x86_64.rpm
```

3. Once the OVS module package is installed, modify the file **/etc/uctv/uctv.conf** to configure and grant permission to monitor ingress and egress traffic and transmit the mirrored packets.

> **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.

```
# br-int mirror-dst

# Changes for OVS Mirroring
# This Value will be used as local Ip in OVS Mirror Config
tunnel-src 172.20.20.11
# This Value will be used as Next Hop for Tunneled Packets
tunnel-gw 172.20.20.1
This Value will be used as local Ipv6 in OVS Mirror Config
tunnel-src-v6 2001::161
This Value will be used as Next Hop ipv6 addr for Tunneled Packets
tunnel-gw-v6 2001::1
# OVS Agent Mode, Values: auto|standard|dpdk|hw-offload
ovs-agent-mode auto
# VLAN Tag value (valid: 0-4094)
ovs-vlan-tag 2020
# Egress Interface for OVS Mirrored Traffic
ovs-egress-if vlan2020
```

4. After modifying the UCT-V OVS config file, start the UCT-V service and verify its status.

```
$ systemctl start uctv.service
$ sudo service uctv status
UCT-V is running
```

- UCT-V OVS module is supported for OpenStack with container-based deployment.
  # Docker name to run the OVS Commands
  docker-name openvswitch_vswitchd
- When you are installing a self-signed RPM package, you must execute the following command to import the signing key into the RPM db.
  sudo rpm --import /path/to/YOUR-RPM-GPG-KEY

To upgrade UCT-V OVS module:

- You must backup the **/etc/uctv/uctv.conf** configuration file before upgrading the UCT-V OVS module and uninstall the old OVS modules.
- Follow the same installation procedure to upgrade the UCT-V OVS modules.
- After upgrading the UCT-V OVS module, copy and modify the **uctv.conf** file, stop the UCT-V, and start the UCT-V. Redeploy the Monitoring Session if required.
  service uctv stop
  service uctv start

Create Monitoring Domain

To create a monitoring domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > OpenStack**. The Monitoring Domain page appears.
2. On the Monitoring Domain page, click **New**. The **Monitoring Domain Configuration** page appears.

3. Enter or select the appropriate information to configure Monitoring Domain for OpenStack. Refer to the following table for field-level details.

   NOTE: For the URL, User Domain Name, Project Domain Name, and Region field values, refer to the RC file downloaded from your OpenStack dashboard.

Field	Description
Monitoring Domain	A name for the monitoring domain. **Note**: You can only view and delete the existing configuration for V Series node 1. You cannot create and perform any other actions on the existing configuration for GigaVUE V Series node 1 as the features are deprecated from GigaVUE-FM fabric manager.
Alias	An alias used to identify the monitoring domain.
URL	The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable. Refer to the OpenStack User Manual for more information on retrieving the authentication URL from the OpenStack. **Note**: To ensure the validity of OpenStack certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. Refer to the Trust Store section in GigaVUE Administration Guide for more detailed information.
User Domain Name	The domain name of your OpenStack authentication domain. **Note**: • If you are using a separate domain for AUTH, enter that domain name as User Domain Name. • If you are not using a separate domain, you can use the same domain for User and Project Domain Name.
Project Domain Name	The domain name of your OpenStack project.
Project Name	The name of the project used for OpenStack authentication.
Region	The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version. **keystone endpoint-list** or **openstack endpoint list** or looking at the RC file in OpenStack to view your credentials.
Username	The username used to connect to your OpenStack cloud. **Note**: If you are using OVS mirroring, you must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to OVS Mirroring Prerequisites for more information.
Password	The password of your OpenStack cloud.
Traffic	Select the type of component used to capture traffic for monitoring:

Field	Description
Acquisition Method	• **OVS Mirroring:** If you select OVS Mirroring as your tapping method, the traffic is acquired from the UCT-Vs installed on the hypervisors. You must configure the UCT-V Controller to monitor the UCT-Vs. **Note**: For software release 6.7.00, only OVS Mirroring is supported on RHOSP 17.1 version.
Projects to Monitor (Only for OVS Mirroring traffic acquisition method)	This field only appears for OVS Mirroring traffic acquisition method. • Click the **Get Project List** to view the list of projects. **Note**: The **Get Project List** button will only work if all the OpenStack credentials have been provided. Refer to OVS Mirroring Prerequisites. • Select projects that you want to monitor from the list. • You can click **Select None** to clear existing selections or **Select All** to add all available projects to the connection configuration.
Traffic Acquisition Tunnel MTU (Maximum Transmission Unit)	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 Cloud Suite V Series node. • For GRE, the default value is 1450. • For VXLAN, the default value is 1400. However, the UCT-V tunnel MTU should be 50 bytes less than the default MTU size.

4. Click **Save**. The **OpenStack Fabric Launch Configuration** page appears. Refer to Configure GigaVUE Cloud Suite for OpenStack Components for detailed information.

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.
- If GigaVUE-FM fails to connect to OpenStack, an error message is displayed specifying the cause of failure. The connection status is also displayed in Audit Logs, refer to About Audit Logs for more information.

Configure GigaVUE Fabric Components in GigaVUE-FM

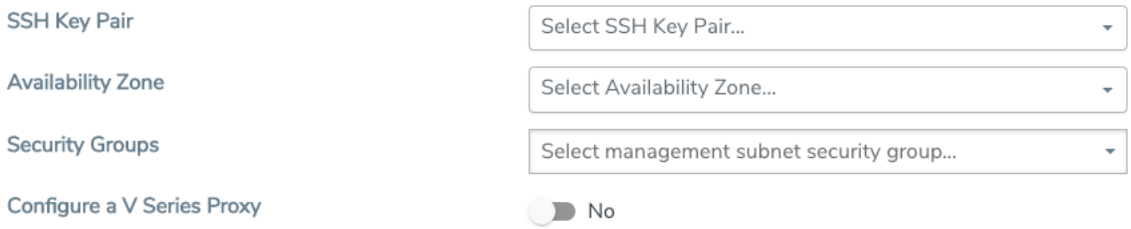
After configuring the Monitoring Domain, you will be navigated to the OpenStack Fabric Launch Configuration page. In the same **OpenStack Fabric Launch Configuration** page, you can configure the following fabric components:

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

In the **OpenStack Fabric Launch Configuration** page, enter or select the required information as described in the following table.

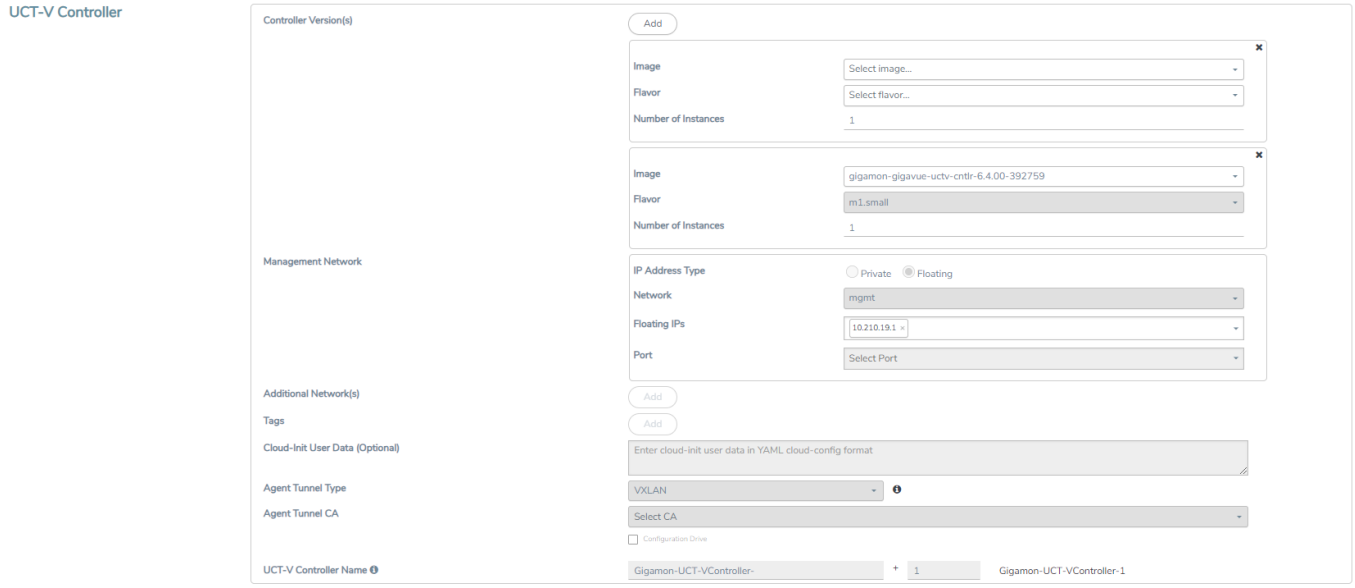
Fields	Description
SSH Key Pair	The SSH key pair for the UCT-V Controller. For more information about SSH key pair, refer to Key Pairs.
Availability Zone	The distinct locations (zones) of the OpenStack region.
Security Groups	The security group created for the UCT-V Controller. For more information, refer to Security Group for OpenStack .
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.
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.
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 more detailed information, refer to .

Select **Yes** to configure a GigaVUE V Series Proxy.



Configure UCT-V Controller

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



- Only if UCT-Vs are used for capturing traffic, then the UCT-V Controllers must be configured in the OpenStack cloud.
- A UCT-V Controller can only manage UCT-Vs that have the same version.

Enter or select the required information in the UCT-V Controller section as described in the following table.

Fields	Description
Controller Version(s)	The UCT-V Controller version that you configure must always have the same version number as the UCT-Vs deployed in the instances. For more detailed information refer GigaVUE-FM Version Compatibility Matrix. **Note**: If there is a version mismatch between the UCT-V Controllers and UCT-Vs, GigaVUE-FM cannot detect the agents in the instances. To add UCT-V Controllers: **a.** Under **Controller Versions**, click **Add**. **b.** 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. **c.** From the **Flavor** drop-down list, select a size for the UCT-V Controller. **d.** In **Number of Instances**, specify the number of UCT-V Controllers to launch. The minimum number you can specify is 1.
Management Network	This segment defines the management network that GigaVUE-FM uses to communicate with UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes. **Network** - Select the management network ID. **Ports** - Select a port, you can choose a port related to the selected management network ID. **IP Address Type**

Fields	Description
	The type of IP address GigaVUE-FM needs to communicate with UCT-V Controllers: o **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy reside inside the same project. o **Floating**—A floating IP is needed only if GigaVUE-FM is not in the same project in the cloud or is outside the cloud. GigaVUE-FM needs a floating IP to communicate with the controllers from an external network.
Additional Network(s)	(Optional) If there are UCT-Vs on networks that are not IP routable from the management network, additional networks or subnets must be specified so that the UCT-V Controller can communicate with all the UCT-Vs. Click **Add** to specify additional networks (subnets), if needed. Also, make sure that you specify a list of security groups for each additional network. **Ports**: Select a port associated with the network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. There is a specific UCT-V Controller Version for OVS Mirroring and OVS Mirroring + DPDK. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name. c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Agent Tunnel Type	The type of tunnel used for sending the traffic from UCT-Vs to GigaVUE V Series nodes. The options are GRE, VXLAN, and Secure tunnels (TLS-PCAPNG).
Agent Tunnel CA	The Certificate Authority (CA) that should be used in the UCT-V Controller for connecting the tunnel.

Configure GigaVUE V Series Proxy

The fields in the GigaVUE V Series Proxy configuration section are the same as those on the UCT-V Configuration page. Refer to Configure UCT-V Controller for the field descriptions.

Configure GigaVUE V Series Node

Creating a GigaVUE V Series node profile automatically launches the V Series node. Enter or select the required information in the GigaVUE V Series Node section as described in the following table.

Prerequisites

Enable **`Host pass through`** by editing the *nova.conf* file and changing the *cpu_mode = host-passthrough*

Parameter	Description
Image	Select the GigaVUE V Series node image file.
Flavor	Select the form of the GigaVUE V Series node.
Management Network	For the GigaVUE V Series Node, the Management Network is what is used by the GigaVUE V Series Proxy to communicate with the GigaVUE V Series Nodes. Select the management network ID. **Ports**— Select a port, you can choose a port related to the selected management network ID. **Note**: When both IPv4 and IPv6 addresses are available, IPv6 address is preferred, however if IPv6 address is not reachable then IPv4 address is used.
Data Network	Click **Add** to add additional networks. This is the network that the GigVUE V Series node uses to communicate with the monitoring tools. Multiple networks are supported. • **Tool Subnet**—Select a tool subnet, this is the default subnet that the GigaVUE-FM use to egress traffic to your tools. This subnet must have proper connectivity to your endpoint. • **IP Address Type** ◦ **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy, or the GigaVUE V Series node 2 reside inside the same project. ◦ **Floating**—A floating IP address specified here will be where V Series node 2x.x can be directly managed by GigaVUE-FM or can optionally managed by controllers. • **Network 1**—Select a network type. • **Ports** —Select a port associated with the network. **Note**: • For OVS Mirroring or OVS Mirroring + DPDK deployments, must select **Floating** in the Data Network section and then specify the IPs in the **Floating IPs** field. You can have multiple Floating IPs. • A network provider that is able to receive the monitored traffic may also be used here for OVS Mirroring and OVS Mirroring + DPDK. In this case, you would not need to provide a floating IP; but could select "private" and choose the provider network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name.

Parameter	Description
	c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Min Instances	The minimum number of GigaVUE V Series nodes to be launched in OpenStack. The minimum number can be 1. • When you deploy an OVS Mirroring or OVS Mirroring + DPDK monitoring session, the V Series nodes will automatically be deployed based on the # of hypervisors being monitored. • When you deploy a UCT-V based monitoring session, the V Series nodes will automatically be deployed based on the # of VMs being monitored and the instance per V Series node ratio defined in the OpenStack Setttings page. **Note**: 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 OpenStack.
V Series Node Name	(Optional) Enter the name of the V Series Node. The V Series Node name must meet the following criteria: o The entire name can be a minimum of 1 to a maximum of 128 characters. o The suffix must only be a numeral and it should range between 0 to 999999999. o When deploying multiple V Series Nodes, the suffix of the consecutive V Series Node name is updated successively. E.g., 000, 001, 002, 003, etc..

Click **Save** to save the OpenStack Fabric Launch Configuration.

To view the fabric launch configuration specification of a visibility node, click on a visibility node or proxy, and a quick view of the Fabric Launch Configuration appears on the Monitoring Domain page.

Setting MTU for OVS Mirroring

When using OVS (Open vSwitch) port mirroring, it's important to configure the MTU (Maximum Transmission Unit) of VM interfaces correctly to avoid packet fragmentation issues in the data path.

- If the original MTU is 1450, reduce it to 1376:
- Calculation: 1450 - (50 bytes VXLAN header + 24 bytes extra headroom)
- If the original MTU is 8950, reduce it to 8816 using the same logic.

For Example: In an OpenStack deployment with a default MTU of 1450, packets are fragmented at the VM based on this MTU. When mirrored via OVS with VXLAN tunneling (which adds ~50 bytes), the total packet size exceeds the path MTU, causing further fragmentation. This can lead to out-of-order IP frame delivery at the VseriesNode.

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

What to do Next

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

Acquire Traffic using Customer Orchestrated Source as Traffic Acquisition Method

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

- Install GigaVUE-FM on OpenStack
- Create Monitoring Domain
- Configure GigaVUE Fabric Components in GigaVUE-FM
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Create Monitoring Domain

To create a monitoring domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > OpenStack**. The Monitoring Domain page appears.
2. On the Monitoring Domain page, click **New**. The **Monitoring Domain Configuration** page appears.

3. Enter or select the appropriate information to configure Monitoring Domain for OpenStack. Refer to the following table for field-level details.

NOTE: For the URL, User Domain Name, Project Domain Name, and Region field values, refer to the RC file downloaded from your OpenStack dashboard.

Field	Description
Monitoring Domain	A name for the monitoring domain. **Note**: You can only view and delete the existing configuration for V Series node 1. You cannot create and perform any other actions on the existing configuration for GigaVUE V Series node 1 as the features are deprecated from GigaVUE-FM fabric manager.
Alias	An alias used to identify the monitoring domain.
URL	The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable. Refer to the OpenStack User Manual for more information on retrieving the authentication URL from the OpenStack. **Note**: To ensure the validity of OpenStack certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. Refer to the Trust Store section in GigaVUE Administration Guide for more detailed information.
User Domain Name	The domain name of your OpenStack authentication domain. **Note**: • If you are using a separate domain for AUTH, enter that domain name as User Domain Name. • If you are not using a separate domain, you can use the same domain for User and Project Domain Name.
Project Domain Name	The domain name of your OpenStack project.
Project Name	The name of the project used for OpenStack authentication.
Region	The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version. **keystone endpoint-list** or **openstack endpoint list** or looking at the RC file in OpenStack to view your credentials.
Username	The username used to connect to your OpenStack cloud. **Note**: If you are using OVS mirroring, you must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to OVS Mirroring Prerequisites for more information.
Password	The password of your OpenStack cloud.
Traffic	Select the type of component used to capture traffic for monitoring:

Field	Description
Acquisition Method	• **Customer Orchestrated Source:** If you selectCustomer Orchestrated Source as the tapping method, you can use tunnels as a source where the traffic is directly tunneled to V Series nodes without deploying UCT-Vs or UCT-V Controllers.
Projects to Monitor (Only for OVS Mirroring traffic acquisition method)	This field only appears for OVS Mirroring traffic acquisition method. • Click the **Get Project List** to view the list of projects. **Note**: The **Get Project List** button will only work if all the OpenStack credentials have been provided. Refer to OVS Mirroring Prerequisites. • Select projects that you want to monitor from the list. • You can click **Select None** to clear existing selections or **Select All** to add all available projects to the connection configuration.
Traffic Acquisition Tunnel MTU (Maximum Transmission Unit)	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 Cloud Suite V Series node. • For GRE, the default value is 1450. • For VXLAN, the default value is 1400. However, the UCT-V tunnel MTU should be 50 bytes less than the default MTU size.

4. Click **Save**. The **OpenStack Fabric Launch Configuration** page appears. Refer to Configure GigaVUE Cloud Suite for OpenStack Components for detailed information.

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.
- If GigaVUE-FM fails to connect to OpenStack, an error message is displayed specifying the cause of failure. The connection status is also displayed in Audit Logs, refer to About Audit Logs for more information.

Configure GigaVUE Fabric Components in GigaVUE-FM

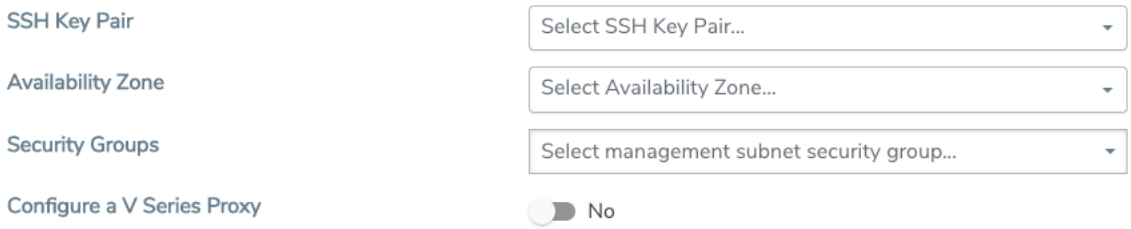
After configuring the Monitoring Domain, you will be navigated to the OpenStack Fabric Launch Configuration page. In the same **OpenStack Fabric Launch Configuration** page, you can configure the following fabric components:

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

In the **OpenStack Fabric Launch Configuration** page, enter or select the required information as described in the following table.

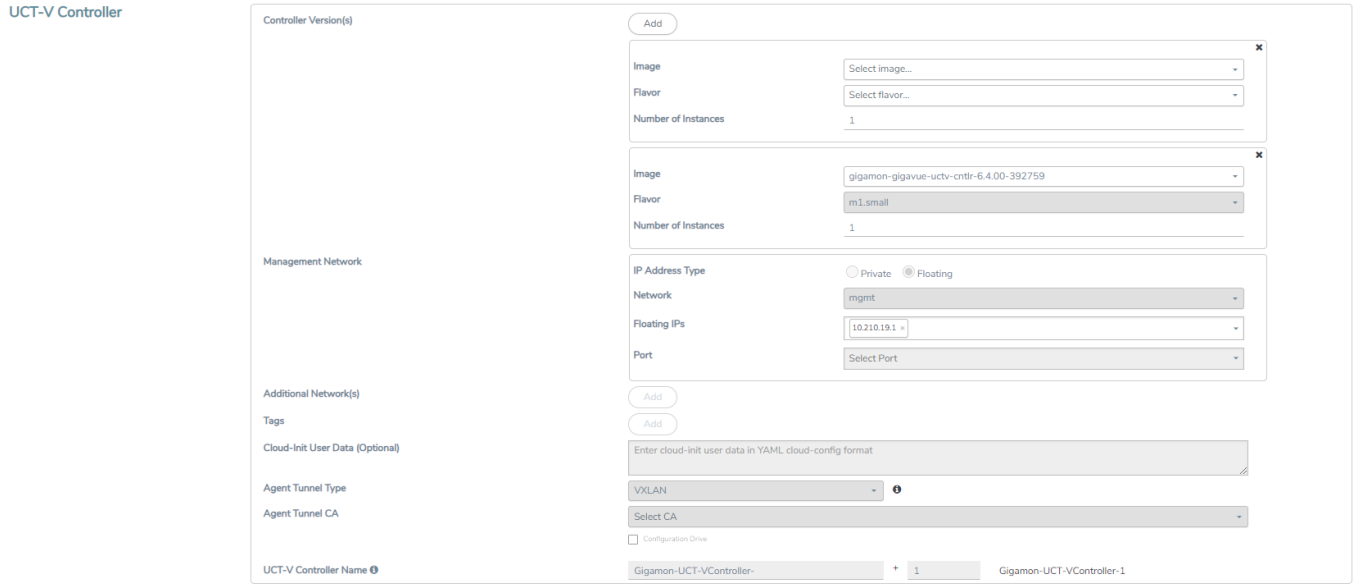
Fields	Description
SSH Key Pair	The SSH key pair for the UCT-V Controller. For more information about SSH key pair, refer to Key Pairs.
Availability Zone	The distinct locations (zones) of the OpenStack region.
Security Groups	The security group created for the UCT-V Controller. For more information, refer to Security Group for OpenStack .
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.
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.
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 more detailed information, refer to .

Select **Yes** to configure a GigaVUE V Series Proxy.



Configure UCT-V Controller

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



- Only if UCT-Vs are used for capturing traffic, then the UCT-V Controllers must be configured in the OpenStack cloud.
- A UCT-V Controller can only manage UCT-Vs that have the same version.

Enter or select the required information in the UCT-V Controller section as described in the following table.

Fields	Description
Controller Version(s)	The UCT-V Controller version that you configure must always have the same version number as the UCT-Vs deployed in the instances. For more detailed information refer GigaVUE-FM Version Compatibility Matrix. **Note**: If there is a version mismatch between the UCT-V Controllers and UCT-Vs, GigaVUE-FM cannot detect the agents in the instances. To add UCT-V Controllers: **a.** Under **Controller Versions**, click **Add**. **b.** 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. **c.** From the **Flavor** drop-down list, select a size for the UCT-V Controller. **d.** In **Number of Instances**, specify the number of UCT-V Controllers to launch. The minimum number you can specify is 1.
Management Network	This segment defines the management network that GigaVUE-FM uses to communicate with UCT-V Controllers, GigaVUE V Series Proxy, and GigaVUE V Series Nodes. **Network** - Select the management network ID. **Ports** - Select a port, you can choose a port related to the selected management network ID. **IP Address Type**

Fields	Description
	The type of IP address GigaVUE-FM needs to communicate with UCT-V Controllers: o **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy reside inside the same project. o **Floating**—A floating IP is needed only if GigaVUE-FM is not in the same project in the cloud or is outside the cloud. GigaVUE-FM needs a floating IP to communicate with the controllers from an external network.
Additional Network(s)	(Optional) If there are UCT-Vs on networks that are not IP routable from the management network, additional networks or subnets must be specified so that the UCT-V Controller can communicate with all the UCT-Vs. Click **Add** to specify additional networks (subnets), if needed. Also, make sure that you specify a list of security groups for each additional network. **Ports**: Select a port associated with the network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. There is a specific UCT-V Controller Version for OVS Mirroring and OVS Mirroring + DPDK. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name. c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Agent Tunnel Type	The type of tunnel used for sending the traffic from UCT-Vs to GigaVUE V Series nodes. The options are GRE, VXLAN, and Secure tunnels (TLS-PCAPNG).
Agent Tunnel CA	The Certificate Authority (CA) that should be used in the UCT-V Controller for connecting the tunnel.

Configure GigaVUE V Series Proxy

The fields in the GigaVUE V Series Proxy configuration section are the same as those on the UCT-V Configuration page. Refer to Configure UCT-V Controller for the field descriptions.

Configure GigaVUE V Series Node

Creating a GigaVUE V Series node profile automatically launches the V Series node. Enter or select the required information in the GigaVUE V Series Node section as described in the following table.

Prerequisites

Enable **`Host pass through`** by editing the *nova.conf* file and changing the *cpu_mode = host-passthrough*

Parameter	Description
Image	Select the GigaVUE V Series node image file.
Flavor	Select the form of the GigaVUE V Series node.
Management Network	For the GigaVUE V Series Node, the Management Network is what is used by the GigaVUE V Series Proxy to communicate with the GigaVUE V Series Nodes. Select the management network ID. **Ports**— Select a port, you can choose a port related to the selected management network ID. **Note**: When both IPv4 and IPv6 addresses are available, IPv6 address is preferred, however if IPv6 address is not reachable then IPv4 address is used.
Data Network	Click **Add** to add additional networks. This is the network that the GigVUE V Series node uses to communicate with the monitoring tools. Multiple networks are supported. • **Tool Subnet**—Select a tool subnet, this is the default subnet that the GigaVUE-FM use to egress traffic to your tools. This subnet must have proper connectivity to your endpoint. • **IP Address Type** ◦ **Private**—A private IP can be used when GigaVUE-FM, the UCT-V Controller, or the GigaVUE V Series Proxy, or the GigaVUE V Series node 2 reside inside the same project. ◦ **Floating**—A floating IP address specified here will be where V Series node 2x.x can be directly managed by GigaVUE-FM or can optionally managed by controllers. • **Network 1**—Select a network type. • **Ports** —Select a port associated with the network. **Note**: • For OVS Mirroring or OVS Mirroring + DPDK deployments, must select **Floating** in the Data Network section and then specify the IPs in the **Floating IPs** field. You can have multiple Floating IPs. • A network provider that is able to receive the monitored traffic may also be used here for OVS Mirroring and OVS Mirroring + DPDK. In this case, you would not need to provide a floating IP; but could select "private" and choose the provider network.
Tag(s)	(Optional) The key name and value that helps to identify the UCT-V Controller instances in your 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 (also known as a tag) that is easy to identify such as us-west-2-uctv-controllers. To add a tag: a. Click **Add**. b. In the **Key** field, enter the key. For example, enter Name.

Parameter	Description
	c. In the **Value** field, enter the key value. For example, us-west-2-uctv-controllers.
Cloud-Init User Data (Optional)	Enter the cloud-init user data in cloud-config format.
Min Instances	The minimum number of GigaVUE V Series nodes to be launched in OpenStack. The minimum number can be 1. • When you deploy an OVS Mirroring or OVS Mirroring + DPDK monitoring session, the V Series nodes will automatically be deployed based on the # of hypervisors being monitored. • When you deploy a UCT-V based monitoring session, the V Series nodes will automatically be deployed based on the # of VMs being monitored and the instance per V Series node ratio defined in the OpenStack Setttings page. **Note**: 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 OpenStack.
V Series Node Name	(Optional) Enter the name of the V Series Node. The V Series Node name must meet the following criteria: o The entire name can be a minimum of 1 to a maximum of 128 characters. o The suffix must only be a numeral and it should range between 0 to 999999999. o When deploying multiple V Series Nodes, the suffix of the consecutive V Series Node name is updated successively. E.g., 000, 001, 002, 003, etc..

Click **Save** to save the OpenStack Fabric Launch Configuration.

To view the fabric launch configuration specification of a visibility node, click on a visibility node or proxy, and a quick view of the Fabric Launch Configuration appears on the Monitoring Domain page.

Setting MTU for OVS Mirroring

When using OVS (Open vSwitch) port mirroring, it's important to configure the MTU (Maximum Transmission Unit) of VM interfaces correctly to avoid packet fragmentation issues in the data path.

- If the original MTU is 1450, reduce it to 1376:
- Calculation: 1450 - (50 bytes VXLAN header + 24 bytes extra headroom)
- If the original MTU is 8950, reduce it to 8816 using the same logic.

For Example: In an OpenStack deployment with a default MTU of 1450, packets are fragmented at the VM based on this MTU. When mirrored via OVS with VXLAN tunneling (which adds ~50 bytes), the total packet size exceeds the path MTU, causing further fragmentation. This can lead to out-of-order IP frame delivery at the VseriesNode.

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 > OpenStack**.
   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 ingress and Egress REPs**:
   a. In the canvas, from the **New** expand menu, drag and drop New Tunnel to the graphical workspace.
   b. On the new tunnel icon, select the ⋮ menu button and select **Details**. The **Add Tunnel Spec** quick view page appears.
   c. Enter the **Alias**, **Description**, and the **Type** details and click **Save**.
3. **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.
4. **Deploy the Monitoring Session**:
   a. From the **Actions** menu, select **Deploy**.
   b. After successful deployment on all the V Series Nodes, the status appears as **Success** on the Monitoring Sessions **Sources** tab.

What to do Next

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

Deploy GigaVUE Cloud Suite for OpenStack 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 sunGroupName 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 AWS:

- Deploy GigaVUE Fabric Components using Azure - Generic Mode
- Deployment GigaVUE Fabric Components using Azure -Integrated Mode

Deploy GigaVUE Fabric Components using OpenStack - 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

Acquire Traffic using UCT-V - Without creating Monitoring Domain

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

- Install GigaVUE-FM on OpenStack
- Configure Role-Based Access for Third Party Orchestration
- Install UCT-Vs
- Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Configure Role-Based Access for Third Party Orchestration

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

Create a Custom User

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

To add users:

1. Go to **Settings** and select **Authentication** > **GigaVUE-FM User Management > Users**.
2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
   - **Name:** Actual name of the user
   - **Username**: User name configured in GigaVUE-FM
   - **Email**: Email ID of the user
   - **Password/Confirm Password**: Password for the user. For more details, refer to Change Your Password.

- **User Group:** Select the desired User Group to associate the user.

  GigaVUE-FM prompts for your password.

4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

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

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

To create a role:

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

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

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

To create a new user group:

1. Go to **Settings**, and then select **Authentication> GigaVUE-FMUser Management >User Groups**.
2. Select **New Group.**
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:
5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Install UCT-Vs

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

- Linux UCT-V Installation
- Windows UCT-V Installation

> **NOTE:** If you want to avoid downloading and installing the UCT-Vs every time there is a new instance to be monitored, you can save the UCT-V running on an instance as a private AMI.
> To save the UCT-V as an AMI from your EC2 console, right click on the instance and navigate to **Image and Templates** > **Create Image**.

Linux UCT-V Installation

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

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

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

Refer to the following sections:

- Linux Network Firewall Requirements
- Install Linux UCT-Vs
- Register Linux UCT-V

Prerequisites

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

Linux Network Firewall Requirements

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

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

After you deploy the UCT-V VM, you can also use the following commands to add the network firewall rule:

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

Linux UCT-V Installation Scenarios

1. **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.
2. **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.
3. **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

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**6.14.02** 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-6.14.02-amd64.deb
    $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
    ```

2. **To install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS:**
    a. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
    $ sudo rpm -i gigamon-gigavue-uctv-6.14.02-x86_64.rpm
    ```

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

```
sudo uctv-wizard
```

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

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

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

Options	Use Command	Description
		If you wish to skip the prompts to add the required firewall policy, enter your option as **y**. The console interface adds the firewall rules automatically. Enter **N** if you wish to configure manually. For details, refer to Install Linux UCT-Vs using Manual Configuration section.
uninstall	sudo uctv-wizard uninstall	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
  ```

Install Linux UCT-Vs using Manual Configuration

- Install UCT-V from Ubuntu/Debian Package
- Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

1. Download the UCT-V**6.14.02** Debian (.deb) 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-6.14.02-amd64.deb
   $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
   ```

4. Modify the file **/etc/uctv/uctv.conf** to configure and register the source and destination interfaces.
   The following examples registers eth0 as the mirror source for both ingress and egress traffic and eth1 as the destination for this traffic:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

1. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
$ sudo rpm -i gigamon-gigavue-uctv-6.14.02-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

After installing the UCT-V package, 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:

- **GigaVUE-FM Orchestration**: Perform the following steps:
  a. Log in to the UCT-V.
  b. Create a local configuration file and enter the following user data. **/etc/gigamon-cloud.conf** is the local configuration file in the Linux platform.

```
Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        remoteIP: <IP address of the UCT-V Controller>
        remotePort: 443
        token: <Token>
```

  c. Restart the UCT-V service.
  Linux platform:

```
$ sudo service uctv restart
```

  d. Verify UCT-V is running:

```
sudo service uctv status
```

For more details on how to create tokens, refer to .

Verify that the service status is Active (running). After the UCT-V VM successfully connects to GigaVUE-FM, it appears in **Inventory** > **Resources** > **UCT-V** in the GigaVUE-FM.

- **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 toGigaVUE-FM using Third Party Orchestration with the following two modes:
    1. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
    2. 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

Rules and Notes

- VXLAN is the only tunnel type supported for Windows UCT-V.
- Loopback interface is not supported on 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 workflows.

Windows UCT-V Installation Scenarios

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

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

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 **`6.14.02`** 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
```

Install Windows UCT-Vs using Manual Configuration

1. Download the Windows UCT-V**6.14.02** 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.

Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode

This section provides information on how to register GigaVUE fabric components using OpenStack or a configuration file.

Recommended Instance Type

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

Fabric Component	Machine type
GigaVUE V Series Node	m1.medium
GigaVUE V Series Proxy	m1.small
UCT-V Controller	m1.small

Deployment Best Practices

Keep in mind the following when deploying the fabric components using generic mode:

- Set the correct MTU for tunnels: Use 1450 for IPv4 tunnels and 1430 for IPv6. To update the Traffic Acquisition Tunnel MTU:

    1. Select the monitoring domain.
    2. Select the **Edit Monitoring Domain** option.
    3. Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

    NOTE: Set the correct MTU before deploying the Monitoring Session. If not, you must undeploy and redeploy the session after updating the MTU.

- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to Create Monitoring Domain.
- GigaVUE V Series Node must have a minimum of two Networks Interfaces (NIC) attached to it, a management NIC and a data NIC. You can add both these interfaces when deploying the GigaVUE V Series Node in OpenStack.
- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to Configure Tokens for Third Party Orchestration.
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section in the 6.9 Documentation.

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

Configure UCT-V Controller in OpenStack

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

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

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

Register UCT-V Controller during Instance Launch

In your OpenStack dashboard, you can launch the UCT-V Controller and register UCT-V Controller using Customization Script.

perform the following steps:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The UCT-V Controller uses this registration data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM>
        sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
        remotePort: 443
```

The UCT-V Controller deployed in OpenStack appears on the Monitoring Domain page of GigaVUE-FM.

Register UCT-V Controller after Instance Launch

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

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

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

and use the details for registering UCT-V Controller.

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

3. Restart the UCT-V Controller service.

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

Assign Static IP address for UCT-V Controller

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

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

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

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

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

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

Example netplan config:

network:
  version: 2
  renderer: NetworkManager
```

```
ethernets:
  ens3:
    addresses:
      - 10.114.53.24/21
    dhcp4: no
    dhcp6: no
    accept-ra: false
    routes:
      - to: 10.114.48.1/32
        scope: link
      - to: default
        via: 10.114.48.1
```

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the UCT-V Controller service.

   $ sudo service uctv-cntlr restart

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

NOTE: When you deploy GigaVUE V Series Nodes or UCT-V Controllers using Third Party orchestration, you cannot delete the monitoring domain without unregistering the V Series Nodes or UCT-V Controllers.

Configure UCT-V in OpenStack

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

To register UCT-V using a configuration file:

1. Install the UCT-V in 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.

3. Edit the local configuration file and enter the following Customization Script.

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

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

   NOTE: 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-agent restart
   - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in OpenStack

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

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

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

Register V Series Nodes or V Series Proxy during Instance Launch

To register V Series nodes or proxy using the Customization Script in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```

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

Register V Series Node or V Series Proxy after Instance Launch

To register V Series node or proxy using a configuration file:

1. Log in to the GigaVUE V Series Node or Proxy.
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.

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

> **NOTE:** You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **`remotePort`** value as 443 and the **`remoteIP`** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using V Series proxy then, enter the **`remotePort`** value as 8891 and **`remoteIP`** as <IP address of the Proxy>. If the V Series Node registration fails with a Connection Refused error, set the remoteIP field to the API network IP of GigaVUE-FM instead of the floating IP.

3. Restart the GigaVUE V Series Node or Proxy service.
   - GigaVUE V Series Node:
     `$ sudo service vseries-node restart`
   - GigaVUE V Series Proxy:
     `$ sudo service vps restart`

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

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

Assign Static IP address for GigaVUE V Series

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the GigaVUE V Series service.

   $ sudo service vseries-node restart

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

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

What to do Next

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

Acquire Traffic using Customer Orchestrated Source - Without creating Monitoring Domain

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

- Install GigaVUE-FM on OpenStack
- Configure Role-Based Access for Third Party Orchestration
- Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Configure Role-Based Access for Third Party Orchestration

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

Create a Custom User

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

To add users:

1. Go to **Settings** and select **Authentication** > **GigaVUE-FM User Management > Users**.
2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
   - **Name:** Actual name of the user
   - **Username**: User name configured in GigaVUE-FM
   - **Email**: Email ID of the user
   - **Password/Confirm Password**: Password for the user. For more details, refer to Change Your Password.
   - **User Group:** Select the desired User Group to associate the user.

   GigaVUE-FM prompts for your password.
4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

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

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

To create a role:

1. Go to **Settings** and select **Authentication> GigaVUE-FM User Management >Roles**.
2. Select **New Role**.
3. In the New Role page, select or enter the following details:
   - **Role Name**: Name of the role.
   - **Description**: Description of the role.

- **Select Permission**: From the **Select Permissions** tab, select **Third Party Orchestration**, and provide write permissions.

4. Select **Apply** to save the configuration.

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

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

To create a new user group:

1. Go to **Settings**, and then select **Authentication> GigaVUE-FMUser Management >User Groups**.
2. Select **New Group.**
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:
5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode

This section provides information on how to register GigaVUE fabric components using OpenStack or a configuration file.

Recommended Instance Type

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

Fabric Component	Machine type
GigaVUE V Series Node	m1.medium
GigaVUE V Series Proxy	m1.small
UCT-V Controller	m1.small

Deployment Best Practices

Keep in mind the following when deploying the fabric components using generic mode:

- Set the correct MTU for tunnels: Use 1450 for IPv4 tunnels and 1430 for IPv6. To update the Traffic Acquisition Tunnel MTU:

    1. Select the monitoring domain.
    2. Select the **Edit Monitoring Domain** option.
    3. Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

    NOTE: Set the correct MTU before deploying the Monitoring Session. If not, you must undeploy and redeploy the session after updating the MTU.

- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to Create Monitoring Domain.
- GigaVUE V Series Node must have a minimum of two Networks Interfaces (NIC) attached to it, a management NIC and a data NIC. You can add both these interfaces when deploying the GigaVUE V Series Node in OpenStack.
- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to Configure Tokens for Third Party Orchestration.
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section in the 6.9 Documentation.

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

Configure UCT-V Controller in OpenStack

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

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

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

Register UCT-V Controller during Instance Launch

In your OpenStack dashboard, you can launch the UCT-V Controller and register UCT-V Controller using Customization Script.

perform the following steps:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The UCT-V Controller uses this registration data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM>
        sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
        remotePort: 443
```

The UCT-V Controller deployed in OpenStack appears on the Monitoring Domain page of GigaVUE-FM.

Register UCT-V Controller after Instance Launch

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

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

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

and use the details for registering UCT-V Controller.

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

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

Assign Static IP address for UCT-V Controller

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

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

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

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

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

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

Example netplan config:

network:
  version: 2
  renderer: NetworkManager
```

```
ethernets:
  ens3:
    addresses:
      - 10.114.53.24/21
    dhcp4: no
    dhcp6: no
    accept-ra: false
    routes:
      - to: 10.114.48.1/32
        scope: link
      - to: default
        via: 10.114.48.1
```

4. Save the file.
5. Apply the configuration.

   ```
   $ sudo netplan apply
   ```

6. Restart the UCT-V Controller service.

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

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

NOTE: When you deploy GigaVUE V Series Nodes or UCT-V Controllers using Third Party orchestration, you cannot delete the monitoring domain without unregistering the V Series Nodes or UCT-V Controllers.

Configure UCT-V in OpenStack

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

To register UCT-V using a configuration file:

1. Install the UCT-V in 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.

3. Edit the local configuration file and enter the following Customization Script.

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

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

   NOTE: 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-agent restart
   - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in OpenStack

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

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

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

Register V Series Nodes or V Series Proxy during Instance Launch

To register V Series nodes or proxy using the Customization Script in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```

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

Register V Series Node or V Series Proxy after Instance Launch

To register V Series node or proxy using a configuration file:

1. Log in to the GigaVUE V Series Node or Proxy.
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.

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

> **NOTE:** You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **`remotePort`** value as 443 and the **`remoteIP`** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using V Series proxy then, enter the **`remotePort`** value as 8891 and **`remoteIP`** as <IP address of the Proxy>. If the V Series Node registration fails with a Connection Refused error, set the remoteIP field to the API network IP of GigaVUE-FM instead of the floating IP.

3. Restart the GigaVUE V Series Node or Proxy service.
   - GigaVUE V Series Node:
     `$ sudo service vseries-node restart`
   - GigaVUE V Series Proxy:
     `$ sudo service vps restart`

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

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

Assign Static IP address for GigaVUE V Series

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the GigaVUE V Series service.

   $ sudo service vseries-node restart

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

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 > OpenStack**.
   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 ingress and Egress REPs**:
   a. In the canvas, from the **New** expand menu, drag and drop New Tunnel to the graphical workspace.
   b. On the new tunnel icon, select the ⋮ menu button and select **Details**. The **Add Tunnel Spec** quick view page appears.
3. **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.
4. **Deploy the Monitoring Session**:
   a. From the **Actions** menu, select **Deploy**.
   b. After successful deployment on all the V Series Nodes, the status appears as **Success** on the Monitoring Sessions **Sources** tab.

What to do Next

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

Acquire Traffic using UCT-V - By creating Monitoring Domain

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

- Install GigaVUE-FM on OpenStack
- Configure Role-Based Access for Third Party Orchestration
- Install UCT-Vs
- Create Monitoring Domain

- Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Configure Role-Based Access for Third Party Orchestration

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

Create a Custom User

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

To add users:

1. Go to **Settings** and select **Authentication** > **GigaVUE-FM User Management > Users**.
2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
   - **Name:** Actual name of the user
   - **Username**: User name configured in GigaVUE-FM
   - **Email**: Email ID of the user
   - **Password/Confirm Password**: Password for the user. For more details, refer to Change Your Password.
   - **User Group:** Select the desired User Group to associate the user.

   GigaVUE-FM prompts for your password.
4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

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

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

To create a role:

1. Go to **Settings** and select **Authentication> GigaVUE-FM User Management >Roles**.
2. Select **New Role**.

3. In the New Role page, select or enter the following details:
   - **Role Name**: Name of the role.
   - **Description**: Description of the role.
   - **Select Permission**: From the **Select Permissions** tab, select **Third Party Orchestration**, and provide write permissions.
4. Select **Apply** to save the configuration.

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

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

To create a new user group:

1. Go to **Settings**, and then select **Authentication> GigaVUE-FMUser Management >User Groups**.
2. Select **New Group.**
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:
5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Install UCT-Vs

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

- Linux UCT-V Installation
- Windows UCT-V Installation

> **NOTE:** If you want to avoid downloading and installing the UCT-Vs every time there is a new instance to be monitored, you can save the UCT-V running on an instance as a private AMI.
> To save the UCT-V as an AMI from your EC2 console, right click on the instance and navigate to **Image and Templates** > **Create Image**.

Linux UCT-V Installation

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

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

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

Refer to the following sections:

- Linux Network Firewall Requirements
- Install Linux UCT-Vs
- Register Linux UCT-V

Prerequisites

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

Linux Network Firewall Requirements

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

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

After you deploy the UCT-V VM, you can also use the following commands to add the network firewall rule:

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

Linux UCT-V Installation Scenarios

1. **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.
2. **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.
3. **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

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**6.14.02** 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-6.14.02-amd64.deb
    $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
    ```

2. **To install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS:**
    a. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
    $ sudo rpm -i gigamon-gigavue-uctv-6.14.02-x86_64.rpm
    ```

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

```
sudo uctv-wizard
```

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

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

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

Options	Use Command	Description
		If you wish to skip the prompts to add the required firewall policy, enter your option as **y**. The console interface adds the firewall rules automatically. Enter **N** if you wish to configure manually. For details, refer to Install Linux UCT-Vs using Manual Configuration section.
uninstall	sudo uctv-wizard uninstall	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
  ```

Install Linux UCT-Vs using Manual Configuration

- Install UCT-V from Ubuntu/Debian Package
- Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

1. Download the UCT-V**6.14.02** Debian (.deb) 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-6.14.02-amd64.deb
   $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
   ```

4. Modify the file **/etc/uctv/uctv.conf** to configure and register the source and destination interfaces.
   The following examples registers eth0 as the mirror source for both ingress and egress traffic and eth1 as the destination for this traffic:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

1. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
$ sudo rpm -i gigamon-gigavue-uctv-6.14.02-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

After installing the UCT-V package, 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:

- **GigaVUE-FM Orchestration**: Perform the following steps:
  a. Log in to the UCT-V.
  b. Create a local configuration file and enter the following user data. **/etc/gigamon-cloud.conf** is the local configuration file in the Linux platform.

```
Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        remoteIP: <IP address of the UCT-V Controller>
        remotePort: 443
        token: <Token>
```

  c. Restart the UCT-V service.
  Linux platform:

```
$ sudo service uctv restart
```

  d. Verify UCT-V is running:

```
sudo service uctv status
```

For more details on how to create tokens, refer to .

Verify that the service status is Active (running). After the UCT-V VM successfully connects to GigaVUE-FM, it appears in **Inventory** > **Resources** > **UCT-V** in the GigaVUE-FM.

- **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 toGigaVUE-FM using Third Party Orchestration with the following two modes:
    1. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
    2. 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

Rules and Notes

- VXLAN is the only tunnel type supported for Windows UCT-V.
- Loopback interface is not supported on 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 workflows.

Windows UCT-V Installation Scenarios

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

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

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 **`6.14.02`** 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
```

Install Windows UCT-Vs using Manual Configuration

1. Download the Windows UCT-V**6.14.02** 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.

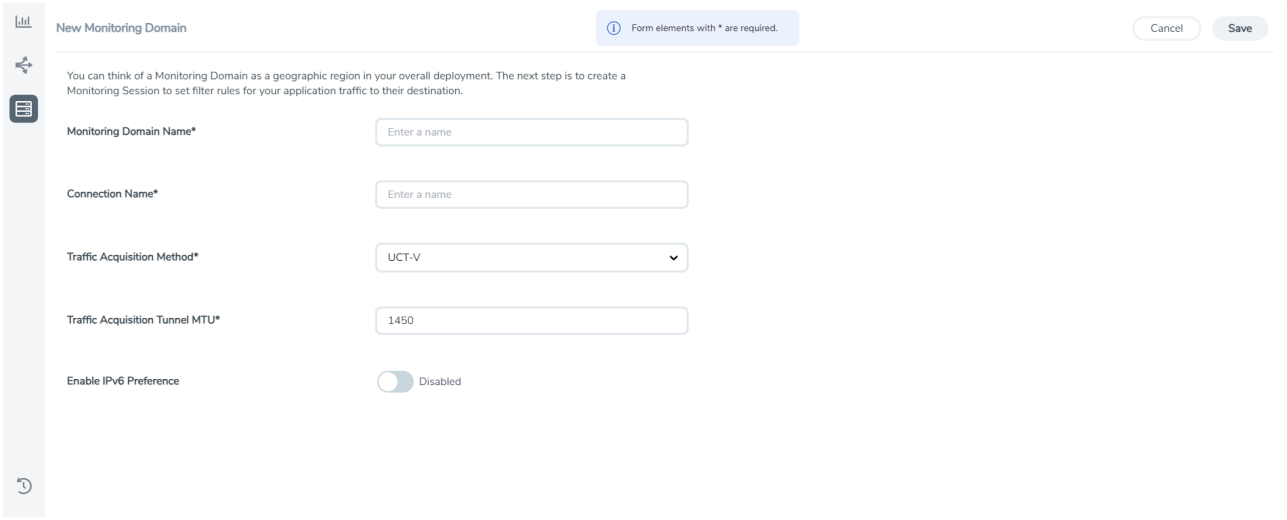
You can register UCT-V in your virtual machine using:

- **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 Monitoring Domain

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



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

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

4. Click **Save**.

Notes:

- Ensure that all V Series Nodes within a single Monitoring Domain are running the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.

- Similarly, when upgrading a V Series Node, ensure that the GigaVUE-FM version is the same or higher than the V Series Node version.

You can perform the following actions in the Monitoring domain page:

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

NOTE: To view and manage the generated sysdump files, select the GigaVUE V SeriesNode and select the **Sysdump** tab in the lower pane.

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

Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode

This section provides information on how to register GigaVUE fabric components using OpenStack or a configuration file.

Recommended Instance Type

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

Fabric Component	Machine type
GigaVUE V Series Node	m1.medium
GigaVUE V Series Proxy	m1.small
UCT-V Controller	m1.small

Deployment Best Practices

Keep in mind the following when deploying the fabric components using generic mode:

- Set the correct MTU for tunnels: Use 1450 for IPv4 tunnels and 1430 for IPv6. To update the Traffic Acquisition Tunnel MTU:

    1. Select the monitoring domain.
    2. Select the **Edit Monitoring Domain** option.
    3. Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

    > **NOTE:** Set the correct MTU before deploying the Monitoring Session. If not, you must undeploy and redeploy the session after updating the MTU.

- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to Create Monitoring Domain.
- GigaVUE V Series Node must have a minimum of two Networks Interfaces (NIC) attached to it, a management NIC and a data NIC. You can add both these interfaces when deploying the GigaVUE V Series Node in OpenStack.
- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to Configure Tokens for Third Party Orchestration.
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section in the 6.9 Documentation.

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

Configure UCT-V Controller in OpenStack

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

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

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

Register UCT-V Controller during Instance Launch

In your OpenStack dashboard, you can launch the UCT-V Controller and register UCT-V Controller using Customization Script.

perform the following steps:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The UCT-V Controller uses this registration data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM>
        sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
        remotePort: 443
```

The UCT-V Controller deployed in OpenStack appears on the Monitoring Domain page of GigaVUE-FM.

Register UCT-V Controller after Instance Launch

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

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

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

and use the details for registering UCT-V Controller.

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

3. Restart the UCT-V Controller service.

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

Assign Static IP address for UCT-V Controller

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

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

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

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

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

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

Example netplan config:

network:
  version: 2
  renderer: NetworkManager
```

```
ethernets:
  ens3:
    addresses:
      - 10.114.53.24/21
    dhcp4: no
    dhcp6: no
    accept-ra: false
    routes:
      - to: 10.114.48.1/32
        scope: link
      - to: default
        via: 10.114.48.1
```

4. Save the file.
5. Apply the configuration.

   ```
   $ sudo netplan apply
   ```

6. Restart the UCT-V Controller service.

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

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

NOTE: When you deploy GigaVUE V Series Nodes or UCT-V Controllers using Third Party orchestration, you cannot delete the monitoring domain without unregistering the V Series Nodes or UCT-V Controllers.

Configure UCT-V in OpenStack

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

To register UCT-V using a configuration file:

1. Install the UCT-V in 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.

3. Edit the local configuration file and enter the following Customization Script.

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

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

   NOTE: 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-agent restart`
   - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in OpenStack

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

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

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

Register V Series Nodes or V Series Proxy during Instance Launch

To register V Series nodes or proxy using the Customization Script in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```

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

Register V Series Node or V Series Proxy after Instance Launch

To register V Series node or proxy using a configuration file:

1. Log in to the GigaVUE V Series Node or Proxy.
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.

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

> **NOTE:** You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **`remotePort`** value as 443 and the **`remoteIP`** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using V Series proxy then, enter the **`remotePort`** value as 8891 and **`remoteIP`** as <IP address of the Proxy>. If the V Series Node registration fails with a Connection Refused error, set the remoteIP field to the API network IP of GigaVUE-FM instead of the floating IP.

3. Restart the GigaVUE V Series Node or Proxy service.
   - GigaVUE V Series Node:
     `$ sudo service vseries-node restart`
   - GigaVUE V Series Proxy:
     `$ sudo service vps restart`

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

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

Assign Static IP address for GigaVUE V Series

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the GigaVUE V Series service.

   $ sudo service vseries-node restart

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

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

What to do Next

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

Acquire Traffic using Customer Orchestrated Source - By creating Monitoring Domain

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

- Install GigaVUE-FM on OpenStack
- Configure Role-Based Access for Third Party Orchestration
- Create Monitoring Domain
- Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

Configure Role-Based Access for Third Party Orchestration

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

Create a Custom User

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

To add users:

1. Go to **Settings** and select **Authentication** > **GigaVUE-FM User Management > Users**.
2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
   - **Name:** Actual name of the user
   - **Username**: User name configured in GigaVUE-FM
   - **Email**: Email ID of the user
   - **Password/Confirm Password**: Password for the user. For more details, refer to Change Your Password.
   - **User Group:** Select the desired User Group to associate the user.

   GigaVUE-FM prompts for your password.
4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

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

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

To create a role:

1. Go to **Settings** and select **Authentication> GigaVUE-FM User Management >Roles**.
2. Select **New Role**.
3. In the New Role page, select or enter the following details:
   - **Role Name**: Name of the role.
   - **Description**: Description of the role.

- **Select Permission**: From the **Select Permissions** tab, select **Third Party Orchestration**, and provide write permissions.

4. Select **Apply** to save the configuration.

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

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

To create a new user group:

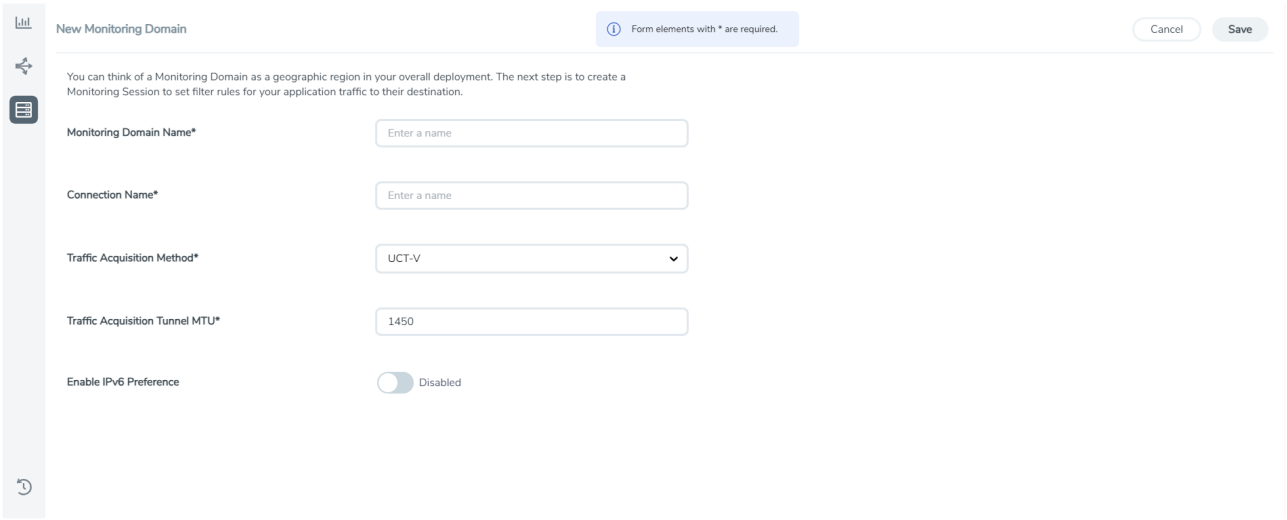
1. Go to **Settings**, and then select **Authentication> GigaVUE-FMUser Management >User Groups**.
2. Select **New Group.**
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:
5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Linux UCT-V Installation

Windows UCT-V Installation

Create Monitoring Domain

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



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

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

4. Click **Save**.

Notes:

- Ensure that all V Series Nodes within a single Monitoring Domain are running the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.

- Similarly, when upgrading a V Series Node, ensure that the GigaVUE-FM version is the same or higher than the V Series Node version.

You can perform the following actions in the Monitoring domain page:

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

NOTE: To view and manage the generated sysdump files, select the GigaVUE V SeriesNode and select the **Sysdump** tab in the lower pane.

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

Configure and Deploy GigaVUE Fabric Components using OpenStack - Generic Mode

This section provides information on how to register GigaVUE fabric components using OpenStack or a configuration file.

Recommended Instance Type

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

Fabric Component	Machine type
GigaVUE V Series Node	m1.medium
GigaVUE V Series Proxy	m1.small
UCT-V Controller	m1.small

Deployment Best Practices

Keep in mind the following when deploying the fabric components using generic mode:

- Set the correct MTU for tunnels: Use 1450 for IPv4 tunnels and 1430 for IPv6. To update the Traffic Acquisition Tunnel MTU:

  1. Select the monitoring domain.
  2. Select the **Edit Monitoring Domain** option.
  3. Enter the **Traffic Acquisition Tunnel MTU** value and select **Save**.

  NOTE: Set the correct MTU before deploying the Monitoring Session. If not, you must undeploy and redeploy the session after updating the MTU.

- You can also create a monitoring domain under Third Party Orchestration and provide the monitoring domain name and the connection name as groupName and subGroupName in the registration data. For details, refer to Create Monitoring Domain.
- GigaVUE V Series Node must have a minimum of two Networks Interfaces (NIC) attached to it, a management NIC and a data NIC. You can add both these interfaces when deploying the GigaVUE V Series Node in OpenStack.
- Create tokens in the **User Management** page in GigaVUE-FM. For details, refer to Configure Tokens for Third Party Orchestration.
- 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 details, refer to the *Configure Role-Based Access for Third-Party Orchestration* section in the 6.9 Documentation.

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

Configure UCT-V Controller in OpenStack

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

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

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

Register UCT-V Controller during Instance Launch

In your OpenStack dashboard, you can launch the UCT-V Controller and register UCT-V Controller using Customization Script.

perform the following steps:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The UCT-V Controller uses this registration data to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM>
        sourceIP: <IP address of UCT-V> (Optional. Required only in multi-NIC deployments to
specify which UCT-V interface to use for communicating with the UCT-V Controller.)
        remotePort: 443
```

The UCT-V Controller deployed in OpenStack appears on the Monitoring Domain page of GigaVUE-FM.

Register UCT-V Controller after Instance Launch

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

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

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

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

3. Restart the UCT-V Controller service.

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

Assign Static IP address for UCT-V Controller

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

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the UCT-V Controller service.

   $ sudo service uctv-cntlr restart

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

NOTE: When you deploy GigaVUE V Series Nodes or UCT-V Controllers using Third Party orchestration, you cannot delete the monitoring domain without unregistering the V Series Nodes or UCT-V Controllers.

Configure UCT-V in OpenStack

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

To register UCT-V using a configuration file:

1. Install the UCT-V in 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.
3. Edit the local configuration file and enter the following Customization Script.

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

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

   NOTE: 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-agent restart
     ```
   - Windows platform: Restart from the Task Manager.

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

Configure GigaVUE V Series Nodes and V Series Proxy in OpenStack

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

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

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

Register V Series Nodes or V Series Proxy during Instance Launch

To register V Series nodes or proxy using the Customization Script in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.

2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```



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

Register V Series Node or V Series Proxy after Instance Launch

To register V Series node or proxy using a configuration file:

1. Log in to the GigaVUE V Series Node or Proxy.
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.

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

> **NOTE:** You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **`remotePort`** value as 443 and the **`remoteIP`** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using V Series proxy then, enter the **`remotePort`** value as 8891 and **`remoteIP`** as <IP address of the Proxy>. If the V Series Node registration fails with a Connection Refused error, set the remoteIP field to the API network IP of GigaVUE-FM instead of the floating IP.

3. Restart the GigaVUE V Series Node or Proxy service.
   - GigaVUE V Series Node:
     `$ sudo service vseries-node restart`
   - GigaVUE V Series Proxy:
     `$ sudo service vps restart`

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

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

Assign Static IP address for GigaVUE V Series

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the GigaVUE V Series service.

   $ sudo service vseries-node restart

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

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 > OpenStack**.
   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 ingress and Egress REPs**:
   a. In the canvas, from the **New** expand menu, drag and drop New Tunnel to the graphical workspace.
   b. On the new tunnel icon, select the ⋮ menu button and select **Details**. The **Add Tunnel Spec** quick view page appears.
3. **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.
4. **Deploy the Monitoring Session**:
   a. From the **Actions** menu, select **Deploy**.
   b. After successful deployment on all the V Series Nodes, the status appears as **Success** on the Monitoring Sessions **Sources** tab.

What to do Next

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

Deployment GigaVUE Fabric Components using OpenStack - Integrated Mode

This section provides detailed information on the multiple ways in which you can configure GigaVUE Cloud Suite for OpenStack to provide visibility for physical and virtual traffic. Using three different options, you can configure GigaVUE Cloud Suite for OpenStack based on the traffic acquisition method and the method in which you want to deploy fabric components. For details, refer to the Deployment Prerequisites.

You can use your own OpenStack Orchestrator to deploy the GigaVUE fabric components instead of using GigaVUE-FM to deploy your fabric components.

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.

For more detailed information and the work flow, refer to the following topics:

- Deployment GigaVUE Fabric Components using OpenStack -Integrated Mode
  - Deployment GigaVUE Fabric Components using OpenStack -Integrated Mode
  - Deployment GigaVUE Fabric Components using OpenStack -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 OpenStack. Refer to the following topics for instructions on configuring traffic acquisition, processing, and forwarding to your desired destination.

- Install GigaVUE-FM on OpenStack
- Install UCT-Vs
- Create Monitoring Domain
- Configure GigaVUE Fabric Components in GigaVUE-FM
- Configure and Deploy Monitoring Session

Install GigaVUE-FM on OpenStack

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

Refer to Install GigaVUE-FM on OpenStack for steps to upload fabric images, install GigaVUE-FM on OpenStack and the steps to start GigaVUE-FM instance and configure it.

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. For details, refer to Network Firewall Requirement.

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

After you deploy the UCT-V VM, you can also use the following commands to add the network firewall rule:

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

Linux UCT-V Installation Scenarios

1. **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.
2. **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.
3. **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

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**6.14.02** 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-6.14.02-amd64.deb
    $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
    ```

2. **To install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS:**
    a. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
    $ sudo rpm -i gigamon-gigavue-uctv-6.14.02-x86_64.rpm
    ```

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

```
sudo uctv-wizard
```

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

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

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

Options	Use Command	Description
		If you wish to skip the prompts to add the required firewall policy, enter your option as **y**. The console interface adds the firewall rules automatically. Enter **N** if you wish to configure manually. For details, refer to Install Linux UCT-Vs using Manual Configuration section.
uninstall	sudo uctv-wizard uninstall	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
  ```

Install Linux UCT-Vs using Manual Configuration

- Install UCT-V from Ubuntu/Debian Package
- Install UCT-V from RPM, Red Hat Enterprise Linux, and CentOS

Install UCT-V from Ubuntu/Debian Package

To install from a Debian package:

1. Download the UCT-V**6.14.02** Debian (.deb) 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-6.14.02-amd64.deb
   $ sudo dpkg -i gigamon-gigavue-uctv-6.14.02-amd64.deb
   ```

4. Modify the file **/etc/uctv/uctv.conf** to configure and register the source and destination interfaces.
   The following examples registers eth0 as the mirror source for both ingress and egress traffic and eth1 as the destination for this traffic:

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

1. Download the UCT-V**6.14.02** 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-6.14.02-x86_64.rpm
$ sudo rpm -i gigamon-gigavue-uctv-6.14.02-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

After installing the UCT-V package, 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:

- **GigaVUE-FM Orchestration**: Perform the following steps:
  a. Log in to the UCT-V.
  b. Create a local configuration file and enter the following user data. **/etc/gigamon-cloud.conf** is the local configuration file in the Linux platform.

```
Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        remoteIP: <IP address of the UCT-V Controller>
        remotePort: 443
        token: <Token>
```

  c. Restart the UCT-V service.
  Linux platform:

```
$ sudo service uctv restart
```

  d. Verify UCT-V is running:

```
sudo service uctv status
```

For more details on how to create tokens, refer to .

Verify that the service status is Active (running). After the UCT-V VM successfully connects to GigaVUE-FM, it appears in **Inventory** > **Resources** > **UCT-V** in the GigaVUE-FM.

- **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 toGigaVUE-FM using Third Party Orchestration with the following two modes:
    1. Generic Mode - Deploy GigaVUE Fabric Components using Generic Mode section in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration
    2. 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

Rules and Notes

- VXLAN is the only tunnel type supported for Windows UCT-V.
- Loopback interface is not supported on 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 workflows.

Windows UCT-V Installation Scenarios

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

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

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 **`6.14.02`** 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
```

Install Windows UCT-Vs using Manual Configuration

1. Download the Windows UCT-V**6.14.02** 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 Monitoring Domain

To create a monitoring domain in GigaVUE-FM:

1. Go to **Inventory > VIRTUAL > OpenStack**. The Monitoring Domain page appears.
2. On the Monitoring Domain page, click **New**. The **Monitoring Domain Configuration** page appears.

3. Enter or select the appropriate information to configure Monitoring Domain for OpenStack. Refer to the following table for field-level details.

NOTE: For the URL, User Domain Name, Project Domain Name, and Region field values, refer to the RC file downloaded from your OpenStack dashboard.

Field	Description
Monitoring Domain	A name for the monitoring domain. **Note**: You can only view and delete the existing configuration for V Series node 1. You cannot create and perform any other actions on the existing configuration for GigaVUE V Series node 1 as the features are deprecated from GigaVUE-FM fabric manager.
Alias	An alias used to identify the monitoring domain.
URL	The authentication URL is the Keystone URL of the OpenStack cloud. This IP address must be DNS resolvable. Refer to the OpenStack User Manual for more information on retrieving the authentication URL from the OpenStack. **Note**: To ensure the validity of OpenStack certificates issued by a trusted Certificate Authority (CA), you must enable the Trust Store. Refer to the Trust Store section in GigaVUE Administration Guide for more detailed information.
User Domain Name	The domain name of your OpenStack authentication domain. **Note**: • If you are using a separate domain for AUTH, enter that domain name as User Domain Name. • If you are not using a separate domain, you can use the same domain for User and Project Domain Name.
Project Domain Name	The domain name of your OpenStack project.
Project Name	The name of the project used for OpenStack authentication.
Region	The region where the Project resides. You can find your region by running one of these commands, depending on your OpenStack version. **keystone endpoint-list** or **openstack endpoint list** or looking at the RC file in OpenStack to view your credentials.
Username	The username used to connect to your OpenStack cloud. **Note**: If you are using OVS mirroring, you must belong to a role that meets the OpenStack minimum requirements for OVS Mirroring. Refer to OVS Mirroring Prerequisites for more information.
Password	The password of your OpenStack cloud.
Traffic	Select the type of component used to capture traffic for monitoring:

Field	Description
Acquisition Method	• **UCT-V:** If you select UCT-V as the tapping method, the traffic is acquired from the UCT-Vs installed on the VMs. You must configure the UCT-V Controller to monitor the UCT-Vs.
Projects to Monitor (Only for OVS Mirroring traffic acquisition method)	This field only appears for OVS Mirroring traffic acquisition method. • Click the **Get Project List** to view the list of projects. **Note**: The **Get Project List** button will only work if all the OpenStack credentials have been provided. Refer to OVS Mirroring Prerequisites. • Select projects that you want to monitor from the list. • You can click **Select None** to clear existing selections or **Select All** to add all available projects to the connection configuration.
Traffic Acquisition Tunnel MTU (Maximum Transmission Unit)	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 Cloud Suite V Series node. • For GRE, the default value is 1450. • For VXLAN, the default value is 1400. However, the UCT-V tunnel MTU should be 50 bytes less than the default MTU size.

4. Click **Save**. The **OpenStack Fabric Launch Configuration** page appears. Refer to Configure GigaVUE Cloud Suite for OpenStack Components for detailed information.

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.
- If GigaVUE-FM fails to connect to OpenStack, an error message is displayed specifying the cause of failure. The connection status is also displayed in Audit Logs, refer to About Audit Logs for more information.

Configure GigaVUE Fabric Components in GigaVUE-FM

You can use your own OpenStack orchestration system to deploy GigaVUE fabric nodes and use GigaVUE-FM to configure the advanced features supported by these nodes. These nodes register themselves with GigaVUE-FM using the information provided by your OpenStack orchestration system. Once the nodes are registered with GigaVUE-FM, you can configure monitoring sessions and related services in GigaVUE-FM. Health status of the registered nodes are determined by the heartbeat messages sent from the respective nodes.

Recommended Instance Type

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

Fabric Component	Machine type
GigaVUE V Series Node	m1.medium
GigaVUE V Series Proxy	m1.small
UCT-V Controller	m1.small

Keep in mind the following when deploying the fabric components using third party orchestration in integrated mode:

- Traffic Acquisition Tunnel MTU is set to the default value of 1500. To edit the Traffic Acquisition Tunnel MTU, select the monitoring domain and click on the **Edit Monitoring Domain** option. Enter the **Traffic Acquisition Tunnel MTU** and click Save.
- When you deploy the fabric components using 3rd party orchestration, you cannot delete the monitoring domain without unregistering the registered fabric components.
- GigaVUE V Series Node must have a minimum of two Networks Interfaces (NIC) attached to it, a management NIC and a data NIC. You can add both these interfaces when deploying the GigaVUE V Series Node in OpenStack.
- Create tokens in the **User Management** page in GigaVUE-FM. Refer to Configure Tokens 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.

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

- Configure V Series Nodes and Proxy in OpenStack
- Configure UCT-V Controller in OpenStack
- Configure UCT-V in OpenStack

Configure V Series Nodes and Proxy in OpenStack

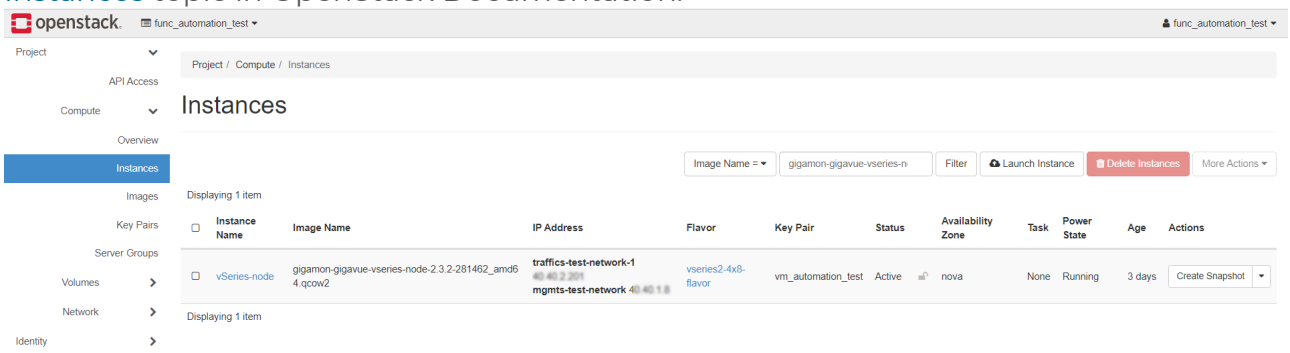
To configure V Series Nodes and V Series Proxy in OpenStack platform:

1. Before configuring GigaVUE fabric components through OpenStack, you must create a monitoring domain in GigaVUE-FM. Refer to Create Monitoring Domain for detailed instructions.
2. In the **Monitoring Domain Configuration** page, select **No** for the **Use FM to Launch Fabric** field as you are going to configure the fabric components in OpenStack Orchestrator.
3. In your OpenStack environment, you can deploy V Series nodes or V Series proxy using the following methods:
   - Register V Series Nodes or V Series Proxy using OpenStack GUI
   - Register V Series Node or V Series Proxy using a configuration file

Register V Series Nodes or V Series Proxy using OpenStack GUI

To register V Series nodes or proxy using the user data in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, click **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.



2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteIP: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```

- You can register your V Series node directly with GigaVUE-FM or you can use V Series Proxy to register your V Series node with GigaVUE-FM. If you wish to register V Series Node directly, enter the **`remotePort`** value as 443 or if you wish to deploy V Series Node using V Series Proxy then, enter the **`remotePort`** value as 8891.

Register V Series Node or V Series Proxy using a configuration file

To register V Series node or proxy using a configuration file:

1. Log in to the V Series Node or proxy.
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 V Series Node or V Series Proxy.

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

NOTE: If you wish to register V Series node using V Series proxy then, enter the **`remotePort`** value as 8891.

3. Restart the V Series node or proxy service.
   - V Series node:
     `$ sudo service vseries-node restart`
   - V Series proxy:
     `$ sudo service vps restart`

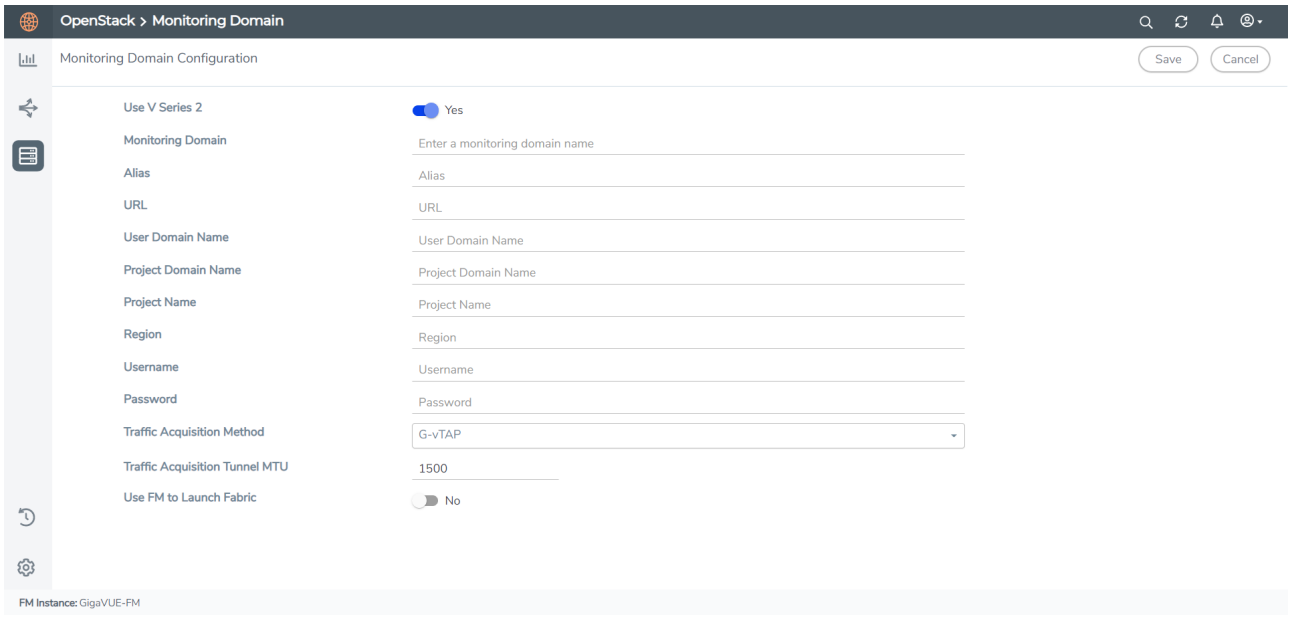
   The deployed V Series node or V Series proxy registers with the GigaVUE-FM. After successful registration the V Series node or proxy sends heartbeat messages to GigaVUE-FM every 30 seconds. If one heartbeat is missing ,the visibility node status appears as 'Unhealthy'. If more than five heartbeats fail to reach GigaVUE-FM, GigaVUE-FM tries to reach the V Series node or proxy and if that fails as well then GigaVUE-FM unregisters the V Series node or proxy and it will be removed from GigaVUE-FM.

Configure UCT-V Controller in OpenStack

To configure GigaVUE fabric components in OpenStack platform:

1. Before configuring GigaVUE fabric components through OpenStack, you must create a monitoring domain in GigaVUE-FM. While creating the monitoring domain, select **UCT-V** as the Traffic Acquisition Method. Refer to Create Monitoring Domain for detailed instructions.

2. In the **Monitoring Domain Configuration** page, select **No** for the **Use FM to Launch Fabric** field as you are going to configure the fabric components in OpenStack Dashboard.

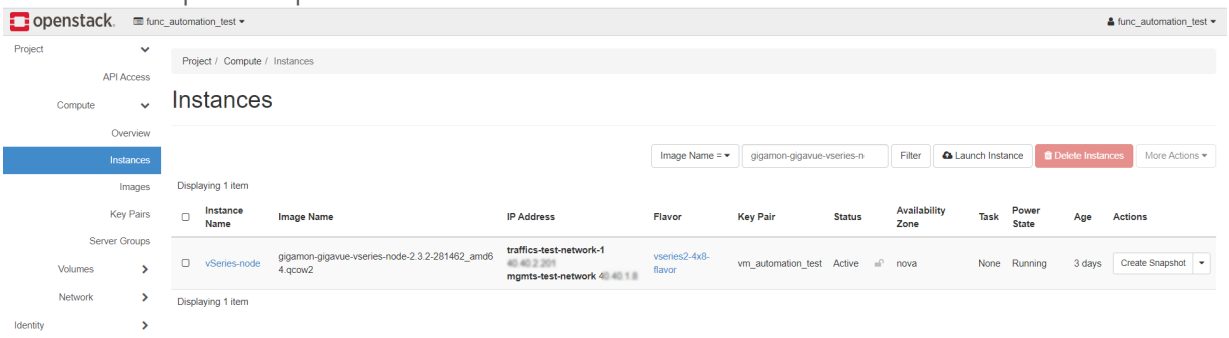


3. In your OpenStack environment, launch the UCT-V Controller using any of the following methods:
   - Register Linux UCT-V
   - Register UCT-V Controller using a configuration file

Register UCT-V Controller using OpenStack GUI

To register UCT-V Controller using the user data in OpenStack GUI:

a. On the Instance page of OpenStack dashboard, click **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.



b. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The UCT-V Controller uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config write_files: - path: /etc/gigamon-cloud.conf owner: root:root
permissions: '0644' content: |Registration: groupName: <Monitoring Domain
Name> subGroupName: <Connection Name> token: <Token> remoteIP: <IP address of
the GigaVUE-FM> sourceIP: <IP address of UCT-V Controller> (Optional Field)
remotePort: 443
```

The UCT-V Controller deployed in OpenStack appears on the Monitoring Domain page of GigaVUE-FM.

Register UCT-V Controller using a configuration file

To register UCT-V Controller using a configuration file:

a. Log in to the UCT-V Controller.

b. Create a local configuration file (**/etc/gigamon-cloud.conf**) and enter the following user data.

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

c. 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 OpenStack

UCT-V should be registered via the registered UCT-V Controller and communicates through PORT 8891.

NOTE: 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.

To register UCT-V using a configuration file:

1. Install the UCT-V in 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.

3. Create a local configuration file and enter the following user data.

   - **/etc/gigamon-cloud.conf** is the local configuration file in Linux platform.
   - **C:\ProgramData\uctv\gigamon-cloud.conf** is the local configuration file in Windows platform.

```
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)
    remotePort: 8891
```

   - 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.

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

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, GigaVUE-FM tries to reach the UCT-V and if that fails as well then GigaVUE-FM unregisters the UCT-V and it will be removed from GigaVUE-FM.

Key Considerations for Upgrade

Consider the following when upgrading the GigaVUE-FM to 6.1.00 or higher version (when using third party orchestration to deploy fabric components):

When upgrading GigaVUE-FM to any version higher than 6.0.00, if the GigaVUE V Series Nodes version deployed in that GigaVUE-FM is lower than or equal to 6.0.00, then, for the seamless flow of traffic, GigaVUE-FM automatically creates **Users** and **Roles** in GigaVUE-FM with the required permission.

For the user created in GigaVUE-FM,

- Username: **orchestration**
- Password: **orchestration123A!**.

  Verify that the username **orchestration** does not belong to another user in GigaVUE-FM.

After the upgrade, we recommend changing the password on the Users page. For details, refer to Configure Role-Based Access for Third Party Orchestration.

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

What to do Next

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

Deploy GigaVUE Cloud Suite using Terraform

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

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

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

License

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

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

Supportability

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

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

Supported Deployment Models

GigaVUE-FM Orchestration

- VMware ESXi-Traffic acquisition using Platform Mirroring.

Third-Party Orchestration (Generic Mode)

Supported platforms

- AWS
- Azure
- Google Cloud Platform (GCP)
- Nutanix
- OpenStack
- VMware ESXi

- NSX-T Federation
- Oracle Cloud Infrastructure (OCI)

Supported traffic acquisition methods

- UCT-V
- Customer-Orchestrated Source

Supported GigaSMART Applications

The Gigamon Terraform provider supports the following GigaSMART applications:

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

Rules and notes

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

Prerequisites

Before starting a Terraform deployment, ensure the following:

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

Get Started

Install Terraform

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

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

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

```
terraform version
```

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

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

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

Create a working directory

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

- which resources are available
- which arguments each resource requires

- which values each argument accepts
- how each resource is used in Terraform

Refer to: Gigamon Terraform provider documentation

Deployment Using Terraform

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

Third-Party Orchestration

- Acquire traffic using UCT-V
- Acquire traffic using Customer-Orchestrated Source

Deploy Third Party Orchestration by using Terraform

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

Supported platforms (Generic Mode):

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

This deployment supports the following traffic acquisition methods:

- UCT-V
- Customer-Orchestrated Source

Before you begin

Complete the following setup before starting the Terraform deployment.

1. **Install GigaVUE-FM.**

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

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

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

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

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

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

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

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

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

The Gigamon provider is available at the Terraform Registry.

Install the Gigamon Terraform Provider Without Internet Access

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

Step 1: Download and install the provider binary

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

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

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

```
mkdir -p ~/.terraform.d/plugins/local/gigamon/gigamon/6.14.0/linux_amd64
```

3. **Extract the zip file:**

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

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

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

5. **Make the binary executable:**

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

Step 2: Update the provider block

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

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

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

Step 3: Continue with the standard deployment flow

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

Prepare the Third Party Orchestration Deployment Files

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

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

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

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

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

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

1. Single-file configurationModule-based configuration

Single-file configuration

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

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

1. Provider block

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

   Example:

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

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

2. Resource blocks

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

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

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

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

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

   Refer to the Terraform Registry documentation for the resource type names, required arguments, and argument descriptions for each resource.

Module-based configuration

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

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

Example main.tf - calling modules

```
module "monitoring_domain" {
```

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

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

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

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

Refer to the Terraform provider documentation for the full list of resources and arguments for each module.

Run the Terraform Deployment File

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

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

   ```
   terraform init
   ```

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

2. Optional: Validate the configuration.

   - Run:

     ```
     terraform validate
     ```

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

3. Review the deployment plan.
   - Run:
     ```
     terraform plan
     ```
   - Review the output carefully. Confirm that Terraform will create or update the expected third-party objects.
   - Do not continue until the plan matches the deployment that you expect.
4. Apply the configuration.
   - Run:
     ```
     terraform apply
     ```
   - When prompted, confirm the deployment.
   - Terraform connects to GigaVUE-FM, creates or updates the third-party monitoring domain and connection, and applies the monitoring session configuration defined in the deployment files.

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

Update an existing Third Party Orchestration deployment

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

```
terraform plan
terraform apply
```

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

Destroy a Third-Party Orchestration Deployment

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

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

UCT-V agents or V Series nodes remain registered. Terraform will destroy the monitoring session, certificate application, CA, and SSL key or certificate in the first run, but will fail when it attempts to delete the connection and monitoring domain.

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

1. **Run the first destroy**
   - Run the following command from the working directory:

     ```
     terraform destroy
     ```

   - Terraform removes the monitoring session in this run:

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

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

     ```
     terraform destroy
     ```

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

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

Verify the deployment

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

1. **Verify the monitoring domain and connection**

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

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

2. **Verify node and tunnel health**

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

3. **Verify the monitoring session configuration**

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

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

4. **Verify traffic flow**

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

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

OpenTofu Support

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

OpenTofu Registry

To use OpenTofu instead of Terraform:

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

For the Gigamon provider, refer to the OpenTofu Registry to download the provider package compatible with OpenTofu.

Best Practices

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

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

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

     Example: main.tf - backend and provider configuration

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

     Example: backend.hcl - backend connection values

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

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

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

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

Verification and Troubleshooting

Verify the deployment

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

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

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

Troubleshoot the deployment

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

Provider issues

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

Authentication or connectivity issues

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

Configuration issues

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

Deployment timeout issues

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

State file issues

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

Recover after a failed deployment

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

Configure

Configure UCT-V Features

Refer to the following sections for more detailed information:

- Configure Precryption in UCT-V
- Create Precryption Template for UCT-V
- Create Prefiltering Policy Template
- Configure Secure Tunnel (OpenStack)

Configure Precryption in UCT-V

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

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

Rules and Notes

- To avoid packet fragmentation, change the option precryption-path-mtu in the UCT-V configuration file (**/etc/uctv/uctv.conf**) within the range 1400-9000 based on the platform path MTU.

- Protocol version IPv4 and IPv6 are supported.
- Using IPv6 tunnels requires GigaVUE-FM and the fabric components version 6.6.00 or above.

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

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

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

Validate Precryption connection

To validate the Precryption connection, follow these steps:

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

Limitations

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

Create Precryption Template for UCT-V

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

Rules and Notes:

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

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

- Create Precryption Template for Filtering based on Applications
- Create Precryption Template for Filtering based on L3-L4 details

Create Precryption Template for Filtering based on Applications

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

To create,

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

   The **Precryption Policies** page appear

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

   The New Precryption Template page appears.

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

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

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

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

Create Precryption Template for Filtering based on L3-L4 details

To create,

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

3. Perform the following steps:

   a. In the **Template** field, enter a name for the template.

   b. In the **Rule Name** field, enter a name for the rule.

   c. For **Action**, select one of the following options:

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

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

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

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

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

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

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

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

- L3
- L4

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

For L3, perform the following:

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

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

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

- Not Equal to
- Equal to

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

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

For L4, perform the following:

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

- Source Port
- Destination Port

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

- Not Equal to
- Equal to

iii. Enter the source or destination port value.

4. Select **Save**.

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

The template is successfully created. To enable Precryption, refer to Configure Monitoring Session Options (OpenStack) section.

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

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

Create Prefiltering Policy Template

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

To create a prefiltering policy template,

1. Go to **Traffic > Resources > Prefiltering** and select **UCT-V**.

2. Select **New**.

3. In the **Template Name** field, enter the name of the template,

4. In the **Rule Name** field, enter the name of a rule.

5. Select one of the following options:

- Pass: Allows the traffic.
- Drop: Blocks the traffic..

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

6. Select one of the following options:

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

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

7. Select a priority value from 1-8.

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

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

8. Select one of the following options as **Filter Type**:
   - L3
   - L4
9. Select one of the following options **Filter Name**:
   - ip4Src
   - ip4Dst
   - ip6Src
   - ip6Dst
   - Proto: Applies toboth ipv4 and ipv6.
10. Select one of the following options for **Filter Relation**:
    - Not Equal to
    - Equal to
11. In the **Value** field, enter the source or destination port.
12. Select **Save**.

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

To enable prefiltering, refer to Monitoring Session Options.

Configure Secure Tunnel (OpenStack)

The Secure tunnels can be configured on:

- Precrypted Traffic
- Mirrored Traffic

Precrypted Traffic

You can send the precrypted traffic through a secure tunnel. When secure tunnels for Precryption is enabled, packets are framed and sent to the TLS socket. The packets are sent in PCAPng format.

When you enable the secure tunnel option for regular and precrypted packets, two TLS secure tunnel sessions are created.

It is recommended always to enable secure tunnels for precrypted traffic to securely transfer the sensitive information.

For more information about PCAPng, refer to PCAPng Application.

Mirrored Traffic

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

Refer to the following sections for Secure Tunnel Configuration:

- Configure Secure Tunnel from UCT-V to GigaVUE V Series Node in UCT-V
- Configure Secure Tunnel between GigaVUE V Series Nodes

Prerequisites

- While creating Secure Tunnel, you must provide the following details:
  - SSH key pair
  - CA certificate
- Port 11443 should be enabled in security group settings. Refer to Security Group for OpenStack for more detailed information on Network Firewall / Security Group.

Notes

- Protocol versions IPv4 and IPv6 are supported.
- If you wish to use IPv6 tunnels, your GigaVUE-FM and the fabric components version must be 6.6.00 or above.
- For UCT-V with a version lower than 6.6.00, if the secure tunnel is enabled in the monitoring session, secure mirror traffic will be transmitted over IPv4, regardless of IPv6 preference.

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

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

S.No	Task	Description
1.	Upload a CA	You must upload a Custom Certificate to UCT-V Controller to establish a connection with the GigaVUE V Series Node. To upload the CA using GigaVUE-FM, follow the steps given below: 1. Go to **Inventory > Resources > Security > CA List**. 2. Select **New** to add a new Custom Authority. The **Add Custom Authority** page appears. 3. Enter or select the following information. • Alias - Alias name of the CA. • File Upload - Choose the certificate from the desired location. 4. Select **Save**. For more information, refer to the section Adding Certificate Authority.
2.	Upload an SSL Key	You must add an SSL key to GigaVUE V Series node. To add an SSL Key, follow the steps in the section Upload SSL Keys.

S.No	Task	Description
3	Enable the secure tunnel	You should enable the secure tunnel feature to establish a connection between the UCT-V and GigaVUE V Series node. To enable the secure tunnel feature follow these steps: **1.** In the Edit Monitoring Session page, click **Options**. The **Monitoring Session options** page appears. **2.** Enable the **Secure Tunnel** button. You can enable secure tunnel for both mirrored and precrypted traffic. **Note**: When GigaVUE V Series is upgraded or deployed to 6.5 version, all the existing monitoring sessions will be redeployed, and individual TLS TEPs are created for each UCT-V agent in GigaVUE V Series node.
4.	Select the SSL Key	You must select the added SSL Key in GigaVUE V Series node Key while creating a monitoring domain configuring the fabric components in GigaVUE-FM. To select the SSL key, follow the steps in the section Configure GigaVUE Fabric Components in GigaVUE-FM. If the existing monitoring domain does not have a SSL key, you can add it by following the given steps: **1.** Select the monitoring domain for which you want to add the SSL key. **2.** Click the **Actions** drop down list and select **Edit SSL Configuration**. An **Edit SSL Configuration** window appears. **3.** Select the CA in the **UCT-V Agent Tunnel CA** drop down list. **4.** Select the SSL key in the **V Series Node SSL key** drop down list. **5.** Click **Save**.
5.	Select the CA	You should select the added Certificate Authority (CA) in UCT-V Controller while creating the monitoring domain configuring the fabric components in GigaVUE-FM. To select the CA certificate, follow the steps in the section Configure GigaVUE Fabric Components in GigaVUE-FM

Configure Secure Tunnel between GigaVUE V Series Nodes

You can create secure tunnel:

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

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

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

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

<table>
<tr><th>S.No</th><th>Task</th><th>Description</th></tr>
<tr><td>1.</td><td>Upload a CA.</td><td>You must upload a CA Certificate to UCT-V Controller to establish a connection between the GigaVUE V Series node.

To upload the CA using GigaVUE-FM follow the steps given below:
1. Go to Inventory > Resources > Security > CA List.
2. Click Add, to add a new Certificate Authority. The Add Certificate Authority page appears.
3. Enter or select the following information.
<table><tr><th>Field</th><th>Action</th></tr><tr><td>Alias</td><td>Alias name of the CA.</td></tr><tr><td>File Upload</td><td>Choose the certificate from the desired location.</td></tr></table>4. Click Save.
5. Click Deploy All.

For more information, refer to the Adding Certificate Authority section.</td></tr>
<tr><td>2.</td><td>Upload an SSL Key.</td><td>You must add an SSL key to GigaVUE V Series node. To add an SSL Key, follow the steps in the section</td></tr>
<tr><td>3</td><td>Create a secure tunnel.</td><td>You should create a secure tunnel to establish a connection between the UCT-V and first GigaVUE V Series Node. To enable the secure tunnel feature follow these steps:
1. In the Edit Monitoring Session page, click Options. The Monitoring Session Options page appears.
2. Enable the Secure Tunnel button. You can enable secure tunnel for both mirrored and precrypted traffic.</td></tr>
<tr><td>4.</td><td>Select the added SSL Key.</td><td>Select the SSL Key added in Step 2, while creating a monitoring domain and configuring the fabric components in GigaVUE-FM for the first GigaVUE V Series Node.
You must select the added SSL Key in the first GigaVUE V Series Node.
To select the SSL key, follow the steps in the section Configure GigaVUE Fabric Components in GigaVUE-FM.</td></tr>
<tr><td>5.</td><td>Select the added CA certificate.</td><td>You should select the added Certificate Authority (CA) in UCT-V Controller while creating the monitoring domain. To select the CA certificate, follow the steps in the section Configure GigaVUE Fabric Components in GigaVUE-FM.</td></tr>
<tr><td>6</td><td>Create an Egress tunnel from the first GigaVUE V Series Node.</td><td>You must create an egress tunnel for traffic to flow out from the first GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session. Refer to Create a Monitoring Session to know about Monitoring Session.

To create the egress tunnel, follow these steps:
1. After creating a new Monitoring Session, or click Actions > Edit on an</td></tr>
</table>

S.No	Task	Description
		existing monitoring session, the GigaVUE-FM canvas appears. 2. In the canvas, select **New** > **New Tunnel**, drag and drop a new tunnel template to the workspace. The **Add** Tunnel Spec quick view appears. 3. On the New Tunnel quick view, enter or select the required information as described in the following table: • Alias - The name of the tunnel endpoint. • Description - The description of the tunnel endpoint. • Type - Select TLS-PCAPNG for creating egress secure tunnel • Traffic Direction - Choose **Out** (Encapsulation) for creating an egress tunnel from the V Series node to the destination. Select or enter the following values: o MTU- The default value is 1500. o Time to Live - Enter the value of the time interval till which the session needs to be available. The value ranges from 1 to 255. The default value is 64. o DSCP - Enter the Differentiated Services Code Point (DSCP) value. o Flow Label - Enter the Flow Label value. o Source L4 Port- Enter the Souce L4 Port value o Destination L4 Port - Enter the Destination L4 Port value. o Flow Label o Cipher- Only SHA 256 is supported. o TLS Version - Select TLS Version1.3. o Selective Acknowledgments - Choose **Enable** to turn on the TCP selective acknowledgments. o SYN Retries - Enter the value for number of times the SYN has to be tried. The value ranges from 1 to 6. o Delay Acknowledgments - Choose **Enable** to turn on delayed acknowledgments. • IP Version - The version of the Internet Protocol. IPv4 and IPv6 are supported. • Remote Tunnel IP - Enter the interface IP address of the second GigaVUE V Series Node (Destination IP). 4. Click **Save**.
7.	Select the added SSL Key in the GigaVUE V Series Node	You must select the added SSL Key while creating a monitoring domain and configuring the fabric components in GigaVUE-FM in the second GigaVUE V Series Node. To select the SSL key, follow the steps in the section Configure GigaVUE Fabric Components in GigaVUE-FM.
8	Create an ingress tunnel in the second GigaVUE V Series node.	You must create a ingress tunnel for traffic to flow in from the first GigaVUE V Series Node with tunnel type as TLS-PCAPNG while creating the Monitoring Session for the second GigaVUE V Series Node. Refer to Create a Monitoring Session to know about monitoring session.

S.No	Task	Description
		To create the ingress tunnel, follow these steps: **1.** After creating a new monitoring session, or click **Actions** > **Edit** on an existing monitoring session, the GigaVUE-FM canvas appears. **2.** In the canvas, select **New** > **New Tunnel**, drag and drop a new tunnel template to the workspace. The **Add** Tunnel Spec quick view appears. **3.** On the New Tunnel quick view, enter or select the required information as described in the following table: • Alias - The name of the tunnel endpoint. • Description - The description of the tunnel endpoint. • Type - Select TLS-PCAPNG for creating egress secure tunnel. **Note**: If you are enabling Secure tunnel in Monitoring Session with traffic acquisition method as UCT-V, you must not create TLS-PCAPNG Tunnel with direction IN, Destination L4 port 11443, and GigaVUE V Series Node version 6.5 and above. • Traffic Direction- Choose **In** (Decapsulation) for creating an ingress tunnel that receives traffic from the first GigaVUE V Series Node. Select or enter the values as described in Step 6: • IP Version - The version of the Internet Protocol. IPv4 and IPv6 are supported. • Remote Tunnel IP - Enter the interface IP address of the first GigaVUE V Series Node (Destination IP). **4.** Click **Save**.

Viewing Status of Secure Tunnel

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

To verify the status of secure tunnel,

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

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

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

Configure Monitoring Session

This chapter describes how to setup ingress and egress tunnel, maps, applications in a monitoring session to receive and send traffic to the GigaVUE 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 (OpenStack)
- Create Ingress and Egress Tunnels (OpenStack)
- Create Raw Endpoint (OpenStack)
- Create a New Map (OpenStack)
- Add Applications to Monitoring Session
- Interface Mapping (OpenStack)
- Deploy Monitoring Session
- View Monitoring Session Statistics
- Visualize the Network Topology (OpenStack)

Create a Monitoring Session (OpenStack)

You must a Create Monitoring Domain before creating a monitoring session.

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 (OpenStack)

The following table outlines the Monitoring Session tabs in OpenStack, each enabling traffic analysis, configuration, and infrastructure visibility.

Tab	Description
Overview	You can view the high level information of the selected Monitoring Session such as, connections, tunnel details, health status, deployment status, and information related to Application Intelligence statistics. You can also view the statistics of the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis. You can filter the statistics based on the elements associated with the Monitoring Session. For more information, refer to View Monitoring Session Statistics.
Sources	Displays the sources and target details monitored by the Monitoring Session. You can view and edit the connection details of the Monitoring Session. You can view the deployment status, number of targets, and targets source health. In the **Selection Status** section, you can view the VM status. The status indicates whether the VM is supported, not supported, selected, or not selected. When you hover over the status, a tooltip displays the reason for that status. **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 template and apply it to the Monitoring Session. Refer to Configure Monitoring Session Options (OpenStack) for more detailed information. **Note**: Traffic Acquisition is only applicable for Monitoring Domain created with UCT-V as Acquisition method.

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

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

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

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

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

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

Configure Monitoring Session Options (OpenStack)

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-d efined 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 User Defined Application.

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.

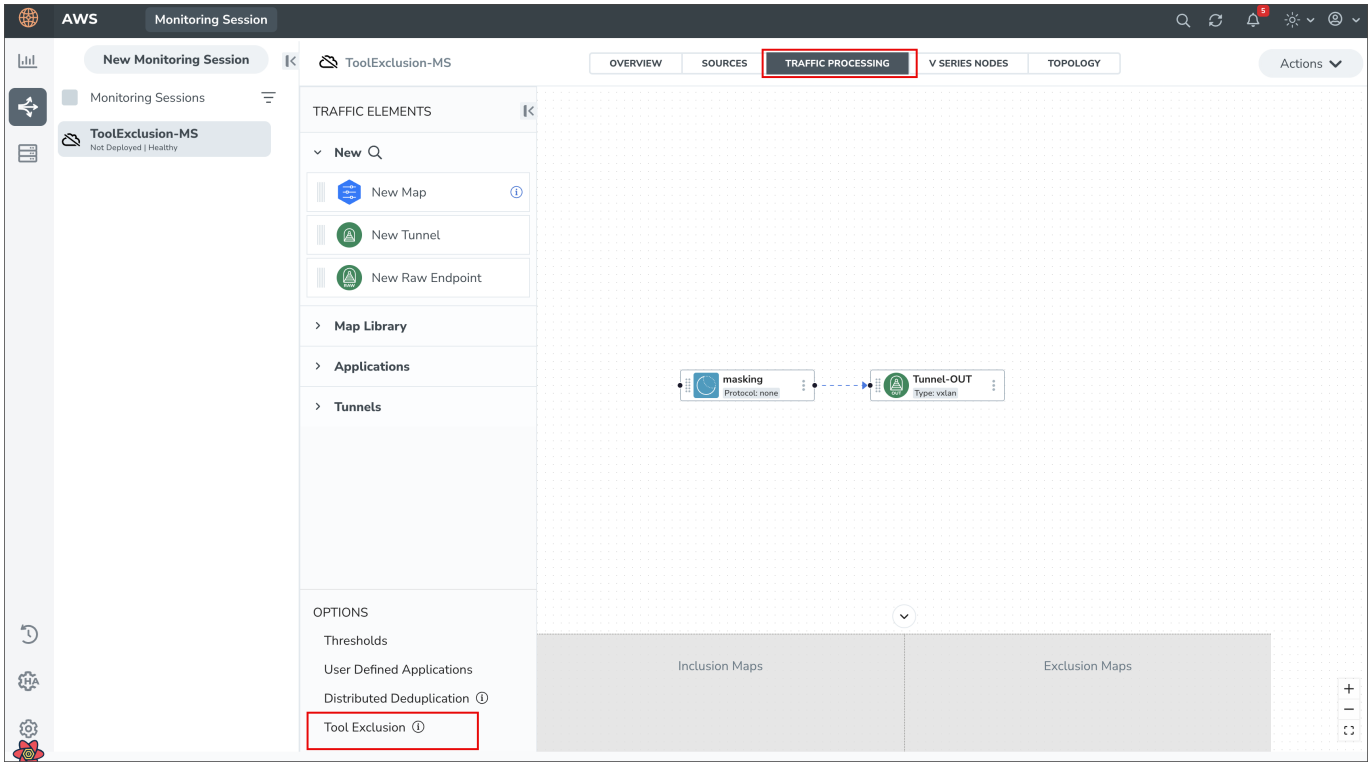
You can exclude tool instances using either of the following methods:

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 (OpenStack)

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, or ERSPAN tunnel.

Create a new tunnel endpoint

To create,

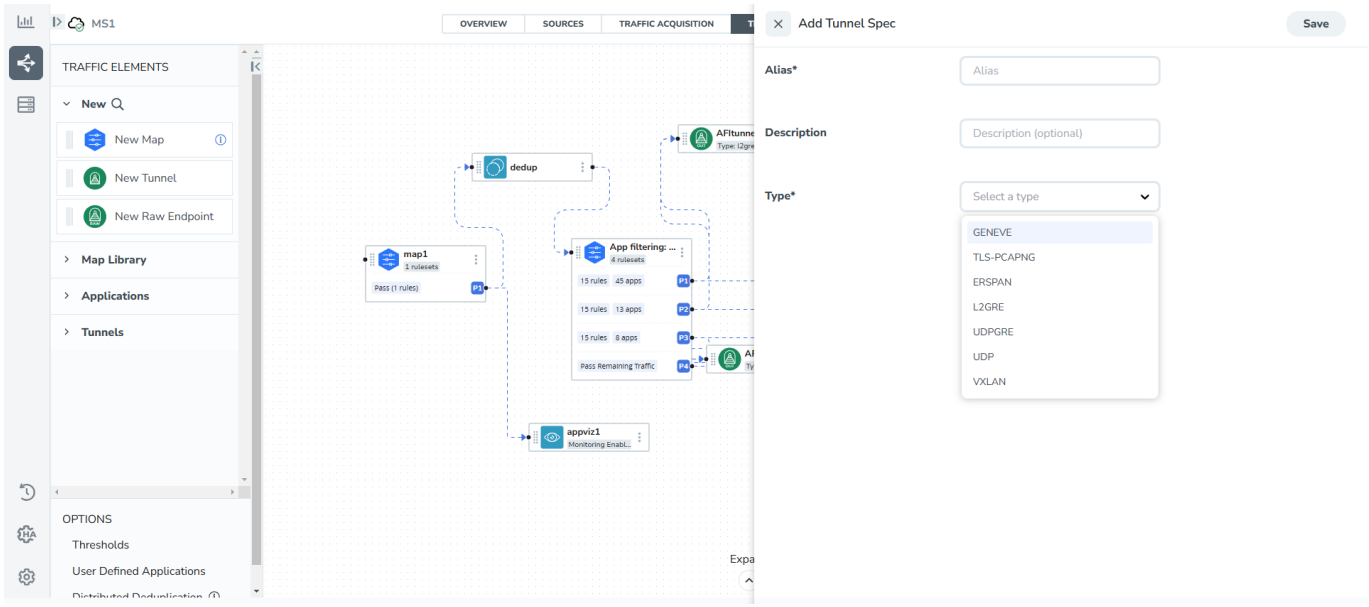
1. Perform one of the following and navigate to the **TRAFFIC PROCESSING** tab:
   - Create a new monitoring session
   - Select **Actions > Edit** on an existing monitoring session.

   The GigaVUE-FM Monitoring Session canvas page appears.

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

   The **Add Tunnel Spec** quick view appears.
4. Enter the **Alias**, **Description**, and **Type** details.

   For details, refer to Details - Add Tunnel Specifications table.
5. Select **Save**.



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

Apply a threshold template to Tunnel End Points

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

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

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

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

Table 1: Details - Add Tunnel Specifications

Field	Description	
Alias	The name of the tunnel endpoint.	
Description	The description of the tunnel endpoint.	
Admin State **Note**: This option appears only after the Monitoring session deployment.	Use this option to send or stop the traffic from GigaVUE-FM to the egress tunnel endpoint. Admin State is enabled by default. You can use this option to stop sending traffic to unreachable or down tools. Each egress tunnel configured on the GigaVUE V SeriesNode has an administrative state that enables GigaVUE-FM to halt the tunnel's traffic flow. GigaVUE-FM only disable the tunnels when it receives a notification via REST API indicating that a tool or group of tools is down. **Note**: This option is not supported for TLS-PCAPNG tunnels.	
Type	The type of the tunnel. Select from the options below to create a tunnel. ERSPAN, L2GRE, VXLAN, TLS-PCAPNG, UDP, or UDPGRE.	
VXLAN		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. **Note**: In the scenario where secure tunnels need to be established between a GigaVUE V Series Node and a GigaVUE HC Series, you can utilize the **Configure Physical Tunnel** option provided in the GigaVUE V Series Secure Tunnel page. This allows you to configure secure tunnels on your physical device conveniently. For details, refer to Secure Tunnels.		
In	Choose **In** (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP/FQDN	For egress tunnels, specify the destination endpoint by using an IPv4 address, IPv6 address, or fully qualified domain name (FQDN). FQDN support applies only to VXLAN and L2GRE egress tunnels. By using an FQDN, you can keep the same tunnel configuration even when the destination IP address changes. If the FQDN resolves to both IPv4 and IPv6 addresses, the system selects the address based on the preferred IP stack configured for the monitoring domain.
	VXLAN Network Identifier	Unique value that is used to identify the VXLAN. The value ranges from 1 to 16777215.
	Source L4 Port	The port used to establish the connection to the target. For example, if A is the source and B is the destination, this port value

Field	Description	
		belongs to A.
	Destination L4 Port	The port used to establish the connection will be established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
Out	Choose **Out** (Encapsulation) for creating an egress tunnel from the GigaVUE V Series Node to the destination endpoint.	
	Remote Tunnel IP	For egress tunnel, the Remote Tunnel IP is the IP address of the tunnel destination endpoint.
	MTU	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry. The default value is 1500.
	Time to Live	Enter the value of the time interval for which the session needs to be available. The value ranges from 1 to 255. The default value is 64.
	DSCP	Differentiated Services Code Point (DSCP) is a value that network devices use to identify traffic to be handled with higher or lower priority. The values ranges from 0 to 63 with 0 being the highest priority and 63 being the lowest priority.
	Flow Label	Unique value, which is used to identify packets that belong to the same flow. A flow is a sequence of packets that need to be treated as a single entity that may require special handling. The accepted value is between 0 and 1048575.
	VXLAN Network Identifier	Unique value which is used to identify the VXLAN. The value ranges from 1 to 16777215.
	Multi Tunnel	Enable the multi-tunnel flag to create multiple tunnels for flow distribution to the 5G-Cloud application. Refer to 5G-Cloud Ericson SCP Support. **Applicable Platforms**: OpenStack, Third Party Orchestration, VMware ESXi **Notes:** • You can configure either a single-tep or multi-tep setup for the egress tunnel. Switching between these configurations is not allowed; to make changes, you must undeploy and redeploy the Monitoring Session. • When you enable Multi-Tunnel on a VXLAN tunnel and set the number of tunnels, GigaVUE-FM automatically creates the additional VXLAN tunnel endpoints. Any later changes to the original VXLAN tunnel, such as disabling Multi-Tunnel or modifying Domain Tagging do not update these auto created endpoints. They continue to retain the configuration

Field	Description	
		that existed at the time they were created. To apply updated settings, you must delete the VXLAN TEP and the associated LB application, then recreate the LB and VXLAN TEP with the new configuration, and re-establish the link between them.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
	Domain Tagging	Enable this option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. **NOTE:** This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
UDPGRE		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	Choose **In** (Decapsulation) for creating an ingress tunnel to carry traffic from the source to the GigaVUE V Series Node.	
	IP Version	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Key	Identifier used to differentiate different UPDGRE/L2GRE tunnels. It routes the encapsulated frames to the appropriate tunnel on the remote endpoint. Enter a value between 0 and 4294967295.
	Source L4 Port	The port from which the connection is established to the target. For example, if A is the source and B is the destination, this port value belongs to A.
	Destination L4 Port	The port to which the connection is established from the source. For example, if A is the source and B is the destination, this port value belongs to B.
L2GRE		
Traffic Direction		

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

Field	Description	
		4294967295.
	Domain Tagging	Enable this option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping. **NOTE:** This setting is available only when Domain Classification is enabled in the associated PCAPng application. Refer to PCAPng Application for details.
ERSPAN		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node.		
In	**IP Version**	The version of the Internet Protocol. Select IPv4 or IPv6.
	Remote Tunnel IP	For ingress tunnel, the Remote Tunnel IP is the IP address of the tunnel source.
	Flow ID	The ERSPAN flow ID is a numerical identifier that distinguishes different ERSPAN sessions or flows. The value ranges from 1 to 1023.
TLS-PCAPNG		
Traffic Direction The direction of the traffic flowing through the GigaVUE V Series Node. **Note**: In the scenario where secure tunnels need to be established between a GigaVUE V Series and a GigaVUE HC Series, you can utilize the **Configure Physical Tunnel** option provided in the GigaVUE V Series Secure Tunnel page. This allows you to conveniently configure secure tunnels on your physical device . For details, refer to Secure Tunnels section.		

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

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

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

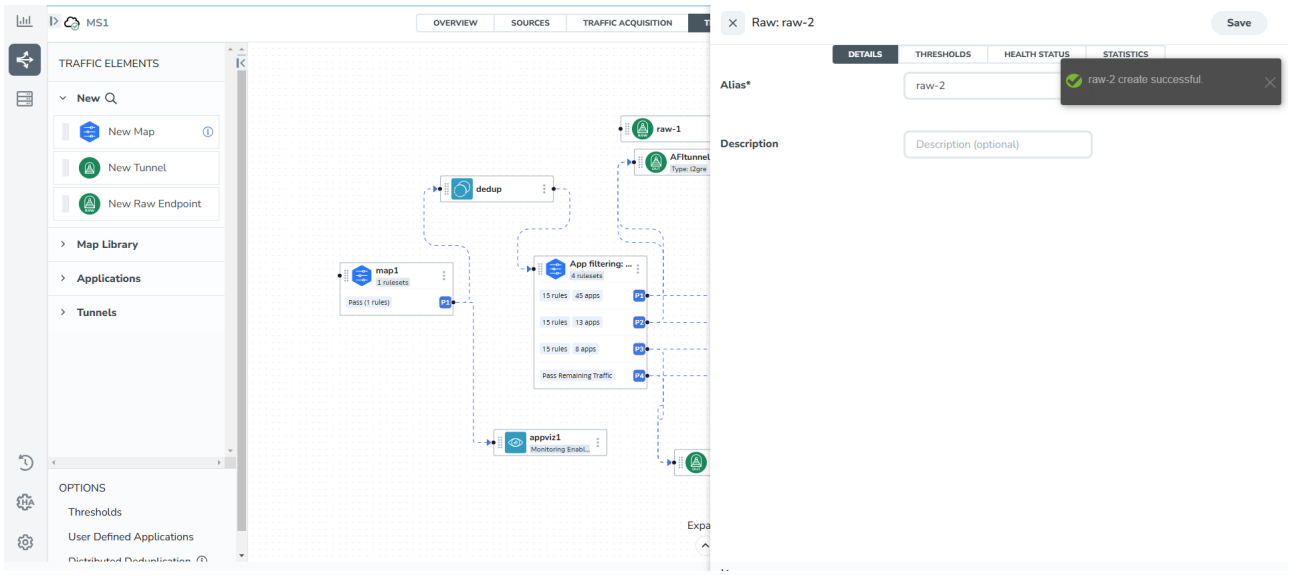
Create Raw Endpoint (OpenStack)

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 (OpenStack)

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 ∩ 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**: • For OpenStack environment, Subnet Name Source and Subnet Name Destination are the exclusion filters available as part of Exclusion Maps with Traffic Acquisition method as OVS Mirroring in the Monitoring Domain. • If no ATS rule filters listed above are used, all VMs and vNICS are selected as targets. When any ATS rule results in a null set, no target is selected and V Series Node does not receive traffic from any VM or vNIC. • 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.

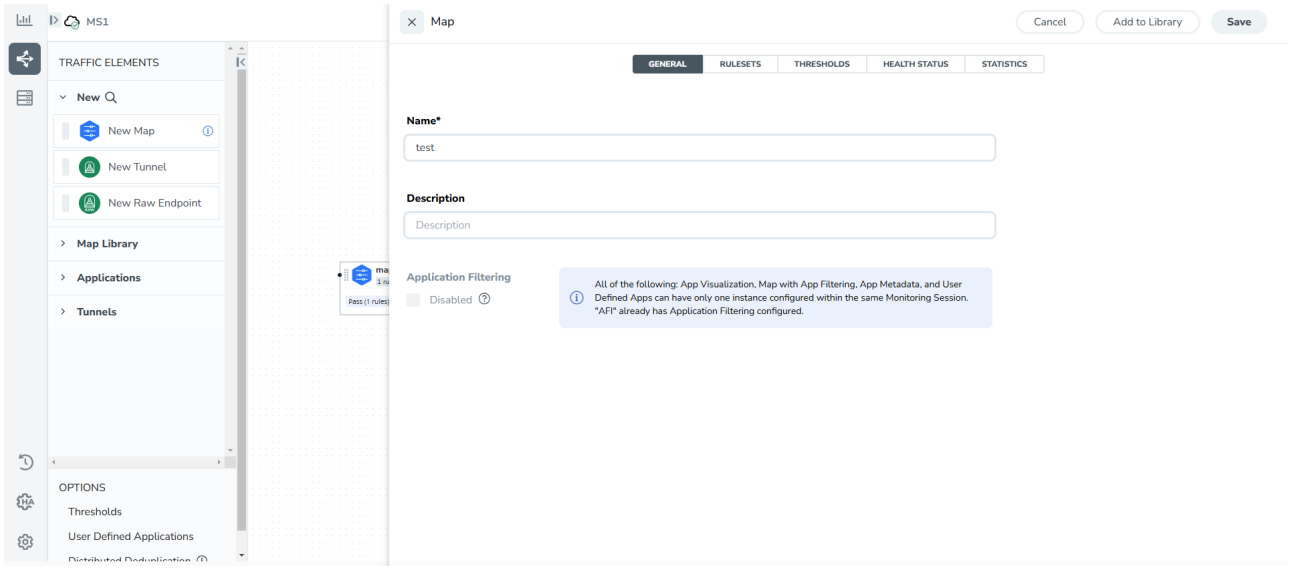
Rules and Notes:

- 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 you are running GigaVUE Cloud Suite on OpenStack, you can add a subnet to the exclusion map. To do this, create an exclusion map and select the Subnet name in the ruleset.
- If a packet is fragmented then all the fragments are destined to the same application end point. You can find the stats of mapped fragmented traffic in GigaVUE-FM.

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.

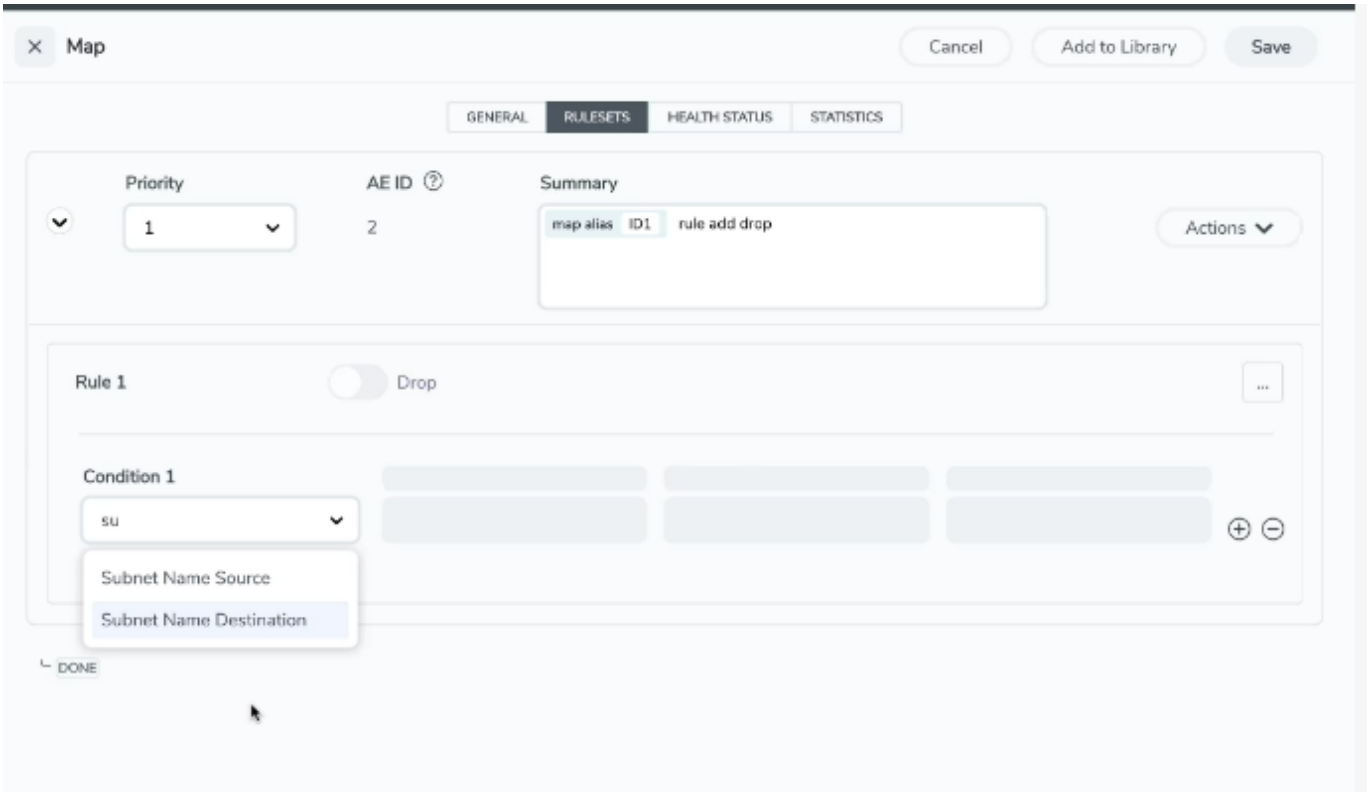
Example- Create a New Map using Inclusion and Exclusion Maps

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.

Starting from software release 6.8 version, to exclude a subnet, a provision to exclude interfaces based on subnet name is added in the Monitoring Domain as part of Exclusion Maps for OpenStack environment with Traffic Acquisition method as OVS mirroring. To add a subnet to the exclusion map, create an exclusion map and select the Subnet name (Subnet Name Source or Subnet Name Destination) in the ruleset.



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 (OpenStack)

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	✓	✓	✓	✓	✓	✓	✓

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
Intelligence							
Application Visualization	✓	✓	✓	✓	✓	✓	✓
De-duplication	✓	✓	✓	✓	✓	✓	✓
L2GRE Tunnel Encapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✓
VXLAN Tunnel Encapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓	✓

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
L2GRE Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✓
VXLAN Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✓	✓	✓	✓	✓	✓
ERSPAN Tunnel Decapsulation	✓	✗	✓	✓	✓	✓	✓

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
Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.							
UDPGRE Tunnel Decapsulation Refer to *Create Ingress and Egress Tunnels* section on the retrospective GigaVUE Cloud Suite Deployment Guide.	✓	✗	✓	✓	✓	✓	✗
GigaSMART NetFlow Generation	✓	✓	✓	✓	✓	✓	✓
GRE-In-UDP Tunnel Decapsulation	✗	✗	✓	✓	✗	✓	✗
GENEVE Decapsulati	✓	✗	✗	✗	✓	✗	✗

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
on							
Header Stripping	✓	✓	✓	✓	✓	✓	✓
Load Balancing	✓	✓	✓	✓	✓	✓	✓
Masking	✓	✓	✓	✓	✓	✓	✓
SSL Decrypt	✓	✓	✓	✓	×	✓	✓
Secure Tunnels Encapsulation	×	✓	×	✓	✓	✓	✓
Secure Tunnels Decapsulation	×	✓	×	✓	✓	✓	✓
Slicing	✓	✓	✓	✓	✓	✓	✓
5G-Cloud Application	×	×	✓	✓	×	✓	×
5G-Service Based Interface Application	×	×	✓	✓	✓	✓	×

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.

<table>
<tr><th>Parameter</th><th>Description</th></tr>
<tr><td colspan="2">Advanced Settings</td></tr>
<tr><td>Number of Flows</td><td>The number of flows supported by the application.
Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms.
<table>
<tr><th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr>
<tr><td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
<tr><td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
</table>

Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.
</td></tr>
<tr><td>Fast Mode</td><td>Enable the Fast Mode option for performance (less CPU cycles and less memory utilization) improvement. When the Fast Mode is enabled, some or all of the attributes of the applications will be disabled. You can view the list of attributes/applications available in the fast mode by navigating to the app editor under AMI feature in the GigaVUE-FM. If all the attributes of the application are disabled then the application itself is disabled.</td></tr>
</table>

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

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

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

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

6. 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.
7. Click the **THRESHOLDS** tab. For more details on how to create and apply threshold template, refer to Traffic Health Monitoring section in the respective cloud deployment guides.

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

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

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

View Application Statistics for Application Filtering

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

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

Application Metadata Exporter

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

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

You can use AMX to:

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

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

- Application Metadata Exporter Deployment Options
- Application Metadata Exporter - Export Scenarios

Application Metadata Exporter Deployment Options

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

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

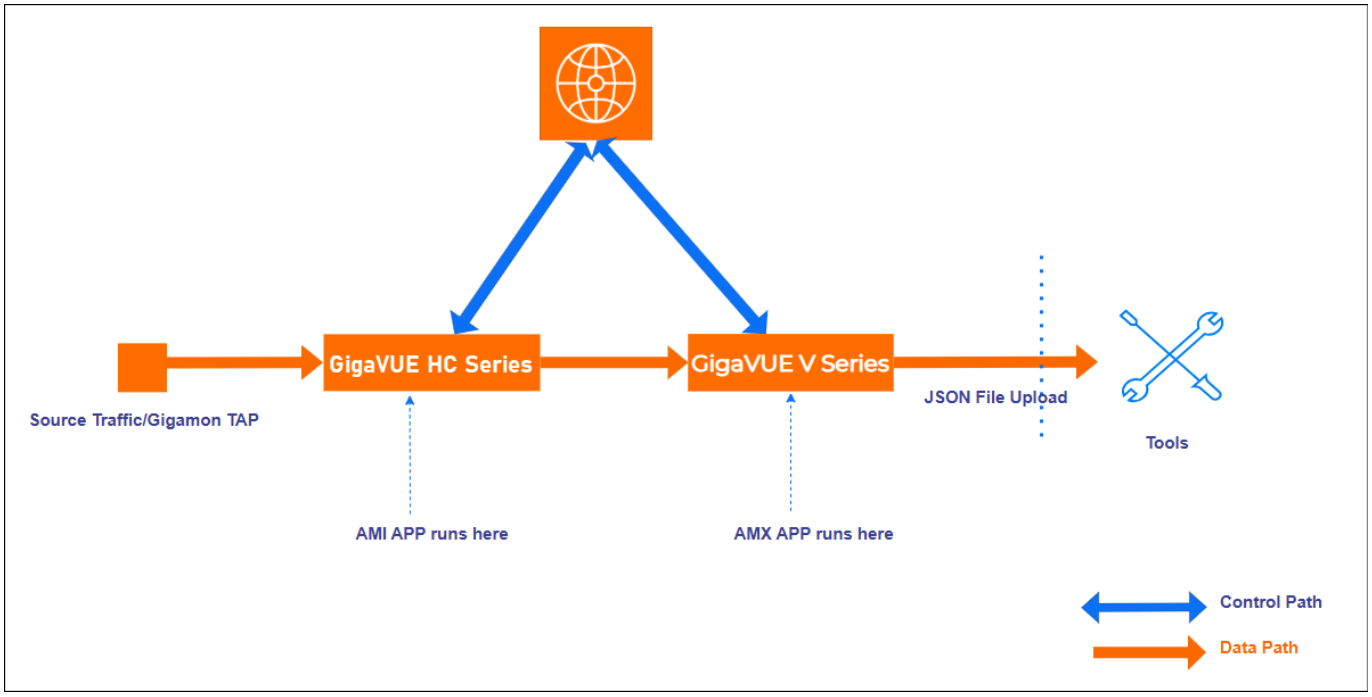
- On-Premises
    - Hardware (AMI)
    - Hardware (Control Plane Metadata)
    - Virtual (VMware, Nutanix, KVM)
- Public Cloud

On-Premises

Hardware (AMI)

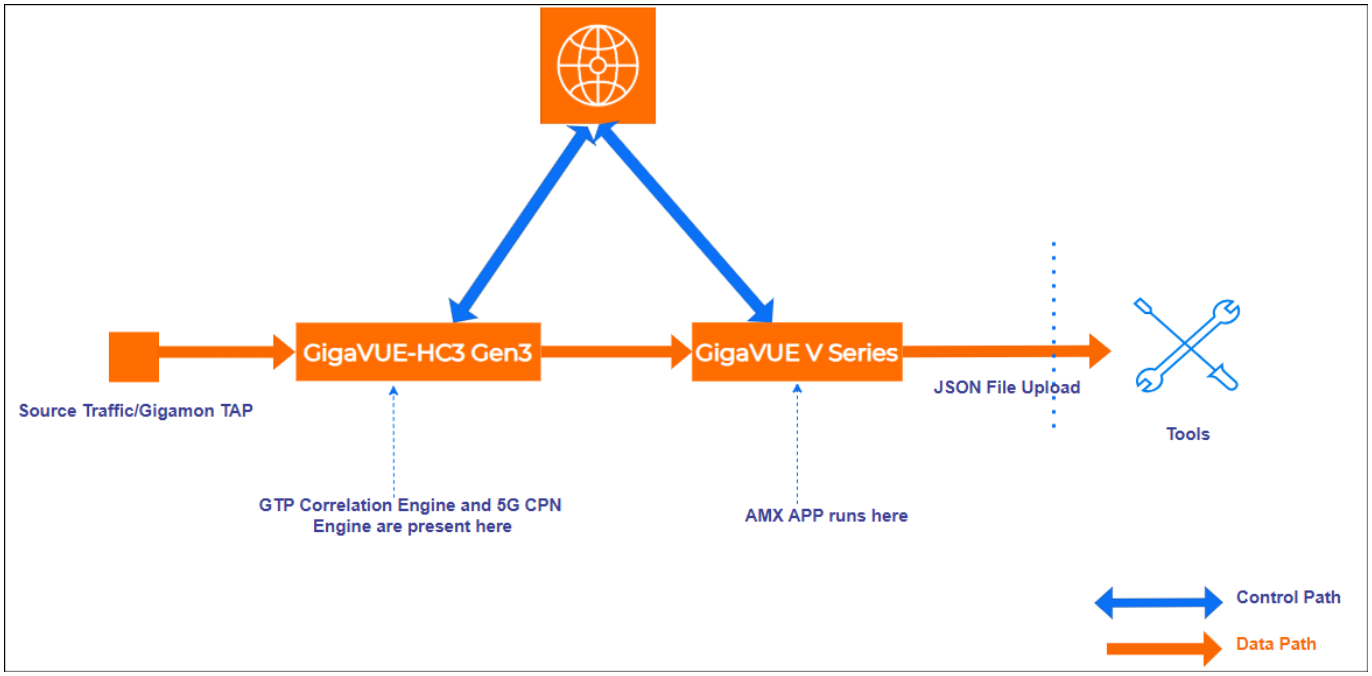
In hardware deployments, the Application Metadata Intelligence (AMI) runs on a physical node/cluster, and the AMX application is deployed on a GigaVUE V Series Node running on VMware ESXi, OpenStack, or Nutanix. The output from the AMI in CEF format is sent to the AMX application in GigaVUE V Series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

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



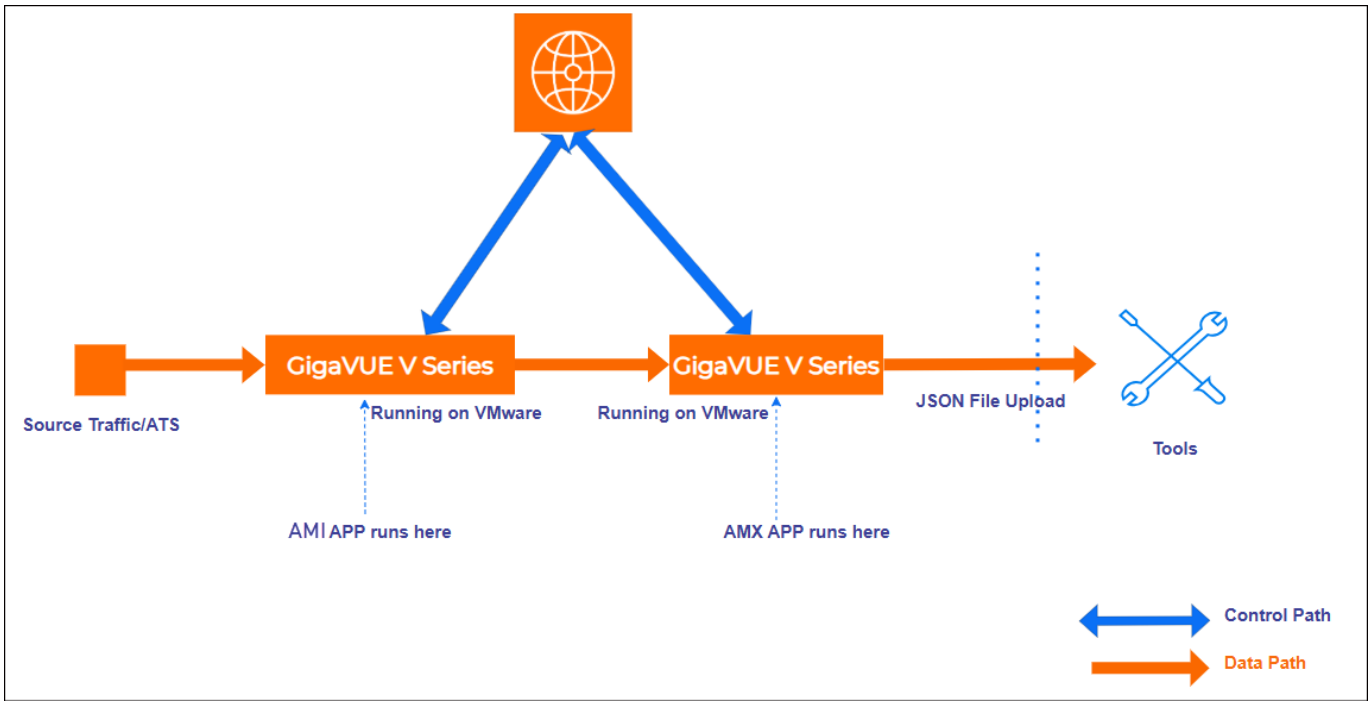
Hardware (Control Plane Metadata)

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



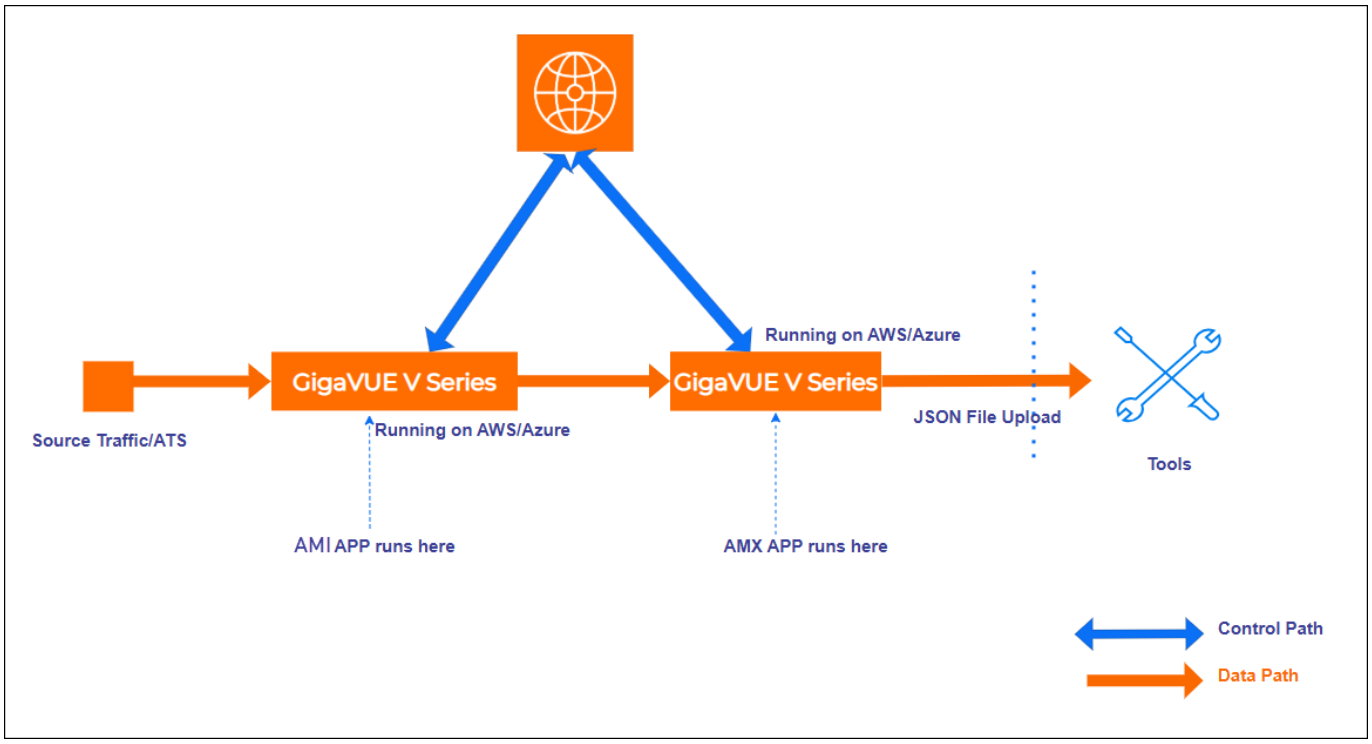
Private Cloud (VMware)

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



Public Cloud

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



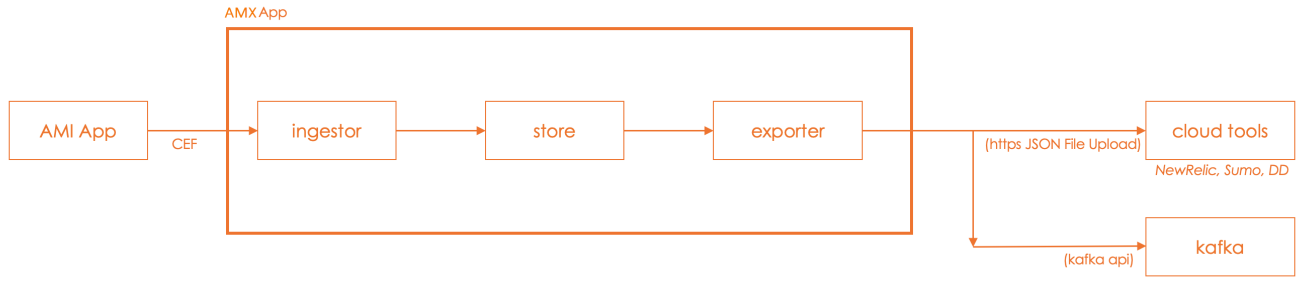
Application Metadata Exporter - Export Scenarios

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

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

Export Application Metadata in JSON format over HTTPS/Kafka

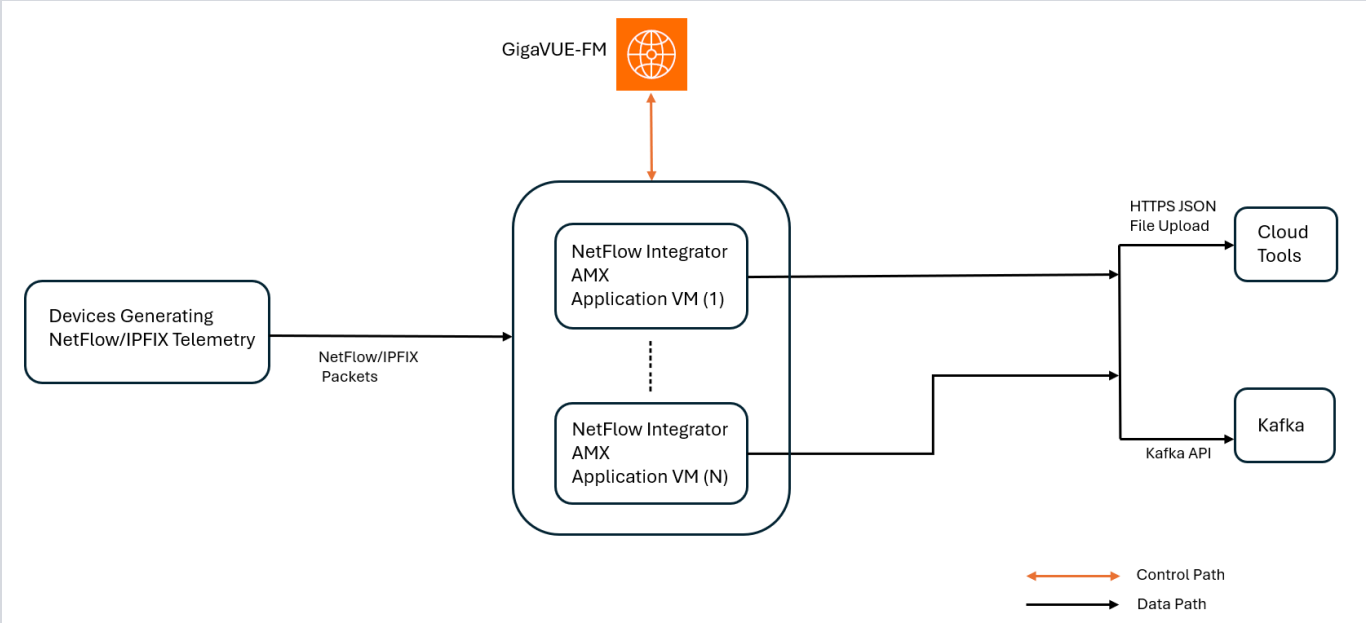
Application Metadata Exporter (AMX) application converts the output from the Application Metadata Intelligence (AMI) in CEF format into JSON format and sends it to the cloud tools and Kafka Consumers. The following image illustrates this end-to-end flow.



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

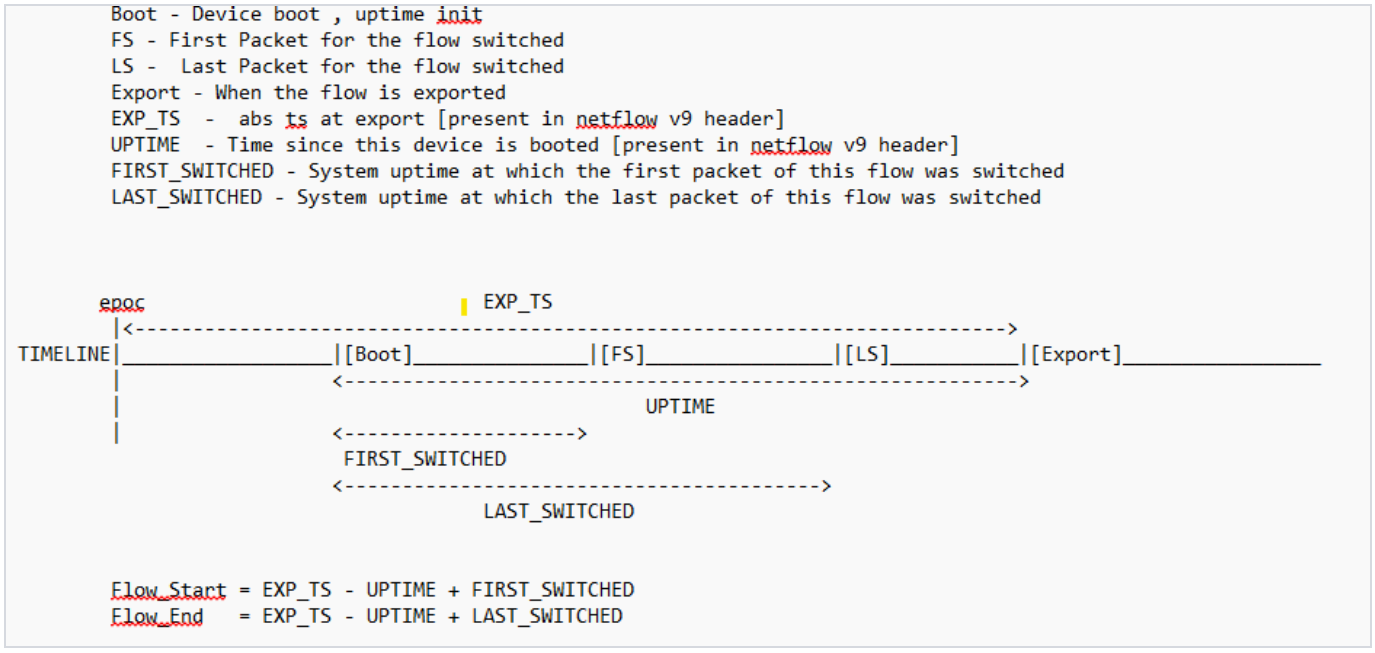
Important Notes

- **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:
  - Flow_Start = Export Timestamp - System Uptime + First_Switched
- Flow end time is derived using the following formula:
  - Flow_End = Export Timestamp - System Uptime + Last_Switched
- The calculated start and end time are converted into nanoseconds.
- Unsupported attributes such as flow start/end in milliseconds, microseconds, or nanoseconds will be ignored. Instead, the export timestamp is used as both the flow start and end times, which are then converted to nanoseconds.

- **NetFlow v10/IPFIX Timestamps**:
  - NetFlow v10/IPFIX directly incorporates timestamps without the need for system uptime.
  - Even if Flow start and end timestamps are available in seconds, milliseconds, microseconds, or nanoseconds, they are always converted to nanoseconds for output.
- **Missing Flow Start/End Attributes** - For records that do not include flow start or end attributes, AMX automatically assigns both values to the export timestamp.

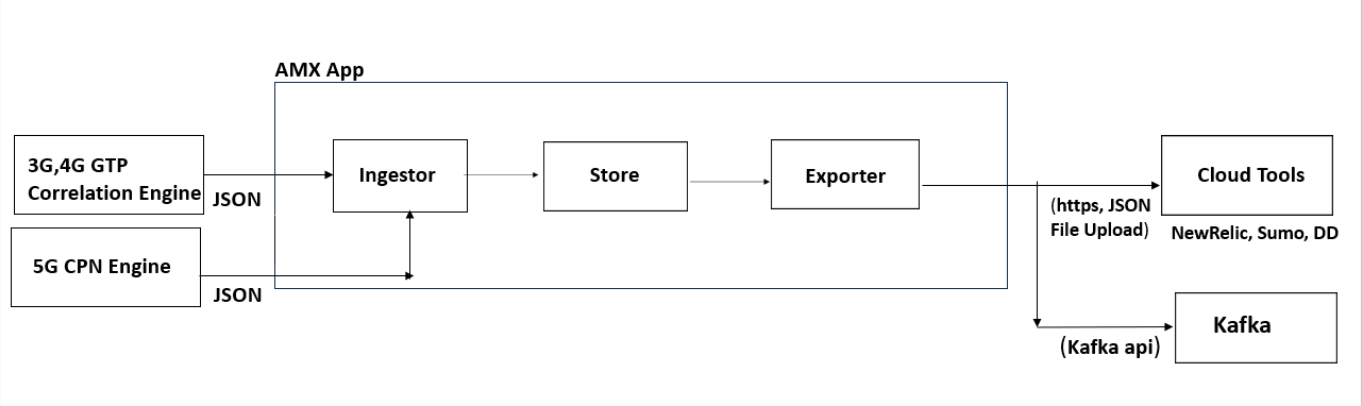
The image below explains how flow start and end times are calculated using system uptime, export timestamp and flow switch times.



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

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

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



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

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

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

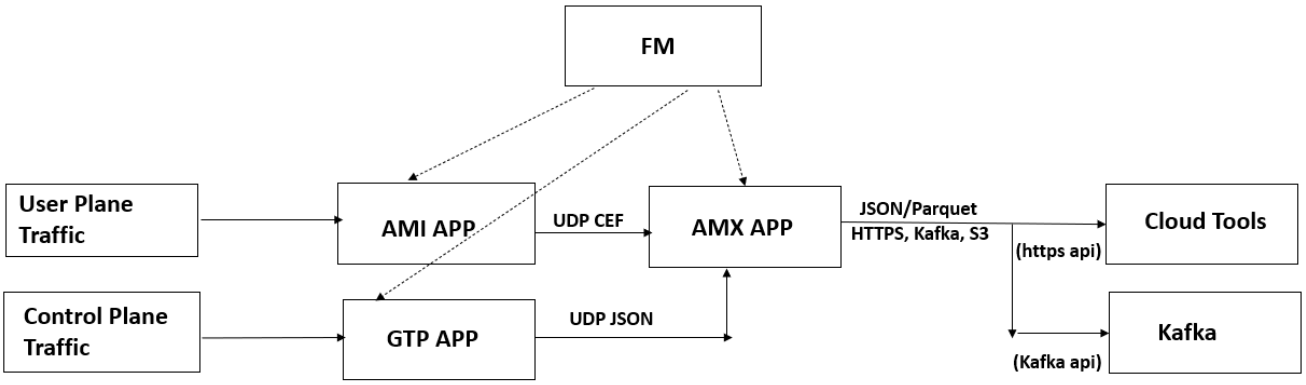
Export of GigaVUE Enriched Metadata for Mobile Networks is supported only for GigaVUE V Series Node deployed using Third Party Orchestration on VMware ESXi.

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

- GigaVUE-HC3
- GigaVUE-HC1-Plus

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

For information on Control Plane Metadata, refer to Control Plane Metadata.



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

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

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

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

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

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

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

- VMware: 300 sec
- AWS: 30 sec

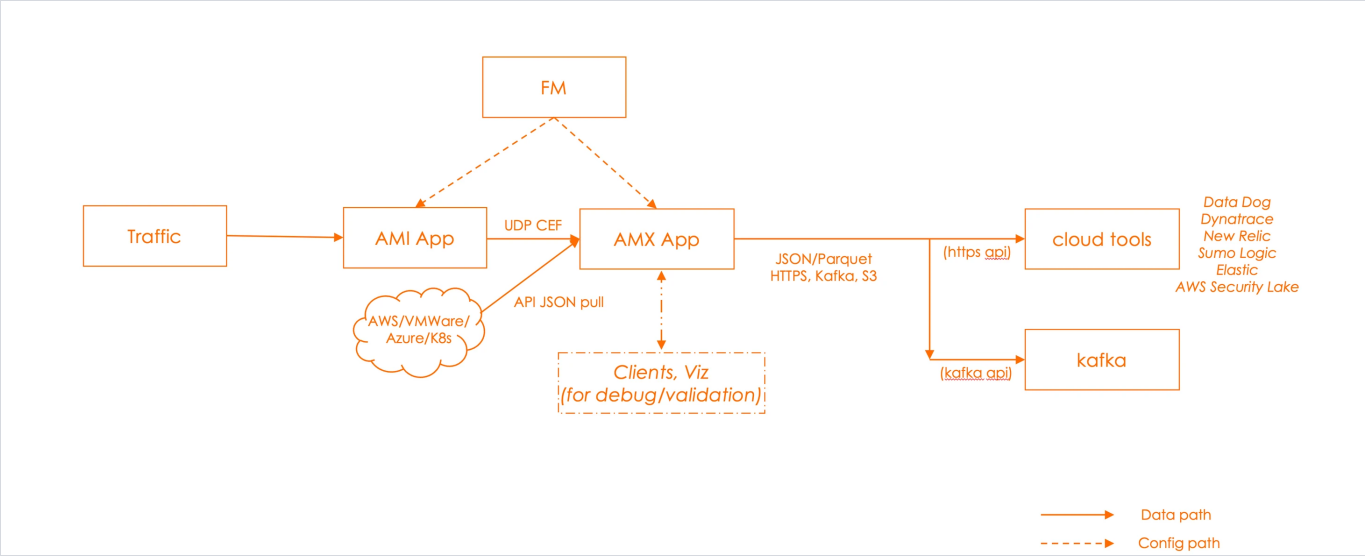
- Azure: 60 sec

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

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

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

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

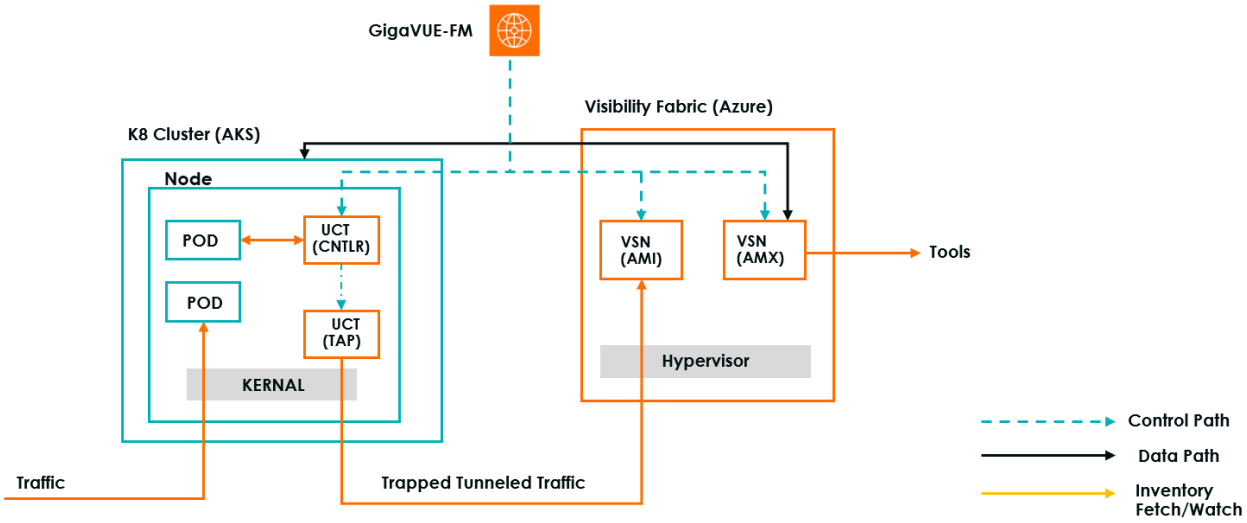


The enrichment supported depends on the type of platform. Refer to Attributes for GigaVUE Enriched Metadata for Cloud Workloads for more details.

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

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

to the AMX node, also deployed in Azure. With the new enhancement, the AMX node enriches the AMI metadata using Kubernetes inventory details fetched directly from the AKS cluster.



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

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

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

Refer to Prerequisites for Export of GigaVUE Enriched Metadata for Cloud Workloads for more detailed information.

Rules and Notes

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

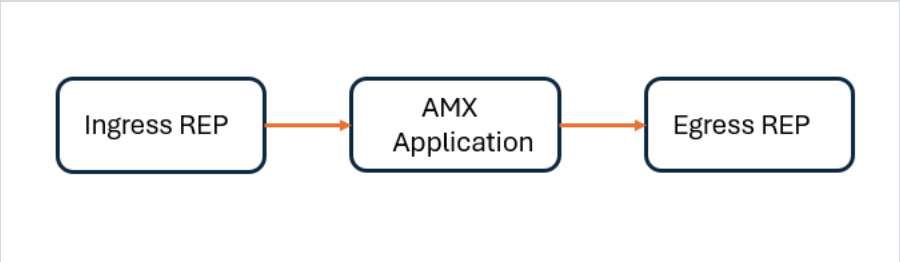
- When using this application for production usage, it is recommended to use large size Virtual Machines.
- If you reload the GigaVUE V Series Node after configuring the AMX application, then the Ingestor in the AMX application fails.
- When using GigaVUE Enriched Metadata for Cloud Workloads, if duplicate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

Configure Application Metadata Exporter

NOTE: Refer to Prerequisites for Application Metadata Exporter topic for all information related to the prerequisites.

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

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



To add AMX application:

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

   Notes:

   - Ensure to create separate Monitoring Domains and Monitoring Sessions for Application Metadata Exporter (AMX) and Application Metadata Intelligence (AMI). Even when AMI is deployed virtually, it operates independently and requires its own Monitoring Session.
   - You can use the Third-Party Orchestration method when deploying fabric components—such as the AMX V Series Node—on your own.
   - When deploying AMX using the Third-Party Orchestration method, Monitoring Session must also be configured in the same work flow. To deploy through Third Party Orchestration, go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**.
   - Before deploying AMX, ensure to change the **Traffic Acquisition Method** to **Customer Orchestrated Source** in the Monitoring Domain. Refer to Create Monitoring Domain.

2. After creating a new Monitoring Session or on an existing Monitoring Session, navigate to the **TRAFFIC PROCESSING** tab. The GigaVUE-FM Monitoring Session canvas page appears.
3. In the canvas, click the icon on the left side of the page to view the traffic processing elements. Drag and drop **AMX** from **APPLICATIONS** to the graphical workspace. The application quick view appears.
4. Enter the **Alias** for the application.

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

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

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

> You can use the **Actions** button to add multiple Metadata Enrichment. Keep in mind the following when configuring multiple Metadata Enrichment:
>
> - You can only configure either Mobility or Workload enrichment
> - You can only configure one Metadata enrichment with either **Mobility** or **Workload** as the **Type**
> - You can only configure a maximum of 5 Metadata enrichment with **Others** as the **Type**.

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

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

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

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable **Secure Endpoint** option to hide the API keys from the URL. Example: • Normal Endpoint: http://examples.com/mykey/key123 • Masked Endpoint: http://examples.com/mykey/****
Headers	Enter the custom header value of the cloud tool instance. Example: Add Content-Type: application/json in a separate header field. Ensure that the **Secure Keys** option is enabled.
Type	Select any one of the following: **AMI** - Select this option to export AMI. **Mobility Control**- Select this option to export control plane metadata. **AMI Enriched** - Select this option to export enriched metadata for cloud workloads. **NetFlow/IPFIX** - Select this option to export NetFlow/IPFIX packets to tools such as HTTPS or Kafka.
Account ID	Enter the account ID number of the selected Cloud Tool.
API Key	Enter the API key of the Cloud Tool.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Cloud Tool. Both IPv4 and IPv6 are supported. This field is optional.
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Format	The default output format is JSON.
Zip	Enable the box to compress the output file. **Note**: Enable this field when using New Relic as the cloud tool.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Export Retries	The number of times the application tries to export the entries to Cloud Tool. The recommended minimum value is 4 and the maximum is 10.
Maximum Entries	The number of JSON entries in a file. The maximum number of allowed entries is 20000 and the minimum is 10, however 1000 is the default value.

Fields	Description
Backoff Reset Window	With a non-zero value, the exporter keeps trying to reach the tool as many times as the retry. The default setting is 0, and the Backoff Reset Window is disabled at this value.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
Labels	Click **Add**. Enter the following details: o Enter the **Key**. o Enter the **Value**. **Note**: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.

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

Fields	Description
Alias	Enter the alias name for the Kafka Export.
Topic	The topic name to push JSON streams to, which is generally given to users part of the Kafka administration.
Type	Select any one of the following: AMI - Select this option to export AMI. Mobility Control - Select this option to export control plane metadata. AMI Enriched - Select this option to export enriched metadata. NetFlow/IPFIX - Select this option to export NetFlow/IPFIX ingestor types.
Brokers	The URL that contains the Kafka cluster endpoints. Click ⊕ to add another broker and click ⊖ to remove an existing broker.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Kafka Broker. Both IPv4 and IPv6 are supported. This field is optional.
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Format	The default output format is JSON.
Zip	Enable the box to compress the output file.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds. The default time interval is 30 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Export Retries	The number of times the application tries to export the entries to Kafka. The recommended minimum value is 4 and the maximum is 10.
Maximum Entries	The number of JSON entries in a file. The maximum number of allowed entries is 20000 and the minimum is 10, however 1000 is the default value.

Fields	Description
Backoff Reset Window	With a non-zero value, the exporter keeps trying to reach the tool as many times as the retry. The default setting is 0, and the Backoff Reset Window is disabled at this value.
Labels	Click **Add**. Enter the following details: o Enter the **Key**. o Enter the **Value**.
Producer Configurations	Click **Add** to enter the authentication details if a Kafka broker needs authentication. For Example: • security.protocol=SASL_SSL • sasl.mechanism=PLAIN or sasl.mechanisms=PLAIN • sasl.username=username • sasl.password=password

9. Click **Save** to deploy the Monitoring Session. The **Select nodes to deploy the Monitoring Session** dialog box appears. Select the GigaVUE V Series Node for which you wish to deploy the Monitoring Session.
10. Click **Deploy** from the **Actions** menu to deploy the Monitoring Session. The status is displayed as **Success** in the Monitoring Sessions page. The session is successfully deployed on all the V Series Nodes.

Switching to outer_ip for Mobility Enrichment

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

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

1. Undeploy the Monitoring Session in GigaVUE-FM.
2. Create a new .yaml file in the GigaVUE V Series Node: vi /etc/amx_settings.yaml
3. Modify the enrichment mode:
```
# Default mobility enrichment mode: inner_ip (default) or outer_ip
mobility_outer_ip_version: <Enter version as v6 or v4>
mobility_enrichment_mode: outer_ip
```
4. Deploy the Monitoring Session.

During Initial AMX Configuration:

1. Create a new .yaml file in the GigaVUE V Series Node: vi /etc/amx_settings.yaml
2. Modify the enrichment mode:
```
# Default mobility enrichment mode: inner_ip (default) or outer_ip
mobility_outer_ip_version: <Enter version as v6 or v4>
mobility_enrichment_mode: outer_ip
```

3. Save the file and proceed with the deployment.

Application Monitoring Options

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

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

You can view the statistics of the application in the **STATISTICS** tab. Refer to View Application Metadata Exporter Statistics for more detailed information.

Deploy Application Metadata Exporter in OpenStack using Third Party Orchestration

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

Predeployment Checklist:

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

   Exporting JSON records involves file read/write operations; using SSD storage is recommended.
2. Ensure you have a valid SecureVUE Plus VBL license to deploy AMX and to export NetFlow/IPFIX records in JSON format. With any AMI entitlement (AMI SKU, ZTA bundle, or SecureVUE PLUS bundle), you are also entitled to deploy AMX.
3. Ensure you have access to the VUE Gigamon Community portal for downloading V Series OVA image.
4. 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.

5. 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.
6. 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.
7. Review the deployment guide for the tool specific configuration needed to export records from AMX. Refer to AMX Tool Integration.
8. 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.
9. Ensure that the standalone GigaVUE-FM deployment meets the recommended requirements shown in the following table. For medium and large scale deployments, refer to Recommended Requirements for GigaVUE-FM Deployments. This step is optional. A standalone GigaVUE-FM instance is required only when an existing GigaVUE-FM is not available for managing AMX.

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

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

- (Optional) Install GigaVUE-FM using ISO Image File on OpenStack
- (Optional) Configure Metadata Export Filter for AMX
- (Optional) Configure Role-Based Access for Third Party Orchestration
- Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source
- Create Tokens to Secure REST API and Third Party Orchestration Access to GigaVUE-FM
- Deploy V Series Node and Register with GigaVUE-FM
- Create a Monitoring Session to Configure Application Metadata Exporter
- Configure Ingestor in AMX to Process Traffic
- Integrate AMX with Third Party Tool
- Configure Exporter to Export Attributes to Third Party tools
- Verify Your Deployment

(Optional) Install GigaVUE-FM using ISO Image File on OpenStack

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

You can subscribe to GigaVUE-FM in the marketplace and launch in AWS.

Refer to the following topics for instruction on installing GigaVUE-FM in AWS:

- Access GigaVUE-FM
- Access GigaVUE-FM

Access GigaVUE-FM

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

(Optional) Configure Metadata Export Filter for AMX

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

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

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

Create a Custom User

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

To add users:

1. Go to **Settings** and select **Authentication** > **GigaVUE-FM User Management > Users**.
2. On the User page, select **New User**.
3. In the Add User page, enter the following details:
   - **Name:** Actual name of the user
   - **Username**: User name configured in GigaVUE-FM
   - **Email**: Email ID of the user
   - **Password/Confirm Password**: Password for the user. For more details, refer to Change Your Password.
   - **User Group:** Select the desired User Group to associate the user.

   GigaVUE-FM prompts for your password.
4. Click **Ok**. The new user is added.

Create a Custom Role and Assign the User

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

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

To create a role:

1. Go to **Settings** and select **Authentication> GigaVUE-FM User Management >Roles**.
2. Select **New Role**.
3. In the New Role page, select or enter the following details:

- **Role Name**: Name of the role.
- **Description**: Description of the role.
- **Select Permission**: From the **Select Permissions** tab, select **Third Party Orchestration**, and provide write permissions.

4. Select **Apply** to save the configuration.

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

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

To create a new user group:

1. Go to **Settings**, and then select **Authentication> GigaVUE-FMUser Management >User Groups**.
2. Select **New Group.**
3. In the New User Group page, select **Next** to progress forward and **Back** to navigate backward.
4. In the **Group Info** tab, enter the Group Name and Description details:
5. In the **Assign Roles** tab, select the role that you want to assign to the user group.
6. In the **Assign Tags** tab, select the required tag key and tag value.
7. In the **Assign Users** tab, select the required users.
8. Select **Apply** to save the configuration. The new user group is added.

Create Monitoring Domain to Acquire Traffic from Customer Orchestrated Source

To create a Monitoring Domain in GigaVUE-FM:

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

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

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

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

To create a token:

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

Deploy V Series Node and Register with GigaVUE-FM

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

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

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

Register V Series Nodes or V Series Proxy during Instance Launch

To register V Series nodes or proxy using the Customization Script in OpenStack GUI:

1. On the Instance page of OpenStack dashboard, select **Launch instance**. The Launch Instance wizard appears. For detailed information, refer to Launch and Manage Instances topic in OpenStack Documentation.
2. On the **Configuration** tab, enter the Customization Script as text in the following format and deploy the instance. The V Series nodes or V Series proxy uses this customization script to generate config file (**/etc/gigamon-cloud.conf**) used to register with GigaVUE-FM.

```
#cloud-config
 write_files:
 - path: /etc/gigamon-cloud.conf
   owner: root:root
   permissions: '0644'
   content: |
     Registration:
        groupName: <Monitoring Domain Name>
        subGroupName: <Connection Name>
        token: <Token>
        remoteAddress: <IP address of the GigaVUE-FM> or <IP address of the Proxy>
        remotePort: 443
```

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

Register V Series Node or V Series Proxy after Instance Launch

To register V Series node or proxy using a configuration file:

1. Log in to the GigaVUE V Series Node or Proxy.
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.

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

> **NOTE:** You can register your GigaVUE V Series Nodes directly with GigaVUE-FM or you can use V Series proxy to register your GigaVUE V Series Nodes with GigaVUE-FM. If you wish to register GigaVUE V Series Nodes directly, enter the **`remotePort`** value as 443 and the **`remoteIP`** as <IP address of the GigaVUE-FM> or if you wish to deploy GigaVUE V Series Nodes using V Series proxy then, enter the **`remotePort`** value as 8891 and **`remoteIP`** as <IP address of the Proxy>. If the V Series Node registration fails with a Connection Refused error, set the remoteIP field to the API network IP of GigaVUE-FM instead of the floating IP.

3. Restart the GigaVUE V Series Node or Proxy service.
   - GigaVUE V Series Node:
     `$ sudo service vseries-node restart`
   - GigaVUE V Series Proxy:
     `$ sudo service vps restart`

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

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

Assign Static IP address for GigaVUE V Series

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

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

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

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

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

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

Example netplan config:

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

4. Save the file.
5. Apply the configuration.

   $ sudo netplan apply

6. Restart the GigaVUE V Series service.

   $ sudo service vseries-node restart

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

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
- Access GigaVUE-FM

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 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.

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 quic k view in the **Cloud Tool Exports** section, enter or select the required details as described in Cloud Tool Exporter - Field References.
2. Click **Save**.

Deploy the Monitoring Session

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

Cloud Tool Exporter - Field References

Fields	Description
Alias	Enter the alias name for the cloud tool export.
Endpoint	Enter the endpoint URL where the cloud tool instance is installed and the ports are configured. Enable **Secure Endpoint** option to hide the API keys from the URL. Example: • Normal Endpoint: http://examples.com/mykey/key123 • Masked Endpoint: http://examples.com/mykey/****
Headers	Enter the custom header value of the cloud tool instance. Example: Add Content-Type: application/json in a separate header field. Ensure that the **Secure Keys** option is enabled.
Type	Select any one of the following: **AMI** - Select this option to export AMI or identity enrichment metadata **Mobility Control**- Select this option to export control plane metadata. **AMI Enriched** - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. **NetFlow/IPFIX** - Select this option to transform NetFlow/IPFIX records to JSON format.
Labels	Click **Add**. Enter the following details: o Enter the **Key**. o Enter the **Value**. **Note**: Refer to the deployment guides of the respective cloud tools for configuring the Key and Value.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Zip	Enable the box to compress the output file. **Note**: Enable this field when using New Relic as the cloud tool.
Interval (sec)	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default value is 5000.
Request Timeout	This is the time exporter waits for a response back from the tool. The default value is 10 seconds.
Source IP Address	Source IP Address is needed when the egress interface is configured with multiple IP addresses. Configure the source IP address which is connected to the Cloud Tool. Both IPv4 and IPv6 are supported. This field is optional.

Configure Kafka Exporter

To configure Kafka Export settings in AMX:

1. On the AMX quick view in the **Kafka Exports** section, enter or select the required details as described in Kafka Exporter - Field References.
2. Click **Save**.

Deploy the Monitoring Session

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

Kafka Exporter - Field References

Fields	Description
Alias	Enter the alias name for the Kafka Export.
Topic	The topic name to push JSON streams to, which is generally given to users part of the Kafka administration.
Type	Select any one of the following: **AMI** - Select this option to export AMI or identity enrichment metadata **Mobility Control** - Select this option to export control plane metadata. **AMI Enriched** - Select this option to export enriched metadata for cloud workloads and identity enrichment using ZPA. **NetFlow/IPFIX** - Select this option to export NetFlow/IPFIX ingestor types.
Brokers	The URL that contains the Kafka cluster endpoints. Click ⊕ to add another broker and click ⊖ to remove an existing broker.
Labels	Click **Add**. Enter the following details: o Enter the **Key**. o Enter the **Value**.
Advanced Settings	
Enable Export	Enable the box to export the Application Metadata Intelligence output in JSON format.
Interval	The time interval (in seconds) in which the data should be uploaded periodically. The recommended minimum time interval is 10 seconds and the maximum time interval is 90 seconds. The default time interval is 30 seconds.
Parallel Writer	Specifies the number of simultaneous JSON exports done.
Max Entries	The number of JSON entries in a file. The valid range is 5000–20000 and the default

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

Extend Your AMI and AMX Deployment

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

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

Verify Your Deployment

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

To troubleshoot any issues in Application Metadata Exporter deployment, refer to Troubleshoot Application Metadata Exporter Deployment.

View Application Metadata Exporter Statistics

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

1. Go to **Traffic > Virtual > Orchestrated Flows > Third Party Orchestration**. The **Monitoring Sessions** Overview page appears.
2. Select the monitoring session you want to view the statistics.
3. Go to the **TRAFFIC PROCESSING** tab. In the canvas page, click the ⋮ menu button on the **AMX** application and select **Details**. The Application quick view appears.
4. Click on **STATISTICS** tab.
5. To view the statistics of a particular GigaVUE V Series Nodes, select the required GigaVUE V Series Node from the **V Series Node IP** dropdown menu to view the **Ingestor**, **Exporter**, and **Attributes Enrichment** statistics.

You can download the Ingestor and Exporter statistics using the following options:

- **Download TCP Dump** (Ingestor)- Select this option to download the PCAP file (for example, cap_ens192_20260407_073719655-run-001.pcap.gz). The TCP dump capture is limited to 100 packets and runs for up to 30 seconds, or until 100 packets are captured (whichever occurs first). You can capture a TCP dump for only one AMX ingestor at a time. The TCP dump download option is available only for the following ingestor types: AMI, NetFlow/IPFIX, Mobility, and Zscaler LSS.
- **Download Records** (Exporter) - Select this option to download the JSON file (for example, Elastic-exporter-records.json.gz) of the exporter records. You can download records for a single AMX exporter or for multiple exporters. The exporter record download waits up to 30 seconds for records to be exported. If no records are produced, the operation times out.

For instructions on how to download the ingestor and exporter records, and troubleshoot AMI/AMX issues, refer to Troubleshoot AMI/AMX Deployment.

For more information on health status of exporters, refer to View Health Status.

You can view the following parameters in the Application Metadata Exporter application statistics page.

Exporter - Statistics to monitor exports from AMX

Parameters	Description
Name	Displays the name of exporters created for this application.
Exported Entries	Displays the number of entries available in the files that will be uploaded from AMX to the cloud tool or the Kafka consumers.
Average File Size	Displays the average size of the file.
File Upload Errors	Displays the number of file uploaded from the AMX to the cloud tool or the Kafka consumers.
File Upload Errors	Displays the number of times the file was not uploaded to the cloud tools or Kafka consumers due to errors.
Re Exported Entries (Only for Kafka Tool Exporter)	Displays the number of records re-exported from the errors directory after the Kafka tool is restored. Errored files are re-exported automatically at a rate of 10 files per 5 seconds.
Exported Entries Rolling 15 Sec	Displays the rolling average of entries exported per second over the past 15 seconds.
Exported Entries Rolling 30 Sec	Displays the rolling average of entries exported per second over the past 30 seconds.
Exported Entries Rolling 60 Sec	Displays the rolling average of entries exported per second over the past 60 seconds.
Exported MBytes	Total number of megabytes exported.
Exported MBytes Rolling 15 Sec	Rolling average of megabytes exported over the past 15 seconds.
Exported MBytes Rolling 30 Sec	Rolling average of megabytes exported over the past 30 seconds.
Exported MBytes Rolling 60 Sec	Rolling average of megabytes exported over the past 60 seconds.

Ingestor - Statistics to monitor ingestion

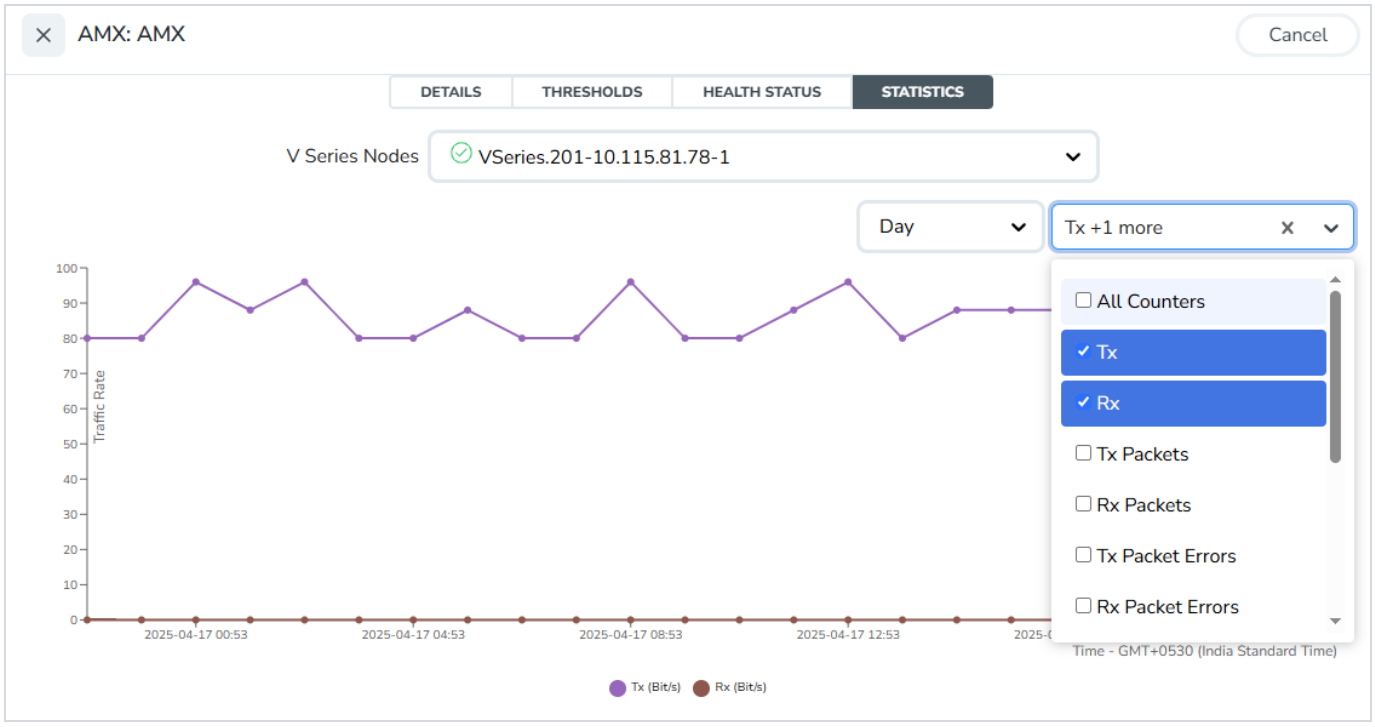
Parameters	Description
Name	Displays the name of the ingestor created for this application.
Packets Received	Displays the packets received from the AMI to AMX ingestor.
Packets Dropped	Displays the packets dropped when the traffic passes from AMI to AMX ingestor.

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

Attributes Enrichment - Statistics to monitor enrichment

Parameters	Description
Name	Displays the name of the Metadata enrichment.
One Min Percent	Displays the percentage of the traffic that is enriched from the last 1 minute.
Five Min Percent	Displays the percentage of the traffic that is being enriched from the last 5 minutes.
Ten Min Percent	Displays the percentage of the traffic that is being enriched from the last 10 minutes.

The following image shows the statistical details for general AMX application.

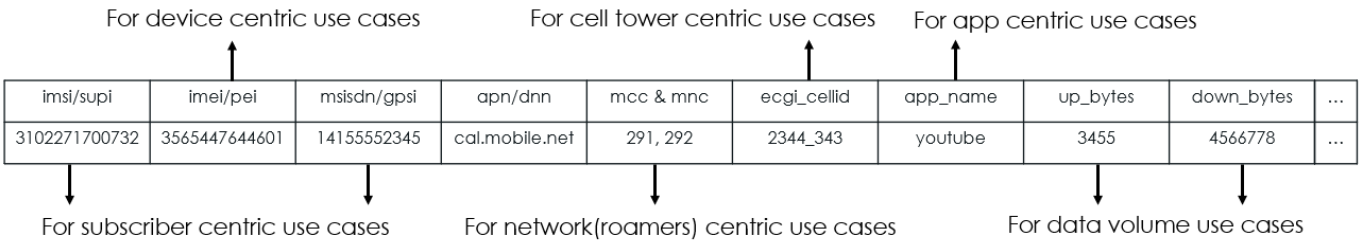


Attributes for GigaVUE Enriched Metadata for Mobile Networks

This section describes the various attributes available within GigaVUE Enriched Metadata for Mobile Networks.

Each flow will generate only one record per export interval; however, in this section, the attributes are separated into the following categories:

- Control Plane Attributes
- User Plane Attributes



Control Plane Attributes

The control plane attributes are generated based on network signaling message transactions with successful responses. These include signaling transactions for session establishment, modification, deletion, and others. These attributes are used for enriching the user plane metadata with the corresponding control plane subscriber, device, and location information.

Attribute	Description
imsi	International Mobile Subscriber Identity. A 3GPP-defined unique private identifier for a mobile user on a 4G network.
imei	International Mobile Equipment Identity. A 3GPP-defined unique identifier for a mobile device on a 4G network.
msisdn	Mobile Station Integrated Services Digital Network. A 3GPP-defined identifier for a mobile device on a 4G network.
supi	Subscription Permanent Identifier. A 3GPP-defined unique identifier for a mobile user on a 5G network.
gpsi	Generic Public Subscription Identifier. A 3GPP-defined public identifier for a mobile user on a 5G network.
pei	Permanent Equipment identifier. A 3GPP-defined unique identifier for a mobile device on a 5G network.
ctrl_tun_ access_ ipv4_addr	Control Plane Access Tunnel IPv4 Address. MME identifier for S11, SGW identifier for S5/S8-C, AMF identifier for N11 and SGSN for Gn, GP
ctrl_tun_ access_ ipv6_addr	Control Plane Access Tunnel IPv6 Address. MME identifier for S11, SGW identifier for S5/S8-C, AMF identifier for N11 and SGSN for Gn, GP
ctrl_tun_ access_teid	Control Plane Access Tunnel Endpoint Identifier
ctrl_tun_ core_ipv4_ addr	Control Plane Core Tunnel IPv4 Address. SGW identifier for S11, PGW identifier for S5/S8-C, SMF identifier for N11 and GGSN for Gn,Gp
ctrl_tun_ core_ipv6_ addr	Control Plane Core Tunnel IPv6 Address. SGW identifier for S11, PGW identifier for S5/S8-C, SMF identifier for N11 and GGSN for Gn,Gp
ctrl_tun_ core_teid	Control Plane Core Tunnel Endpoint Identifier

user_tun_up_link_ipv4_addr	User Plane Uplink Tunnel IPv4 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_ipv6_addr	User Plane Uplink Tunnel IPv6 Address. Identifier of the SGW/UPF that receives the uplink user plane data from the eNodeB/gNodeB
user_tun_up_link_teid	User Plane Uplink Tunnel Endpoint Identifier
user_tun_down_link_ipv4_addr	User Plane Downlink Tunnel IPv4 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_ipv6_addr	User Plane Downlink Tunnel IPv6 Address. Identifier of the eNodeB/gNodeB that receives the uplink user plane data from the SGW/UPF
user_tun_down_link_teid	User Plane Downlink Tunnel Endpoint Identifier
if_name	3GPP Control Plane Interface Name
name	Name of the event on the 3GPP control plane interface
cause_code	Outcome of the event on the 3GPP control plane interface. Mostly set to 16 to indicate a successful event
sm_context_ref	Session Management Context
ebi	EPS(Evolved Packet System) Bearer ID
lbi	Linked Bearer Identity
pdu_session_id	PDU(Packet Data Unit) Session ID
apn	Access Point Name
dnn	Data Network Name
ue_ipv4_addr	User Equipment IPv4 Address
ue_ipv6_addr	User Equipment IPv6 Address
ue_ipv4v6_addr	User Equipment IPv4v6 Address
ue_non_ip_addr	User Equipment non-IP Address
ue_addr_type	User Equipment IP address type (IPv4 or IPv6)
qci	QoS Class Identifier
qfi	QoS Flow Identifier
five_qi	5G QoS Identifier

cgi_mcc	Mobile Country Code from Cell Global Identity
cgi_mnc	Mobile Network Code from Cell Global Identity
cgi_lac	Local Area Code from Cell Global Identity
cgi_cell_id	Cell Identification from Cell Global Identity
sai_mcc	Mobile Country Code from Service Area Identifier
sai_mnc	Mobile Network Code from Service Area Identifier
sai_lac	Local Area Code from Service Area Identifier
sai_sac	Service Area Code from Service Area Identifier
tai_mcc	Mobile Country Code from Tracking Area Identity
tai_mnc	Mobile Network Code from Tracking Area Identity
tai_tac	Tracking Area Code from Tracking Area Identity
ecgi_mcc	Mobile Country Code from E-UTRAN Cell Global Identifier
ecgi_mnc	Mobile Network Code from E-UTRAN Cell Global Identifier
ecgi_cell_id	Cell Identification from E-UTRAN Cell Global Identifier
lai_mcc	Mobile Country Code from Location Area Identity
lai_mnc	Mobile Network Code from Location Area Identity
lai_lac	Local Area Code from Location Area Identity
enode_id_mcc	Mobile Country Code from Evolved Node
enode_id_mnc	Mobile Network Code from Evolved Node
macro_enode_id	Evolved Node ID
ncgi_mcc	Mobile Country Code from NR Cell Global Identity
ncgi_mnc	Mobile Network Code from NR Cell Global Identity
ncgi_nr_cell_id	Cell Identification from NR Cell Global Identity
rat_type	Type of Radio Access Technology
snssai_sst	Slice/Service Type from Single Network Slice Selection Assistance Information
snssai_sd	Slice Differentiator id from Single Network Slice Selection Assistance Information

User Plane Attributes

The user plane attributes are generated by Application Metadata Intelligence (AMI), and include information such as IP addresses, protocols, timestamps, and byte/packet counters. These attributes are generated from the network flows and do not require DPI capabilities. The attributes for the mobile network traffic are generated for the inner IP flow by bypassing the outer headers

Attribute	Description
ts	Timestamp of the metadata generated in UTC
vendor	Identifying Gigamon as the vendor providing the metadata
version	Version number of the GigaSMART software release from the AMI application
generator	Identifies the Gigamon device that generates the user plane metadata. For example - HC3
src_ip	Source IPv4 address of the inner flow
dst_ip	Destination IPv4 address of the inner flow
src_mac	Source MAC address of the inner flow
dst_mac	Destination MAC address of the inner flow
mpls	MPLS Label if the flow has the MPLS header
protocol	Layer 4 protocol of the inner flow
src_port	Layer 4 protocol source port of the inner flow
dst_port	Layer 4 protocol destination port of the inner flow
device_inbound_ interface	Traffic receiving port on the Gigamon device example HC3
tcp-sport	Source TCP port
tcp_dport	Destination TCP port
udp_sport	Source UDP port
udp_dport	Destination UDP port
src_bytes	Source bytes from client to server in the flow including the outer and the inner IP header
dst_bytes	Destination bytes from client to server in the flow including the outer and the inner IP header
src_inner_bytes	Source bytes from the client to the server in the flow with the inner IP header (without the outer header)
dst_inner_bytes	Destination bytes from the server to the client in the flow with the inner IP header (without the outer header)
src_packets	Source packets from client to server in the flow including the header
dst_packets	Destination packets from client to server in the flow including header
total_bytes	Total bytes is sum of source and destination bytes
total_packets	Total packets is sum of source and destination packets
id	Unique flow id assigned to a flow
ingress_vlan_id	Vlan id in the packet received for the flow
src_ipv6	Source IPv6 address of the inner flow
dst_ipv6	Destination IPv6 address of the inner flow
ip_version	IP version v4, v6
ip_cos_id	QOS value in the IPv4 header
ip_dscp	DSCP value in the IPv4 header

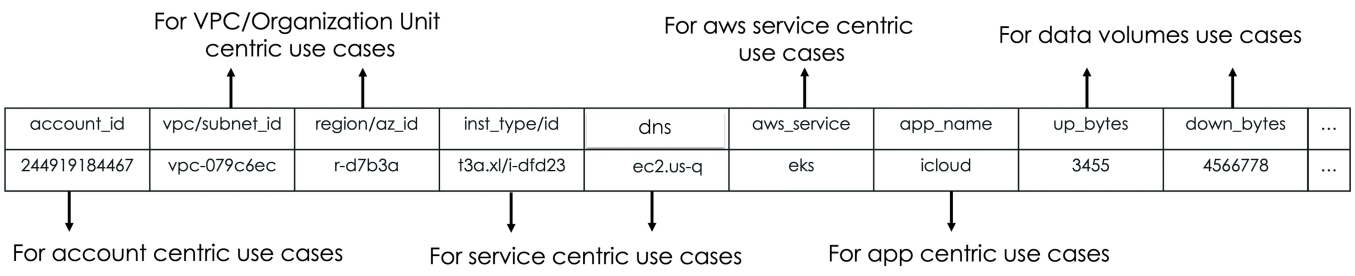
ip_ttl	Time to live in the IPv4 header
ip6hop_limit	Hop limit in the IPv6 header
flow_labelv6	flow label in the IPv6 header
ip6_ds	DSCP value in the IPv6 header
tcp_flags	TCP flag in the TCP header
gre_key	Generic routing encapsulation header key in the GRE header
ip_hdr_len	IPv4 header length
ip_frag_flags	IPv4 fragmentation flags
ipv6frag_flags	IPv6 fragmentation flags
ip_frag_id	IPv4 fragment id
ipv6frag_id	IPv6 fragment id
ip_frag_offset	IPv4 fragment offset
ipv6frag_offset	IPv6 fragment offset
ipv4opt	IPv4 options
ip_precendence_id	IPv4 precedence id
ip6precendence_id	IPv6 precedence id
ip_tot_len	IPv4 total length
ip6tot_len	IPv6 total length
ipv6hdr_len	IPv6 header length
payload_len_id	Payload length for the flow excluding the L3 IPv4 header
next_hdr_v6	Layer 4 protocol
icmp_type_v4	ICMP message type IPv4
icmp_code_v4	ICMP response code IPv4
icmp_code_v6	ICMP response code IPv6
icmp_type_v6	ICMP message type IPv6
tcp_ack_id	TCP ack id
tcp_hdr_len	TCP header length
tcp_seq_no	TCP sequence number
tcp_urgent_ptr	TCP urgent pointer
tcp_win_size	TCP window size
udp_msg_len	UDP message length
ip6traffic_class	IPv6 traffic class same as QoS in IPv4
flow_start_usec	Inner flow start time in microseconds in UTC
flow_end_usec	Inner flow end time in microseconds in UTC
flow_start_sec	Inner flow start time in seconds in UTC
flow_end_sec	Inner flow end time in seconds in UTC
start_time	Inner flow start time in milliseconds in UTC
end_time	Inner flow end time in milliseconds in UTC

egress_intf_id	egress IP interface-id for the Gigamon device sending the metadata
sys_up_time_first	Difference between the flow start time and the gigaSMART uptime in milliseconds
sys_up_time_last	Difference between the flow end time and the gigaSMART uptime in milliseconds
end_reason	Inner flow end reason – TCP ack, reset, inactive, etc
tcpflagsyn	TCP flag SYN from TCP header
tcpflagsynack	TCP flag SYNACK from TCP header
tcpflagfin	TCP flag FIN (finish)from TCP header
tcpflagrst	TCP flag RST (reset) from TCP header
Labels	Labels are added to the metadata by configuration in AMX. There can be multiple labels configured and added. Each label is configured as a key and a string value.
Labels configured in the JSON example in the section below	
label_event_type	Label identifying the vendor generating the event in the AMX exporter
label_deployment	Label identifying the deployment address of the AMX exporter
label_traffic	Label identifying the traffic region where the AMX is deployed
MSS	Maximum payload size a TCP segment can carry without the headers. The MSS value is exported at the beginning of a flow. Its exported only when bidirectional flows are configured.
aggr-window-size	The actual TCP window size minimum, mean, and maximum values, calculated separately for both sender and receiver packets.
zero-window	This captures TCP zero window event statistics for both sender and receiver, indicating when the TCP receive buffer is full and data transmission is temporarily paused due to flow control.

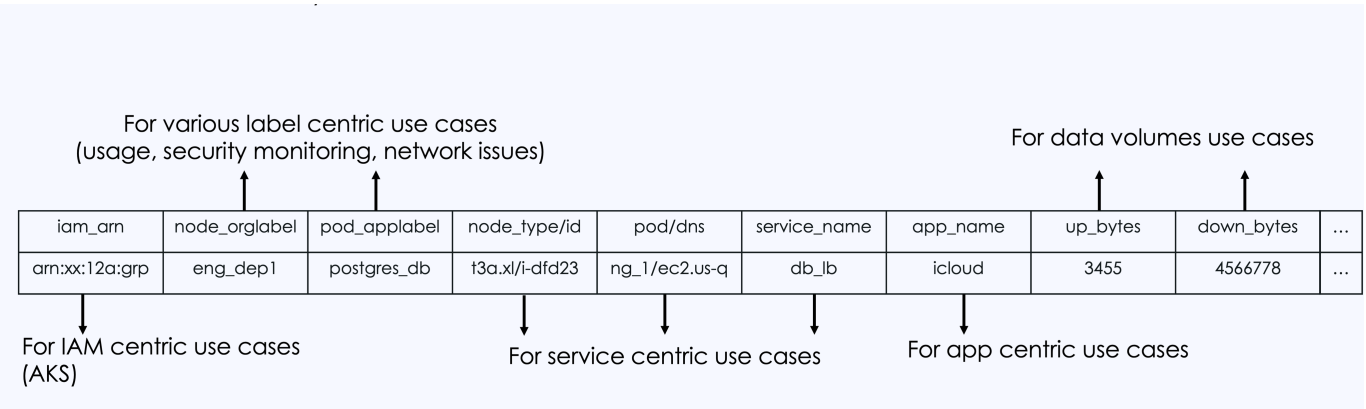
Attributes for GigaVUE Enriched Metadata for Cloud Workloads

This section describes the various attributes available within GigaVUE Enriched Metadata for Cloud Workloads for the following platforms:

- AWS
- Azure
- VMware vCenter
- Azure Kubernetes Service (AKS)
- Zscaler Private Access (ZPA) Attributes



- Azure Kubernetes Service (AKS)



The following table defines the purpose and scope of each referenced dataset, detailing configuration, resource, and operational parameters across AWS, Azure, and Kubernetes environments to support multi-cloud infrastructure documentation.

AWS

Attribute	Description	Source	Destination
aws_instance_id	ID of the EC2 instance	✓	✓
aws_instance_type	Type of EC2 instance	✓	✓
aws_availability_zone	Availability Zone of the EC2 instance	✓	✓
aws_image_id	ID of the base image use for EC2 instance	✓	✓
aws_private_dns_name	Private DNS hostname name assigned to the EC2 instance	✓	✓
aws_private_ip	Private IPv4 address assigned to the EC2 instance	✓	✓
aws_public_dns_name	Public DNS name assigned to the EC2 instance	✓	✓
aws_public_ip	Public IPv4 address assigned to the EC2 instance	✓	✓
aws_instance_state	Current state of the EC2 instance	✓	✓
aws_subnet_id	ID of the subnet in which the EC2 instance is running	✓	✓
aws_vpc_id	ID of the VPC in which the EC2 instance is running	✓	✓

aws_iam_instance_profile_arn	Amazon Resource Name (ARN) of the IAM instance profile	✓	✓
aws_iam_instance_profile_id	ID of the IAM instance profile	✓	✓
aws_security_group_name	Security group name assigned for the EC2 instance	✓	✓
aws_security_group_id	Security group ID assigned for the EC2 instance	✓	✓
aws_network_if_id	ID of the network interface	✓	✓
aws_network_if_attach_id	ID of the network interface attachment	✓	✓
aws_tags	Tags assigned to the EC2 instance in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
aws_flat_tags	Tags assigned to the EC2 instance extracted to top level keys and values	✓	✓
aws_owner_id	ID of the AWS account that created the network interface	✓	✓
workload_platform	Platform Name - AWS	✓	✓

Azure

Attribute	Description	Source	Destination
azure_subscription_id	Subscription ID of the Azure account that created the VM	✓	✓
azure_resource_group	Resource Group of the VM	✓	✓
azure_vm_id	Complete path of the Azure VM. This includes subscription	✓	✓
azure_instance_id	Instance ID of the Azure VM	✓	✓
azure_instance_type	Instance type of the Azure VM	✓	✓
azure_availability_zone	Availability zone of the Azure VM	✓	✓
azure_image_id	Image ID of the Azure VM	✓	✓

azure_private_dns_name	Private DNS name assigned to the VM instance	✓	✓
azure_private_ip	Private IP address assigned to the VM instance	✓	✓
azure_public_dns_name	Public DNS name assigned to the VM instance	✓	✓
azure_public_ip	Public IP address assigned to the VM instance	✓	✓
azure_instance_state	Current state of the VM instance	✓	✓
azure_subnet_id	Subnet ID of the VM instance	✓	✓
azure_vnet_id	ID of the VNET in which the VM instance is running	✓	✓
azure_network_if_id	ID of the network interface	✓	✓
azure_network_if_attach_id	ID of the network interface attachment	✓	✓
azure_tags	Tags assigned to the VM instance in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
azure_flat_tags	Tags assigned to the VM instance extracted to top level keys and values	✓	✓
azure_owner_id	Owner of the VM instance	✓	✓
azure_location	Resource Location	✓	✓
azure_scaleset_name	Scale set name	✓	✓
azure_scaleset_mode	Scale set orchestration mode	✓	✓
workload_platform	Platform Name - Azure	✓	✓

VMware vCenter

Attribute	Description	Source	Destination
vmware_vm_name	Name of the guest VM	✓	✓
vmware_vm_network	Network name of the guest VM	✓	✓
vmware_vm_status	Current status of the guest VM	✓	✓
vmware_vm_tags	Tags assigned to the guest VM in kv format with kv delimiter ':' and pair delimiter ','	✓	✓
vmware_vm_flat_tags	Tags assigned to the guest VM extracted to top level keys and values	✓	✓

vmware_vm_os_family_name	OS type of the guest VM (eg: Linux)	✓	✓
vmware_vm_os_distro_name	OS distribution name of the guest VM	✓	✓
vmware_vm_os_distro_version	OS distribution version of the guest VM	✓	✓
vmware_vm_os_distro_pretty_name	OS distribution pretty name of the guest VM	✓	✓
vmware_vm_dns_host_name	Host DNS name of the guest VM	✓	✓
vmware_vm_dns_host_domain_name	Host DNS domain name of the guest VM	✓	✓
vmware_host_name	Host name of the guest VM	✓	✓
vmware_host_ip	Host IP address of the guest VM	✓	✓
vmware_datacenter_name	Datacenter name of the guest VM	✓	✓
vmware_cluster_name	Cluster name of the guest VM	✓	✓
vmware_vcenter_name	vCenter name of the guest VM	✓	✓
workload_platform	Platform Name - VMware	✓	✓

Azure Kubernetes Service (AKS)

Attribute	Description	Source	Destination
k8s_cluster_name	Name identifying the Kubernetes cluster	✓	✗
k8s_cluster_version	Version of Kubernetes running on the cluster	✓	✗
k8s_cluster_status	Operational state of the Kubernetes cluster	✓	✗
k8s_node_name	Identifier for the node within the cluster	✓	✗
k8s_node_internal_ip	Internal IP address assigned to the node	✓	✗
k8s_node_external_ip	External IP address of the node	✓	✗
k8s_node_os_version	Operating system version installed on the node	✓	✗
k8s_node_os_image	Name of the OS image used by the node	✓	✗
k8s_node_kernel_version	Kernel version of the node	✓	✗
k8s_node_cri_version	Version of the container runtime interface used on the node	✓	✗
k8s_node_status	Health and availability status of the node	✓	✗
k8s_node_labels	Node labels parsed into key-value	✓	✗

	pairs for metadata enrichment		
k8s_pod_name	Name assigned to the Kubernetes pod	✓	✗
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✓	✗
k8s_pod_namespace	Namespace in which the pod is deployed	✓	✗
k8s_pod_status	Operational state of the pod	✓	✗
k8s_pod_ip	IP address assigned to the pod	✓	✗
k8s_service_name	Name of the Kubernetes service	✓	✗
k8s_service_type	Type of service	✓	✗
k8s_container_name	Name of the container running inside the pod	✓	✗
k8s_container_image	Image used to create the container	✓	✗
k8s_container_port	Container port details	✓	✗
k8s_cluster_name	Name identifying the Kubernetes cluster	✗	✓
k8s_cluster_version	Version of Kubernetes running on the cluster	✗	✓
k8s_cluster_status	Operational state of the Kubernetes cluster	✗	✓
k8s_node_name	Identifier for the node within the cluster	✗	✓
k8s_node_internal_ip	Internal IP address assigned to the node	✗	✓
k8s_node_external_ip	External IP address of the node	✗	✓
k8s_node_os_version	Operating system version installed on the node	✗	✓
k8s_node_os_image	Name of the OS image used by the node	✗	✓
k8s_node_kernel_version	Kernel version of the node's operating system	✗	✓
k8s_node_cri_version	Version of the container runtime interface used on the node	✗	✓
k8s_node_status	Health and availability status of the node	✗	✓
k8s_node_labels	Node labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_name	Name assigned to the Kubernetes pod	✗	✓
k8s_pod_labels	Pod labels parsed into key-value pairs for metadata enrichment	✗	✓
k8s_pod_namespace	Namespace in which the pod is	✗	✓

	deployed		
k8s_pod_status	Operational state of the pod	✖	✓
k8s_pod_ip	IP address assigned to the pod	✖	✓
k8s_service_name	Name of the Kubernetes service	✖	✓
k8s_service_type	Type of service	✖	✓
k8s_container_name	Name of the container running inside the pod	✖	✓
k8s_container_image	Docker image used to create the container	✖	✓
k8s_container_port	Container port details	✖	✓

Zscaler Private Access (ZPA) Attributes

Attribute	Description
z_lss_username	Username of Zscaler Private Access user.
z_lss_client_private_ip	Private IP address of the Zscaler Private Access client
z_lss_client_public_ip	Public IP address of the Zscaler Private Access. client.
z_lss_customer	Name of Zscaler Private Access customer.
z_lss_connection_id	ID of the Zscaler Private Access connection.
z_lss_connection_status	Status of the Zscaler Private Access connection.
z_lss_connection_status_code	Status code of the Zscaler Private Access connection.
z_lss_connector	Name of the Zscaler Private Access application connector.
z_lss_hostname	Name of the host device as given by the Zscaler client connector.
z_lss_client_country_code	Country code of the Zscaler Private Access client location.
z_lss_policy	Name of the Zscaler Private Access policy.
z_lss_app_group	Group name of the Zscaler Private Access application.
z_lss_policy_processing_time	Processing time of the Zscaler Private Access policy.
z_lss_idp	Identity provider used by Zscaler Private Access.
z_lss_client_city	City of the Zscaler Private Access client location.
z_lss_platform	Client platform of Zscaler Private Access.
z_lss_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[1]	Supported for NetFlow v5	Supported for NetFlow v9	Supported for NetFlow v10/IPFIX
type	Type of flow message	✓	✓	✓
sampler_ address	Address of the device that generated the packet	✓	✓	✓
sampling_ rate	Sampling rate of the flow	✓	✓	✓
export_time_ ns	Export time in Nano seconds	✓	✓	✓
time_ received_ns	Timestamp in nanoseconds of when the message was received	✓	✓	✓
sequence_ num	Sequence number of the flow packet	✓	✓	✓
source_id	An identifier that uniquely identifies the flow source. For IPFIX it refers to the observation domain ID.	✗	✓	✓
etype	Ethernet type (0x86dd for IPv6...)	✓	✓	✓
src_addr	Source address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_SRC_ADDR(8), IPV6_ SRC_ADDR(27) • NetFlow v10/IPFIX - sourceIPv4Address(8), sourceIPv6Address(27)	✓	✓	✓
dst_addr	Destination address (IP) • NetFlow v5 - IPv4 only • NetFlow v9 - IPV4_DST_ADDR (12), IPV6_ DST_ADDR (28) • NetFlow v10/IPFIX - destinationIPv4Address(12), destinationIPv6Address(28)	✓	✓	✓
ip_tos	IP Type of Service • NetFlow v5 – tos • NetFlow v9 - SRC_TOS (5) • NetFlow v10/IPFIX - ipClassOfService (5)	✓	✓	✓
ip_dscp	Differentiated Services Code Point • ipDiffServCodePoint (195)	✗	✓	✓
ip_ precedence	The value of the IP Precedence • ipPrecedence (196)	✗	✓	✓
proto	Protocol (UDP, TCP, ICMP...) • NetFlow v5 – proto • NetFlow v9 - PROTOCOL (4)	✓	✓	✓

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

	• NetFlow v10/IPFIX - protocolIdentifier (4)			
ip_ttl_max	Maximum TTL value observed for packets of the flow • NetFlow v9 - MAX_TTL (53) • NetFlow v10/IPFIX - maximumTTL (53)	✗	✓	✓
src_port	Source port (when UDP/TCP/SCTP) • NetFlow v5 – srcport • NetFlow v9 - L4_SRC_PORT (7) • NetFlow v10/IPFIX - sourceTransportPort (7)	✓	✓	✓
dst_port	Destination port (when UDP/TCP/SCTP) • NetFlow v5 – dstport • NetFlow v9 - L4_DST_PORT (11) • NetFlow v10/IPFIX - destinationTransportPort (11)	✓	✓	✓
bytes	Number of bytes in a flow • NetFlow v5 – dOctets • NetFlow v9 - IN_BYTES (1), OUT_BYTES (23) • NetFlow v10/IPFIX - octetDeltaCount (1), postOctetDeltaCount (23)	✓	✓	✓
packets	Number of packets in a flow • NetFlow v5 – dPkts • NetFlow v9 - IN_PKTS (2), OUT_PKTS (24) • NetFlow v10/IPFIX - packetDeltaCount (2), postPacketDeltaCount (24)	✓	✓	✓
bytes_total	Running byte counter for a permanent flow • NetFlow v9 - IN_PERMANENT_BYTES (85) • NetFlow v10/IPFIX - octetTotalCount(85)	✗	✓	✓
packets_total	Running packet counter for a permanent flow • NetFlow v9 - IN_PERMANENT_PKTS (86) • NetFlow v10/IPFIX - packetTotalCount(86)	✗	✓	✓
flow_ direction	Flow direction • NetFlow v9 - DIRECTION (61) • NetFlow v10/IPFIX - flowDirection(61)	✗	✓	✓
in_if	Input interface • NetFlow v5 – input • NetFlow v9 - INPUT_SNMP (10) • NetFlow v10/IPFIX - ingressInterface (10)	✓	✓	✓
out_if	Output interface • NetFlow v5 – output • NetFlow v9 - OUTPUT_SNMP (14) • NetFlow v10/IPFIX - egressInterface (14)	✓	✓	✓

time_flow_start_ns	Time the flow started in nanoseconds. Refer to Important Notes. • NetFlow v5 – System uptime and first • NetFlow v9 - System uptime and FIRST_SWITCHED (22) • NetFlow v10 - flowStartXXX (150, 152, 154, 156)	✓	✓	✓
time_flow_end_ns	Time the flow ended in nanoseconds. Refer to Important Notes. • NetFlow v5 – System uptime and last • NetFlow v9 - System uptime and LAST_SWITCHED (21) • NetFlow v10/IPFIX - flowEndXXX (151, 153, 155, 157)	✓	✓	✓
tcp_flags	TCP flags • NetFlow v5 – tcp_flags • NetFlow v9 - TCP_FLAGS (6) • NetFlow v10/IPFIX - tcpControlBits (6)	✓	✓	✓
ip_length_min	Minimum IP packet length • NetFlow v9 - MIN_PKT_LNGTH (25) • NetFlow v10/IPFIX - minimumIpTotalLength (25)	✗	✓	✓
ip_length_max	Maximum IP packet length • NetFlow v9 - MAX_PKT_LNGTH (26) • NetFlow v10/IPFIX - maximumIpTotalLength (26)	✗	✓	✓
next_hop	Nexthop IP address • NetFlow v5 – nexthop • NetFlow v9 - IPV4_NEXT_HOP (15), IPV6_NEXT_HOP (62) • NetFlow v10/IPFIX - ipNextHopIPv4Address (15), ipNextHopIPv6Address (62)	✓	✓	✓
ipv4_next_hop	Nexthop IPv4 address • NetFlow v9 - IPV4_NEXT_HOP (15) • NetFlow v10/IPFIX - ipNextHopIPv4Address (15)	✗	✓	✓
ipv6_next_hop	Nexthop IPv6 address • NetFlow v9 - IPV6_NEXT_HOP (62) • NetFlow v10/IPFIX - ipNextHopIPv6Address (62)	✗	✓	✓
flow_id	An identifier of a flow that is unique within an Observation Domain • NetFlow v9 - flowId (148) • NetFlow v10/IPFIX - flowId (148)	✗	✓	✓

firewall_event	Indicates a firewall event. Allowed values are listed in the firewall Event registry. • firewallEvent(233)	✖	✓	✓
icmp_type	Type of the ICMP message • NetFlow v9 - ICMP_TYPE (32) • NetFlow v10/IPFIX - icmpTypeXXX (176, 178), icmpTypeCodeXXX (32, 139)	✖	✓	✓
icmp_code	Code of the ICMP message • NetFlow v9 - ICMP_TYPE (32) • NetFlow v10/IPFIX - icmpCodeXXX (177, 179) icmpTypeCodeXXX (32, 139)	✖	✓	✓
flow_end_reason	The reason for flow termination. Values are listed in the flowEndReason registry.	✖	✓	✓
application_id	Specifies an Application ID • NetFlow v9 - APPLICATION TAG (95) • NetFlow v10/IPFIX - applicationId (95)	✖	✓	✓
application_description	Specifies the description of an application • NetFlow v9 - APPLICATION DESCRIPTION (94) • NetFlow v10/IPFIX - applicationDescription (94)	✖	✓	✓
application_name	Specifies the name of an application • NetFlow v9 - APPLICATION NAME (96) • NetFlow v10/IPFIX - applicationName (96)	✖	✓	✓
tcp_window_size	The window field in the TCP header • tcpWindowsSize (186)	✖	✓	✓
viptela_vpn_id	Private enterprise number for vIPtela Inc. (41916)	✖	✖	✓

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.

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

Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.

</td></tr>
<tr><td>Fast Mode</td><td>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.</td></tr>
<tr><td>Aggregate Round-trip Time</td><td>Enable this option to export the minimum, maximum, and mean of RTT values for the following list of supported protocols and attributes and also the aggregate of TCP Lost byte values collected per export time interval.
<table>
<tr><th>Protocol</th><th>Arrtibute</th></tr>
<tr><td>http</td><td>rtt</td></tr>
<tr><td>icmp</td><td>rtt</td></tr>
<tr><td>icmp6</td><td>rtt</td></tr>
<tr><td>ssh</td><td>rtt</td></tr>
<tr><td>tcp</td><td>rtt</td></tr>
<tr><td>tcp</td><td>rtt_app</td></tr>
<tr><td>telnet</td><td>rtt</td></tr>
<tr><td>wsp</td><td>connect_rtt</td></tr>
<tr><td>wsp</td><td>query_rtt</td></tr>
</table>
</td></tr>
</table>

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 *Tool Templates*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 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:

<table>
<tr><th>Parameter</th><th>Description</th></tr>
<tr><td>Name</td><td>Enter a name for the application.</td></tr>
<tr><td>Description</td><td>Enter the description.</td></tr>
<tr><td>Export Interval</td><td>The time interval in seconds at which the export must be done. The export interval is set to 300 seconds. It cannot be modified.</td></tr>
<tr><td colspan="2">Advanced Settings</td></tr>
<tr><td>Number of Flows</td><td>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>
<tr><th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr>
<tr><td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
<tr><td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
</table>

Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.
</td></tr>
<tr><td>Monitoring</td><td>You can use this option, enable or disable the Application Visualization application functionality.</td></tr>
<tr><td>Fast Mode</td><td>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.</td></tr>
</table>

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 sectionfor 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 with in the past 7 days.

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

De-duplication

De-duplication application targets, identifies, and eliminates duplicate packets, blocking unnecessary duplication and sending optimized flows to your security and network monitoring tools. De-duplication lets you detect and choose the duplicate packets to count or drop in a network analysis environment.

Duplicate packets are common in network analysis environments where both the ingress and egress data paths are sent to a single output. They can also appear when packets are gathered from multiple collection points along a path. The de-duplication application lets you eliminate these packets, only forwarding a packet once and thus reducing the processing load on your tools.

Refer to the following topics:

- Feature Overview
- Configure De-duplication Application
- Distributed De-duplication

Feature Overview

There are two actions that can be specified for handling the duplicate packets detected:

- drop, which drops the duplicate packets
- count, which counts the duplicate packets, but does not drop them

A time interval can be configured within which an identical packet will be considered a duplicate. The greater the interval over which traffic can be checked for duplicates, the higher the accuracy of the de-duplication detection and subsequent elimination.

For example, if two of the same packets are seen in the specified time interval, the packets will be detected as duplicates. If one packet is seen in the time interval and another packet is seen in a later time interval, the packets will not be detected as duplicates.

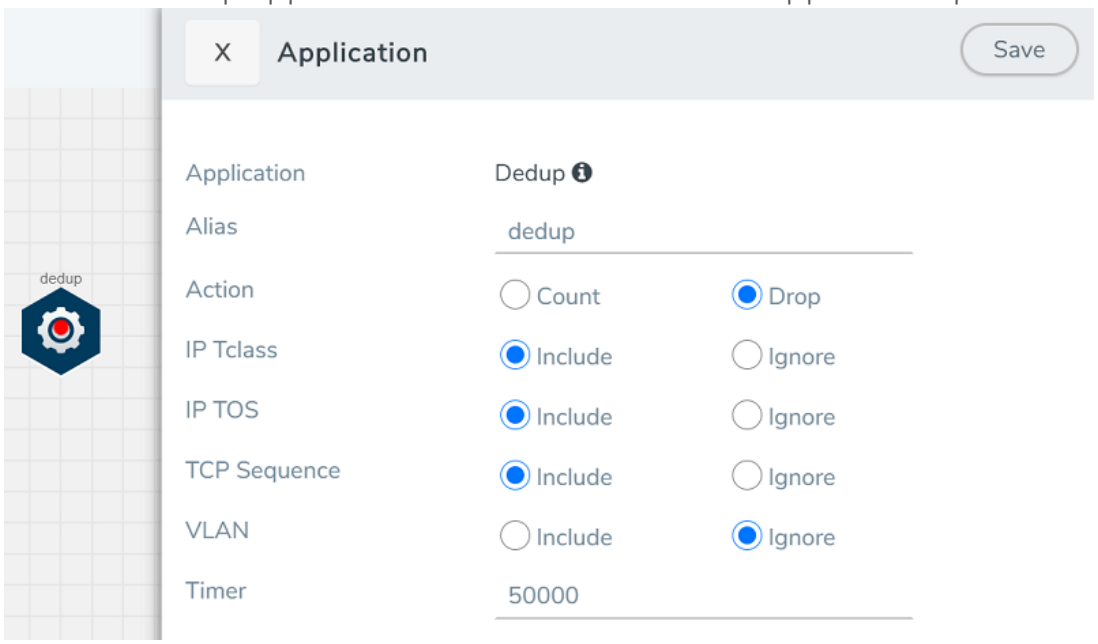
For IPv4 and IPv6 packets, to determine if a packet is considered to be a duplicate, parts of the IP headers (Layer 3 and Layer 4), as well as part of the payload are compared.

For non-IP packets, a packet is considered to be a duplicate if it is identical.

Configure De-duplication Application

To add a de-duplication application:

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



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.

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:

<table>
<tr><th>Parameter</th><th>Description</th></tr>
<tr><td>Name</td><td>Enter a name for the application.</td></tr>
<tr><td>Description</td><td>Enter the description.</td></tr>
<tr><td colspan="2">Application Metadata Settings</td></tr>
<tr><td>Flow Direction</td><td>Enable or Disable Bi-Directional Flow behavior. Bi-Directional is enabled by default. Disable this option for Uni-Directional Flow behavior.</td></tr>
<tr><td>Timeout</td><td>Specify the traffic flow inactivity timeout, in seconds. The session will be removed due to inactivity when no packets match.</td></tr>
<tr><td>Data Link</td><td>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.</td></tr>
<tr><td>Observation ID</td><td>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.</td></tr>
<tr><td colspan="2">Advanced Settings</td></tr>
<tr><td>Number of Flows</td><td>The number of flows supported by the application.
Refer to the following table for the maximum number of flows supported for VMware, AWS, and Azure platforms.
<table>
<tr><th>Cloud Platform</th><th>Instance Size</th><th>Maximum Number of Flows</th></tr>
<tr><td>VMware</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
<tr><td>Nutanix</td><td>Large (8 vCPU and 16GB RAM)</td><td>200k</td></tr>
</table>

Medium Form Factor is supported for VMware ESXi only when secure tunnels option is disabled. The maximum Number of Flows for VMware ESXi when using a medium Form Factor is 50k..
If the rate of unique UDP sessions per second exceeds the threshold—calculated as maximum number of flows per second divided by the UDP timeout value—the system may fail to classify applications correctly. In such cases, AFI may not filter packets accurately, resulting in incorrect packet passes or drops. However, this limitation does not apply to DNS flows.
</td></tr>
</table>

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**.

GRE-In-UDP Tunnel Decapsulation

The GRE-In-UDP tunnel is a new tunnel type that supports decapsulating the GRE-in-UDP headers in the vTAP captured data packets. The decapsulation of the tunnel headers is common for both control and data packets. However, the control packets (with JSON metadata format) are dropped based on the GRE Key value '0'. The data packets (either fragmented or unfragmented) are handled as follows:

S. No	Process	Description
1	**Reassembly of fragmented data packets (if fragmented packets are received)**	The fragmented packets are reassembled based on standard IP header-based reassembly using the following fields: • More Fragments (MF) field in IPv4 packets • Identification field in IPv6 packets
2	**Decap tunnel lookup**	The ingress packet fields are parsed to find a matching Tunnel End Point (TEP) based on the following header fields:

S. No	Process	Description
		• Destination IP • Source IP • UDP port numbers • GRE key
3	**Decapsulation of headers**	If a matching TEP is found, the encapsulation headers in the packet are removed and the packet is forwarded to the PCAPng application. Refer to the PCAPng Application section for more details.

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 striping:

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

Note:

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

- VMware vCenter
- VMware NSX-T
- Nutanix

Limitation

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

Configure Header Stripping Application

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

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

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

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

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

4. Click **Save**.

Create Custom Port Template

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

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

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

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

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

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

What's Next

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

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

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

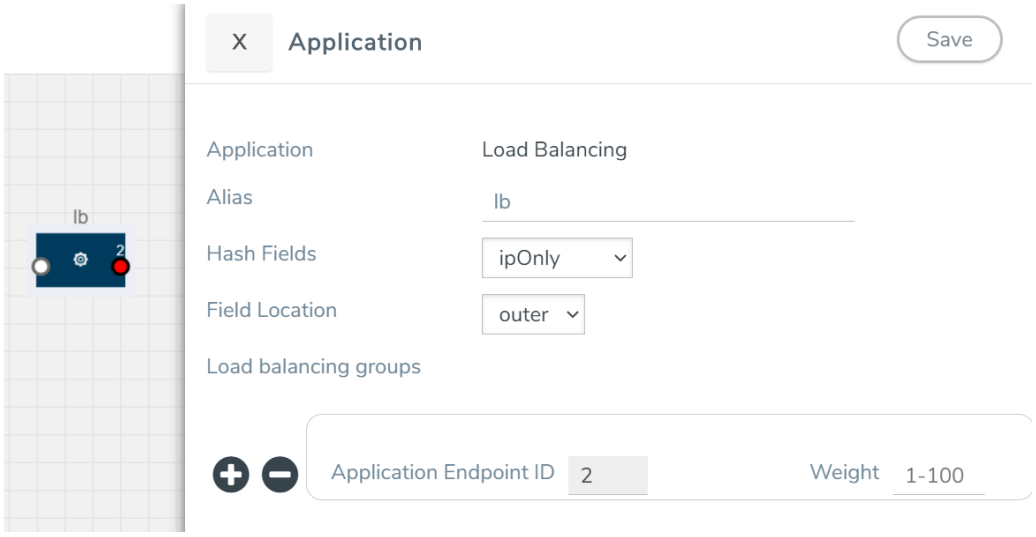
Load Balancing

Load balancing application performs stateless distribution or Enhanced Load Balacing 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.



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. • **fiveTupl**e: 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

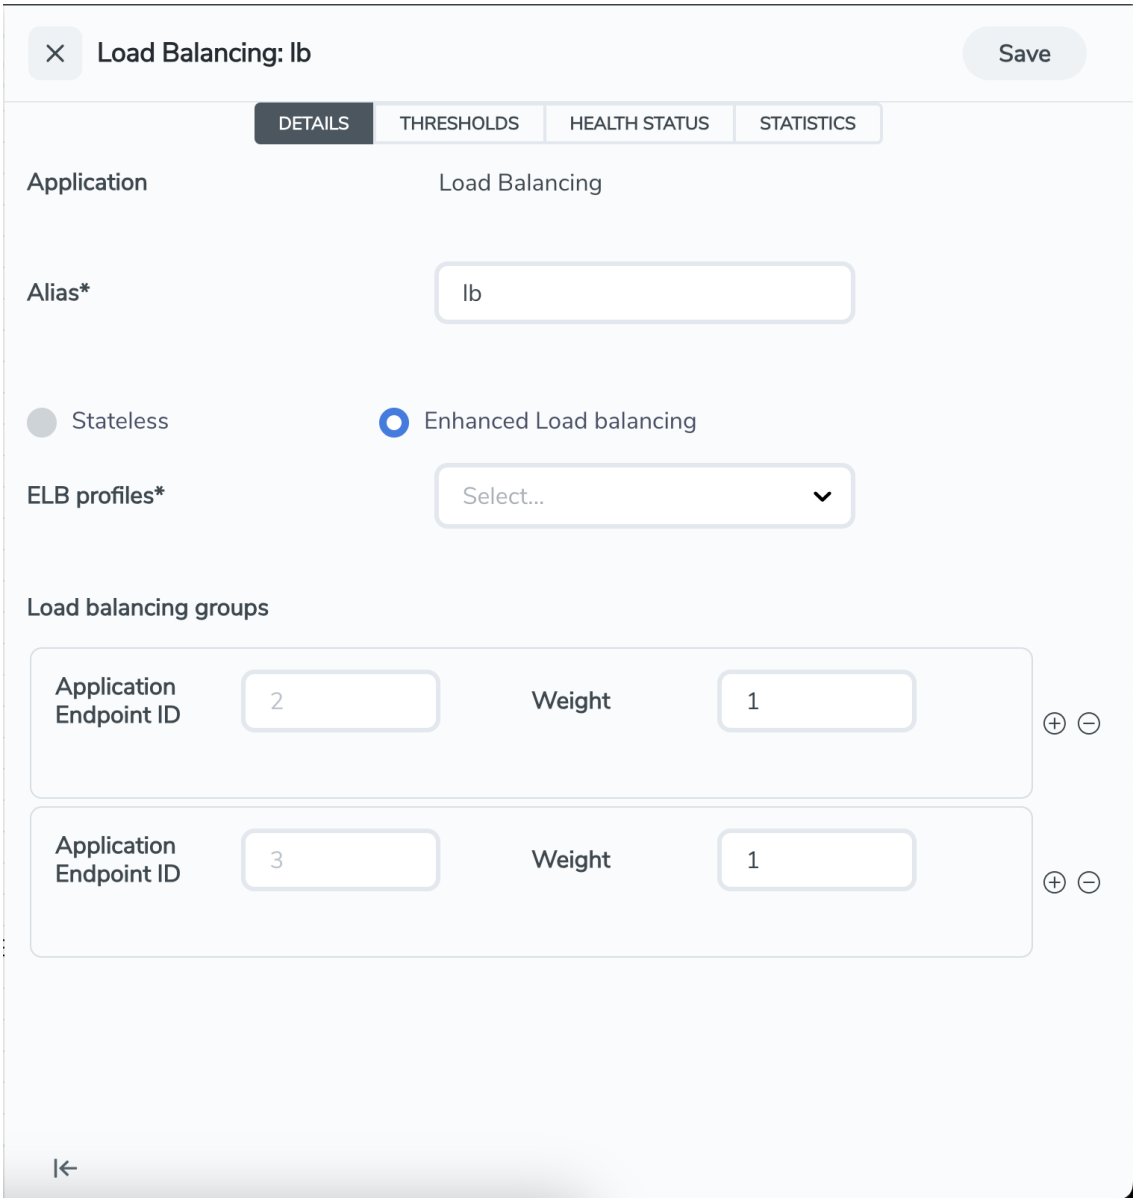
Enhanced Load Balancing redistributes the traffic from a failed endpoint to another endpoint which is a part of the ELB group. When the failed endpoint recovers, the redistributed traffic is restored to the recovered endpoint. The traffic across other endpoints remain undisturbed during this process. 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 gtputeid
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.



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

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

3. Click **Save**.

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

What's Next

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

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

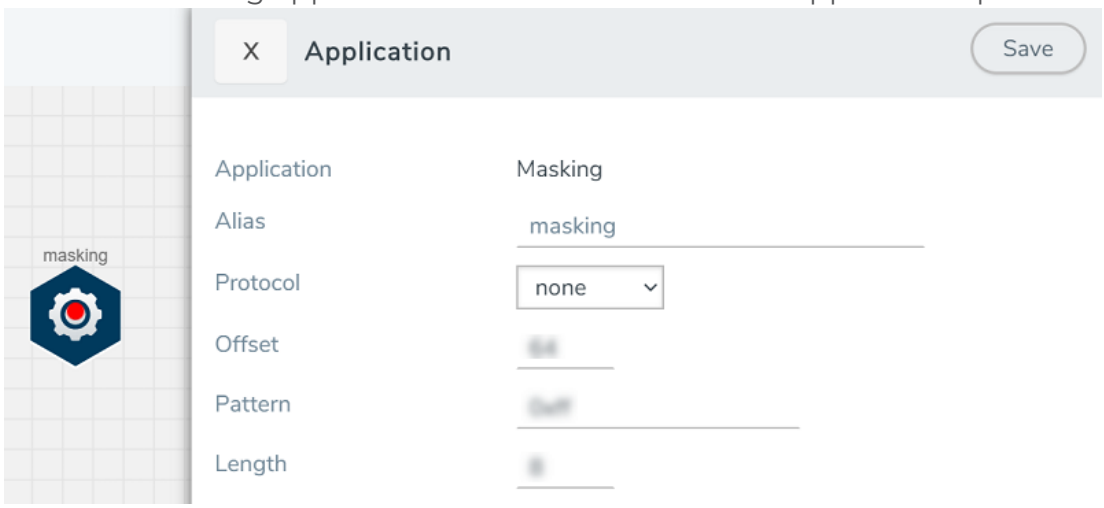
You can view the statistics of the application in the **STATISTICS** tab.

Masking

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

To add a Masking application:

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



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

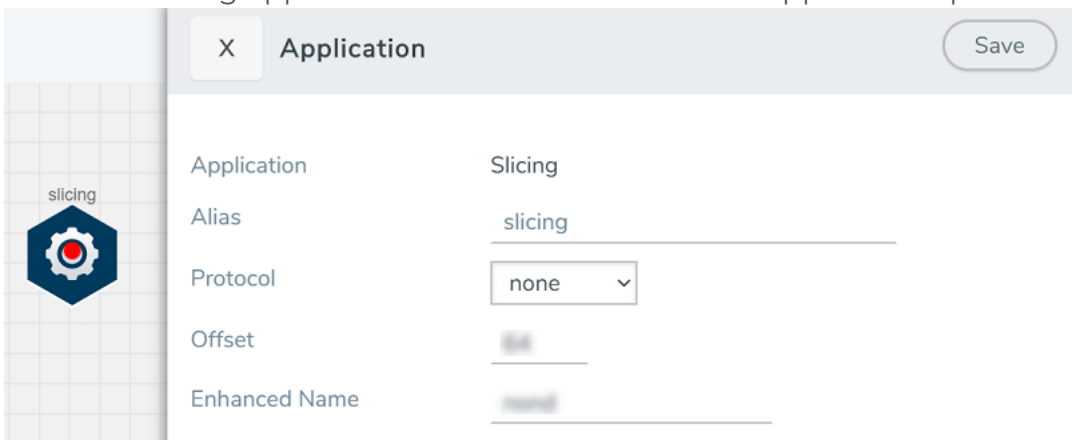
Packet Slicing allows you to truncate packets after a specified header and slice length, preserving the portion of the packet required for monitoring purposes. Slicing operations are typically configured to preserve specific packet header information, allowing effective network analysis without the overhead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



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

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

4. Click **Save**.

What's Next

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

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

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

SSL Decrypt

License: For information on licensing, refer to the *)* section of respective GigaVUE Cloud Suite Deployment Guide.

SSL Decrypt application delivers decrypted traffic to out-of-band tools that can then detect threats entering the network. Secure Socket Layer (SSL) is a cryptographic protocol that adds security to TCP/IP communications such as Web browsing and email. The protocol allows the transmission of secure data between a server and client who both have the keys to decode the transmission and the certificates to verify trust between them.

SSL encryption secures traffic between a client and a server, such as a Web server. SSL decryption uses keys to decode the traffic between the client and server.

SSL and Transport Layer Security (TLS) protocols consist of a set of messages exchanged between a client and server to set up and tear down the SSL connection between them. To set up the connection, the client and server use the Public Key Infrastructure (PKI) to exchange the bulk encryption keys needed for data transfer.

IMPORTANT: To use SSL Decrypt application in GigaVUE-FM 6.3.00, install new GigaVUE-FM 6.3.00 image. Refer to *GigaVUE-FM Installation and Upgrade Guide* for step-by-step instructions on how to install GigaVUE-FM. SSL Decrypt application does not work if you upgrade from any previous GigaVUE-FM version to GigaVUE-FM 6.3.00.

Keep in mind the following when using SSL Decrypt application:

1. On updating the keys, service, or key maps which are already used in a monitoring session, the monitoring session is dynamically updated, and you need not re-deploy the monitoring session. You can also see if the updated keys, services, or key maps were successfully updated to the monitoring session and the respective GigaVUE V Series Nodes on the **All Events** page. Refer to Overview of Events section in the *GigaVUE Administration Guide* for detailed information on Events.
2. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the key will be removed from the Key Map. If that key is the only available entry in the Key Map, then it will not be removed.
3. When deleting a key that is part of a Key Map and that Key Map is used in a monitoring session that is not deployed, then the key will be removed from the Key Map and if that key is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
4. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is already deployed, then the service will be removed from the Key Map. If that service is the only available entry in the Key Map, then it will not be removed.

5. When deleting a service that is part of a Key Map and that Key Map is used in a monitoring session which is not deployed, then the service will be removed from the Key Map and if that service is the only available entry in the Key Map, the whole key map will be removed from the monitoring session.
6. In VMware NSX-T platform, the throughput of SSL Decrypt application is improved to 480 Mbps.

Refer to the following topics for more detailed information:

- Supported Protocols, Algorithms, and Ciphers for SSL Decrypt
- Configure SSL Decrypt

Supported Protocols, Algorithms, and Ciphers for SSL Decrypt

The supported protocols are as follows:

- SSL 3.0
- TLS 1.0
- TLS 1.1
- TLS 1.2

The supported authentication (Au) is as follows:

- RSA

The supported key exchange (Kx) is as follows:

- RSA

The supported encryption algorithms (Enc) are as follows:

- NULL
- RC4
- DES
- 3DES
- AES (including GCM mode)
- CAMELLIA
- SEED
- IDEA

The supported compression algorithm is as follows:

- NULL

The supported digest algorithms are as follows:

- MD5
- SHA1
- SHA2

The supported key sizes are 128, 256, 512, 1024, 2048, and 4096.

The supported TLS extensions are as follows:

- Extended Master Secret, RFC 7627
- Encrypt-then-MAC, RFC 7366

The following table lists the supported ciphers:

Table 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

The supported applications are as follows:

- HTTPS
- FTPS
- SMTP, IMAP, and POP3 with StartTLS

Configure SSL Decrypt

To configure SSL Decrypt Application using GigaVUE-FM follow the steps given below:

- Upload SSL Key Pairs
- Create SSL Service
- Key Mapping
- SSL Key Pair Store
- Add SSL Decrypt to Monitoring Session

Upload SSL Key Pairs

To upload an SSL private key, do the following:

1. Go to **Inventory > Resources > Security > SSL Keys**.
2. Click **Add**. The **Create SSL Key Pair** page appears.
3. Enter the following details:

Field	Description
Key Alias	Enter a name for the key.
Comment	Enter a description
Key Type	Select the either of the key type:
PEM	
PassPhrase (optional)	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.
Private key	Enter the Private Key using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
Certificate	Enter the Certificate using any of the following options: • Copy and Paste • Install from URL • Install from Local Directory
PKCS12	
PassPhrase	SSH passphrases allows you to protect your private key from being used with out the passphrase. Enter the passphrase created with the private key.
SSL Key Store	Enter the SSL Key Pair Store in which the Key is stored.

4. Click **Save**.

NOTE: SSL Decrypt application does not support HSM.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Key pair, select the key from the list in the SSL Key Pairs page and click the **Edit** button.
Delete	To delete a SSL Key pair, select the key from the list in the SSL Key pairs page and click the **Delete** button.
Delete all	Use this button to delete all the SSL Key pairs in the SSL Key pairs page.
View certificate	To view the certificate associated with the particular SSL Key Pair, select the key from the list in the SSL Key Pairs page and click the **View Certificate** button.

Create SSL Service

After uploading a private key, you can add a service. A service maps to a physical server, such as an HTTP server. One server can run multiple services. A service is a combination of an IP address and a server port number.

Prerequisite

Before creating a service, upload a private key as described in Upload SSL Key Pairs

To create a SSL service, do the following:

1. Go to **Inventory > Resources > Security > SSL Service**. The SSL Services page appears.
2. Click **Add**. The **Create SSL Service** page appears.
3. On the Create SSL Service page, enter the following details:

Field	Description
Alias	Enter a name for the SSL Service.
Default Service	Enable this to use default service.
Server IP Address	Enter the IP address of the server in which the service runs.
Server Port	Enter the port number of the server.

4. Click **Save**.

The following actions can also be performed from the SSL Service Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the **Edit** button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the **Delete** button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

Key Mapping

After adding the SSL Service, now you map the private key with the service using Key Mapping.

To map a key with the service, follow the steps given below,

1. Go to **Inventory > Resources > Security > SSL Key Mapping**. The SSL Key Pair Mapping page appears.
2. Click **Add**.

3. Enter the Key Mapping Alias.
4. Select the SSL Service and Key Alias from the drop-down.
5. Click **Save**.

The following actions can also be performed from the SSL Key Pairs Page:

Field	Description
Edit	To edit a SSL Service, select the service from the list in the SSL Service page and click the **Edit** button.
Delete	To delete a SSL Service, select the service from the list in the SSL Service page and click the **Delete** button.
Delete all	Use this button to delete all the SSL Service in the SSL Service page.

SSL Key Pair Store

SSL Key Pair Store is a repository, that allows you to save all the key under a single location. You can create multiple key stores and in each key store you can store multiple keys.

1. Go to **Inventory > Resources > Security > SSL Key Store**. The SSL Key Pair Store page appears.
2. Click **Add**.
3. Enter the **Key Store Alias** and **Comment**.
4. Click **Save**.

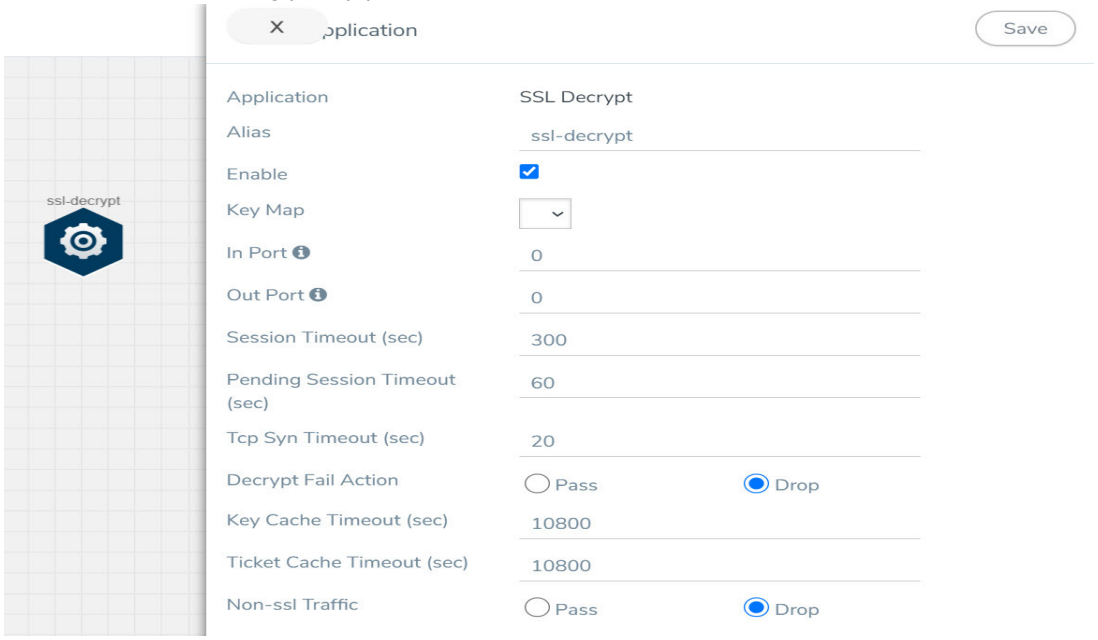
The following actions can also be performed from the SSL Key Pair Store Page:

Field	Description
Edit	To edit a SSL Key pair Store, select the Key Store from the list in the SSL Key Pair Store page and click the **Edit** button.
Delete	To delete a SSL SSL Key Pair Store, select the SSL Key Pair Store from the list in the SSL Key Pair Store page and click the **Delete** button.
Delete all	Use this button to delete all the SSL Key Pair Store in the SSL Key Pair Store page.

Add SSL Decrypt to Monitoring Session

After mapping your keys with service, to add GigaSMART applications to GigaVUE V Series Node, follow the steps given below,

1. Drag and drop **SSL Decrypt** from APPLICATIONS to the graphical workspace.
2. Click the SSL Decrypt application and select **Details**.



3. Enter the following details in the Application quick view:

Fields	Description
Alias	Enter the alias name for the application.
Enable	Enable the box to enable SSL Decryption.
Key Map	Select the Key Map from the list of available Key Maps. Refer to Key Mapping for more details on how to map the key to SSL Service.
In Port	Enter the source port number from which the traffic should be fetched.
Out Port	Enter the destination port number to which the decrypted traffic should be delivered.
Session Timeout	Enter the value in seconds after which the session should be timeout. The default value is 300 seconds.
Pending Session Timeout	Enter the value in seconds after which the session must timeout if the session is in pending state
Tcp Syn Timeout	Enter the value in seconds after which the session must timeout when the session does not synchronize TCP.
Decrypt Fail Action	Select **Pass** to allow the traffic to pass through the application when the decryption fails and select **Drop** to drop the traffic before passing through the application when the decryption fails.
Key Cache Timeout (sec)	Enter the value in seconds until which the key cache information can be reused for resumption.
Ticket Cache Timeout (sec)	Enter the value in seconds until which the ticket cache information can be reused for resumption.
Non-ssl Traffic	Select **Pass** to allow the non-SSL traffic to pass through the application and select **Drop** to drop the non- SSL traffic before passing through the application.

4. Click **Save**.
5. Click **Deploy**. The Select nodes to deploy the monitoring session page appears.
6. Select the GigaVUE V Series Nodes you want to deploy and select an interface for each GigaVUE V Series Node. Then, click **Deploy**.

5G-Service Based Interface Application

The 5G-Service Based Interface (5G-SBI) application synthesizes HTTP2 transactions with proper L2, L3, and L4 headers from JSON-encoded data received from Ericsson vTAPs via UDP-GRE. The 5G-SBI application takes the raw encoded data that monitoring tools cannot understand and reconstructs complete, tool-facing packets with valid TCP flows.

The 5G-SBI application operates on a DPDK-based V Series Node and supports the following traffic types:

- **SBI traffic (L7-JSON)** - AMF, SMF signaling from Ericsson PCC vTAPs (key value 1). The 5G-SBI application synthesizes this traffic into HTTP2 packets.
- **Non-SBI traffic** - PSVP, GTPC, Diameter, and other protocols from Ericsson PCC vTAPs (key values 5, 7, 9, 11, 15, 21, 23). The PCAPng application reads the data from the Enhanced Packet Block and forwards it directly to the egress tunnel without synthesis.

Ericsson vTAP

Ericsson virtual TAPs (vTAPs) mirror traffic from 5G Core network functions. Each network function module has its own vTAP. Ericsson PCC (Packet Core Controller) contains vTAPs for AMF, SMF, and other protocols. Ericsson CCx (Cloud Core) contains vTAPs for SCP traffic.

Each vTAP assigns a key value to identify the protocol type. The vTAP wraps the mirrored packet in PCAPng format (Enhanced Packet Block) and sends it over a UDP-GRE tunnel to the GigaVUE V Series Node.

The following table lists the Ericsson vTAP key values and their corresponding traffic types:

Key Value	Traffic Type	Category
0,2	Metadata (domain classification)	Control
1	AMF, SMF	SBI (L7-JSON)
3, 17, 19	SCP protocols	SBI (L7-Native)
5, 7, 9, 11, 15, 21, 23	PSVP, GTPC, Diameter	Non-SBI

NOTE: Key values 0 and 2 carry domain metadata only. This traffic is processed by the PCAPng application in Primary mode to identify which flow belongs to which domain. This traffic does not go to monitoring probes.

Ericsson vTAP Domain Tagging

Ericsson vTAP components mirror control plane and selected user plane traffic, for example, from Ericsson PCC nodes, and send it to the GigaVUE V Series using UDP-GRE with

PCAPng/JSON metadata. Previously, the 5G-SBI application treated all mirrored flows uniformly, requiring monitoring tools to independently infer the Ericsson domain (for example, pc-mm, pc-sm, sc-scp) and the associated 3GPP interface.

With Domain Tagging, the PCAPng application on V Series reads the Ericsson vTAP metadata, identifies the Ericsson domain for each flow, and assigns a VLAN ID based on a configurable domain to VLAN map. When you enable domain classification in the PCAPNG app and VLAN tagging on the VXLAN or L2GRE egress tunnel, V Series tags each packet with the appropriate VLAN so probes can separate domains simply by filtering on VLAN.

You can also reserve VLANs for:

- Traffic mapped to the **others** domain (for example, unsupported or unexpected Ericsson domains)
- Traffic mapped to the **unclassified** domain (for example, traffic for which domain information is not yet available)

Encoded messages from the Ericsson vTAP contain Kubernetes Pod ephemeral IP addresses for NF instances, which are internal to the Ericsson deployment and not accessible outside the Kubernetes cluster. Because these IPs cannot be used externally, the 5G SBI application always maps the NF identifiers (NF instance ID and FQDN) in the encoded message to the configured ingress IP address using the string-to-IP mapping table.

The NF instance ID and FQDN ID must be provided in the form of CSV file. You can upload the CSV file through GigaVUE-FM.

How SBI Application works

In the GigaVUE V Series Node, the SBI application receives the HTTP2 transaction messages as JSON encoded data from the PCAPng application, which in turn receives the data from the UDP-GRE TEP.

The SBI application:

1. Parses the JSON data to extract source and destination information, ports, and message type (request/response).
2. Synthesizes a complete HTTP/2 transaction with L2, L3, and L4 headers, along with the HTTP/2 headers and body from the original transactions.
3. Sends the synthesized packets to the egress TEP, which forwards them to the appropriate virtual probes.

The following image shows the block diagram of the data flow in the V Series Node containing the SBI application.

Supported Platforms:

The application is supported on the following platforms:

- VMware ESXi
- OpenStack
- Third Party Orchestration

Rules and Notes

- The maximum number of HTTP2 headers (in the synthesized HTTP2 transactions) that is supported is 64.
- The PCAPng application that is linked to 5G-SBI application (on the right side) should only be linked to UDP-GRE TEP with key value 1 on the left side. If it is linked to other UDP-GRE TEPs (key values other than 1), then the behavior cannot be defined and leads to unexpected results.
- The maximum number of NF and FQDN entries supported is 4K.
- In 5G-SBI application, the V Series node can log the following details to CSV files:
  - **Transaction details** - Represents the transaction or flow of request and response packets into the application. The details of the flow or transaction are recorded in the CSV file for 5 minutes or 60 minutes based on the configuration.
  - **Flow statistics details** - Represents the packet and flow statistics in 60 seconds time interval.

  These files help you to understand the records and traffic efficiently. The files are named as per the date and time in which the files were created. When the number and size of files grow, the application automatically detects the old files and deletes them.
- Ericsson vTAP domain tagging can be configured only in primary mode when configuring Pcapng application. Refer to 5G-Service Based Interface Application.

5G-Ericsson vTAP Packet Processing (5G-EVP)

The 5G-Ericsson vTAP Packet Processing (5G-EVP) application orders packets received from Ericsson vTAP sources and optionally decrypts DTLS(Datagram Transport Layer Security)-encrypted traffic before forwarding it to the downstream 5G-SBI and 5G-Cloud applications.

When multiple Ericsson vTAPs send traffic simultaneously, packets can arrive at the V Series Node out of order. Out-of-order packets can prevent monitoring probes from building session tables correctly, resulting in lost traffic visibility. The 5G-EVP application addresses this by ordering packets based on the timestamp at which each packet was originally observed on the network, ensuring that downstream applications and monitoring probes receive traffic in the correct sequence.

The 5G-EVP application runs on a dedicated V Series Node and sits in front of the existing 5G-SBI and 5G-Cloud nodes. It does not modify or replace those applications.

NOTE: The 5G-SBI and 5G-Cloud applications can operate without the 5G-EVP application. If packet ordering is not required, use the existing two-node deployment. However, the 5G-EVP application requires a downstream 5G-SBI or 5G-Cloud application to function.

Solution Architecture

The complete Ericsson vTAP solution with the 5G-EVP application uses three V Series Nodes managed by GigaVUE-FM. Each V Series Node requires a separate Monitoring Domain and a separate Monitoring Session.

V Series Apps	GigaVUE-FM
5G-EVP V Series Node	Orders packets from Ericsson vTAPs based on embedded timestamps and forwards ordered traffic to the downstream V Series Nodes over UDP-GRE. Optionally decrypts DTLS-encrypted traffic.
5G-SBI V Series Node	SBI and Non-SBI pipeline: Processes SBI (L7-JSON) traffic through the 5G-SBI application for synthesis. Forwards non-SBI traffic directly to egress tunnels.
	SCP pipeline: Distributes SCP (L7-Native) traffic through the LB application to multi-tap egress tunnels for downstream processing by the 5G-Cloud Node.
5G-Cloud V Series Node	Processes SBI (L7-Native) traffic from Ericsson SCP using the 5G-Cloud application.
GigaVUE-FM	Configures and monitors all three V Series Nodes. Manages NTP server configuration for the 5G-EVP Node.

Recommended VM Specifications

V Series Node	vCPUs	Memory	Disk Space
5G-EVP	20	32 GB	80 GB
5G-SBI	8	16 GB	80 GB
5G-Cloud	24	32 GB	80 GB

Recommended VM Specifications (DTLS)

V Series Node	vCPUs	Memory	Disk Space
5G-EVP	32	32 GB	80 GB
5G-SBI	8	16 GB	80 GB
5G-Cloud	24	32 GB	80 GB

Supported Deployment Combinations

The following table lists the supported deployment combinations. SBI and non-SBI traffic always coexist on the same 5G-SBI monitoring session. SCP traffic can be deployed independently or alongside SBI/non-SBI traffic.

Combination	Monitoring Sessions Required
SBI + Non-SBI traffic	EVP + 5G-SBI
SCP traffic	EVP + LB App + Multi-TEP VM + 5G-Cloud
SBI + Non-SBI + SCP	EVP + 5G-SBI + 5G-Cloud

Supported Platforms

- VMware ESXi (FM orchestration and third-party orchestration)
- OpenStack (FM orchestration and third-party orchestration)

Licensing

No new license is required. The 5G-EVP application is included in the existing V Series 5GC add-on license.

Overview of Packet Ordering

When multiple Ericsson vTAPs send traffic simultaneously, packets can arrive at the V Series Node out of order. Out-of-order packets can prevent monitoring probes from building session tables correctly, resulting in lost traffic visibility.

The 5G-EVP application orders packets based on the timestamp at which each packet was originally observed on the network. The application maintains a configurable time window and uses this window to hold incoming packets, order them by observed timestamp, and forward them in the correct sequence to the downstream 5G-SBI or 5G-Cloud Node.

Packets that arrive with timestamps outside the current ordering window are handled based on the configured policy.

NOTE: The maximum sorting window is 100 ms (numBuckets × bucketInterval must not exceed 100). The application validates this constraint and returns an error if the values exceed the limit.

DTLS Support

Ericsson can send vTAP traffic over DTLS (Datagram Transport Layer Security) for encrypted delivery. When you enable DTLS on the 5G-EVP application:

- Ericsson sends encrypted traffic on UDP port 4755 (instead of 4754 for non-secure traffic).
- The 5G-EVP application decrypts the traffic using the configured private key and certificate.
- The application forwards cleartext packets to the downstream 5G-SBI application over UDP-GRE on port 4754. The downstream applications see no difference.

Supported Modes

NOTE: You cannot disable both packet ordering and DTLS simultaneously. At least one must be enabled for the 5G-EVP application to function.

NTP Time Synchronization

The 5G-EVP application requires NTP time synchronization to maintain accurate clock alignment for packet ordering. When you enable packet ordering, the Primary NTP Server and Secondary NTP Server fields appear in the 5G-EVP application configuration.

GigaVUE-FM sends the NTP server configuration to the V Series Node during each deployment.

- If NTP synchronization fails after deployment, the 5G-EVP application health is impacted and GigaVUE-FM generates an event.
- If the chronyd configuration fails during deployment, GigaVUE-FM still proceeds with the deployment. However, packet ordering accuracy may be affected.

NOTE: The deployment may take additional time to complete due to NTP synchronization. Allow sufficient time for the NTP configuration to take effect before verifying application health.

Multiple Egress IP Addresses

When the 5G-EVP application forwards traffic to a downstream V Series Node using a single source IP, all traffic hashes to a single CPU core on the downstream node. To distribute traffic evenly across cores, configure up to 8 source IP addresses in the 5G-EVP egress tunnel. The application internally creates multiple tuples using these IPs and port combinations. Each flow always maps to a single tuple, which preserves per-flow ordering.

You must configure the additional IP addresses in NetPlan on the 5G-EVP VM before specifying them in the application configuration. Refer to Configuring Multiple Egress IP Addresses.

Rules and Notes

- The 5G-EVP application operates on a dedicated V Series Node. Do not deploy it on the same node as the 5G-SBI or 5G-Cloud applications.
- A maximum of 8 source IP addresses can be configured for the 5G-EVP egress tunnel.
- The ordering window (Number of Buckets × Bucket Interval) must not exceed 100 ms.
- The Ingress REP and Egress REP on the 5G-EVP and 5G-Cloud monitoring session canvases are placeholder elements, they do not require configuration.
- Only one PCAPng application instance can be set to Primary mode per monitoring session. Any instance can be designated as Primary, it is not tied to specific key values.
- The 5G-SBI and 5G-Cloud applications can operate without the 5G-EVP application, but the 5G-EVP application cannot operate without a downstream 5G-SBI or 5G-Cloud application.
- Domain Tagging can coexist with packet ordering. There are no limitations on using both features together.

PCAPNG Processing and Flow Tracking

In this Ericsson vTAP design, PCAPng classifies flows by Ericsson domain and inserts the corresponding VLAN tag on the egress traffic.

The PCAPNG application on the V Series serves as the configuration point for Ericsson domain classification and VLAN mapping (app mode, domainClassification, flowTimeout, and domainTable). The 5G SBI application performs the actual parsing of Ericsson vTAP metadata, flow and domain determination, and SBI failure detection. The dataplane and egress tunnel configuration apply the VLAN tags and forward the traffic.

> **NOTE:** In the PCAPng configuration, the Domain VLAN Mapping and Flow Timeout fields are displayed only when App Mode is set to Primary and the Domain Classification check box is enabled.

Egress Tunnels and VLAN Tag Insertion

When the V Series classifies Ericsson vTAP traffic by domain in the PCAPNG application, it forwards packets to the probes through egress tunnels (VXLAN or L2GRE).

When creating an Egress tunnel in the Monitoring Session Canvas, enable the **Domain Tagging** option to tag packets on the egress tunnel with the Ericsson domain-specific VLAN IDs derived from the PCAPng Domain VLAN Mapping.

> **NOTE:** This setting is available only when **Domain Classification** is enabled in the associated PCAPng application. Refer to 5G-Service Based Interface Application for details.

When **Domain Classification** is enabled in the PCAPNG app and **Domain Tagging** is enabled on the egress tunnel, the V Series does the following:

- Looks up the Ericsson domain for each packet.
- Maps it to a VLAN ID from the Domain Table (including any failed SBI or special case VLANs)
- Inserts that VLAN tag on the egress packet

This ensures each probe interface only receives the Ericsson domains it is configured to monitor.

> **Notes:**
>
> - After you deploy a Monitoring Session, you cannot edit its associated 5G SBI or PCAPNG application configurations. To modify them, undeploy the Monitoring Session, update the configuration, and redeploy.
> - If a PCAPNG application instance is set as the Primary for a Monitoring Session, you cannot delete it while it holds that role. To delete or replace it, undeploy the Monitoring Session or reassign the Primary.

End-to End Solution Overview and Configuration:

For end-to-end configuration workflows, refer to the scenario-specific pages based on your deployment:

- Configure PCC vTAP Traffic Flow - SBI (L7-JSON) and non-SBI traffic from Ericsson PCC vTAPs. This is the most common deployment.
- Configure CCx vTAP Traffic Flow - SCP (L7-Native) traffic from Ericsson CCx vTAPs. Can be deployed independently or alongside SBI traffic.

For a full deployment that includes all traffic types, combine the SBI and SCP workflows into a single monitoring session on the 5G-SBI Node.

Upgrade and Migration

If you have an existing 5G-SBI and 5G-Cloud deployment, you can add the 5G-EVP application by inserting a new EVP Node between the Ericsson vTAPs and the 5G-SBI Node.

When you add the EVP Node, re-assign the ingress IP address from the 5G-SBI Node to the EVP Node. This ensures that the vTAPs send traffic to the EVP Node without requiring reconfiguration on the Ericsson side. Provision new IP addresses to connect the EVP egress to the 5G-SBI ingress. Both interfaces must be in the same subnet.

The 5G-SBI to 5G-Cloud pipeline remains unchanged during migration.

For step-by-step instructions, refer to the Upgrade and Migration Steps in Configure PCC vTAP Traffic Flow or Configure CCx vTAP Traffic Flow.

NOTE: If the existing deployment has a source IP configured in the tunnel endpoint (TEP) configuration, remove it before deploying the EVP application.

Configure PCC Traffic Flow

This deployment scenario applies when your environment includes traffic from Ericsson PCC (Packet Core Controller) vTAPs. The Ericsson PCC sends SBI traffic (L7-JSON) and non-SBI traffic simultaneously through its vTAPs, these two traffic types always coexist on the same monitoring session. This is the most common deployment.

The 5G-EVP application receives all PCC traffic, reorders packets based on the embedded observed timestamp, and forwards sorted traffic to the 5G-SBI Node over UDP-GRE. On the 5G-SBI Node, the monitoring session processes traffic through three parallel paths:

Metadata (keys 0, 2) -The PCAPng application processes metadata for domain classification. This traffic stays internal and does not reach monitoring probes.

SBI traffic (key 1) - The PCAPng application processes the traffic and passes it to the 5G-SBI application. The processed traffic is forwarded to the egress tunnel.

Non-SBI traffic (keys 5, 7, 9, 11, 15, 21, 23) The PCAPng application processes the traffic and forwards it directly to the egress tunnel.

Monitoring Sessions Required

Monitoring Session	V Series Node	Traffic Handled
5G-EVP Monitoring Session	5G-EVP	All PCC traffic, packet ordering and optional DTLS decryption
5G-SBI Monitoring Session	5G-SBI	SBI (key 1) + Non-SBI (keys 5,7,9,11,15,21,23) + Metadata (keys 0,2)
5G-Cloud Monitoring Session	5G-Cloud	SCP synthesis (if applicable)

Prerequisites

Before you begin, complete the following:

- Deploy three V Series Nodes: 5G-EVP , 5G-SBI, and 5G-Cloud.
- If you plan to use multiple egress IPs on the 5G-EVP Node, configure the additional IP addresses in NetPlan on the VM. Refer to Configuring Multiple Egress IP Addresses.
- Create a separate monitoring domain for each V Series Node in GigaVUE-FM.

- Verify that the V Series 5GC add-on license is active.
- Configure NTP servers for the 5G-EVP Node (required when packet ordering is enabled).
- If you plan to use DTLS, upload the private key and certificate in GigaVUE-FM via Inventory > Security > SSL Key.

End to End Configuration

Monitoring Session 1: 5G-EVP

Create this monitoring session on the 5G-EVP V Series Node.

1. In GigaVUE-FM, go to the monitoring domain for the 5G-EVP V Series Node.
2. **Create** a new monitoring session.
3. Configure Multiple Egress IP Addresses.
   (missing or bad snippet)
4. On the canvas, drag and drop an **Ingress REP**.
5. Drag and drop the **5G-EVP** application. Select **Details** and configure the Ingress UDP Tunnel, Egress UDP Tunnel, Packet Ordering Configuration, Diagnostics, and Logging sections.

Configure 5G-EVP Application

To configure the 5G-EVP application, drag and drop it onto the monitoring session canvas and select Details. The 5G-EVP application quick view appears. Enter or select the required information as described in the following tables.

Field	Description
Application	The name **5g-EVP** appears by default.
Alias	Enter the alias name as **5gEVP**.
Packet Ordering Config	
Packet Ordering	Toggle to enable or disable timestamp-based packet reordering. When enabled, the 5G-EVP application sorts packets based on the observed timestamp in the PCAPng Enhanced Packet Block before forwarding them to the downstream application.
Primary NTP Server*	Enter the IP address (IPv4 or IPv6) of the primary NTP server. The 5G-EVP application requires NTP time synchronization through chronyd for millisecond-accurate clock alignment. GigaVUE-FM sends this configuration to the V Series Node during each deployment. This field is mandatory when packet ordering is enabled. **NOTE:** Visible when Packet Ordering is enabled.
Secondary NTP Server	Enter the IP address (IPv4 or IPv6) of the secondary NTP server. This field is optional. The secondary NTP server provides failover if the primary NTP server becomes unreachable. **NOTE:** Visible when Packet Ordering is enabled.
Number of Buckets	Enter the number of time buckets for sorting. Default: 100. The value of Number of Buckets × Bucket Interval must not exceed 100. For example, if Bucket Interval is 5 ms, the maximum Number of Buckets is 20. If the configured values exceed this limit, the application returns an error.
Packets per Bucket	Enter the maximum number of packets each bucket can hold. Default: 20,000.
Bucket Interval (ms)	Enter the time span (in milliseconds) for each bucket. Minimum: 1 ms. Maximum: 5 ms. Default: 1 ms. The sorting window equals Number of Buckets × Bucket Interval. For example, 100 buckets × 1 ms = a 100 ms sorting window.
Packet Outside Bucket Interval	Select the action for packets that arrive with timestamps outside the current sorting window. Forward (recommended), sends the packet to the downstream application without reordering. Discard, drops the packet at the 5G-EVP application.
Monitored Flows	
Number of Egress Flows	Enter the maximum number of egress flows. Default: 512. Maximum: 16,384.
Egress Flow Timeout Value	Enter the timeout (in seconds) for egress flow entries. Default: 660.
Ingress UDPGRE Tunnel	

Listening IP	Specify the tunnel's local listen IP address to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. Enter the UDP port number. Default: 4754 (non-secure) or 4755 (DTLS). You can enter any valid port number.
RX Thread	Specify the RX value to receive the packet. Enter a value between 1 and 16. The default value is 8.
DTLS	Select Enable to receive DTLS-encrypted traffic, or Disable for plaintext UDP-GRE traffic.
DTLS Key	Select the key alias associated with the SSL certificate used for DTLS decryption.
Egress UDP Tunnel	
Tool IP	Egress Tunnels Endpoint, Destination IP address.
Tx Destination Port	Enter the destination port. This value must match the listen port on the downstream application. Default: 4754.
Tx Source IP List	Enter up to 8 source IP addresses for the egress tunnel. Configure these IPs in the VM's NetPlan before adding them here.
Tx Source Port	Enter the source port for the egress UDP-GRE tunnel.
Tx Thread	Enter the number of transmit threads. Range: 1–8.
Diagnostics	
Packet Disable	Enter a protocol key value to disable processing for a specific Ericsson protocol. Default: 0.
Tx Disable	Select true to disable packet transmission (for debugging), or false for normal transmission. Default: false.
Logging	
Log Level	Select the minimum severity of system events to log. Higher levels generate more detailed logs.
Packet Capture Log Level	Select the log level: all, receive, transmit, or none.
Message Log Level	Select the log level: error, warn, info, or debug. Default: error.
Log Directory	Thebpath where log files should be stored.

Configure 5G-EVP Application

To configure the 5G-EVP application, drag and drop it onto the monitoring session canvas and select Details. The 5G-EVP application quick view appears. Enter or select the required information as described in the following tables.

Field	Description
Application	The name **5g-EVP** appears by default.
Alias	Enter the alias name as **5gEVP**.
Packet Ordering Config	
Packet Ordering	Toggle to enable or disable timestamp-based packet reordering. When enabled, the 5G-EVP application sorts packets based on the observed timestamp in the PCAPng Enhanced Packet Block before forwarding them to the downstream application.
Primary NTP Server*	Enter the IP address (IPv4 or IPv6) of the primary NTP server. The 5G-EVP application requires NTP time synchronization through chronyd for millisecond-accurate clock alignment. GigaVUE-FM sends this configuration to the V Series Node during each deployment. This field is mandatory when packet ordering is enabled. **NOTE:** Visible when Packet Ordering is enabled.
Secondary NTP Server	Enter the IP address (IPv4 or IPv6) of the secondary NTP server. This field is optional. The secondary NTP server provides failover if the primary NTP server becomes unreachable. **NOTE:** Visible when Packet Ordering is enabled.
Number of Buckets	Enter the number of time buckets for sorting. Default: 100. The value of Number of Buckets × Bucket Interval must not exceed 100. For example, if Bucket Interval is 5 ms, the maximum Number of Buckets is 20. If the configured values exceed this limit, the application returns an error.
Packets per Bucket	Enter the maximum number of packets each bucket can hold. Default: 20,000.
Bucket Interval (ms)	Enter the time span (in milliseconds) for each bucket. Minimum: 1 ms. Maximum: 5 ms. Default: 1 ms. The sorting window equals Number of Buckets × Bucket Interval. For example, 100 buckets × 1 ms = a 100 ms sorting window.
Packet Outside Bucket Interval	Select the action for packets that arrive with timestamps outside the current sorting window. Forward (recommended), sends the packet to the downstream application without reordering. Discard, drops the packet at the 5G-EVP application.
Monitored Flows	
Number of Egress Flows	Enter the maximum number of egress flows. Default: 512. Maximum: 16,384.
Egress Flow Timeout Value	Enter the timeout (in seconds) for egress flow entries. Default: 660.
Ingress UDPGRE Tunnel	

Listening IP	Specify the tunnel's local listen IP address to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. Enter the UDP port number. Default: 4754 (non-secure) or 4755 (DTLS). You can enter any valid port number.
RX Thread	Specify the RX value to receive the packet. Enter a value between 1 and 16. The default value is 8.
DTLS	Select Enable to receive DTLS-encrypted traffic, or Disable for plaintext UDP-GRE traffic.
DTLS Key	Select the key alias associated with the SSL certificate used for DTLS decryption.
Egress UDP Tunnel	
Tool IP	Egress Tunnels Endpoint, Destination IP address.
Tx Destination Port	Enter the destination port. This value must match the listen port on the downstream application. Default: 4754.
Tx Source IP List	Enter up to 8 source IP addresses for the egress tunnel. Configure these IPs in the VM's NetPlan before adding them here.
Tx Source Port	Enter the source port for the egress UDP-GRE tunnel.
Tx Thread	Enter the number of transmit threads. Range: 1–8.
Diagnostics	
Packet Disable	Enter a protocol key value to disable processing for a specific Ericsson protocol. Default: 0.
Tx Disable	Select true to disable packet transmission (for debugging), or false for normal transmission. Default: false.
Logging	
Log Level	Select the minimum severity of system events to log. Higher levels generate more detailed logs.
Packet Capture Log Level	Select the log level: all, receive, transmit, or none.
Message Log Level	Select the log level: error, warn, info, or debug. Default: error.
Log Directory	Thebpath where log files should be stored.

6. Drag and drop an **Egress REP**. This is a placeholder, no configuration is required.
7. Create links: **Ingress REP -> 5G-EVP -> Egress REP**.
8. From the **Actions** menu, select **Deploy**.

NOTE: You cannot disable both packet ordering and DTLS simultaneously. At least one must be enabled for the 5G-EVP application to function.

Monitoring Session 2: 5G-SBI

Create this monitoring session on the **5G-SBI** V Series Node (DPDK-based).

1. In GigaVUE-FM, go to the monitoring domain for the **5G-SBI** V Series Node.
2. Create a new monitoring session.
3. Create one UDP-GRE ingress tunnel for each Ericsson vTAP key value that you want to process. The following table lists the key values and their traffic types:

Key Value	Traffic Value
0,2	Metadata (domain classification)
1	SBI (L7-JSON) - AMF, SMF
5, 7, 9, 11, 15, 21, 23	Non-SBI

Path 1: Metadata (Keys 0 and 2)

4. Drag and drop the first PCAPng application for metadata processing. Select Details and configure the following:

Field	Description
Application	The name **Pcapng** appears by default.
Alias	Enter an Alias as required.
App Mode	Set App Mode to **Primary** to view and configure domain related options. **NOTE:** Primary mode is required to configure domain classification. When you set App Mode to Secondary, Domain Classification, Domain VLAN Mapping, and Flow Timeout fields are not visible.
Domain Classification	Enable the Domain Classification option to classify Ericsson vTAP traffic by domain.
Domain VLAN Mapping	• From the **Domain VLAN Mapping** list, select the VLAN mapping profile that maps Ericsson domains to VLAN IDs for per-domain tagging on the egress tunnel. • Click **Add New** to create a new 5G-Apps configuration. Refer to Adding CSV file for IP Mapping section in 5G-Service Based Interface Application for details. When **Domain Classification** is enabled, the selected Domain VLAN Mapping is used to tag packets per Ericsson domain with the configured VLAN ID (and, if applicable, a separate VLAN ID for failed SBI transactions) before sending traffic to the probes.
Flow Timeout	Set the **Flow Timeout** value in seconds to define how long an inactive Ericsson vTAP flow is kept before it is closed and its resources are released. Valid range: 360-1860 Default value: 660

5. Link the ingress tunnels for key values 0 and 2 to this PCAPng instance.

   NOTE: This PCAPng instance does not require an egress tunnel. It processes metadata only and does not forward traffic to probes. Only one PCAPng instance can be set to Primary mode per monitoring session.

   Path 2: SBI Traffic (Key 1)

6. Drag and drop the second PCAPng application for SBI traffic. Select Details and set App Mode to Secondary.
7. Link the ingress tunnel for key value 1 to this PCAPng instance.
8. Drag and drop the 5G-SBI application. Select Details and configure the following:

Field	Description
Application	The name **5g-sbi** appears by default.
Alias	The name **sbi5gAppTemplate** appears by default.
Vendor Integrations	Select the option **ericssonVTap** from the drop-down list.
SBI Operational Mode	Select the SBI Monitoring Mode for 5G-SBI Network traffic. L7json is selected by default.
Destination IP	The Destination IP toggle determines how synthesized SCP mediated SBI transactions are tagged. You can choose from the following options: • **SCP** - Sets the synthesized transaction's destination IP to the SCP IP address. This option is selected by default. • **Destination Network Function** - Sets the synthesized transaction's destination IP to the IP address of the destination network function.
Ericsson vTAP Version	Select 1 or 2 from the drop-down list box.
MTU Packet Size	Select how packets should be sent to the tool: • **Full Packet** - Sends the entire synthesized packet as a single packet without segmentation. The packet is fragmented at the egress TEP. • **MTU Size** - Segments packets based on the configured tool port MTU value before sending them to the tool. You can specify a value between 1200 and 9200 bytes.
Incomplete SBI Transaction Domain Tag	Enable this option to tag unclassified SBI transactions with the dedicated 'Failed SBI Transaction' VLAN ID defined in the Domain table. The probe uses this tag to identify and troubleshoot these transactions more effectively. **NOTE:** This option appears only when at least one Primary PCAPng file is configured on a V Series node.
Indexed Headers	Enable the check box to index the HTTP2 headers in the 5G-SBI application.
Compressed Headers	Enable the check box to compress the HTTP2 headers in the 5G-SBI application.
TCP Flows	Specify the maximum number of TCP flows based on the VM form factor: • **Small** - Supports 128 to 2048 TCP flows

9.

	• **Medium** - Supports 128 to 8192 TCP flows • **Large** - Supports 128 to 16384 TCP flows
TCP Flow Timeout	Specify the flow range value (in seconds) that determines how long a TCP flow remains valid in the application. • Minimum value: 0 • Maximum Value: 7200 • Default value: 1800
Number of Transaction Flows	Specify the maximum number of streams or flows based on the VM form factor: • **Small** - Supports 1 to 20480 Transaction flows • **Medium** - Supports 1 to 51200 Transaction flows • **Large** - Supports 1 to 122880 Transaction flows
Request Timeout	Specify the time (in seconds) for which the request packet waits for the response packet in a stream. • Minimum value: 1 • Maximum Value: 300 • Default value: 10
Response Timeout	Specify the time (in seconds) for which the response packet waits for the request packet in a stream. • Minimum value: 1 • Maximum Value: 300 • Default value: 2
IP Mapping	
Network Function Instance Mapping	Select the required CSV file from the drop-down list with required Network Function Instance ID (NFID) instance mapping. • **Select an alias**: Choose the alias name associated with the uploaded NFID table (CSV file). • **Add a new alias**: If no alias exists, click **Add New** to create one. Refer to Adding CSV file for IP Mapping for details. **NOTE:** If duplicate entries are detected during mapping table upload, a validation error is displayed. After all other validation errors are resolved, select **Save Anyway** to save the file.
FQDN Mapping	Select the required CSV file from the drop-down list with required FQDN mapping. • **Select an alias**: Choose the alias name

	associated with the uploaded FQDN table (CSV file). • **Add a new alias**: If no alias exists, click **Add** to create one. Refer to Adding CSV file for IP Mapping section for details. **NOTE:** If duplicate entries are detected during mapping table upload, a validation error is displayed. After all other validation errors are resolved, select **Save Anyway** to save the file.
Logging	
Transaction Log	Enable or disable 5G-SBI transaction logging. When enabled, transaction logs are written and rotated based on the configured Transaction Log File Interval.
Transaction Log Interval(Minutes)	Set the transaction log file interval in minutes. This defines how often a new Transaction Log CSV file is created. • Valid options: 5 or 60 minutes • Default: 60 This setting is used only when Transaction Log is enabled.
Stats Log	Enable or disable statistics logging for the 5G-SBI application. When enabled, statistics are written to the configured log folder and managed based on the Log Folder Size setting.
Log Folder Size(MB)	Specify the memory size (in MB) allocated to the log folder. This sets the maximum allowable size for the 5G-SBI log folder and applies to both transaction and statistics logs stored in the configured Log Folder Location. Maximum Value: 40960 MB.
Log Folder Location	Enter the file system path where 5G-SBI log files should be stored. This is the location for both transaction and statistics logs. **Example**: /home/user/log

10. Link the SBI PCAPng instance (from Step 6) to the 5G-SBI application.

    Path 3: Non-SBI Traffic (Keys 5, 7, 9, 11, 15, 21, 23)

11. Drag and drop the third **PCAPng** application for non-SBI traffic. Select **Details** and set **App Mode** to **Secondary**.

12. Link the ingress tunnels for non-SBI key values (5, 7, 9, 11, 15, 21, 23) to this PCAPng instance.

    NOTE: Non-SBI traffic does not require a 5G-Cloud application. The PCAPng application reads the data from the Enhanced Packet Block and forwards it directly to the egress tunnel.

    Create Egress Tunnels and Deploy

13. Create one or more egress tunnels (VXLAN or L2GRE). If Domain Classification is enabled, enable Domain Tagging on the egress tunnel.
14. Link the **5G-SBI** application output to the egress tunnel (for SBI traffic).
15. Link the non-SBI **PCAPng** application output to the egress tunnel (for non-SBI traffic).
16. Verify all links on the canvas:
    - Ingress tunnels (key 0, 2) -> PCAPng (Primary), no egress
    - Ingress tunnel (key 1) -> PCAPng (Secondary) -> 5G-Cloud -> Egress tunnel
    - Ingress tunnels (non-SBI keys) -> PCAPng (Secondary) -> Egress tunnel
17. From the **Actions** menu, select **Deploy**.

Monitoring Session 3: 5G-Cloud

Create this monitoring session on the 5G-Cloud V Series Node (Linux-based).

1. In GigaVUE-FM, navigate to the monitoring domain for the 5G-Cloud V Series Node.
2. Create a new monitoring session.
3. Drag and drop an Ingress REP (placeholder, no configuration required).
4. Drag and drop the 5G-Cloud application. Select Details and configure the following: (missing or bad snippet)
5. Drag and drop an Egress REP (placeholder, no configuration required).
6. Create links: Ingress REP -> 5G-Cloud -> Egress REP.
7. From the Actions menu, select Deploy.

Notes:

- After you deploy a Monitoring Session, you cannot edit its associated 5G SBI or PCAPng application configurations. To modify them, undeploy the Monitoring Session, update the configuration, and redeploy.

- If a PCAPng application instance is set as the Primary for a Monitoring Session, you cannot delete it while it holds that role. To delete or replace it, undeploy the Monitoring Session or reassign the Primary.

Adding CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field.
3. From the **Type** drop-down list, select one of the following. You can also choose to download the template using the **Download Template** option and add the required entries.
   - **5G-SBI Network Function Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping in 5G-Ericsson.

     Header details with an example:
     - **instanceID** - e4d909c2-5d2f-4f7f-9bcb-12a3b6c4d789
     - **ipAddress** - 1.1.1.1
   - **5G-SBI FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

     Header details with an example:
     - **fqdnID** - amf.local
     - **ipAddress** - 1.1.1.1
     - **nfType** - SCP
   - **Domain VLAN Mapping** - Add the CSV file containing a valid domain and a valid VLAN ID for IP mapping.

     Refer the table below for example header details of the CSV file.

domainName	vlanID	failedSBITransactionVlanID
pc-mm	100	200
pc-sm	300	400
ccsm-5gc	500	
ccsm-cnhss	600	
ccrc-nrf	700	
nrfagent-nrf	800	

domainName	vlanID	failedSBITransactionVlanID
ccrc-nssf	900	
ccpc-pcf	1000	
ccpc-provisioning	1100	
ccdm-udr-5g	1200	
ccdm-provisioning	1300	
ccdm-udr-legacy	1400	
cces-nef	1500	
cces-ees	1600	
cces-apigm	1700	
sc-bsf	1800	
sc-scp	1900	
sc-sepp	2000	
others	2100	
unclassified	2200	

4. Click **Choose File** in **File Name** field to upload the CSV file into GigaVUE-FM.
5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

Notes:

- Failed SBI Transactions are applicable only for **pc-mm** and **pc-sm** domains.
- Each domain in the Domain table must use a unique VLAN ID and Failed SBI Transaction ID to allow the probes to distinguish traffic for each Ericsson domain.

Recommended VM Specifications

The configuration of the V Series application should be performed on large VMs.

V Series App	vCPUs	Memory	Disk Space
5G-EVP	20	32 GB	80 GB
5G-SBI	8	16 GB	80 GB
5G-Cloud	24	32 GB	80 GB

- Supported Ports: 1 Management interface, 1 ingress, and 1 egress interfaces.

- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to Configure GigaVUE Fabric Components using VMware ESXi topic for more details.

FHA Dashboards for 5G Applications

After configuring the 5G-SBI application, you can monitor the statistics for Ericsson vTAP by the reports displayed in the Dashboard. To access the details, refer to FHA Dashboards for 5G Applications.

Upgrade and Migration Steps

Follow these steps to add the 5G-EVP application to an existing deployment. If you are deploying as part of a new setup, skip this section.

1. Upgrade GigaVUE-FM to the current release.
2. Upgrade the existing 5G-SBI and 5G-Cloud V Series to the current release.
3. Deploy a new 5G-EVP V Series Node running the current release.
4. Reassign the ingress IP address from the 5G-SBI Node to the EVP Node ingress interface.
5. Provision a new IP address for the EVP Node egress interface.
6. Provision a new IP address for the 5G-SBI Node ingress interface. This must be in the same subnet as the EVP egress IP.
7. If the existing deployment has a source IP in the TEP configuration, remove it.
8. Configure the 5G-EVP monitoring session. Refer to 5G-EVP for field descriptions.
9. Update the 5G-SBI monitoring session with the new ingress IP address.
10. Deploy both monitoring sessions from the Actions menu in GigaVUE-FM.

NOTE: The 5G-SBI to 5G-Cloud pipeline remains unchanged. No modifications are required on the 5G-Cloud Node.

Configure SCP Traffic Flow

Use this workflow when your deployment includes SCP (L7-Native) traffic from Ericsson CCx vTAPs. You can deploy this workflow independently or alongside the SBI traffic flow. The SCP traffic passes through a Load Balancing (LB) application for distribution before reaching the 5G-Cloud Node for L7-Native synthesis.

Monitoring Sessions Required

Monitoring Session	V Series Node	Traffic Handled
5G-EVP	5G-EVP	All traffic, packet ordering and optional DTLS decryption
SCP	5G-SBI	SCP ingress (keys 3,17,19) + Metadata (keys 0,2) + LB + multi-tap egress
5G-Cloud	5G-Cloud	SCP synthesis (L7-Native)

Prerequisites

Before you begin, complete the following:

- Deploy three V Series Nodes: 5G-EVP, 5G-SBI, and 5G-Cloud.
- Create a separate monitoring domain for each V Series Node in GigaVUE-FM.
- Verify that the V Series 5GC add-on license is active.
- Configure NTP servers for the 5G-EVP Node (required when packet ordering is enabled).
- If you plan to use DTLS, upload the private key and certificate in GigaVUE-FM via Inventory > Security > SSL Key.
- If you plan to use multiple egress IPs on the 5G-EVP Node, configure the additional IP addresses in NetPlan on the VM. Refer to Configuring Multiple Egress IP Addresses.
- Upload CSV Files for the 5G-Cloud Application
    - In GigaVUE-FM, go to Inventory > VIRTUAL.
    - Select your cloud platform.
    - Click Settings, and then click 5G-Apps.
    - Upload the FQDN, NF Instance, and User Agent CSV files.

End to End Configuration

Configure 5G-EVP Application

To configure the 5G-EVP application, drag and drop it onto the monitoring session canvas and select Details. The 5G-EVP application quick view appears. Enter or select the required information as described in the following tables.

Field	Description
Application	The name **5g-EVP** appears by default.
Alias	Enter the alias name as **5gEVP**.
Packet Ordering Config	
Packet Ordering	Toggle to enable or disable timestamp-based packet reordering. When enabled, the 5G-EVP application sorts packets based on the observed timestamp in the PCAPng Enhanced Packet Block before forwarding them to the downstream application.
Primary NTP Server*	Enter the IP address (IPv4 or IPv6) of the primary NTP server. The 5G-EVP application requires NTP time synchronization through chronyd for millisecond-accurate clock alignment. GigaVUE-FM sends this configuration to the V Series Node during each deployment. This field is mandatory when packet ordering is enabled. **NOTE:** Visible when Packet Ordering is enabled.
Secondary NTP Server	Enter the IP address (IPv4 or IPv6) of the secondary NTP server. This field is optional. The secondary NTP server provides failover if the primary NTP server becomes unreachable. **NOTE:** Visible when Packet Ordering is enabled.
Number of Buckets	Enter the number of time buckets for sorting. Default: 100. The value of Number of Buckets × Bucket Interval must not exceed 100. For example, if Bucket Interval is 5 ms, the maximum Number of Buckets is 20. If the configured values exceed this limit, the application returns an error.
Packets per Bucket	Enter the maximum number of packets each bucket can hold. Default: 20,000.
Bucket Interval (ms)	Enter the time span (in milliseconds) for each bucket. Minimum: 1 ms. Maximum: 5 ms. Default: 1 ms. The sorting window equals Number of Buckets × Bucket Interval. For example, 100 buckets × 1 ms = a 100 ms sorting window.
Packet Outside Bucket Interval	Select the action for packets that arrive with timestamps outside the current sorting window. Forward (recommended), sends the packet to the downstream application without reordering. Discard, drops the packet at the 5G-EVP application.
Monitored Flows	
Number of Egress Flows	Enter the maximum number of egress flows. Default: 512. Maximum: 16,384.
Egress Flow Timeout Value	Enter the timeout (in seconds) for egress flow entries. Default: 660.
Ingress UDPGRE Tunnel	

Listening IP	Specify the tunnel's local listen IP address to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. Enter the UDP port number. Default: 4754 (non-secure) or 4755 (DTLS). You can enter any valid port number.
RX Thread	Specify the RX value to receive the packet. Enter a value between 1 and 16. The default value is 8.
DTLS	Select Enable to receive DTLS-encrypted traffic, or Disable for plaintext UDP-GRE traffic.
DTLS Key	Select the key alias associated with the SSL certificate used for DTLS decryption.
Egress UDP Tunnel	
Tool IP	Egress Tunnels Endpoint, Destination IP address.
Tx Destination Port	Enter the destination port. This value must match the listen port on the downstream application. Default: 4754.
Tx Source IP List	Enter up to 8 source IP addresses for the egress tunnel. Configure these IPs in the VM's NetPlan before adding them here.
Tx Source Port	Enter the source port for the egress UDP-GRE tunnel.
Tx Thread	Enter the number of transmit threads. Range: 1–8.
Diagnostics	
Packet Disable	Enter a protocol key value to disable processing for a specific Ericsson protocol. Default: 0.
Tx Disable	Select true to disable packet transmission (for debugging), or false for normal transmission. Default: false.
Logging	
Log Level	Select the minimum severity of system events to log. Higher levels generate more detailed logs.
Packet Capture Log Level	Select the log level: all, receive, transmit, or none.
Message Log Level	Select the log level: error, warn, info, or debug. Default: error.
Log Directory	Thebpath where log files should be stored.

Monitoring Session 2: SCP Traffic (5G-SBI Node)

Create this monitoring session on the 5G-SBI V Series Node. This session handles the ingress, PCAPng processing, and load balancing for SCP traffic before forwarding it to the 5G-Cloud Node.

1. In GigaVUE-FM, navigate to the monitoring domain for the 5G-SBI V Series Node.
2. Create a new monitoring session.
3. Create one UDP-GRE ingress tunnel for each SCP key value:

Key Value	Traffic Value
0,2	Metadata (domain classification)
3, 17, 19	SCP (L7-Native)

4. Drag and drop the first PCAPng application for metadata processing. Select Details and configure the following:
(missing or bad snippet)
5. Link the ingress tunnels for key values 0 and 2 to this PCAPng instance.

   NOTE: This PCAPng instance does not require an egress tunnel. Only one PCAPng instance can be set to Primary mode per monitoring session.

6. Drag and drop the second PCAPng application for SCP traffic. Select Details and set App Mode to Secondary.
7. Link the ingress tunnels for key values 3, 17, 19 to this PCAPng instance.
8. Drag and drop the LB (Load Balancing) application. Select Details and configure the hash field to greFlowid.
9. Link the SCP PCAPng instance (from Step 6) to the LB application.
10. Create a VXLAN egress tunnel with the Multi-Tunnel flag enabled. If Domain Classification is enabled, enable Domain Tagging on the egress tunnel.

    NOTE: You create one egress tunnel. The system internally creates multiple taps for load distribution (minimum 1, maximum 16).

11. Link the LB application output to the egress tunnel.
12. Verify all links on the canvas:
    - Ingress tunnels (key 0, 2) -> PCAPng (Primary), no egress
    - Ingress tunnels (key 3, 17, 19) -> PCAPng (Secondary) -> LB -> Egress tunnel (multi-tap)
13. From the Actions menu, select Deploy.

Monitoring Session 3: 5G-Cloud

Create this monitoring session on the 5G-Cloud V Series Node (Linux-based).

1. In GigaVUE-FM, navigate to the monitoring domain for the 5G-Cloud V Series Node.
2. Create a new monitoring session.
3. Drag and drop an Ingress REP (placeholder, no configuration required).
4. Drag and drop the 5G-Cloud application. Select Details and configure the following: (missing or bad snippet)
5. Drag and drop an Egress REP (placeholder, no configuration required).
6. Create links: Ingress REP -> 5G-Cloud -> Egress REP.
7. From the Actions menu, select Deploy.

> **Notes:**
>
> - After you deploy a Monitoring Session, you cannot edit its associated 5G SBI or PCAPNG application configurations. To modify them, undeploy the Monitoring Session, update the configuration, and redeploy.
> - If a PCAPNG application instance is set as the Primary for a Monitoring Session, you cannot delete it while it holds that role. To delete or replace it, undeploy the Monitoring Session or reassign the Primary.

Adding CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field.
3. From the **Type** drop-down list, select one of the following. You can also choose to download the template using the **Download Template** option and add the required entries.
   - **5G-SBI Network Function Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping in 5G-Ericsson.

     Header details with an example:
     - **instanceID** - e4d909c2-5d2f-4f7f-9bcb-12a3b6c4d789
     - **ipAddress** - 1.1.1.1
   - **5G-SBI FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

     Header details with an example:
     - **fqdnID** - amf.local

- **ipAddress** - 1.1.1.1
- **nfType** - SCP

- **Domain VLAN Mapping** - Add the CSV file containing a valid domain and a valid VLAN ID for IP mapping.

  Refer the table below for example header details of the CSV file.

domainName	vlanID	failedSBITransactionVlanID
pc-mm	100	200
pc-sm	300	400
ccsm-5gc	500	
ccsm-cnhss	600	
ccrc-nrf	700	
nrfagent-nrf	800	
ccrc-nssf	900	
ccpc-pcf	1000	
ccpc-provisioning	1100	
ccdm-udr-5g	1200	
ccdm-provisioning	1300	
ccdm-udr-legacy	1400	
cces-nef	1500	
cces-ees	1600	
cces-apigm	1700	
sc-bsf	1800	
sc-scp	1900	
sc-sepp	2000	
others	2100	
unclassified	2200	

4. Click **Choose File** in **File Name** field to upload the CSV file into GigaVUE-FM.
5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

Notes:

- Failed SBI Transactions are applicable only for **pc-mm** and **pc-sm** domains.
- Each domain in the Domain table must use a unique VLAN ID and Failed SBI Transaction ID to allow the probes to distinguish traffic for each Ericsson domain.

Recommended VM Specifications

The configuration of the V Series application should be performed on large VMs.

V Series App	vCPUs	Memory	Disk Space
5G-EVP	20	32 GB	80 GB
5G-SBI	8	16 GB	80 GB
5G-Cloud	24	32 GB	80 GB

- Supported Ports: 1 Management interface, 1 ingress, and 1 egress interfaces.
- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to Configure GigaVUE Fabric Components using VMware ESXi topic for more details.

FHA Dashboards for 5G Applications

After configuring the 5G-SBI application, you can monitor the statistics for Ericsson vTAP by the reports displayed in the Dashboard. To access the details, refer to FHA Dashboards for 5G Applications.

Upgrade and Migration Steps

Follow these steps to add the 5G-EVP application to an existing deployment. If you are deploying as part of a new setup, skip this section.

1. Upgrade GigaVUE-FM to the current release.
2. Upgrade the existing 5G-SBI and 5G-Cloud V Series to the current release.
3. Deploy a new 5G-EVP V Series Node running the current release.
4. Reassign the ingress IP address from the 5G-SBI Node to the EVP Node ingress interface.
5. Provision a new IP address for the EVP Node egress interface.
6. Provision a new IP address for the 5G-SBI Node ingress interface. This must be in the same subnet as the EVP egress IP.
7. If the existing deployment has a source IP in the TEP configuration, remove it.
8. Configure the 5G-EVP monitoring session. Refer to Adding CSV file for IP Mapping for field descriptions.
9. Update the 5G-SBI monitoring session with the new ingress IP address.
10. Deploy both monitoring sessions from the Actions menu in GigaVUE-FM.

> **NOTE:** The 5G-SBI to 5G-Cloud pipeline remains unchanged. No modifications are required on the 5G-Cloud Node.

5G-Cloud Application

5G-Core is a service-based architecture in which many control plane network functions are available and communication across these network functions happens through HTTP2 protocol. These HTTP2 transactions are mirrored using some specific network functions, which are in JSON-encoded format.

The 5G Cloud application supports synthesizing of L2, L3, L4, and HTTP2 headers from the JSON-encoded data of the 5G core network. Currently, Ericsson VTAP 5G core input format is supported in DPDK based platforms and in non-DPDK based platforms, Casa vTAP, Oracle SCP, Nokia SCP Inbound, and Nokia SCP Inbound/Outbound input formats are supported.

In non-DPDK mode, V Series Node supports processing of the HTTP2/HTTP2S input formats and extracts the actual packet in the 5G core, and then sends it out on the egress tunnel to probe using either VXLAN or L2GRE or UDP-GRE Tunnel. To achieve this, V Series supports two new applications:

- GVHTTP2
- 5G-Cloud

Supported Platforms

The application is supported on the following platforms:

- VMware ESXi
- OpenStack
- Third Party Orchestration

To know more about the configuration of GVHTTP2 and 5G-Cloud application for Casa vTAP, Oracle SCP, and Nokia SCP, refer to the following sections:

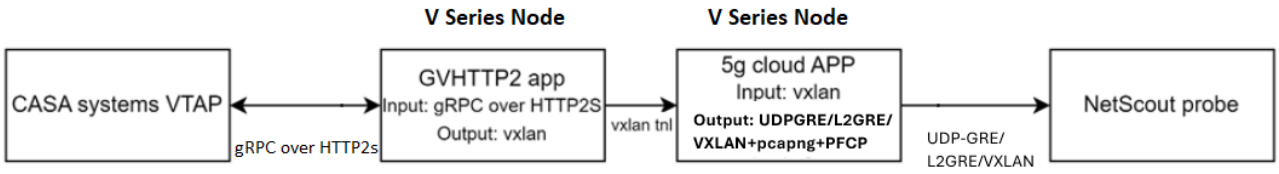
- Configuration of GVHTTP2 Application for Casa vTAP
- 5G-Cloud Oracle and Nokia SCP Support
- Add 5G-Cloud Application in Ericsson SCP
- How Nokia HEP3 IMS Solution works

5G-Cloud Casa vTAP Support

V Series Node supports receiving data from 5G Casa network functions in the Casa vTAP input format. This augments the currently existing functionality of receiving data from Ericsson network functions in Ericsson vTAP format and receiving data from Nokia network

functions in Nokia format. 5G-Cloud Casa vTAP is a non-service-based interface architecture in which input from the Casa vTAP system is processed, decoded, and sent out through an egress tunnel.

Casa systems vTAP network function mirrors the traffic in the 5G core and sends it out in gRPC encoded protobuf format. To convert the encoded protobufs to actual network traffic in 5G core, V Series Node supports two applications, GVHTTP2 and the 5G-Cloud.



GVHTTP2 Application

The first application, GVHTTP2, acts as an HTTP2 server. It receives the gRPC data over HTTP2 and establishes a TCP connection with the Casa vTAP system by acknowledging it. The application will extract the HTTP2 header and HTTP2 body, which are gRPC encoded, then encapsulate the data into a VXLAN header and send it out over the VXLAN tunnel to the next application, 5G-Cloud.

5G-Cloud Application

The 5G-Cloud application receives the VXLAN packet, decapsulates, and decodes the gRPC protobuf content. The application will extract the PFCP packets (N4 traffic) and convert the traffic into PCAPng format. It then sends the extracted packets to the NetScout tool probe over the UDP-GRE/L2GRE/VXLAN tunnels.

NOTE: Since the NetScout probe only understands Ericson vTAP format, which is UDP-GRE encapsulated in PCAPng format, the same is followed in the Casa vTAP system. As part of the 6.7 release, only N4 non-SBI traffic (PFCP) is supported.

Configuration of 5G-Cloud Casa vTAP

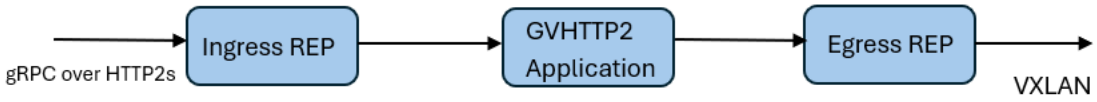
To configure the GVHTTP2 and 5G-Cloud applications in GigaVUE-FM, refer to the following sections:

- Configuration of GVHTTP2 Application for Casa vTAP
- Configuration of 5G-Cloud Application for Casa vTAP

Configuration of GVHTTP2 Application for Casa vTAP

In GigaVUE-FM, you must do the following to add the GVHTTP2 application in the Monitoring Session of a Monitoring Domain:

1. Create an ingress REP (Raw End Point) to receive the gRPC data over TCP on top of HTTP2. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the GVHTTP2 application in the Monitoring Session.
3. Create a link between ingress REP and the GVHTTP2 application.
4. Create egress REP.
5. Create a link between the GVHTTP2 application and egress REP.

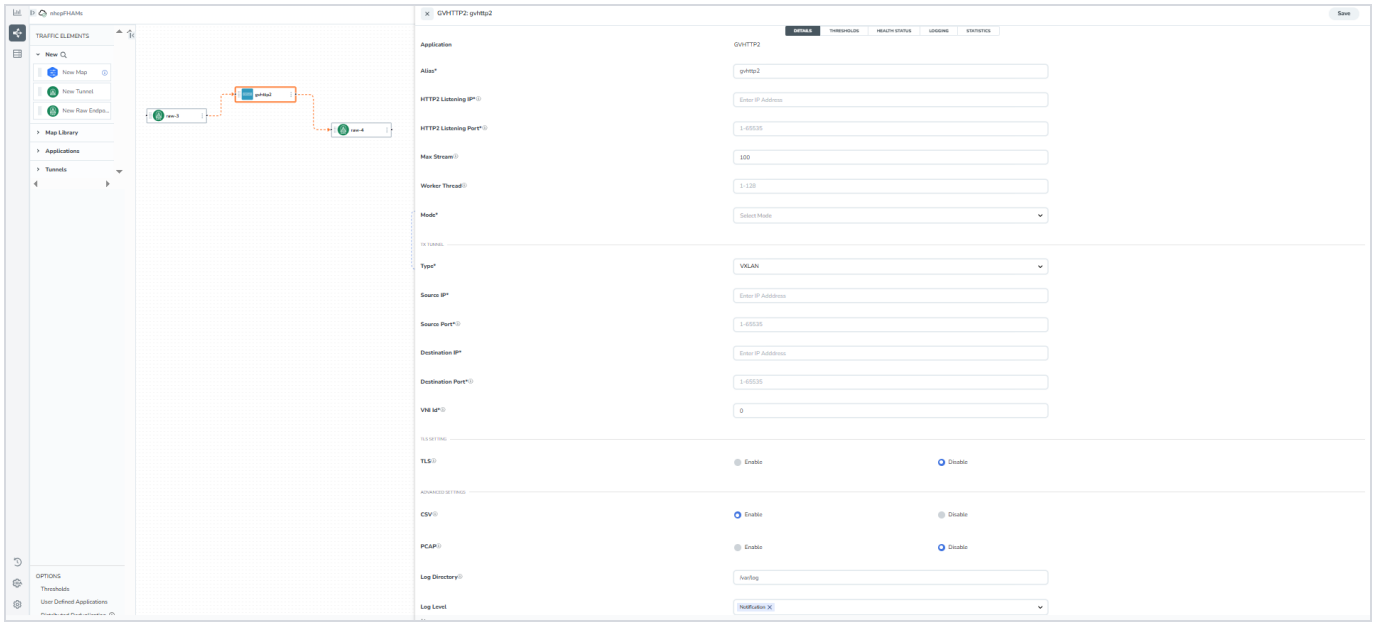


You can add a GVHTTP2 application to:

- **New Monitoring Session** - Add the GVHTTP2 application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
- **Existing session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

To add a GVHTTP2 application:

1. In the canvas, drag and drop the GVHTTP2 application and select **Details**. The GVHTTP2 quick view appears.



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

Field	Description
Application	The name **GVHTTP2** appears by default.
Alias	Enter the alias name as **gvhttp2**.
HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port. Enter a value between 1 and 65535.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specifies the number of worker threads. Enter a value between 1 and 128. The default value is one.
Mode	Select **Casa vTAP** as mode from the drop-down menu.
TX Tunnel (Specifies the tunnel configuration from GVHTTP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the source port to bind for sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the destination IP address to which packets will be sent.
Destination Port	Specify the destination port where packets will be sent.
VNI id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TLS Setting	
TLS	Select a TLS/SSL Monitor Mode from one of the following options:

	▪ Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. ▪ Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on.
Advanced Setting	
Log Directory	Specify the path to store the log files.
Log Level	In the log level drop-down list, select the severity log level of the events from the following options: • PCAP • CSV • Info • Detail • Full-Parse

Downloading Logs for GVHTTP2 application

You can view the log files of a V Series Node or download them as CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and click the **APPS LOGGING** link.
3. The **GVHTTP2 Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Configuration of 5G-Cloud Application for Casa vTAP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain:

1. Create ingress REP to extract the HTTP2 body (gRPC content). Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi)for more information on creating a REP.
2. Add the 5G-Cloud application in the Monitoring Session.
3. Create a link between ingress REP and the 5G-Cloud application.
4. Create egress REP.
5. Create a link between the 5G-Cloud application and egress REP.

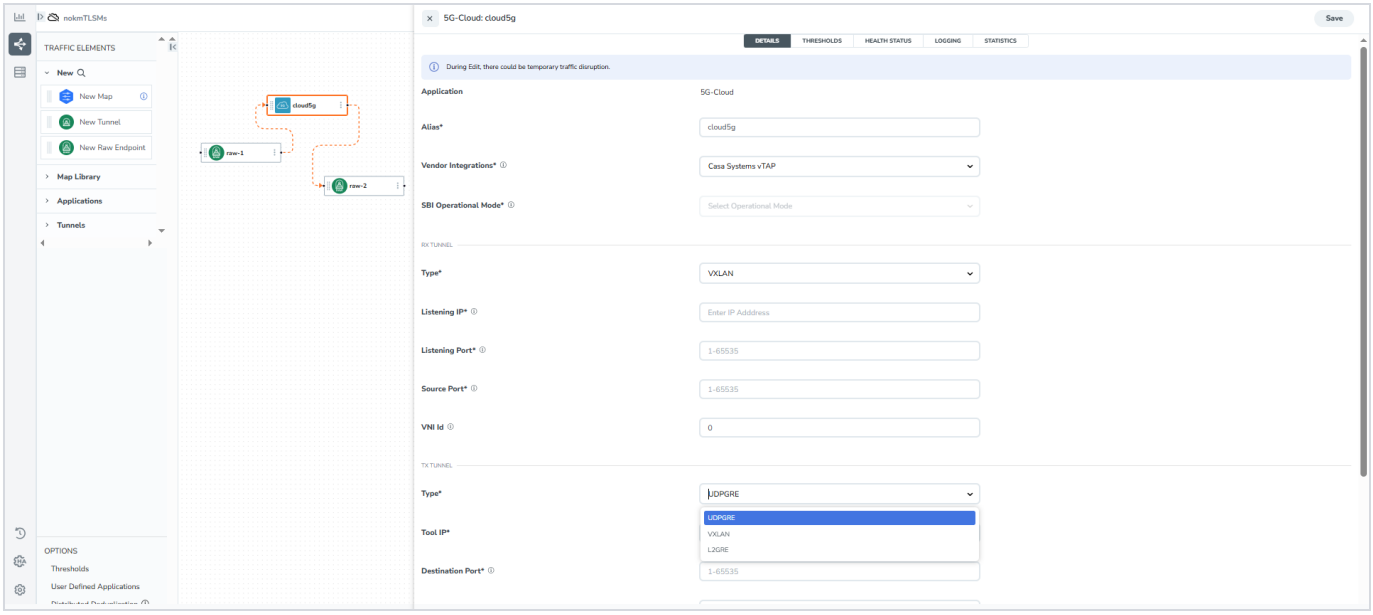


You can add a 5G-Cloud application to:

- **New Monitoring Session** - Add the 5G-Cloud application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
- **Existing session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

To add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



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

Field	Description
Application	The name **5g-Cloud** appears by default.
Alias	Enter the alias name as **cloud5g**.
Vendor Integrations	From the drop-down list, select **Casa Systems vTAP**.

RX Tunnel (Specifies the tunnel configuration from GVHTTP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Listening IP	Specify the tunnel's local listen IP address (TX tunnel destination IP address used in GVHTTP2) to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is 0.
TX Tunnel	
Type	Select the tunnel type from the drop-down list. The available tunnel types are: • UDPGRE • VXLAN • L2GRE
Tool IP	Specify the remote IP address to which the packet will be sent.
Destination Port	Specify the tunnel destination port where the packet will be sent. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind for sending the packet. Enter a value between 1 and 65535.
VNI ID **(Applicable only when the selected tunnel type is VXLAN**)	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key **(Applicable only when the selected tunnel type is L2GRE**)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU. **Note**: For V Series Node version 6.8.00 and above, the range should be between 1400 and 8800. The default value is 8800. For V Series Node version below 6.8.00, range should be between 1500 and 8800. The default value is 8800.
Log Directory	Specify the path to store the log files.
Log Level	In the tunnel log level drop-down list, select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug

Downloading Logs for 5G-Cloud application

You can view the log files of a V Series Node or download them as CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g** application, click the ⋮ menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
3. The **CLOUD5G Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Limitations

- The 5G-Cloud application does not validate the VNI ID configured by the user in the GVHTTP2 application when processing the VXLAN packets.
- The VNI ID configured by the user will not be validated while processing the packets received on the ingress interface. So, the 5G-Cloud application will not drop the VXLAN received on the ingress interface if the VNI ID received in the packet does not match the VNI ID configured by the user.
- After deploying the Monitoring Session from GigaVUE-FM, users can only change the following configuration parameters for the GVHTTP2 application:
  - Server (Listen IP Address and Port Number)
  - VXLAN (Source IP and Port Number)
  - VXLAN (Destination IP and Port Number)
- After deploying the Monitoring Session from GigaVUE-FM, users can only change the following configuration parameters for the 5G-Cloud application:
  - Ingress VXLAN (Listen IP, Listen Port Number, and Source Port Number)
  - Egress UDP-GRE/L2GRE/VXLAN Tunnels (Destination IP Address, Destination Port, Source IP, and Source Port)

Rules and Notes

- On a V Series Node, only one application (GVHTTP2 or 5G-Cloud or AMX) can be deployed across Monitoring Sessions.
- GVHTTP2, 5G-Cloud, and AMX applications should not be combined with general GigaSMART applications in a Monitoring Session or across Monitoring Sessions.
- For GVHTTP2, 5G-Cloud, and AMX applications, the source and destination links should be Raw Endpoints.

- Support for GVHTTP2, 5G-Cloud, and AMX applications is limited to 'Customer Orchestrated Source' tapping type.
- The 5G-Cloud and GVHTTP2 applications require the "VBL-ADD-ADV5GC" or "VBL-ADD-5GC" license.
- To deploy a Monitoring Session, ensure that the 'HTTP2 Listening IP address' and 'TX Tunnel Source IP address' in GVHTTP2 corresponds to the IP addresses of one of the V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the 'RX Tunnel Listening IP address' and 'TX Tunnel Source IP address' in 5G-Cloud correspond to the IP addresses of one of the V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the TX Tunnel Destination IP address used in GVHTTP2 and the RX Tunnel Listening IP address used in 5G-Cloud are the same.

NOTE: For information on statistics details of Casa vTAP for GVHTTP2 and 5G-Cloud applications, refer to Appendix - Casa vTap Statistics.

5G-Cloud Oracle and Nokia SCP Support

GigaVUE V Series Node supports receiving data from 5G network functions in Oracle SCP and Nokia SCP input formats. In both Oracle and Nokia SCP, 3GPP standard method is followed, which has all the standard network components implemented.

Oracle SCP

Oracle systems input format is HTTP2/HTTP2S POST over TCP with TLS enabled. In Oracle SCP, the 5G-Cloud application synthesizes the packets and egresses them to the tool port after it has received all the 4-leg transactions. Oracle SCP processing in the 5G-Cloud application supports converting the discovery and the 4-legged (consumer, SCP, and producer) transactions to packets as per 3GPP Model C and Model D.

Model C

The Consumer knows the producer and sends messages through SCP. Model C comprises of 4-legged transactions between the consumer, SCP, and the producer. The 5G-Cloud application converts the 4-legged transactions to packets and sends them to the probe.

Model D

SCP sends discovery messages to discover the producer after receiving the messages from the consumer. Model D comprises of 2-legged transactions between the consumer and the SCP, 4-legged transactions between the SCP, SCP discovery, and the NRF to find the producer, and the 2-legged transactions between the producer and the SCP. The 5G-Cloud application converts the 4-legged transactions between the consumer, SCP, and the producer to packets and sends them to the probe. The 5G-Cloud application also converts the discovery transactions between the SCP and the NRF and sends them to the probe.

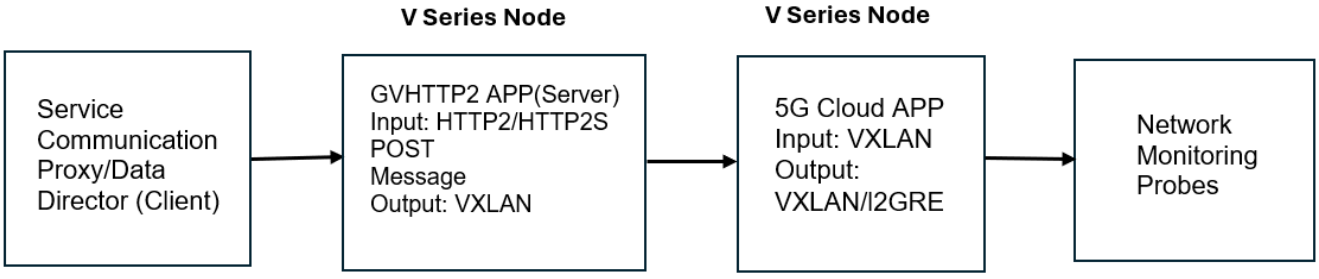
5G- Cloud application will ensure that both the discovery transaction and the 4-legged transaction are transmitted in order, to the connected tool or probe.

Nokia SCP

In Nokia SCP, traffic is mirrored between control functions using the HTTP2 protocol. This traffic is mirrored from a service called SCP (Service Control Proxy), a centralized point through which all the communications between all the control plane functions pass. SCP encrypts all the messages that are going across the control plane network functions and forwards that information in JSON-encoded format to GVHTTP2.

Both Model C and Model D transactions can be received from the SCP. However, Model D cannot be identified as Model D, because a separate message id is used for the discovery messages. In the case of Model D, SCP will include the Model D discovery transactions that happened with NRF along with the network function communication. Therefore, 6 transactions will be forwarded from SCP for a single request and response pair between any two network functions. The Model C transactions will be processed in 2 modes:

- Transparent (2 legged): The 5G-Cloud application will synthesize a direct transaction between the network functions. The 4-legged transactions received between the consumer, SCP, and the producer are synthesized to look like 2-legged transactions between the consumer and the producer.
- Inbound and outbound (4 Legged): The 5G-Cloud application will synthesize indirect transactions, meaning all transactions between any network function will go via SCP. All the transactions received in Model C or Model D are synthesized and sent to the probes.



GVHTTP2 Application

The first application, GVHTTP2, acts as an HTTP2 server, receives the data over HTTP2 frame and establishes a TCP connection with multiple Oracle SCP/DD or Nokia SCP pods. It receives the HTTP2 headers and stores them in a temporary buffer. Once it receives all the headers and body of that stream, it acknowledges back to the client. The application will extract the HTTP2 header and HTTP2 body, encapsulate the data into a VXLAN header, and send it out over the VXLAN tunnel to the next application, 5G-Cloud.

5G-Cloud Application

The 5G-Cloud application receives and decapsulates the VXLAN packet. It extracts the HTTP2 body, parses the SCP/DD JSON mirrored message, and associates the mirrored message transaction ID to track the complete transaction. It then synthesizes TCP Flow and HTTP2 Session representing full 5G CNF / SCP / 5G CNF transaction data and sends it out over L2GRE/VXLAN tunnel to the probe.
The 5G-Cloud application runs in Oracle SCP, Nokia SCP Transparent, and Nokia SCP Inbound-Outbound modes.

- In Nokia SCP Inbound-Outbound mode, both Inbound (between Consumer and SCP) and Outbound (between SCP and Producer) traffic is synthesized and sent to the connected Tool port. In Nokia SCP Transparent mode, only Inbound (between Consumer and SCP) traffic is processed and sent to the Tool Port.
- In the Nokia SCP Transparent mode, the 5G-Cloud Application will drop the Outbound traffic from the SCP (if received in the input).
- In the Nokia SCP Transparent mode, the Producer replaces the SCP in transactions between the Consumer and the SCP.
- In the Nokia SCP Transparent mode, authority header value will be replaced (if the replace authority configuration is enabled).

Configuration of 5G-Cloud Oracle/Nokia SCP

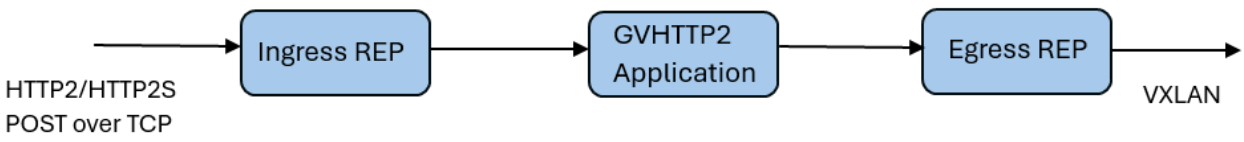
To configure the GVHTTP2 and 5G-Cloud applications in GigaVUE-FM, refer to the following sections:

- Configuration of GVHTTP2 Application for Oracle SCP/Nokia SCP
- Configuration of 5G-Cloud Application for Oracle SCP/Nokia SCP

Configuration of GVHTTP2 Application for Oracle SCP/Nokia SCP

In GigaVUE-FM, you must do the following to add the GVHTTP2 application in the Monitoring Session of a Monitoring Domain:

1. Create an ingress REP (Raw End Point) to receive the HTTP2/HTTP2S data over TCP. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the GVHTTP2 application in the Monitoring Session.
3. Create a link between ingress REP and the GVHTTP2 application.
4. Create egress REP.
5. Create a link between the GVHTTP2 application and egress REP.

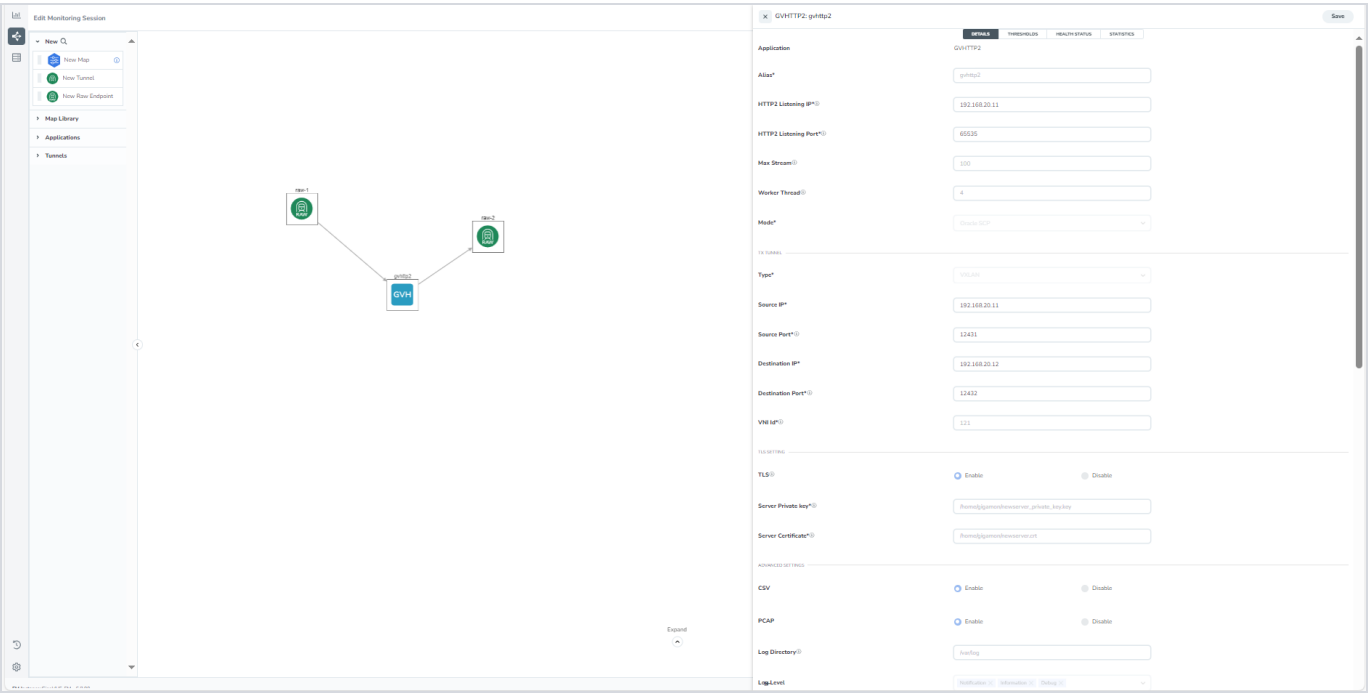


You can add a GVHTTP2 application to:

- New Monitoring Session - You can add the GVHTTP2 application after creating a new Monitoring Session and when the canvas appears.
- Existing Monitoring Session - Click **Edit** on the existing Monitoring Session, and the GigaVUE-FM canvas appears.

To add a GVHTTP2 application:

1. In the canvas, drag and drop the GVHTTP2 application and select **Details**. The GVHTTP2 quick view appears.



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

Field	Description
Application	The name **GVHTTP2** appears by default.
Alias	Enter the alias name as **gvhttp2**.

HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specifies the number of worker threads. Enter a value between 1 and 128. The default value is four.
Mode	Select Oracle SCP or Nokia SCP as mode from the drop-down menu.
TX Tunnel (Specifies the tunnel configuration from GVHTPP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the source port to bind for sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the destination IP address to which packets will be sent.
Destination Port	Specify the destination port where packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TLS Setting	
TLS	Select a TLS/SSL Monitor Mode from one of the following options: ▪ Enable—When the monitor mode is enabled, the TLS/SSL decryption or encryption is off. ▪ Disable—This is the default value. When the monitor mode is disabled, the TLS/SSL decryption or encryption is on.
Server Private key*	Specify the server private key only when TLS is enabled.
Server Certificate*	Specify server certificate only when TLS is enabled.
Advanced Setting	
CSV	Enable or disable the CSV log level.
PCAP	Enable or disable the PCAP log level.
Log Directory	Specify the path to store the log files.
Log Level	In the log level drop-down list, select the severity log level of the events from the following options. The default value level is Notification. • Notification • Information • Debug

Download Logs - GVHTTP2 application

You can view the log files of a V Series Node or download them as .CSV or .txt file.

To download the log files to a local environment:

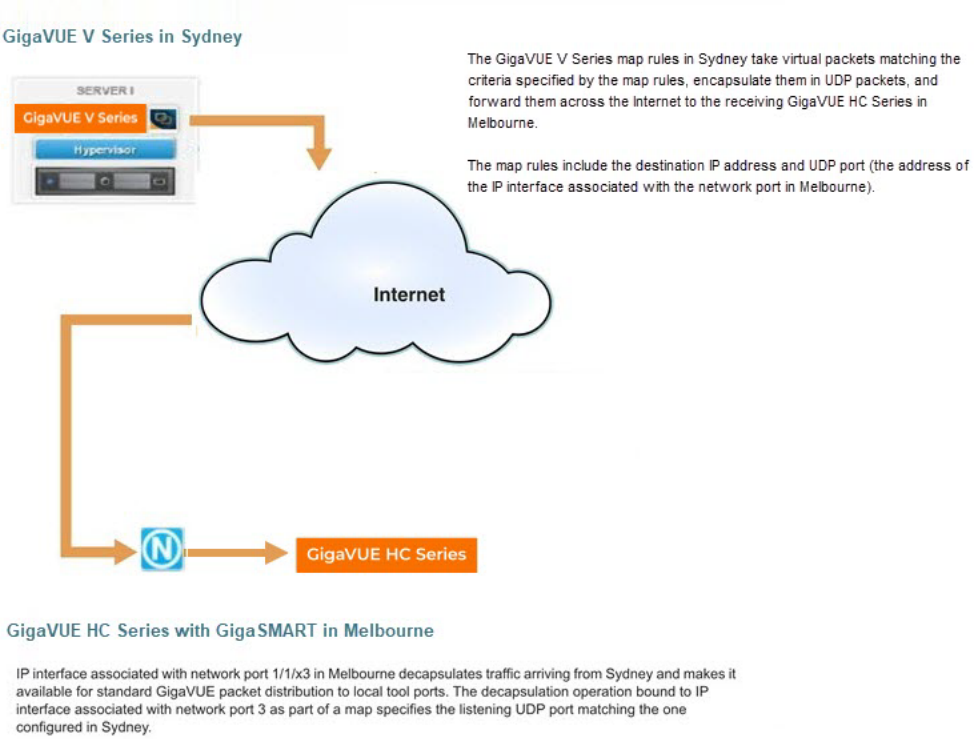
1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.

2. Select the required Monitoring Session and click the **APPS LOGGING** link.
3. The **GVHTTP2 Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Configuration of 5G-Cloud Application for Oracle SCP/Nokia SCP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain:

1. Create ingress REP to extract the HTTP2 body. Refer to *Create Raw Endpoint* section in GigaVUE Cloud Suite Deployment Guide - VMware (ESXi) for more information on creating a REP.
2. Add the 5G-Cloud application in the Monitoring Session.
3. Create a link between ingress REP and the 5G-Cloud application.
4. Create egress REP.
5. Create a link between the 5G-Cloud application and egress REP.



Adding 5G-Cloud Application in Oracle SCP/Nokia SCP

Prerequisites

The pre-requisite to add a 5G-Cloud Application in Oracle SCP/Nokia SCP is:

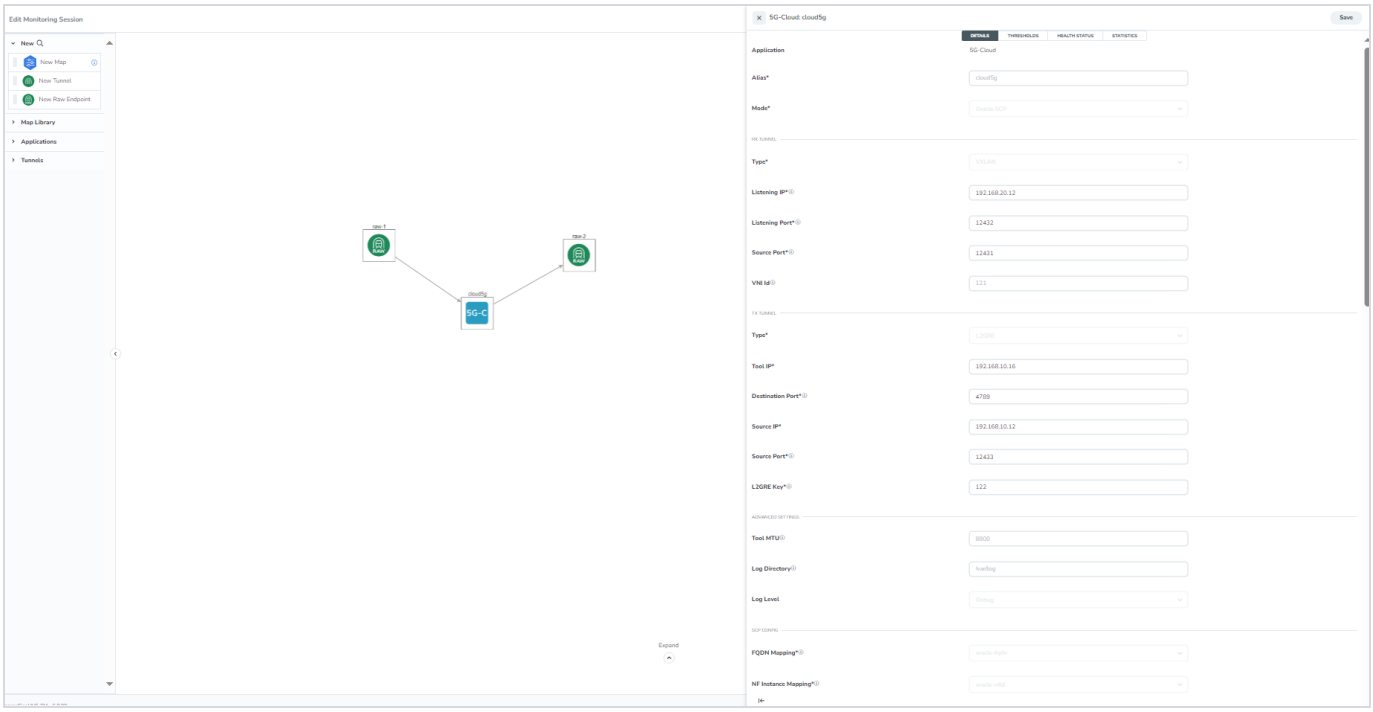
- You must upload CSV files containing a valid FQDN, NF Instance, and User Agent ID's and a valid IPv4/IPv6 addresses. To upload the CSV file, refer to Adding CSV file for IP Mapping.

You can add a 5G-Cloud application to:

- New Monitoring Session - You can add the 5G-Cloud application after creating a new Monitoring Session and when the canvas appears.
- Existing session - Click **Edit** on the existing Monitoring Session, and the GigaVUE-FM canvas appears.

To add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



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

Field	Description
Application	The name **5g-Cloud** appears by default.
Alias	Enter the alias name as **cloud5g**.
Mode	From the drop-down list, select the required mode from the following options: • Oracle SCP • Nokia SCP Transparent • Nokia SCP Inbound/Outbound
RX Tunnel (Specifies the tunnel configuration from GVHTPP2 to 5G-Cloud)	
Type	Specify the tunnel type. VXLAN is selected by default.
Listening IP	Specify the tunnel's local listen IP address (TX tunnel destination IP address used in GVHTTP2) to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
TX Tunnel	
Type	Specify the tunnel type.
Tool IP	Specify the remote IP address to which the packet will be sent.

Destination Port	Specify the tunnel destination port where the packet will be sent. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind for sending the packet. Enter a value between 1 and 65535.
VNI Id **(Applicable only when the selected tunnel type is VXLAN**)	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key **(Applicable only when the selected tunnel type is L2GRE**)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU. **Note**: For V Series Node version 6.8.00 and above, the range should be between 1400 and 8800. The default value is 8800. For V Series Node version below 6.8.00, range should be between 1500 and 8800. The default value is 8800.
Log Directory	Specify the path to store the log files.
Log Level	In the tunnel log level drop-down list, select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug
SCP Config	
FQDN Mapping	Specify the alias name created for the uploaded FQDN table CSV file. You can also add new FDQN by utilizing the **ADD** button. **Note**: **Header details**: gigamon@test-5gc-new:/var/log/cloud5g_tabledir$ cat FQDN.txt FQDN,IngressIP1,IngressIP2,IngressPort,EgrIP,NFType,NFLocation (The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional) **Examples - 5G-Cloud FQDN Entries**: **Oracle SCP**: FQDN.txt:scp1.pp2.scp.5gc.mnc001.mcc208.3gppnetwork.org,10.64.104.30,10.64.104.36,9443,10.64.104.45,SCP,plolp

	Nokia SCP: s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,172.23.13.187,,8080,,SCP,plolp
NF Instance Mapping	Specify the alias name created for the uploaded NFID table CSV file. You can also add new NF Instance by utilizing the **ADD** button. **Note**: **Header details**: gigamon@test-5gc-new:/var/log/cloud5g_tabledir$ cat NFID.txt NFID,IngressIP1,IngressIP2,IngressPort,EgrIP,NFType,NFLocation (The fields NFID and ingressIP1 are mandatory. All other fields are optional) **Examples - 5G-Cloud NFID Entries:** • **Oracle SCP**: NFID.txt:04fba086-6ac1-4472-9ee3-cf10e57849f8,10.64.107.85,10.64.107.79,8080,,SMF,plolp • **Nokia SCP**: 8b94c7ee-d750-4bcd-95be-634b31a5fa0d,172.23.14.4,,8080,,AMF,jeolp
User Agent Mapping	Specify the alias name created for the uploaded User Agent table CSV file. You can also add new User Agent by utilizing the **ADD** button. **Note**: **Header details**: gigamon@test-5gc-new:/var/log/cloud5g_tabledir$ cat UA.txt UserAgent,IngressIP1,NFType,NFLocation (The fields useragentID and ingressIPAddress are mandatory. All other fields are optional) **Examples - 5G-Cloud User Agent Entries:** • **Oracle SCP**: pp5gs2nd15gcam01,10.64.106.67,AMF,jeolp • **Nokia SCP**: pp5gs2nd15gcam01,10.64.106.67,AMF,plolp pp5gs2nd15gcam02,10.64.106.107,AMF,plolp ofrpp25gcsmsf01,10.64.106.138,SMSF,jeolp
SCP Advanced Config	
No. of TCP Flows	Specify the total number of concurrent TCP flows. Enter a value between 128 and 2048. The default value is 1024.
No. of Transaction Flows	Specify the total number of Transaction Stream Flows allocated for tracking active Transaction Streams. Enter a value between 128 and 5500. The default value is 2048.
TCP Flow Timeout	Specify the TCP flow timeout ranging between 0 to 7200. The default value is 900.
SCP Transaction Timeout	Specify the SCP transaction timeout value ranging between 1 and 300. The default value is 10. **Note**: The recommended timeout value is between 10 and 20 seconds. An increase in the transaction timeout value will impact the total number of transactions that can be processed in parallel.
Minimum TCP Client Port	Specify the minimum value to be used for TCP client port ranging between 1023 and 65535. The default value is 32768.
Maximum TCP Client Port	Specify the maximum value to be used for client port ranging between 1023 and 65535. The default value is 36863.
Header	Enable or disable the header indexing flag from the available options. The default value is

Index	disable.
Header Compression Code **Note**: HTTP/2 requests and responses typically contain a large number of redundant headers. HTTP/2 uses header compression to compress headers significantly.	Enable or disable the header compression code from the available options. The default value is disable.
Gigamon Header **Note**: Helps in the debugging process to identify all transactions.	Enable or disable the Gigamon Header from the available options. The default value is Enable.
Nokia Use 3gpp Target API Root **(Applicable only when the selected mode is Nokia SCP Transparent)**	Specify the value which upon detecting producer IP/FQDN, treats the 3gpp Target API Root to be predictive of the Producer IP if value is non-zero as either enable or disable. The default value is disable.
Nokia Replace Authority	Specify the value which upon detecting producer IP/FQDN, treats the 3gpp Target API Root to be predictive of the Producer IP if the value is non-zero as either enable or disable. The default value is disable.

(Applicable only when the selected mode is Nokia SCP Transparent)	
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
5G-Cloud Log Level	Select the required 5G-Cloud CSV log level from the drop down list. The default value is None. • All • Flow • Message • Transaction • None

Download Logs - 5G-Cloud application

You can view the log files of a V Series Node or download them as .CSV or .txt file.

To download the log files to a local environment:

1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and click the **APPS LOGGING** link.
3. The **CLOUD5G Logging** page displays the logs currently available. You can download the preferred logs by selecting the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download Files** from the **Actions** menu.
5. The system downloads the selected files to your local environment.

Adding CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field
3. Select any of the following from the **Type** as per the requirement:

- **5G-Cloud FQDN** - Adding the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping. Adding FQDN is mandatory.
- **5G-Cloud NF Instance** - Adding the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.
- **5G-Cloud UA** - Adding the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.

4. Click **Choose File** in **FileName** field to upload the CSV file into GigaVUE-FM.
5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

NOTE: For Oracle SCP, configuration of FQDN Mapping, NF Instance Mapping, and User Agent Mapping are mandatory. For Nokia SCP Transparent and Nokia SCP Inbound/Outbound, configuration of FQDN Mapping is mandatory.

Limitations

- The following configuration parameters for the GVHTTP2 application are non-editable and can be configured only during the initial configuration:
  - Mode
  - TX Tunnel – txType
  - Worker thread
  - RX VNI Id
  - Max Streams
  - TLS configuration
  - Log Level
  - Log Directory
- The following configuration parameters for the 5G-Cloud application are non-editable and can be configured only during the initial configuration:
  - Mode
  - TX Tunnel – txType
  - RX Tunnel - rxType
  - RX VNI Id
  - Tool MTU
  - Log Directory
  - Log Level
  - Header Index
  - Header Compression Code
  - Packet capture log level
  - SCP log level
  - Gigamon Header

- FQDN Mapping
- NF Instance Mapping
- User Agent Mapping

Rules and Notes

- The 5G-Cloud and GVHTTP2 applications require the "VBL-250T-ADD-5GC" license.
- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to **Procedure to deploy V Series Node in VMware ESXi with SR-IOV Adapte** topic in GigaVUE Cloud Suite Deployment Guide - Third Party Orchestration guide for more details.
- On a V Series Node, only one application (GVHTTP2 or 5G-Cloud or AMX) can be deployed across Monitoring Sessions.
- GVHTTP2, 5G-Cloud, and AMX applications should not be combined with general GigaSMART applications in a Monitoring Session or across Monitoring Sessions.
- For GVHTTP2, 5G-Cloud, and AMX applications, the source and destination links should be Raw Endpoints.
- Support GVHTTP2, 5G-Cloud, and AMX applications is limited to 'Customer Orchestrated Source' tapping type.
- Configuration of GVHTTP2 and 5G-Cloud in V Series Node will fail in versions less than v6.7.00.
- To deploy a Monitoring Session, ensure that the 'HTTP2 Listening IP address' and 'TX Tunnel Source IP address' in GVHTTP2 correspond to the IP addresses of one of the GigaVUE V Series Nodeinterfaces.
- To deploy a Monitoring Session, ensure that the 'RX Tunnel Listening IP address' and 'TX Tunnel Source IP address' in 5G-Cloud correspond to the IP addresses of one of the GigaVUE V Series Node interfaces.
- To deploy a Monitoring Session, ensure that the TX Tunnel Destination IP address used in GVHTTP2, and the RX Tunnel Listening IP address used in 5G-Cloud are the same.
- When you have already deployed the 5G-Cloud application with all the table files uploaded, you should undeploy and redeploy the Monitoring Session again with the updated files to update the table.
- In the GVHTTP2 application, when the TLS/SSL configuration is enabled, the private key and certificate path must be configured.
- Before upgrading GigavUE V Series Node, user should undeploy the Monitoring Session where the GVHTTP2 application is deployed and take the backup of the SSL certificate and private key files. Following the upgrade, user should manually install/copy the private key and the certificate file in GigaVUE GigaVUE V Series Node and redeploy the GVHTTP2 application.
- To upload the table configuration file in CSV format, go to Inventory -> Virtual ->Settings -> 5G-Apps, enter the alias, type the CSV file name, and upload it.
- To download the files, go to Inventory -> Virtual ->Settings -> 5G-Apps and select the alias name to be downloaded.

- From release version 6.9.00, the maximum Transaction Flow value for GigaVUE V Series Nodes with memory less than 32 GB is 5500. For V Series Nodes classified as extra-large with memory greater than 32 GB, the maximum Transaction Flow value is 26000.
- From release version 6.9.00, to make use of the multithreading support for egress/ingress tunnels in GVHTTP2/5G-Cloud applications, you should update the thread count after the GigaVUE V Series Node is upgraded.
- The configuration of the 5G-Cloud application should be performed on extra-large VMs.
  - Configuration specifications for extra large VM:
    - vCPUs: 12-20
    - RAM: 32 GB
    - Disk: 160 GB
  - Supported Ports: Management Ports, Tunneling Ports, and 2 Data Ports

FHA Dashboards for 5G-Cloud Applications

After configuring the GVHTTP2 and 5G-Cloud applications, you can monitor the statistics for Oracle SCP and Nokia SCP by the reports displayed in the Dashboard.

To access the details, refer to **FHA Dashboards for 5G-Cloud Applications** topic in GigaVUE Fabric Management Guide.

5G-Cloud Ericsson SCP Support

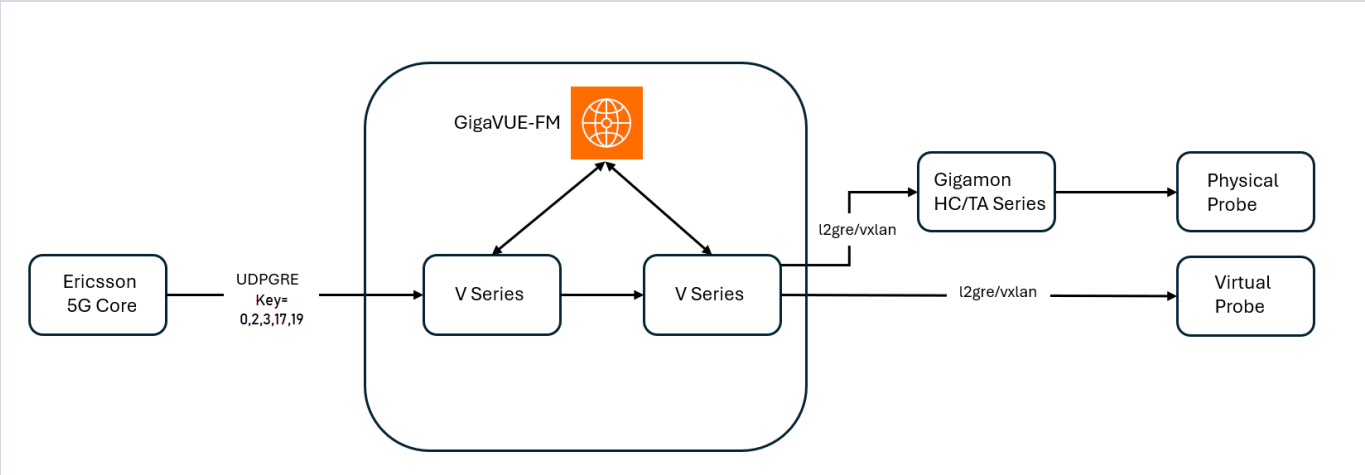
The 5G Core follows a service-based architecture where control plane network functions communicate using the HTTP2 protocol. In Ericsson SCP Mode, HTTP2 transactions are mirrored through the Service Communication Proxy (SCP) and encapsulated using UDP-GRE for further processing.

Ericsson SCP

In **Ericsson vTAP**, traffic is tapped from the Packet Core Controller (PCC), including AMF and SMF, in Layer 7 HTTP SBI JSON format. The 5G SBI application supports this PCC traffic, ensuring efficient processing and monitoring. Refer to 5G-Service Based Interface Application.

In **Ericsson SCP** mode, the Service Communication Proxy (SCP) mirrors 5G SBI traffic within the Cloud Core Signaling Controller, capturing communication between multiple network functions via SCP.

The 5G-Cloud application in GigaVUE V Series Nodes is enhanced to receive and process traffic from the Ericsson SCP and forward it to the monitoring tools. This setup supports both configuration and visualization through GigaVUE-FM integration.



Traffic from the Ericsson 5G Core is encapsulated using UDP-GRE tunnels, with specific keys 3, 17, and 19 to identify different traffic flows from SCP. GigaVUE-FM serves as the centralized management platform orchestrating the visibility fabric, which includes the V Series Nodes. The V Series Nodes inspect, aggregate, and process the incoming packets.

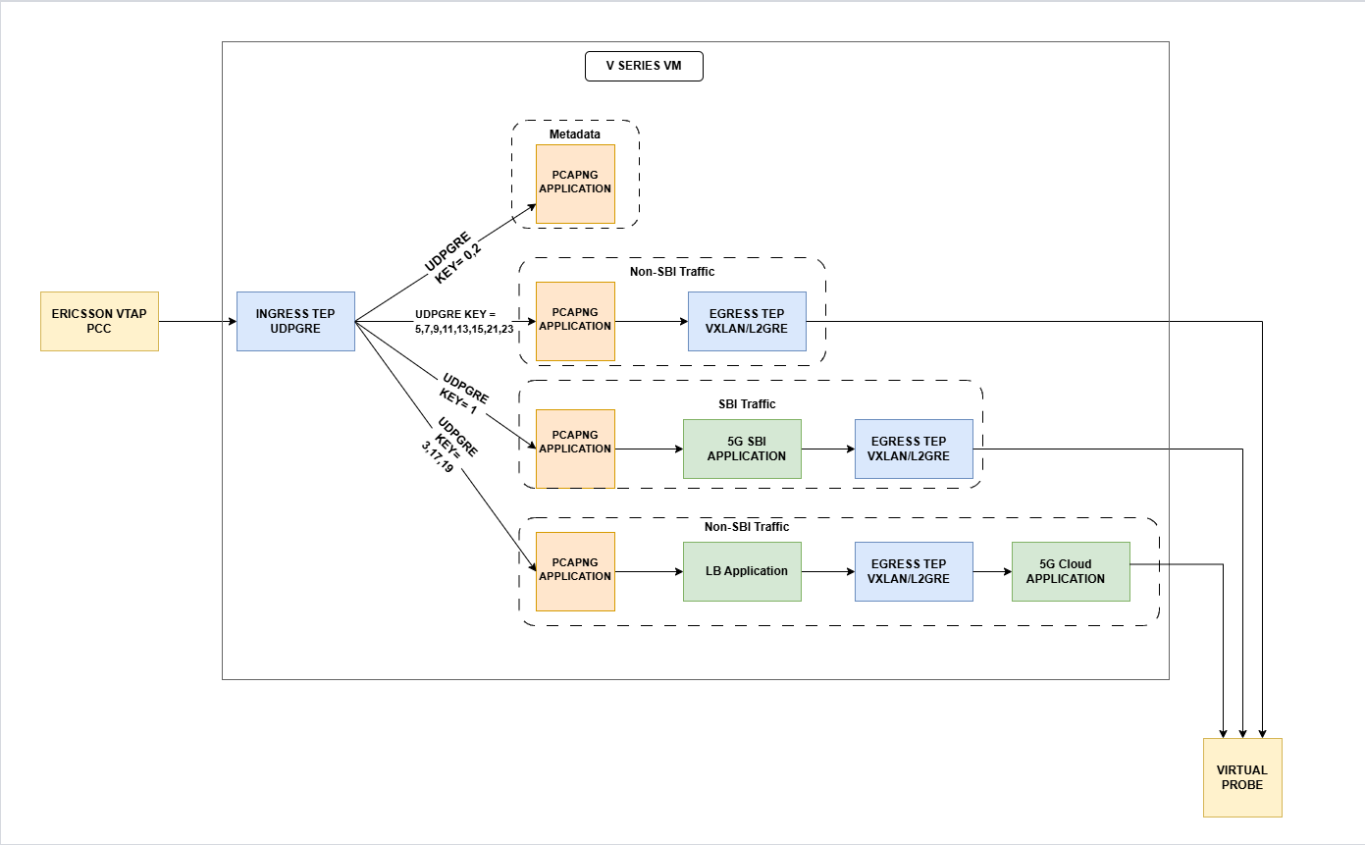
After processing, traffic is encapsulated using L2GRE or VXLAN tunneling protocols and forwarded to either physical or virtual probes as follows:

- **Physical Probe**: If IP connectivity is not available, the traffic is sent to a GigaVUE HC/GigaVUE TA Series device, which serves as a traffic aggregator and distributor. The device decapsulates and distributes the traffic to the probe.
- **Virtual Probe**: If IP connectivity is available, the processed traffic is forwarded directly to virtual probe for detailed traffic inspection and analytics.

This end-to-end solution ensures comprehensive visibility into 5G network traffic, facilitating efficient monitoring and performance analysis across both physical and cloud-based infrastructures.

How Ericsson SCP Solution works

The following image shows the block diagram of the data flow in the V Series Nodes.



The encapsulated traffic enters the system, and based on the key values, is passed through a PCAP-NG, 5G SBI, and LB application to capture and monitor the data.

For traffic mirroring and flow processing in Ericsson VTap 5G Core network with key value 1, refer to 5G-Service Based Interface Application

When traffic is processed through keys 3, 17, and 19, the **PCAPng** application reads and validates various blocks within received PCAPng files before forwarding them to the Load Balancing application. It processes data based on the packet type, which includes a combination of the following blocks.

1. Mandatory Blocks:
    - Section Header Block (SHB)
2. Optional Blocks:
    - Interface Description Block (IDB)
    - Enhanced Packet Block (EPB)
    - Simple Packet Block
    - Name Resolution Block
    - Interface Statistics Block

Also, the PCAPng application classifies domains by using Ericsson vTAP metadata and a configured domain table. Each Ericsson domain (for example, pc-mm, pc-sm, ccsm-5gc, ccrc-nrf, nrfagent-nrf, ccpc-pcf, sc-scp, sc-sepp) maps to a VLAN ID. The V Series applies this VLAN ID on the egress tunnel so downstream probes can distinguish traffic by Ericsson domain.

The domain table can also include others and unclassified to handle traffic that does not match known domains or cannot be classified. Refer to PCAPng Application for more details.

The **Load Balancing** application efficiently distributes traffic with different flow IDs across multiple cores, enabling parallel processing of packets with distinct stream IDs. It ensures that request and response frames with the same stream ID are processed sequentially within the same core, preventing race conditions where these frames might otherwise be handled in parallel across different cores. Refer to Load Balancing.

The extracted data is encapsulated within a VXLAN and routed to the 5G-Cloud application. The 5G-Cloud application processes VXLAN packets by first decapsulating them to extract the HTTP2 headers and body. It then parses the EEVTAP TCP mirrored message and maps the transaction ID to track the full transaction.

After synthesizing the TCP flow and HTTP2 session to represent the complete 5G CNF/SCP/5G CNF transaction data, it forwards the data over an L2GRE/VXLAN tunnel to the probe.

You should enable VLAN Tagging on the egress tunnel (VXLAN or L2GRE) used with the domain-classification PCAPng instance so the tunnel applies the VLAN ID from the Domain VLAN Mapping table for each Ericsson domain. Refer to Create Ingress and Egress Tunnels.

Notes:

- Ericsson SCP will process only Model C traffic.
- You can deploy traffic ingress applications together in a single V Series Node and Monitoring Session. However, for the 5G-Cloud and egress applications, you must deploy a separate Monitoring Session and V Series Node.
- You can configure either a single-tep or multi-tep setup for the egress tunnel. Switching between these configurations is not allowed; to make changes, you must undeploy and redeploy the Monitoring Session.
- Ensure that all V Series Nodes within a single Monitoring Domain are running the same version. Mixing different versions in the same Monitoring Domain may lead to inconsistencies when configuring Monitoring Session traffic elements.

Configuration of 5G-Cloud Ericsson SCP

In GigaVUE-FM, you must do the following to configure the 5G-Cloud application in the Monitoring Session of a Monitoring Domain in the V Series:

V Series Node 1:

S.No	Steps	Refer to
1	Configure UDP-GRE Ingress TEP to receive the HTTP2 messages. **Note**: You should configure individual UDP-GRE tunnels for each of the following key values: • Key Value 1: Ericsson vTAP • Key Values 3 and 17: Ericsson SCP	Create Ingress and Egress Tunnel (VMware vCenter) Create Ingress and Egress Tunnels (OpenStack) Create Ingress and Egress Tunnels (OpenStack)
2	Configure PCAPng application and link ingress TEP and PCAPng application instances in the Monitoring Session. **NOTE:** For Domain tagging configuration, refer to **Configure the PCAPng Application for Domain Tagging** section in PCAPng Application.	PCAPng Application
3	Configure LB application and link LB application and PCAPng application instances in the Monitoring Session.	Load Balancing
4	Create a VXLAN egress TEP. **NOTE:** Enable the **Domain Tagging** option when creating an egress tunnel to tag packets with the Ericsson domain specific VLAN IDs.	Create Ingress and Egress Tunnel (VMware vCenter) Create Ingress and Egress Tunnels (OpenStack) Create Ingress and Egress Tunnels (OpenStack)
5	Create a link between egress TEP and LB application.	

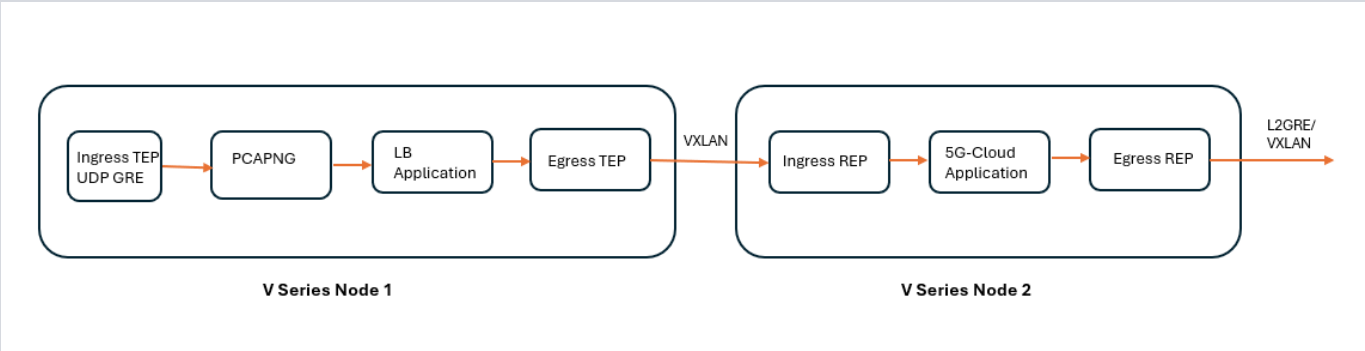
Note: When upgrading GigaVUE-FM with a Monitoring Session in V Series Node 1 to version 6.12 from an earlier release, follow these steps to ensure a successful migration:

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

V Series Node 2:

S.No	Steps	Refer to
1	Create a ingress REP to extract the HTTP2 body.	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (OpenStack) Create Raw Endpoint (Third Party Orchestration)
2	Create a link between the VXLAN egress TEP and ingress REP.	
3	Add the 5G-Cloud application in the Monitoring	5G-Cloud Application

	Session.	
4	Create a link between ingress REP and the 5G-Cloud application.	
5	Create egress REP.	Create Raw Endpoint (VMware vCenter) Create Raw Endpoint (OpenStack) Create Raw Endpoint (Third Party Orchestration)
6	Create a link between the 5G-Cloud application and egress REP.	



Add 5G-Cloud Application in Ericsson SCP

Pre-requisite:

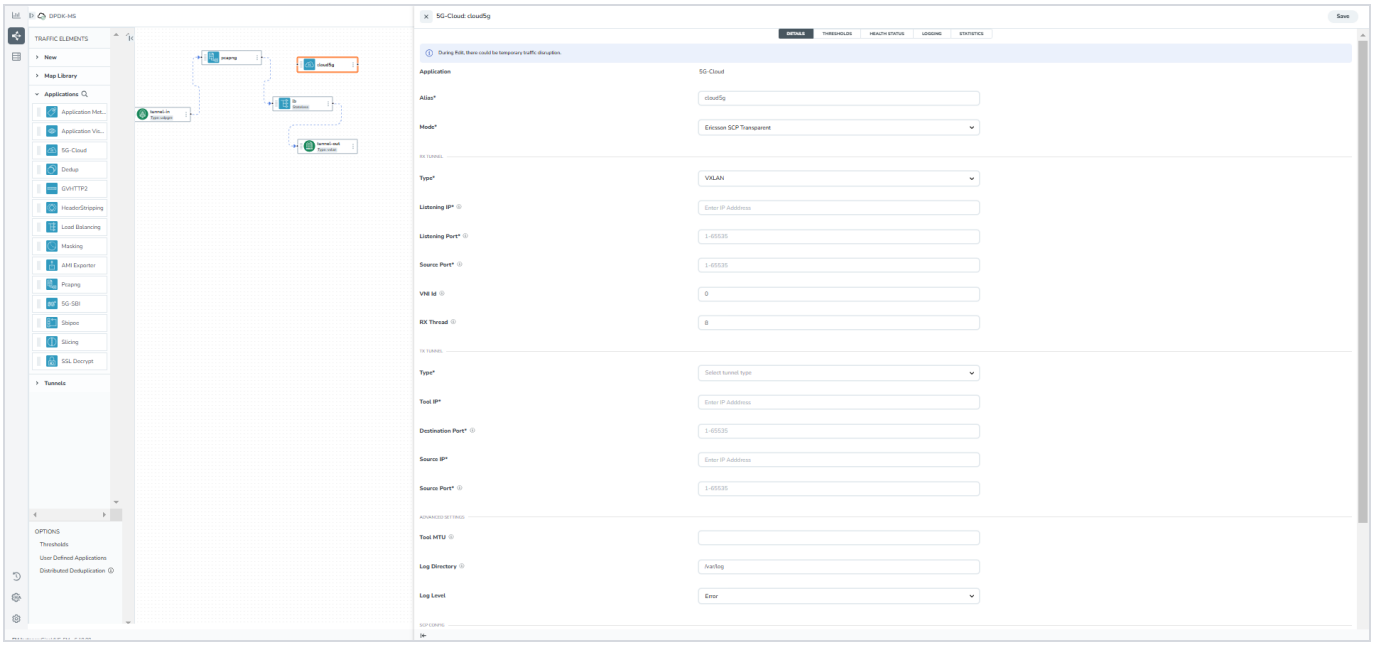
- You must upload CSV files containing a valid FQDN, NF Instance, and User Agent IDs and a valid IPv4/IPv6 address. Refer to Add CSV file for IP Mapping.

You can add a 5G-Cloud application to:

- New Monitoring Session - Add the 5G-Cloud application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
- Existing session - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

To add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



2. On the application quick view, enter or select the required information as described in the 5G Cloud Ericsson SCP Configuration - Field References.

NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

Download Logs - 5G-Cloud application

You can view the log files of a V Series Node or download them as .CSV or .txt files.

To download the log files to a local environment:

1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and click **Details** on the 5G-Cloud application. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
3. Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.

2. Click **New**. Enter the name for the CSV file in the **Alias** field
3. From the Type drop-down list, select one of the following:
   - **5G-Cloud FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.
   - **5G-Cloud NF Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.
   - **5G-Cloud UA** - Add the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.
4. Click **Choose File** in the **FileName** field to upload the CSV file to GigaVUE-FM.

Notes - Header details of the CSV file:

- **FQDN Mapping**: gigamon@vseries:/var/log/cloud5g_tabledir$ cat FQDN.txt (The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional)
  - Header details: FQDNid,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation
  - Example: s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,170.00.13.187,,8080,,SCP,plolp
- **NF Instance Mapping**: gigamon@vseries:/var/log/cloud5g_tabledir$ pwd NFID.txt (The fields NFID and ingressIP1 are mandatory. All other fields are optional)
  - Header details: NFID,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation
  - Example: 8b94c7xx-d700-4xcd-90be-634x11a5xx0d,100.23.14.4,,8080,,AMF,plolp
- **User Agent Mapping**: gigamon@vseries:/var/log/cloud5g_tabledir$ pwd UA.txt (The fields useragentID and ingressIPAddress are mandatory. All other fields are optional)
  - Header details: UserAgent,IngressIP1,NFType,NFLocation
  - Example: xx5gs20015gcay01,192.64.10.67,AMF,plolp

The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.

5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

FHA Dashboards for 5G-Cloud Applications

After configuring the 5G-Cloud application, you can monitor the statistics for Ericsson SCP by the reports displayed in the Dashboard. To access the details, refer to FHA Dashboards for 5G-Cloud Applications.

5G Cloud Ericsson SCP Configuration - Field References

Field	Description
Application	The name **5g-Cloud** appears by default.
Alias	Enter the alias name as **cloud5g**.
Mode	From the drop-down list, select **Ericsson SCP Transparent**.
RX Tunnel	
Type	Specify the tunnel type. The default is VXLAN.
Listening IP	Specify the tunnel's local listen IP address to receive the packet.
Listening Port	Specify the tunnel's local listening port to bind to receive the packet. The application will listen to the traffic coming to the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port from where the packets will be sent. Enter a value between 1 and 65535.
VNI Id	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215. The default value is zero.
RX Thread	Specify the RX value to receive the packet. Enter a value between 1 and 16. The default value is 8.
TX Tunnel	
Type	Specify the tunnel type. The available tunnel types are: • VXLAN • L2GRE
Tool IP	Specify the remote IP address to send the packet.
Destination Port	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending the packet. Enter a value between 1 and 65535.
VNI Id **(Applicable only when the selected tunnel type is VXLAN)**	Specify the VNI to use for the VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key **(Applicable only when the selected tunnel type is L2GRE)**	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Advanced Setting	
Tool MTU	Specify the tool port MTU.
Log Directory	Specify the path to store the log files.
Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification

	• Information • Debug
SCP Config	
FQDN Mapping	Specify the alias name created for the uploaded FQDN table CSV file.
NF Instance Mapping	Specify the alias name created for the uploaded NFID table CSV file.
User Agent Mapping	Specify the alias name created for the uploaded User Agent table CSV file.
TCP Server Ports	Specify the TCP server port or port range to allow TCP communication endpoints.
SCP Advanced Config	
No. of TCP Flows	Specify the total number of concurrent TCP flows. Enter a value between 128 and 2048. The default value is 1024.
No. of Transaction Flows	Specify the total number of Transaction Stream Flows allocated for tracking active Transaction Streams. Enter a value between: • 128 and 26000 for V Series Node of extra large form factor (total memory ≥ 32GB). • 128 and 5500 for V Series Node of large form factor (total memory < 32GB). The default value is 2048.
TCP Flow Timeout	Specify the TCP flow timeout ranging between 0 to 7200. The default value is 900.
SCP Transaction Timeout	Specify the SCP transaction timeout value ranging between 1 and 300. The default value is 10. **Note**: The recommended timeout value is between 10 and 20 seconds. An increase in the transaction timeout value will impact the total number of transactions that can be processed in parallel.
Minimum TCP Client Port	Specify the minimum value to be used for TCP client port ranging between 1023 and 65535. The default value is 32768.
Maximum TCP Client Port	Specify the maximum value to be used for client port ranging between 1023 and 65535. The default value is 36863.
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. Enter a value between 1 and 16. The default value is 8.
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. Enter a value between 100 and 8000. The default value is 1000.
Header Index	Enable or disable the header indexing. The default value is disabled.
Header Compression Code	Enable or disable the header compression code. The default value is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default value is Enable.
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None

5G-Cloud Log Level	Select the required 5G-Cloud CSV log level from the drop down list. The default value is None. • All • Flow • Message • Transaction • None
HTTP2 Monitored Flows	
Number of Streams	Specify the number of monitored stream flows. Enter a value between 1024 and 16384. The default is 1024.
Request Timeout	Specify the HTTP/2 request timeout in seconds. Enter a value between 1 and 300. The default is 15 seconds.
Response Timeout	Specify the HTTP/2 response timeout in seconds. Enter a value between 1 and 300. The default is 2 seconds.
TCP Monitored Flows	
Number of Flows	Specify the number of monitored TCP flows. Enter a value between: • 1024 and 32768 for V Series Node of extra large form factor (total memory > 32GB). • 1024 and 2048 for V Series Node of lesser form factor. The default value is 2048.
Flow Timeout	Specify the timeout for TCP flows in seconds. Enter a value between 1 and 7200. The default value is 60 seconds.
Flow Reassembly Timeout	Specify the timeout for TCP flow reassembly in seconds. Enter a value between 1 and 7200. The default value is 500 seconds.

5G Cloud Nokia HEP3 Support

Nokia HEP3 (Homer Encapsulation Protocol Version 3) method is supported by Nokia for network functions across 3G,4G, 5G and the IMS core. HEP3 efficiently organizes and encodes data for various network protocols and enhances the functionality of Nokia's core network solutions by accommodating diverse packet requirements.

The 5G-Cloud application supports processing the Nokia HEP3 format from the **IMS** network functions and 3G, 4G network functions. Non-SBI traffic is received as a TCP stream with HEP3 encoding by the 5G Cloud application. The 5G-Cloud application receives encapsulated traffic from the IMS vTAP in HEP3 over TCP, using either clear text or mTLS. Non-SBI control traffic is decoded from HEP3, emphasizing data handling without IP translation.

Nokia Cloud Mobility Manager (**CMM**), including AMF and MME, generates SBI and Non-SBI traffic over HTTP2 using clear text / TLS. The gvHTTP2 application receives this traffic and forwards it to the 5G-Cloud application. The 5G-Cloud synthesizes SBI payloads into L2GRE/VXLAN packets, enabling seamless monitoring across 5G core interfaces.

In **Nokia SCP** mode, the traffic is mirrored between the 5G control functions using HTTP2 protocol, which is mirrored from a service called SCP (Service Communication Proxy). The 5G-Cloud application receives the JSON encoded data from the VXLAN ingress TEPs. Refer to Nokia SCP.

The system supports the following Nokia source functions and traffic types:

- 5G Cloud Nokia HEP3 IMS Integration (Non-SBI)
- 5G Cloud Nokia HEP3 CMM Integration (SBI and Non-SBI)

Supportability Table

The supportability table provides a quick reference to the supported Nokia functions:

Source Functions	Traffic Type	Protocols	Reference
3G, 4G, IMS	Non-SBI	HEP3 over TCP with Clear Text / mTLS	5G Cloud Nokia HEP3 IMS Support
CMM (AMF, MME)	SBI and Non-SBI	HEP3 over HTTP2 over TCP with Clear Text / TLS (via gvHTTP2)	5G Cloud Nokia HEP3 CMM Support

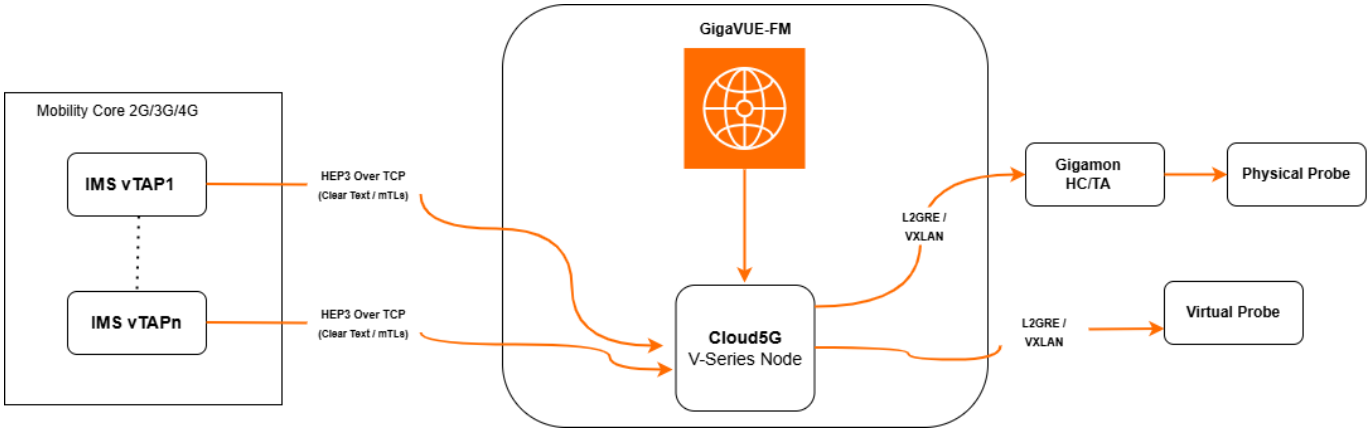
5G Cloud Nokia IMS Solution (Non-SBI)

Nokia HEP3 (Homer Encapsulation Protocol Version 3) is a binary encapsulation format that Nokia uses for control-plane traffic in the IMS core and across 3G and 4G network functions. HEP3 efficiently encodes protocol metadata and payloads so that Nokia core network solutions can handle diverse packet types.

In this solution, the 5G-Cloud application processes non-SBI IMS control traffic that Nokia vTAPs export as HEP3 over TCP, in clear text or protected with mTLS. 5G-Cloud also supports non-SBI HEP3 traffic from 3G and 4G network functions. It terminates the TCP session, decodes the HEP3 payloads, and exposes the non-SBI control traffic for monitoring without performing IP address translation.

How Nokia HEP3 IMS Solution works

The following image shows the ingress and egress data flow through the V Series Node for Nokia HEP3 IMS solution.



Traffic from client vTAP systems (in clear text or secured with mTLS) travels over TCP to the V Series node. The 5G-Cloud application receives this TCP traffic, terminates the TCP connections, decodes the HEP3 payloads, and synthesizes the data for monitoring.

GigaVUE-FM serves as the centralized management platform that orchestrates the visibility fabric, including the 5G-Cloud V Series nodes. The V Series nodes inspect, aggregate, and process the incoming traffic. After processing, the nodes encapsulate the traffic using L2GRE or VXLAN tunneling protocols and forward it to probes as follows:

Physical probe- If IP connectivity is not available, the node sends traffic to a GigaVUE HC or GigaVUE TA Series device, which aggregates and distributes traffic to the probe after decapsulation.

Virtual probe- If IP connectivity is available, the node forwards the processed traffic directly to the virtual probe for detailed inspection and analytics.

This end-to-end solution provides comprehensive visibility into 5G IMS control-plane traffic and supports efficient monitoring and performance analysis across both physical and cloud-based infrastructures.

Configure Nokia IMS HEP3 Solution (Non-SBI)

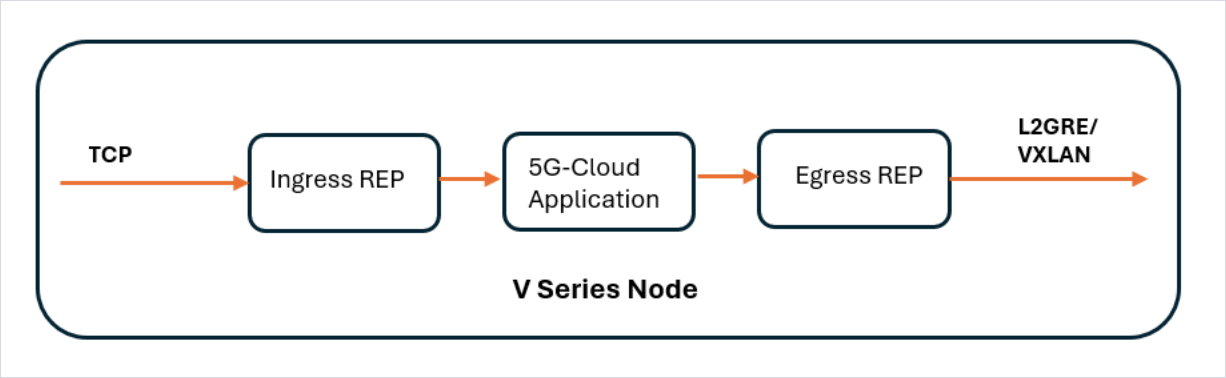
In GigaVUE-FM, you must do the following to add the 5G-Cloud IMS integration to either a new or existing monitoring session of a Monitoring Domain:

Pre-requisite:

Before starting, ensure that:

- You have deployed the V Series nodes and added them to a monitoring domain in GigaVUE-FM.

- IMS vTAPs can reach the V Series node over TCP on the HEP3 collector IP address and port.
- You must upload CSV files containing a valid FQDN and a valid IPv4/IPv6 address. Refer to Add CSV file for IP Mapping.



1. **Create New or Use Exisiting Monitoring Session**

In GigaVUE-FM, go to the monitoring domain that will carry IMS HEP3 traffic. You can add the 5G-Cloud aplication, to a new or existing monitoring session of a Monitoring Domain.

- **New Monitoring Session** - Add the 5G-Cloud application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
- **Existing Monitioring Session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

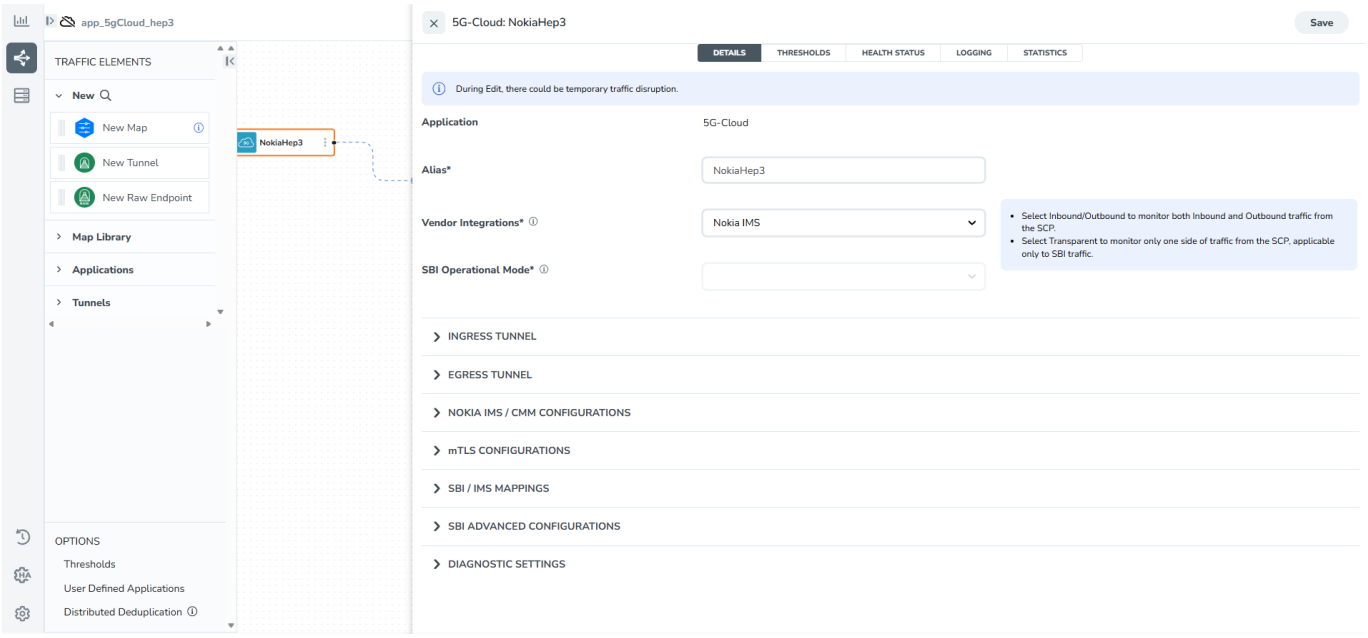
2. **Create an ingress REP to receive the data over TCP**

- On the canvas, create an ingress raw endpoint (REP) on the V Series node to receive HEP3/TCP from the IMS vTAPs.
- Configure the listening IP address and TCP port that the vTAPs use as the HEP3 collector.
- Refer to Create Raw Endpoint (VMware vCenter), Create Raw Endpoint (OpenStack), or Create Raw Endpoint (Third Party Orchestration) for platform-specific steps.

3. **Configure Nokia HEP3 in 5G-Cloud Application**

To add a 5G-Cloud application:

1. In the canvas, drag and drop the 5G-Cloud application and select **Details**. The 5G-Cloud quick view appears.



2. On the application quick view, enter or select the required information as described in the Reference - Configuration Table .

NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

Reference - Configuration Table

Field	Description
Application	The name **5g-Cloud** appears by default.
Alias	Enter the required alias name. (for example: NokiaHep3)
Vendor Integrations	From the drop-down list, select **Nokia IMS**.
SBI Operational Mode	The Operational Mode is grayed out when you select Nokia CMM or Nokia IMS.
Ingress Tunnel	
Type	Specify the tunnel type. The default is TCP.
Listening IP	Specify the tunnel's local listen IP address to receive packets.
Listening Port	Specify the tunnel's local listening port to bind for receiving packets. The application listens to the traffic on the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port to send packets. Enter a value between 1 and 65535.
Egress Tunnel	
Type	Specify the tunnel type. Available tunnel types are: • VXLAN

	• L2GRE
Tool IP	Specify the remote IP address to which you want to send the packets.
Destination Port	Specify the tunnel destination port to send the packets. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending packets. Enter a value between 1 and 65535.
Tool MTU	Specify the tool port MTU: • For V Series Node version 6.8.00 and above, enter a value between 1400 and 8800. The default value is 8800. • For V Series Node version below 6.8.00, enter a value between 1500 and 8800. The default value is 8800.
VNI ID **(Applicable only when the selected tunnel type is VXLAN**)	Specify the ID for VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key **(Applicable only when the selected tunnel type is L2GRE**)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Nokia IMS / CMM Configuration	
Number of Ingress TCP Connections	Specify the number of concurrent TCP connections VTAP can establish. Enter a value between 128 and 2048. The default value is 1024.
Ingress TCP Timeout	Specify the timeout value for an Ingress TCP connection in seconds. Enter a value between 30 and 3600 seconds. The default is 60 seconds. **NOTE:** If no packets are received within the configured time, the TCP connection will terminate due to a timeout.
Number of Egress TCP Flows	Specify the total number of TCP flows allocated for tracking non-SBI traffic. • For V Series Node (extra-large form factor, total memory ≥ 32 GB): 1024 – 16000000 • For V Series Node (large form factor, total memory < 32 GB): 1024 – 16384. • Default value: 4096.
Egress TCP Flow Timeout	Specify the timeout value for an Egress TCP flow in seconds. Enter a value between 30 and 7200 seconds. The default is 900 seconds. **NOTE:** If no packets are received within the configured time, the TCP flow will terminate due to a timeout.
Number of Receiver Threads	Specify the number of receiver threads for processing incoming packets. Enter a value between 1 and 128. The default value is 8.
Number of Egress SCTP Flows	Specify the total number of SCTP flows allocated to track non-SBI traffic. Enter a value between 128 and 2000000. The default value is 1024.

Egress SCTP Flow Timeout	Specifies the timeout value for an Egress SCTP flow in seconds. Enter a value between 30 and 7200. The default value is 900 seconds. **NOTE:** If no packets are received within the configured time, the SCTP flow will terminate due to a due to timeout.
mTLS Configuration	
mTLS	Use the toggle button to enable or disable mTLS encryption and decryption in Monitor Mode.
mTLS key	• If mTLS encryption is enabled, select the mTLS key alias for the SSL certificate from the drop-down list. • To create a key alias, click **Add** next to the text box. For further details, refer to Configure SSL Decrypt.
SBI / IMS Mappings Refer to Add CSV file for IP Mapping section for more details.	
FQDN Mapping	Use this option to provide mappings between Fully Qualified Domain Names and IP addresses for all Network Functions in the network. • **Select an alias**: Choose the alias name associated with the uploaded FQDN table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
Network Function Instance Mapping (Optional)	Use this option to provide mappings between NF Instance IDs and IP addresses for all Network Functions in the network. • **Select an alias**: Choose the alias name associated with the uploaded NFID table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
User Agent Mapping (Optional)	Use this option to provide mappings between User Agents and IP addresses for all Network Functions in the network. • **Select an alias**: Choose the alias name associated with the uploaded User Agent table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
Service Mapping (Optional)	Use this option to map ephemeral attribute values to standard service attributes and define service-specific actions. • **Select an alias**: Choose the alias name associated with the uploaded Service Mapping table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one. • **None**: Select None to clear any previously selected Service Map entry. Refer to Reference - Service Map Field Definitions section for example details.
SBI Advanced Configuration **NOTE:** The below configurations apply only to SBI traffic and will not be used for non-SBI traffic.	
TCP Server Ports	Specify the TCP server port or port range to enable TCP communication between endpoints. • Valid port range: 1-65535. • Maximum number of ports: 1024.

	NOTE: TCP communication is allowed only on the configured ports.
Number of TCP Flows	Specify the total number of TCP Flows allocated for tracking active TCP Flows. • Valid range: 128-2048 • Default value: 1024
Number of Transaction Flows	Specify the total number of Transaction Stream flows allocated for tracking active Transaction. Refer to Recommend Form Factor for VMware vCenter (Instance Types) for more details. **Streams:** • Large Form Factor VMs (total memory < 32GB): 128-5500. • Extra Large Form Factor VMs (total memory ≥ 32GB): 128-26000 • Default value: 2048
TCP Flow Timeout	Specify the TCP flow timeout value. • Valid range: 0-7200 • Default value: 900
SCP Transaction Timeout	Specify the Transaction timeout value. • Valid range: 1-300 • Default value: 10
Minimum TCP Client Port	Specify the starting port number for Synthesize Flow. • Valid range: 1023-65535 • Default value: 32768
Maximum TCP Client Port	Specify the end port number for the Synthesize Flow. • Valid range: 1023-65535 • Default value: 36863
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. • Valid range: 1-16 • Default value: 8
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. • Valid range: 100-8000 • Default value: 1000
Header Index	Enable or disable the header indexing. The default option is disabled.
Header Compression Code	Enable or disable the header compression code. The default option is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default option is Enable.
Diagnostic Settings	
Log Directory	Specify the path to store the log files.
Application Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information

	• Debug
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
Packet Capture Timestamp	Select the required timestamp used to save the Egress PCAPs. The available timestamps are: • Observed Timestamp • System Timestamp **NOTE:** Timestamp option is disabled if the Packet capture level is set to None.
Message Trace Log Level	Select the 5G-Cloud CSV log level from the drop-down list. • All • Flow • Message • Transaction • None (Default Value)

1. **Create links between the ingress REP, 5G-Cloud, and egress REP**

- On the canvas, draw a link from the ingress raw endpoint to the 5G-Cloud application.
- Create an egress raw endpoint (REP) for the target probe or GigaVUE HC/TA device (using the same raw endpoint procedures as in Step 2).
- Draw a link from the 5G-Cloud application to the egress raw endpoint.

2. **Deploy the Monitoring Session**

- After adding all the required elements to the canvas, to deploy the session do the following:
- From the **Actions** menu, select **Deploy**.

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

3. **View the Deployment Report**

   - You can view the Monitoring Session Deployment Report in the **SOURCES** and **V SERIES NODES** tab.
   - When you select the **Status** link, the Deployment Report is displayed.
   - When the deployment is incorrect, the Status column displays one of the following errors:

     a. **Success:** Not deployed on one or more instances due to V Series Node failure.
     b. **Failure:** Not deployed on all V Series Nodes or Instances.

4. **Download Logs - 5G-Cloud application**

   You can view the log files of a V Series Node or download them as .CSV or .txt files.

   To download the log files to a local environment:

   - Go to **Traffic** > **VIRTUAL** > select your cloud platform.
   - Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g** application, click the ⋮ menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
   - Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
   - Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field.
3. From the **Type** drop-down list, select one of the following. You can also choose to download the template using the **Download Template** option and add the required entries.
   - **5G-Cloud FQDN** - Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
FQDNID	abc.xyz.com
IngressIP1	1.1.1.1
IngressIP2	2.2.2.2
IngressPort	100
EgressIP	3.3.3.3
NFType	SMF
NFLocation	US
IMSIP	1001:0008:85a2:0000:0000:8f1e:0340:7110
IMSIntfName	100 200
IMSMask	96
IMSType	nef

Notes:

- The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.
- The fields imsipaddress, imsinterfacename, imsmask and imstype are applicable only when the Integrated solution Vendor is "Nokia HEP3".

- **5G-Cloud NF Instance** - Add the CSV file containing a valid NF instance ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
NFID	amf.local
ingressIPAddress1	1.1.1.1
ingressIPAddress2	2.2.2.2
IngressPort	100
egressIPAddress	3.3.3.3
NFType	AMF
NFLocation	US

NOTE: The fields NFID and ingressIP1 are mandatory. All other fields are optional.

- **5G-Cloud UA** - Add the CSV file containing a valid user agent ID and a valid IPv4/IPv6 address for IP mapping.

Header Details	Example
useragentid	nef.nb.5gc.mc005.mcc580.yournetwork
ingressIPAddress	1.1.1.1
NFType	AMF
NFLocation	US

NOTE: The fields useragentID and ingressIPAddress are mandatory. All other fields are optional.

- **5G-Cloud Service Map** - Add the CSV file containing valid ephemeral port details and a valid IPv4/IPv6 address for IP mapping. Refer to Reference - Service Map Field Definitions for header details and examples.

4. Click **Choose File** in the **File Name** field to upload the CSV file into GigaVUE-FM.
5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

Recommended VM Specifications

The configuration of the 5G-Cloud Nokia HEP3 application should be performed on extra-large and large VMs.

- Configuration specifications for Extra Large VM:
  - vCPUs: 26
  - Memory: 32 GB
  - Disk Space: 200 GB

NOTE: Nokia HEP3 configuration on extra-large VMs is supported only through the Third-Party Orchestration method.

- Configuration specifications for Large VM:
  - vCPUs: 8
  - Memory: 16 GB
  - Disk Space: 80 GB
- Supported Ports: 1 Management interface, 1 ingress, and 1 egress interfaces.
- It is recommended to use SRIOV enabled drivers for better performance and higher throughput. Refer to Configure GigaVUE Fabric Components using VMware ESXi topic for more details.

Reference - Service Map Field Definitions

Service mapping normalizes SIP traffic by converting vendor-specific or ephemeral ports to standard SIP service ports (such as 5060). GigaVUE-FM matches incoming SIP control messages against configurable rules and, when a rule applies, rewrites the source and/or destination port before forwarding the traffic. This ensures that monitoring tools that rely on standard SIP ports continue to work without any changes.

Refer below for field specific details:

NOTE: All fields except sipMessageType, srcPortMap, and dstPortMap are mandatory.

1. **srcIP** - Source IP address of the SIP traffic, as carried in HEP3. You can specify a single IPv4 / IPv6 address, a prefix, or ipany to match any address.

   Example:
   - ipany – matches SIP traffic from any source IP.
   - 1.1.1.1 – matches only traffic coming from 1.1.1.1.
2. **dstIP** - Destination IP address of the SIP traffic. You can specify a single IPv4 / IPv6 address, a prefix, or ipany to match any address.

   Example:
   - ipany – any destination IP.
   - 2.2.2.2 – match traffic going to 2.2.2.2.
3. **srcPort** - Source port number of the SIP traffic. You can use a specific port to target a known source, or portany to match all source ports.

   Example:
   - portany – match SIP traffic regardless of the source port.
   - 5060 – match only when the source port is 5060.
4. **dstPort** - Destination port number of the SIP traffic.

   Example:
   - portany – any destination port.
   - 5060 – only traffic with destination port 5060.
5. **protocolFamily** - Identifies the IP family of the traffic.

   Example:
   - IPv4 SIP traffic.
   - IPv6 SIP traffic.
6. **transportProtocol** - Identifies the L4 transport protocol (TCP or UDP).

   NOTE: This field accepts only tcp or udp; it does not support integer values.

7. **applicationProtocol** - Encodes the application protocol type carried in the HEP3 encapsulation.

   Example:
   - 1 – match SIP messages.

8. **interfaceName** - Identifies the Nokia interface number on which the traffic is received.

   Example:
   - 90 – match traffic received on Nokia HEP3 vTAP interface 90 (IMS Non-SBI).

9. **Direction** - Traffic direction from the node's point of view.

   Example:
   - ingress – apply this rule only to traffic arriving at the node.
   - egress – apply this rule only to traffic sent out from the node.

10. **sipMessageType** - Tells whether the SIP message is a request or a response.

    Example:
    - request – map ports only for SIP requests.
    - response – map ports only for SIP responses.

11. **action** - Specifies what the 5G Cloud application does with matching traffic.

    Example:
    - portmap – change source and/or destination ports as specified in srcPortMap / dstPortMap.
    - discard – drop matching traffic.
    - ignore – match but do not modify the ports.

12. **srcPortMap** - Applies only to server response frames. Use a valid port number in response frames (sipMessageType = response) and set it to portany for request frames.

    Example:
    - 5060 – set the source port to 5060 for matching SIP responses.
    - portany – keep the original source port.

13. **dstPortMap** - Applies onto client-to-server request frames. Use a valid port number in request frames (sipMessageType = request) and set it to portany for response frames.

    Example:
    - 7777 – set the destination port to 7777 for matching SIP requests.
    - portany – keep the original destination port.

FHA Dashboards for 5G-Cloud Applications

After configuring the 5G-Cloud application, you can monitor the statistics for Nokia HEP3 by the reports displayed in the Dashboard. To access the details, refer to FHA Dashboards for 5G-Cloud Applications.

5G Cloud Nokia HEP3 CMM Solution (SBI and Non-SBI)

Nokia 5G network functions use the HEP3 binary format to encode and send packets. HEP3 supports both HTTP2 and TCP protocols, enabling efficient traffic transmission across core network functions. This solution supports traffic processing from Nokia Cloud Mobility Manager (CMM), and the Service Communication Proxy (SCP).

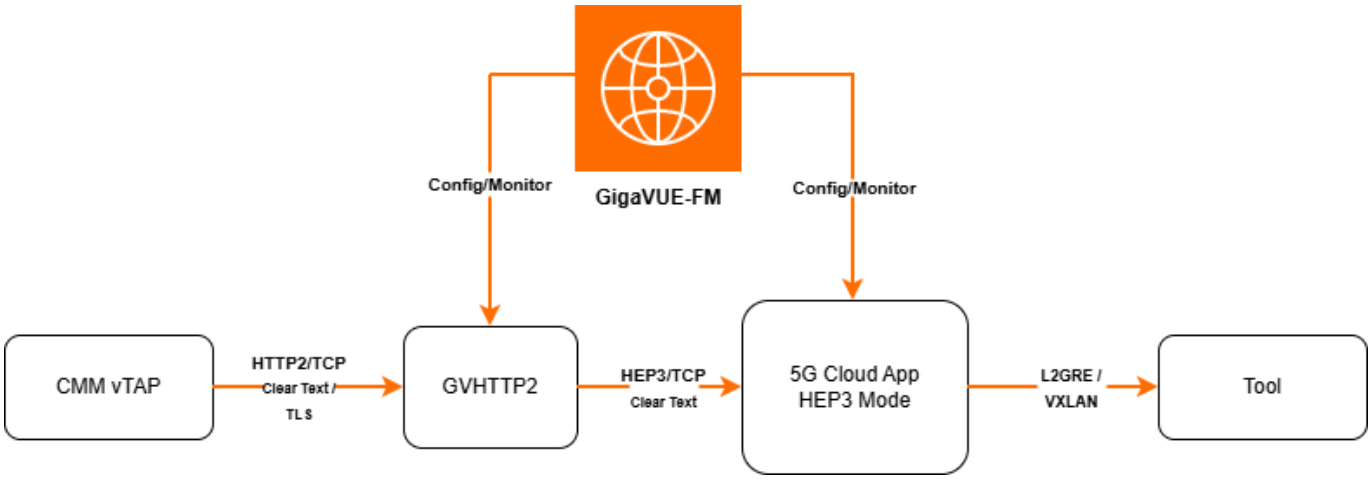
Cloud Mobility Manager (CMM) includes AMF and MME network functions. It generates both Service-Based Interface (SBI) and non-SBI traffic and sends it over HTTP2. The gvHTTP2 application receives CMM traffic and forwards it to the 5G Cloud application, which synthesizes SBI and non-SBI payloads into L2GRE/VXLAN packets before sending them to probes. This approach provides extended visibility across 5G core interfaces and enables operators to monitor both SBI and non-SBI traffic seamlessly.

How Nokia HEP3 CMM Solution works

Nokia vTAP systems in the mobility core send HEP3 data over both HTTP2 and TCP.

Nokia CMM Solution

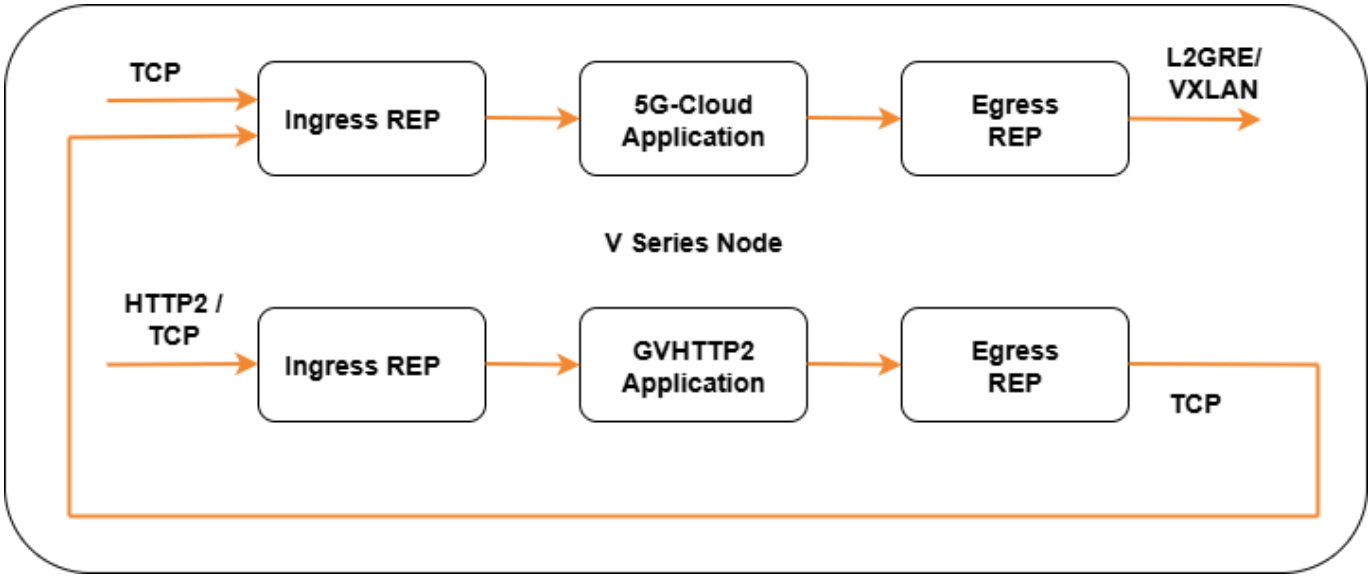
The following image shows the ingress and egress data flow through the V Series node for the Nokia HEP3 CMM solution.



The Cloud Mobility Manager (CMM) function includes key control-plane elements such as AMF and MME. CMM generates both SBI and non-SBI traffic, making it central to mobility management in the 5G Core.

- CMM vTAPs export SBI and non-SBI control-plane traffic as HEP3 over HTTP2/TCP, in either clear text or TLS-protected sessions.

- The gvHTTP2 application on the V Series node listens on the configured HTTP2 IP address and port, terminates HTTP2 sessions from CMM, and extracts the encapsulated HEP3 payloads.
- gvHTTP2 sends the processed HEP3 payloads to the 5G-Cloud application over TCP by using a dedicated Nokia HEP3 over TCP tunnel between the gvHTTP2 and 5G-Cloud V Series nodes.
- 5G-Cloud receives HEP3/TCP from gvHTTP2, terminates the TCP connections, decodes the HEP3 payloads, and synthesizes SBI and non-SBI flows into a format that monitoring tools can consume without IP address translation.
- The V Series node encapsulates the synthesized traffic using L2GRE or VXLAN and then forwards it to:
  - Physical probe, through a GigaVUE HC or GigaVUE TA Series device that aggregates and decapsulates the traffic, or
  - Virtual probe, directly over IP when connectivity is available.



Configure Nokia CMM HEP3 Solution (SBI and Non-SBI)

In GigaVUE-FM, you must do the following to add the 5G-Cloud CMM integration to either a new or existing monitoring session of a Monitoring Domain:

Pre-requisite:

- You must upload CSV files containing a valid FQDN and a valid IPv4/IPv6 address. Refer to Add CSV file for IP Mapping.

1. **Create New or Use Exisiting Monitoring Session**

   In GigaVUE-FM, go to the monitoring domain that will carry IMS HEP3 traffic. You can add the 5G-Cloud aplication, to a new or existing monitoring session of a Monitoring Domain.

   - **New Monitoring Session** - Add the 5G-Cloud application after creating a new Monitoring Session and when the GigaVUE-FM canvas appears. Refer to Create a Monitoring Session section in the respective GigaVUE Cloud Suite Deployment Guide.
   - **Existing Monitioring Session** - Select any existing Monitoring Session and go to **TRAFFIC PROCESSING** tab. The GigaVUE-FM canvas appears.

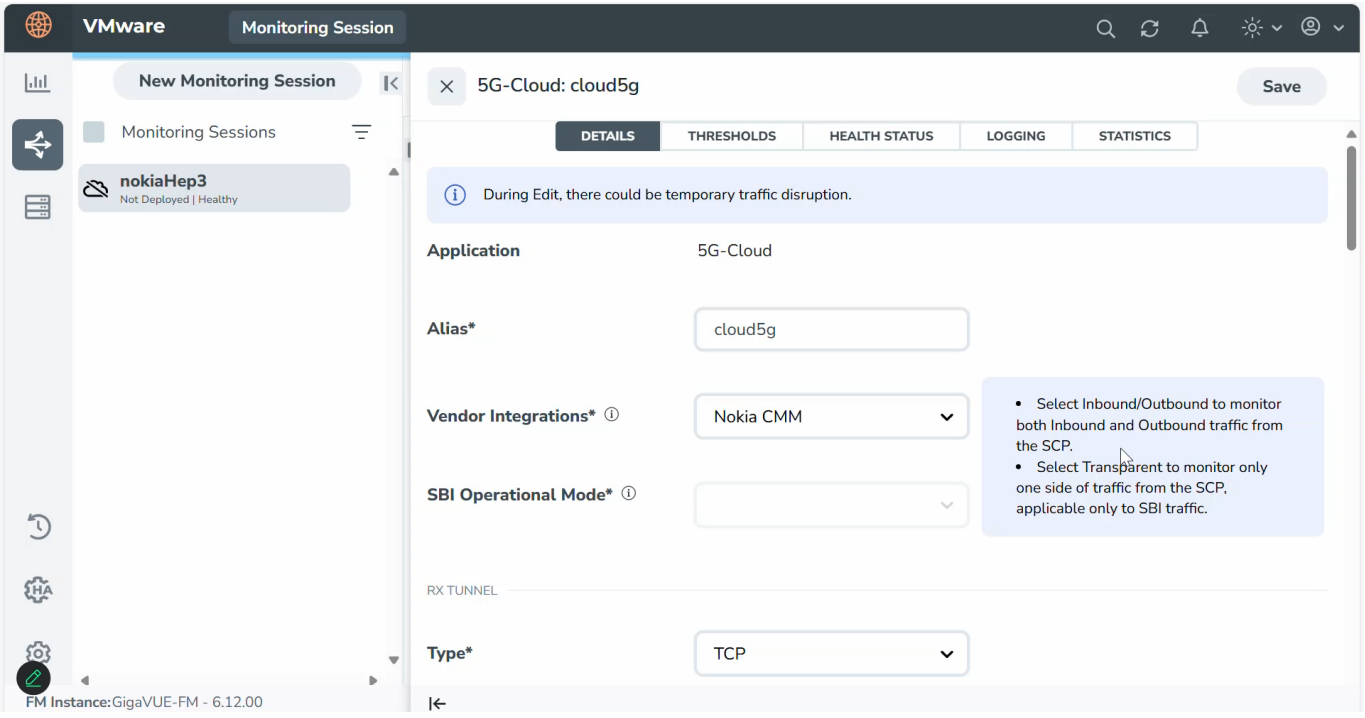
2. **Add gvHTTP2 and connect it to CMM vTAPs.**

   - On the canvas, drag and drop the gvHTTP2 application onto the V Series node.
   - Select the gvHTTP2 node and click Details to open the quick view.
   - Configuration of GvHTTP2 Application for CMM vTAP

Field	Description
Application	The name **gvhttp2** appears by default.
Alias	Enter the alias name as **gvhttp2**.
Vendor Integrations	From the drop-down list, select **Nokia CMM**.
HTTP2 Listening IP	Specify the HTTP2 IP address. The application will listen to the traffic on the specified IP address.
HTTP2 Listening Port	Specify the HTTP2 Port. The application will listen to the traffic on the specified port.
Max Stream	Specify the maximum number of concurrent streams in one HTTP2 session. Enter a value between 1 and 100. The default value is 100.
Worker Thread	Specify the number of parallel threads that can be sued to process the requests. Enter a value between 1 and 16. The default value is 1.
TX Tunnel	
Type	Specify the tunnel type as Nokia HEP3 over TCP from the dropdown for Nokia CMM.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending the packet. Enter a value between 1 and 65535.
Destination IP	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Destination Port	Specify the tunnel destination port to which the packet will be sent. Enter a value between 1 and 65535.
Number of Egress TCP Connections	Specify the number of concurrent TCP connections GvHTTP2 can establish. Enter the value from 1 to 200. The default value is 4.
TLS Config	
TLS	Enable or disable TLS encryption or decryption using the toggle button.
Server Private key	Enter the path where the private key is stored.
Server Certificate	Enter the path where the certificate is stored.
Advanced Setting	
CSV	Enable this option to capture all stats counters for GvHTTP2.

PCAP	Enable this option to capture the packets at egress.
Log Directory	Specify the path to store the log files.
Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug

3. Add the 5G-Cloud application and connect it to gvHTTP2
   - On the canvas, drag and drop the 5G-Cloud application onto the V Series node.
   - Place 5G-Cloud downstream of gvHTTP2.
   - Draw a link from gvHTTP2 to 5G-Cloud so that gvHTTP2 forwards HEP3/TCP to 5G-Cloud.
   - Select the 5G-Cloud node and click Details to open the quick view.



NOTE: It is recommended to maintain a 60-90 second delay when undeploying and deploying a Monitoring Session in GigaVUE-FM.

4. Configuration of 5G-Cloud Application for CMM vTAP

Field	Description
Application	The name **5g-Cloud** appears by default.
Alias	Enter the required alias name. (for example: NokiaHep3)
Vendor Integrations	From the drop-down list, select **Nokia CMM** .
SBI Operational Mode	The Operational Mode is grayed out when you select Nokia CMM or Nokia IMS.
Ingress Tunnel	
Type	Specify the tunnel type. The default is TCP.
Listening IP	Specify the tunnel's local listen IP address to receive packets.
Listening Port	Specify the tunnel's local listening port to bind for receiving packets. The application listens to the traffic on the specified port. Enter a value between 1 and 65535.
Source Port	Specify the tunnel destination port to send packets. Enter a value between 1 and 65535.
Egress Tunnel	
Type	Specify the tunnel type. Available tunnel types are: • VXLAN • L2GRE
Tool IP	Specify the remote IP address to which you want to send the packets.
Destination Port	Specify the tunnel destination port to send the packets. Enter a value between 1 and 65535.
Source IP	Specify the source IP address to use when sending the packet.
Source Port	Specify the tunnel source port to bind when sending packets. Enter a value between 1 and 65535.
Tool MTU	Specify the tool port MTU: • For V Series Node version 6.8.00 and above, enter a value between 1400 and 8800. The default value is 8800. • For V Series Node version below 6.8.00, enter a value between 1500 and 8800. The default value is 8800.
VNI ID **(Applicable only when the selected tunnel type is VXLAN**)	Specify the ID for VXLAN traffic. Enter a value between 0 and 16777215.
L2GRE Key **(Applicable only when the selected tunnel type is L2GRE**)	Specify the key for the L2GRE tunnel type. Enter a value between 0 and 4294967295.
Nokia IMS / CMM Configuration	
Number of Ingress TCP Connections	Specify the number of concurrent TCP connections VTAP can establish. Enter a value between 128 and 2048. The default value is 1024.
Ingress TCP Timeout	Specify the timeout value for an Ingress TCP connection in seconds. Enter a value between 30 and 3600 seconds. The default is 60 seconds.

	NOTE: If no packets are received within the configured time, the TCP connection will terminate due to a timeout.
Number of Egress TCP Flows	Specify the total number of TCP flows allocated for tracking non-SBI traffic. • For V Series Node (extra-large form factor, total memory ≥ 32 GB): 1024 – 16000000 • For V Series Node (large form factor, total memory < 32 GB): 1024 – 16384. • Default value: 4096.
Egress TCP Flow Timeout	Specify the timeout value for an Egress TCP flow in seconds. Enter a value between 30 and 7200 seconds. The default is 900 seconds. **NOTE:** If no packets are received within the configured time, the TCP flow will terminate due to a timeout.
Number of Receiver Threads	Specify the number of receiver threads for processing incoming packets. Enter a value between 1 and 128. The default value is 8.
Number of Egress SCTP Flows	Specify the total number of SCTP flows allocated to track non-SBI traffic. Enter a value between 128 and 2000000. The default value is 1024.
Egress SCTP Flow Timeout	Specifies the timeout value for an Egress SCTP flow in seconds. Enter a value between 30 and 7200. The default value is 900 seconds. **NOTE:** If no packets are received within the configured time, the SCTP flow will terminate due to a due to timeout.
mTLS Configuration	
mTLS	Use the toggle button to enable or disable mTLS encryption and decryption in Monitor Mode.
mTLS key	• If mTLS encryption is enabled, select the mTLS key alias for the SSL certificate from the drop-down list. • To create a key alias, click **Add** next to the text box. For further details, refer to Configure SSL Decrypt.
SBI / IMS Mappings Refer to Add CSV file for IP Mapping section for more details.	
FQDN Mapping	Use this option to provide mappings between Fully Qualified Domain Names and IP addresses for all Network Functions in the network. • **Select an alias**: Choose the alias name associated with the uploaded FQDN table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
Network Function Instance Mapping (Optional)	Use this option to provide mappings between NF Instance IDs and IP addresses for all Network Functions in the network. • **Select an alias**: Choose the alias name associated with the uploaded NFID table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
User Agent Mapping (Optional)	Use this option to provide mappings between User Agents and IP addresses for all Network Functions in the network.

	• **Select an alias**: Choose the alias name associated with the uploaded User Agent table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one.
Service Mapping (Optional)	Use this option to map ephemeral attribute values to standard service attributes and define service-specific actions. • **Select an alias**: Choose the alias name associated with the uploaded Service Mapping table (CSV file) from the drop-down list. • **Add a new alias**: If no alias exists, click **Add** to create one. • **None**: Select None to clear any previously selected Service Map entry.
SBI Advanced Configuration **NOTE:** The below configurations apply only to SBI traffic and will not be used for non-SBI traffic.	
TCP Server Ports	Specify the TCP server port or port range to enable TCP communication between endpoints. • Valid port range: 1-65535. • Maximum number of ports: 1024. **NOTE:** TCP communication is allowed only on the configured ports.
Number of TCP Flows	Specify the total number of TCP Flows allocated for tracking active TCP Flows. • Valid range: 128-2048 • Default value: 1024
Number of Transaction Flows	Specify the total number of Transaction Stream flows allocated for tracking active Transaction. Refer to Recommend Form Factor for VMware vCenter (Instance Types) for more details. **Streams:** • Large Form Factor VMs (total memory < 32GB): 128-5500. • Extra Large Form Factor VMs (total memory ≥ 32GB): 128-26000 • Default value: 2048
TCP Flow Timeout	Specify the TCP flow timeout value. • Valid range: 0-7200 • Default value: 900
SCP Transaction Timeout	Specify the Transaction timeout value. • Valid range: 1-300 • Default value: 10
Minimum TCP Client Port	Specify the starting port number for Synthesize Flow. • Valid range: 1023-65535 • Default value: 32768
Maximum TCP Client Port	Specify the end port number for the Synthesize Flow. • Valid range: 1023-65535 • Default value: 36863
SCP Processing Threads	Specify the number of worker threads for processing In-Out packets. • Valid range: 1-16

	• Default value: 8
TCP Client Ports per Thread	Specify the number of TCP client ports for each worker thread. • Valid range: 100-8000 • Default value: 1000
Header Index	Enable or disable the header indexing. The default option is disabled.
Header Compression Code	Enable or disable the header compression code. The default option is disable.
Gigamon Header	Enable or disable the Gigamon Header. The default option is Enable.
Diagnostic Settings	
Log Directory	Specify the path to store the log files.
Application Log Level	Select the severity log level of the events from the following options: • None • Fatal • Error • Notification • Information • Debug
Packet Capture Level	Select the packet capture level from the drop-down list: • All • Receive • Transmit • None
Packet Capture Timestamp	Select the required timestamp used to save the Egress PCAPs. The available timestamps are: • Observed Timestamp • System Timestamp **NOTE:** Timestamp option is disabled if the Packet capture level is set to None.
Message Trace Log Level	Select the 5G-Cloud CSV log level from the drop-down list. • All • Flow • Message • Transaction • None (Default Value)

5. Create the egress raw endpoint and tool link

- On the canvas, create an egress raw endpoint (REP) on the V Series node that connects to your probe or GigaVUE HC/TA device.
- Draw a link from the 5G-Cloud application to the egress REP so that 5G-Cloud forwards encapsulated L2GRE/VXLAN traffic to the tool.

6. Deploy the monitoring session
    - On the Monitoring Sessions page, select the session.
    - From the Actions menu, select Deploy.
    - Wait for deployment to complete on all V Series nodes.
    - Confirm that the session Status shows Success.
    - If the status shows Failure or partial success, open the SOURCES and V SERIES NODES tabs and click the Status link to view the deployment report.

Download Logs - 5G-Cloud / GvHTTP2 Application

You can view the log files of a V Series Node or download them as .CSV or .txt files.

To download the log files to a local environment:

1. Go to **Traffic** > **VIRTUAL** > select your cloud platform.
2. Select the required Monitoring Session and go to **TRAFFIC PROCESSING** tab. On the **cloud5g** / **gvhttp2**application, click the ⋮ menu button and select **Details**. Go to **LOGGING** in the quick view. The **Logging** page displays the logs currently available.
3. Select the required Days, Timestamps, File Name, and Type (TextLog and FlowStats) details.
4. Select the log files to download, and then click **Download > Files**. The system downloads the selected files to your local environment.

Add CSV file for IP Mapping

To add the CSV file for IP mapping:

1. Go to **Inventory > VIRTUAL >** select your cloud platform, and then click **Settings > 5G-Apps**. The 5G-Apps Configuration page appears.
2. Click **New**. Enter the name for the CSV file in the **Alias** field
3. From the Type drop-down list, select **5G-Cloud FQDN**. You can also choose to download the template using the **Download Template** option and add the required entries.
4. Click **Choose File** in the **FileName** field to upload the CSV file into GigaVUE-FM. Add the CSV file containing a valid FQDN ID and a valid IPv4/IPv6 address for IP mapping.

Notes - Header details of the CSV file:

FQDN Mapping: gigamon@vseries:/var/log/cloud5g_tabledir$ cat FQDN.txt

- Header details: FQDNid,IngressIP1,IngressIP2,IngressPort,EgressIP,NFType,NFLocation,IMSIP,IMSIntfName,IMSMask,IMSType
- Example: s25scp01.scp.5gc.mnc003.mcc525.3gppnetwork.org,170.00.13.187,,8080,,SCP,plolp,2000:4000:124:a0d9:ac1:4::,92,303 314 315,sbc

The fields fqdnID and ingressIPAddress1 are mandatory. All other fields are optional.

The fields imsipaddress, imsinterfacename, imsmask and imstype are applicable only when the Integrated solution Vendor is "Nokia HEP3".

5. Click **Validate** to validate the CSV file.
6. Click **Save** to add the CSV file.

Appendix - Casa vTap Statistics

This section provides the stats details for GVHTTP2 and 5G-Cloud applications. For more information on counter details, refer to the respective application on the GigaVUE API Reference page.

NOTE: Currently, the following stats are not supported in GigaVUE-FM. You should login to the corresponding V Series Node to execute the commands.

GVHTTP2 Application

The following table provides information on the counter details and their usage for the GVHTTP2 application:

S.No	Counter Name	Counter Description
1	http2TransactionsSuccessful	Total HTTP2 Transaction successes that occurred in the GVHTTP2 Server.
2	http2StreamInitializeCount	Total HTTP2 New Streams that have been observed by the GVHTTP2 Server.
3	http2StreamCloseCount	Total HTTP2 Streams that have been successfully completed by the GVHTTP2 Server.
4	http2StreamPrepareResponseCount	Total calls to ask GVHTTP2 to prepare a response for inbound stream transactions.
5	http2StreamOnHeaderCount	Total calls to GVHTTP2 when a header value for a given stream has been observed.
6	http2StreamOnDataChunkCount	Total calls to GVHTTP2 when a data chunk for a given stream has been observed.
7	http2StreamOnRecvFrameCount	Total calls to GVHTTP2 when Receive Frame has

S.No	Counter Name	Counter Description
		occurred for any HTTP2 frame.
8	http2StreamOnRecvFrameEndStreamCount	Total calls to GVHTTP2 when the last HTTP2 Frame has been received indicating that the end of the stream has occurred.
9	http2GETRequests	Total HTTP2 GET Requests received by the GVHTTP2 Server.
10	http2POSTRequests	Total HTTP2 POST Requests received by the GVHTTP2 Server.
11	http2PUTRequests	Total HTTP2 PUT Requests received by the GVHTTP2 Server.
12	http2OtherRequests	Total HTTP2 Requests not GET, POST, or PUT received by the GVHTTP2 Server.
13	http2ValidURI	Total HTTP2 Requests that have valid PATH for the current GVHTTP2 Mode.
14	http2NewSessions	Total TCP Sessions newly initiated by Client applications connecting to the GVHTTP2 Server.
15	http2TerminateSessions	Total TCP Sessions that have been terminated by the Client applications connecting to the GVHTTP2 Server.
16	http2OpenSessions	Total TCP Open / Accept Sessions received by Client Applications connecting to the GVHTTP2 Server.
17	http2CloseSessions	Total TCP Sessions that have closed either by Client Applications or GVHTTP2 Server.
18	vxlanTXPackets	Total number of packets that have been successfully transmitted to the VXLAN TX Tunnel.
19	vxlanTXBytes	Total number of bytes that have been successfully transmitted to the VXLAN TX Tunnel.
20	vxlanThreadCount	Total number of thread cycles that occurred during the interval.
21	http2TransactionsFailed	Total HTTP2 Transaction failures that occurred in the GVHTTP2 Server.
22	http2StreamTimeouts	Total times when the start of a new HTTP2 Stream has occurred but is not completed within the predefined timeout.
23	http2StreamHeaderOverflows	Total events where header has been received, but there is no space in the GVHTTP2 server code to accept the Header Value.
24	http2StreamBufferOverflows	Total events where data frames have been received, but there is no space in the GVHTTP2 server code to accept the data frame.
25	http2InvalidURI	Total HTTP2 Requests that have invalid PATH for

S.No	Counter Name	Counter Description
		the current GVHTTP2 Mode.
26	http2SSLConnectionFailures	Total SSL Connection Failures that occur when the GVHTTP2 Server receives a TCP Connection but fails SSL negotiation.
27	http2SocketFailures	Total Socket Failures either in mTLS or Clear Text mode.
28	vxlanTXSocketSendErrors	Total number of packets that could not be transmitted to the VXLAN TX Tunnel because of socket error.
29	vxlanTXSocketSizeErrors	Total number of packets that could not be transmitted to the VXLAN TX Tunnel because the size of the packet is too large.
30	vxlanThreadTimeouts	Total number of calls to select that timed out waiting for received data. This is not an error.
31	vxlanRXPackets	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been created in the VXLAN thread.
32	vxlanRXBytes	Total number of RX Bytes received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been created in the VXLAN thread.
33	vxlanRXDrops	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been dropped by the VXLAN thread because there are no more RX descriptors in the ingress queue of the VXLAN thread.
34	vxlanRXOversized	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that the VXLAN thread has dropped because the RX message size is too large.
35	vxlanQueueCurrBufferIn	Buffer index of the queue slot for a current packet to be inserted into the VXLAN Thread RX Queue.
36	vxlanQueueCurrBufferOut	Buffer index of the queue slot for a current packet to be removed from the VXLAN Thread RX Queue.
37	vxlanQueueNumBuffers	Total number of buffers in the VXLAN RX Queue.
38	vxlanQueueMaxQueueSize	Total number of packets that were in the VXLAN RX Queue during the current interval.
39	vxlanQueueBufferSize	Size of each RX Packet buffer in the VXLAN RX Queue.
40	vxlanQueuePacketsQueuedIn	Total number of packets that have been put into the VXLAN RX Queue.

S.No	Counter Name	Counter Description
41	vxlanQueuePacketsQueuedOut	Total number of packets that have been removed from the VXLAN RX Queue.
42	vxlanQueuePacketsQueuedFull	Total number of packets that attempted to be queued when the VXLAN RX Queue was full.
43	nghttp2Errors	Total errors that occur within the nghttp2 library used by the GVHTTP2 server application.
44	pcapRXPackets	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been created in the PCAP thread.
45	pcapRXBytes	Total number of RX Bytes received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been created in the PCAP thread.
46	pcapRXDrops	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that have been dropped by the PCAP thread because there are no more RX descriptors in the ingress queue of the PCAP thread.
47	pcapRXOverSized	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that the PCAP thread has dropped because the RX message size is too large.
48	pcapRXErrors	Total number of RX Packets received from inbound HTTP2 Stream transactions received by the GVHTTP2 server from the Client VTAP that the PCAP thread has dropped because of an error.
49	pcapQueueCurrBufferIn	Buffer index of the queue slot for a current packet to be inserted into the PCAP Thread RX Queue.
50	pcapQueueCurrBufferOut	Buffer index of the queue slot for a current packet to be removed from the PCAP Thread RX Queue.
51	pcapQueueNumBuffers	Total number of buffers in the PCAP RX Queue.
52	pcapQueueMaxQueueSize	Total number of packets that were in the PCAP RX Queue during the current interval.
53	pcapQueueBufferSize	Size of each RX Packet buffer in the PCAP RX Queue.
54	pcapQueuePacketsQueuedIn	Total number of packets that have been put into the PCAP RX Queue.
55	pcapQueuePacketsQueuedOut	Total number of packets that have been removed from the PCAP RX Queue.
56	pcapQueuePacketsQueuedFull	Total number of packets that attempted to be queued when the PCAP RX Queue was full.

S.No	Counter Name	Counter Description
57	pcapFiles	Total number of PCAP Files written during this interval.
58	pcapOpenErrors	Total number of file errors that occurred while trying to create new PCAP Files during this interval.
59	pcapWriteErrors	Total number of file errors that occurred while trying to write data to PCAP Files during this interval.
60	pcapFlushErrors	Total number of file errors that occurred while trying to flush cached data to PCAP Files during this interval.
61	pcapCloseErrors	Total number of file errors that occurred while trying to close PCAP Files during this interval.
62	pcapPackets	Total number of PCAP Packets written during the interval.
63	pcapBytes	Total number of PCAP Bytes written during the interval.
64	diskLoggingThreadTimeouts	Total number of calls to delay waiting to receive data in the PCAP file thread. This is not an error.
65	diskLoggingThreadCount	Total number of thread cycles that occurred during the interval.

5G-Cloud Application

The following table provides information on the counter details and their usage for the 5G-Cloud application:

S.No	Counter Name	Counter Description
1	rxPkts	Total number of RX Packets / Messages from the GVHTTP2 Application.
2	txPkts	Total number of packets that have been successfully transmitted to the UDPGRE TX Tunnel.
3	rxPktBytes	Total number of RX Bytes for messages from the GVHTTP2 Application.
4	txPktBytes	Total number of bytes that have been successfully transmitted to the UDPGRE TX Tunnel.
5	rxDrops	Total number of RX Packets received from inbound VXLAN Tunnel by the 5G-Cloud Application from the GVHTTP2 Server APP that have been dropped by the PCAP thread because

S.No	Counter Name	Counter Description
		there are no more RX descriptors in the ingress queue of the PCAP thread.
6	txDrops	Total number of Tx Packets that have been dropped by the 5G-Cloud APP because there are no more RX descriptors in the egress queue.
7	rxThreadCount	Total number of receive thread cycles that occurred during the interval.
8	txThreadCount	Total number of Transmit thread cycles that occurred during the interval.
9	rxSelectTimeouts	Total number of calls to select that timed out waiting for received data. This is not an error.
10	txThreadTimeout	Total number of calls to egress Queue (Timed Semaphore wait) that timed out waiting for received data. This is not an error.
11	rxIPv4Pkts	Total number of IPv4 Packets received by the 5G-Cloud Application (this counter is not used currently).
12	rxIPv6Pkts	Total number of IPv6 Packets received by the 5G-Cloud Application (this counter is not used currently).
13	rxTcpPkts	Total number of TCP Packets received by the 5G-Cloud Application (this counter is not used currently).
14	rxNonTcpPkts	Total number of Non-TCP Packets received by the 5G-Cloud Application (this counter is not used currently).
15	rxVlanPkts	Total number of VLAN Packets received by the 5G-Cloud Application (this counter is not used currently).
16	rxPktsTooShort	Total Number of Incomplete Packets received by the 5G-Cloud Application (this counter is not used currently).
17	rxSelectNoEvents	Total Number of Errors while reading the packet data from the ingress Socket.
18	rxOversizedBuffers	Total number of RX Packets received by the 5G-Cloud Application from the GVHTTP2 Server that have been dropped because the RX message size is too large.
19	rxRecvFromNonCritical	Total number of socket timeout (timed out without Receiving any packet) for the ingress Interface.
20	rxErrors	Total number of Rx Packets dropped by the 5G-Cloud Application because of an error in Processing the Packet.

S.No	Counter Name	Counter Description
21	rxProtocolIDErrors	Total number of Rx Packets dropped by the 5G-Cloud Application because of error in IP Protocol ID (this counter is used only when L2GRE is used as ingress Tunnel Type hence it is not used currently).
22	rxIPProtocolErrors	Total number of Rx Packets dropped by the 5G-Cloud Application because of an error in IP Protocol (this counter is used only when L2GRE is used as ingress Tunnel Type; hence, it is not used currently).
23	rxIPPayloadLengthErrors	Total number of Rx Packets dropped by the 5G-Cloud Application because of an error in IP payload length (this counter is not used currently).
24	txSocketErrors	Total number of packets that could not be transmitted to the UDPGRE TX Tunnel because of socket error.
25	elPkts	Total number of Packets successfully written into the log files.
26	elFiles	Total number of files opened successfully to log the packets.
27	elErrors	Total number of errors while opening the file to write the packets.

Appendix-5G Cloud Nokia HEP3 Support (SBI)

The following table provides information on the counter details and their usage for the 5G-Cloud application:

S.No	Counter Name	Counter Description
1	sslAccept	Total TCP Open / Accept Sessions received by Client Applications connecting to the Server.
2	sslAcceptSSLPending	Total number of times SSL Connection is temporary suspension of the handshake, awaiting further action from the application's client hello callback.
3	sslAcceptSSLWantedRead	Total number of times SSL Connection needs to read more data from the network to proceed, but no data is currently available for reading.
4	sslAcceptSSLWantWrite	Total number of times there was error in writing complete data to the socket.
5	sslAcceptSSLAcceptFailure	Total number of times a new SSL connection is not accepted because of certificate, protocol or other errors.

S.No	Counter Name	Counter Description
6	sslRecvSuccess	Total number of times a new SSL connection has been accepted successfully.
7	sslRecvBytes	Total number of RX Bytes received from the vTAP in SSL Mode.
8	sslRecvSSLPending	Total number of times SSL Connection establishment is which indicates that TLS handshake routine is interrupted by client hello callback.
9	sslRecvSSLWantRead	Total number of times read operation is retried when data is available for reading.
10	sslRecvSSLWantWrite	Total number of times write operation is retried when data is available for writing.
11	sslRecvSSLReadFailures	Total number of times TLS/SSL connection has been closed cleanly by the peer while handshake is in progress.
12	sslBadCert	Total number of times there is error in establishing SSL connection with client because certificates are rejected by the peer.
13	sslCertVerifyFailure	Total number of times where client attempting to establish a secure connection (using SSL/TLS) with a server could not verify the server's SSL certificate.
14	sslUnknownProtocol	Total number of times there is error in establishing SSL connection with client because when the server cannot recognize or negotiate SSL/TLS protocol version.
15	sslHandshakeFailure	Total number of times there is error in establishing SSL connection with client because Handshake could not negotiate a cipher/protocol.
16	sslAlertUnknownCA	Total number of times there is error in establishing SSL connection with client because Certificate authority not recognized.
17	sslNoSharedCipher	Total number of times there is error in establishing SSL connection with client because no cipher overlap with peer.

Interface Mapping (OpenStack)

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

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

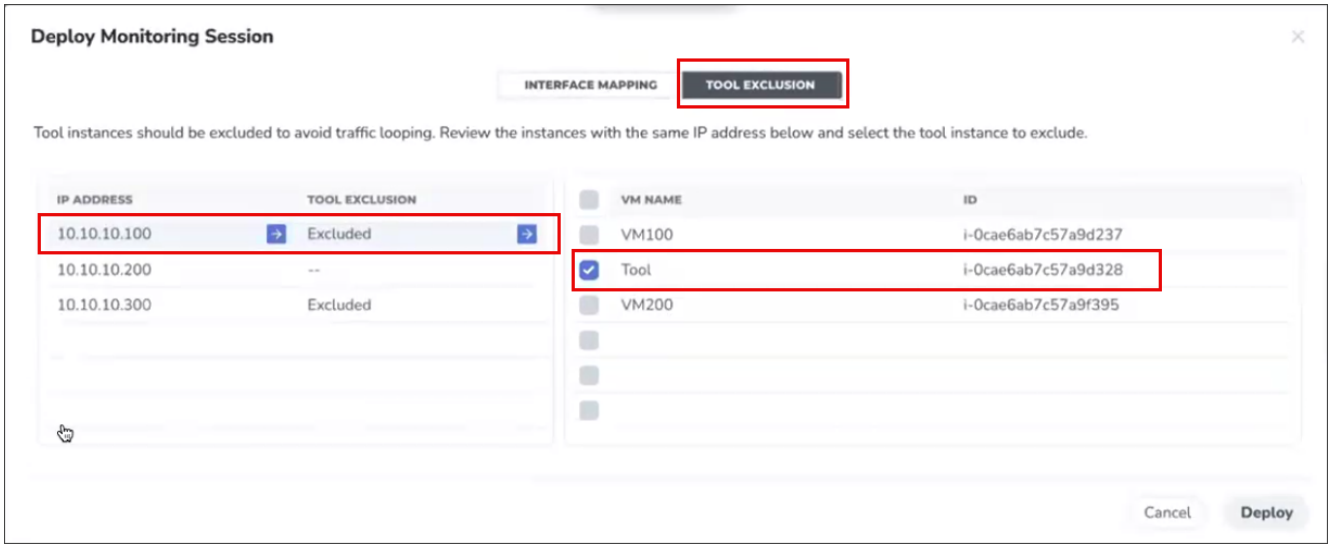
Interface Mapping

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**.



- 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 the configuration health status of each monitoring session and its components.

You can view:

- The overall health status of a monitoring session.
- The health status of individual components.

For details, refer to the following topics:

- Configuration Health Monitoring
- Traffic Health Monitoring
- View Health Status
- View Monitoring Session Statistics
- Visualize the Network Topology (OpenStack)

View Health Status

You can view the health status of the Monitoring Session on the Monitoring Session details page. The health status of the Monitoring Session is healthy only if both the configuration health and traffic health are healthy.

View Health Status of an Application

To view the health status of an application across an entire Monitoring Session,

1. Go to **Traffic > Virtual > Orchestrated Flows** and select your cloud platform.
2. Select a Monitoring Session and navigate to the **TRAFFIC PROCESSING** tab.
3. Select the application for which you wish to see the health status and select **Details**. The quick view page appears.
4. Select the **HEALTH STATUS** tab.

This displays the application's **Configuration Health**, **Traffic Health**, and the **Operational Health**, along with the thresholds applied to each.

NOTE: The secure tunnel status is refreshed every 5 minutes, and the GigaVUE-FM does not display UCT-V secure tunnel status that is older than 7 minutes. If the secure tunnel in the UCT-V is removed, it takes up to 7 minutes to reset the status on the GigaVUE-FM.

View Operational Health Status of an Application

When you configure the Application Metadata Exporter to use **Kubernetes** as the workload platform, the V Series Node transmits failure and error events to GigaVUE-FM, which processes them and updates the node's health status on the Monitoring Session page. When interacting with Kubernetes workloads, the system may encounter errors while retrieving resources such as pods, services, nodes, or endpoints. Refer to Errors for additional error details.

Operational events for Exporter:

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

Format: <Server Type>_<Message>

Server Types: CLOUD EXPORT, KAFKA

Message	Description
UPLOAD_MAX_TRIES_EXCEED	Upload retries exceeded the maximum limit **Example**: CLOUDEXPORTER_UPLOAD_MAX_TRIES_EXCEED
REACHABILITY_FROM_AMX_TO_TOOLS	AMX failed to reach the tool (Cloud Exporter server or Kafka server) **Example**: CLOUDEXPORTER_REACHABILITY_FROM_AMX_TO_TOOLS
NO_IP_ADDRESS	No IP address was configured on the interface **Example**: CLOUDEXPORTER_NO_IP_ADDRESS
EXPORTER_UPLOAD_ERROR	Upload to the exporter failed **Example**: CLOUDEXPORTER_EXPORTER_UPLOAD_ERROR
EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED	The exporter health-monitor probe or upload to the tool timed out. **Example:** CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_DEADLINE_EXCEEDED
EXPORTER_ERR_EVENT_CONTEXT_CANCELED	The exporter health-monitor probe or upload was canceled before completion, typically due to shutdown or reconfiguration. **Example:** CLOUDEXPORTER_EXPORTER_ERR_EVENT_CONTEXT_

Message	Description
	CANCELED KAFKA_EXPORTER_ERR_EVENT_CONTEXT_CANCELED
EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL	The hostname of the configured Kafka broker or Cloud Exporter endpoint could not be resolved due to a DNS failure. **Example:** CLOUDEXPORTER_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL KAFKA_EXPORTER_ERR_EVENT_DNS_LOOKUP_FAIL
EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE	The exporter probe failed at the network layer, for example, due to no route to the host or an unreachable network. **Example:** CLOUDEXPORTER_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE KAFKA_EXPORTER_ERR_EVENT_NETWORK_UNREACHABLE
EXPORTER_ERR_EVENT_TRASPORT_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_TRASPORT_LAYER_ERR KAFKA_EXPORTER_ERR_EVENT_TRASPORT_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_AUTHENTICAION_FAILURE
K8S_UNHANDLED_ERROR	An unspecified error occurred. Check the error description **Example**: GETPODS_K8S_UNHANDLED_ERROR

View Health Status for Individual GigaVUE V Series Nodes

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

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

View Application Health Status for Individual V Series Nodes

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

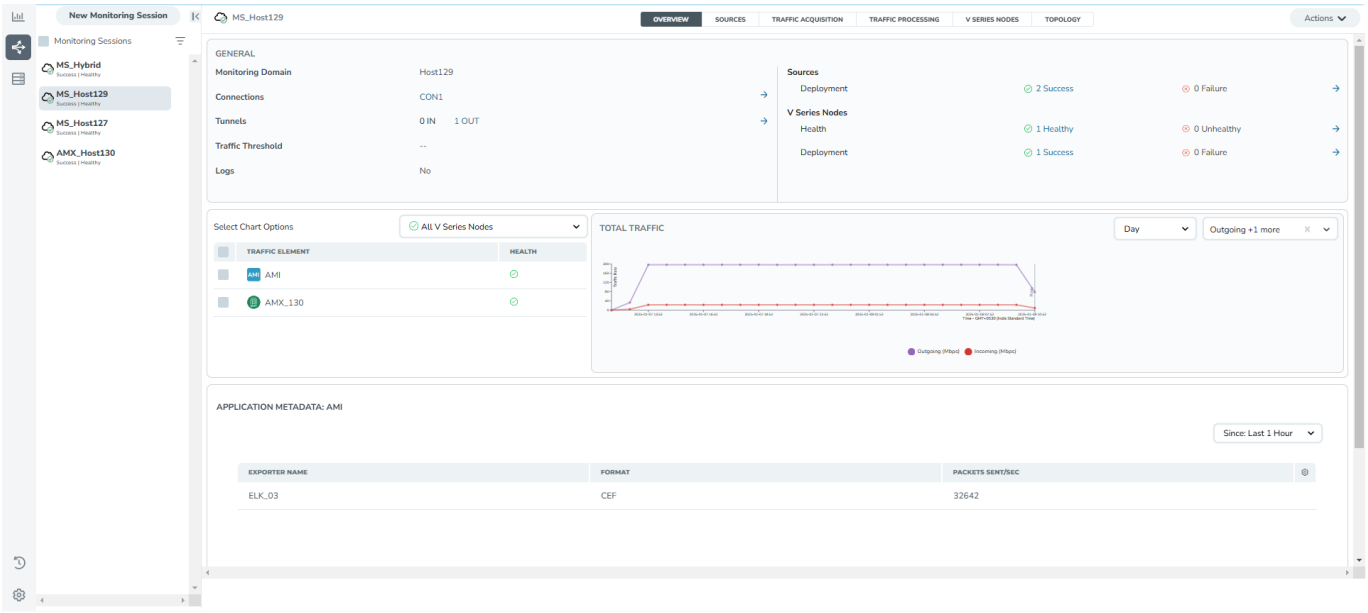
1. On the Monitoring Session page, select the required Monitoring Session from the list view.
2. On the Overview tab, select the GigaVUE V Series Node from the All V Series Nodes drop-down menu.

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

View Monitoring Session Statistics

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

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



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

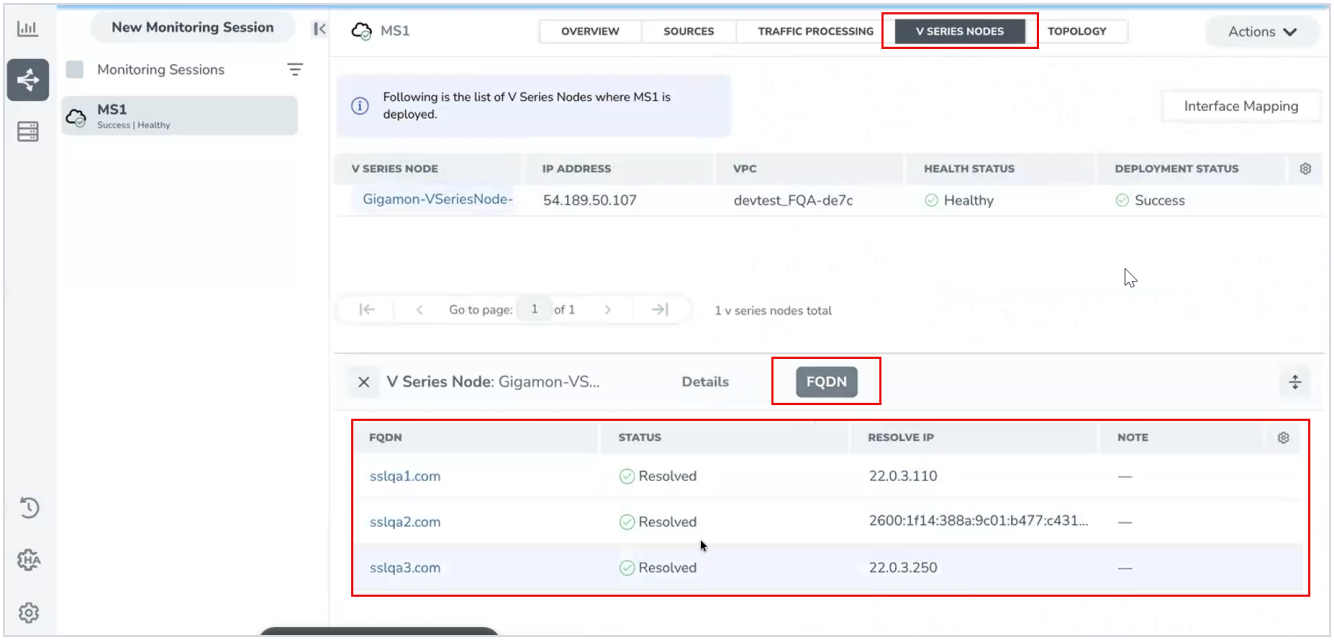
- You can view the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis.
- You can filter the traffic and view the statistics based on factors such as **Incoming**, **Outgoing**, **Ratio (Out/In), Incoming Packets**, **Outgoing Packets**, **Ratio (Out/In) Packets**. You can select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

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.



Visualize the Network Topology (OpenStack)

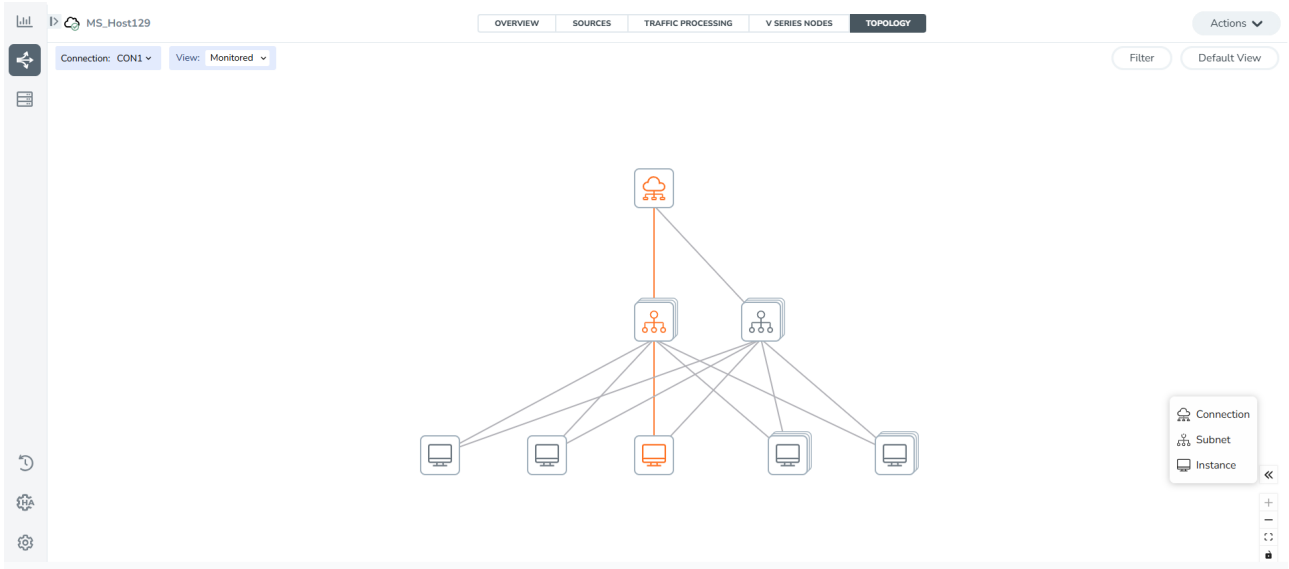
You can have multiple connections in GigaVUE-FM. Each connection can have multiple Monitoring Sessions configured within it. The Topology tab provides a visual representation of the monitored elements within a selected connection and Monitoring Session.

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 OpenStack

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

- Configure the OpenStack Settings
- Role Based Access Control
- About Audit Logs
- About Events

Manage Monitoring Domain

You can view the details of the monitoring domain that are created in the list view. The list view details can be viewed based on:

- Monitoring Domain
- Connections Domain
- Fabric
- 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 - Click the Filter button on the right to filter the monitoring domain based on a specific criterion.
- Left filter - Click the to filter the monitoring domain based on the 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 when using the Secure Tunnels.
Filter	Filters the Monitoring Domain based on the list view options that are configured: • **Tunnel MTU** • **Acquisition Method** • **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
- Credential
- 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

You can use the Actions button to perform the following actions:

- **Edit Fabric** - You can select one fabric or multiple fabrics of the same Monitoring Domain to edit a fabric. You cannot choose different fabrics of multiple Monitoring Domains at the same time and edit their fabric components.
- **Delete Fabric** - You can delete all the fabrics associated with the Monitoring Domain of the selected fabric.
- **Upgrade Fabric** - You can select a Monitoring Domain or multiple Monitoring Domains to upgrade the fabric. You can upgrade the GigaVUE V Series Nodes using this option.
- **Restart OVS Traffic** - You can restart the OVS traffic. You can view the **Restart OVS Traffic** option only when you enable the check box **OVS Agent Traffic when V Series unreachable** in **Advanced Settings**. For more information on settings, refer to Configure the OpenStack Settings.
- **Shut Down OVS Traffic** - You can shut down the OVS traffic. You can view the **Shut down OVS Traffic** option only when you enable the check box **OVS Agent Traffic when V Series unreachable** in **Advanced Settings**. For more information on settings, refer to Configure the OpenStack Settings.

- **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 TroubleshootingDebuggability and Troubleshooting.

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.

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 hearbeat time
- Agent mode
- Status

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.

Refer to Configure the OpenStack Settings, for information regarding **Settings**.

Configure Certificate Settings

To configure certificate settings:

1. Go to **Inventory > VIRTUAL**.
2. Select your cloud platform.
3. Select **Settings > Certificate Settings**.
   The **Certificate Settings** page appears.

4. From the **Algorithm** drop-down list, select the algorithm that determines the cryptographic security of the certificate.

   NOTE: If selecting RSA 8192, the certificate generation may take longer due to the increased key size.

5. In the **Validity** field, enter the total validity period of the certificate.
6. In the **Auto Renewal** field, enter the number of days before expiration of the auto-renewal process should begin.
7. Select **Save**.

Configure the OpenStack Settings

To configure the OpenStack Settings:

1. Go to **Inventory > VIRTUAL > OpenStack,** and then click **Settings**.
2. Click the **Settings** drop-down, and then select **Advanced Settings**.
3. Click **Edit** to edit the Advanced Settings fields.

Refer to the following table for descriptions of the Settings fields.

Settings	Description
Refresh interval for VM target selection inventory (secs)	Specifies the frequency for updating the inventory of VMs in OpenStack.
Refresh interval for fabric deployment inventory (secs)	Specifies the frequency for updating the inventory of GigaVUE fabrics in OpenStack.
Number of UCT-Vs per V Series Node (applicable only for UCT-V based connections)	Specifies the maximum number of instances that can be assigned to the V Series node.
Number of hypervisors per V Series Node (applicable only for OVS mirroring)	Specifies the maximum number of hypervisors that can be assigned to the V Series node.
Refresh interval for UCT-V inventory (secs)	Specifies the frequency for discovering the UCT-Vs available in the project. This is applicable for UCT-Vs only. **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.
OVS Mirror tunnel	Specifies the startup range value of the OVS mirror tunnel ID. This is applicable for UCT-

Settings	Description
range start	V OVS modules only.
OVS Mirror tunnel range end	Specifies the closing range value of the OVS mirror tunnel ID. This is applicable for UCT-V OVS modules only.
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.
OVS Agent Traffic when V Series is down	Enable this option to stop the OVS module from sending the traffic to the V Series node. You can stop the traffic either manually or automatically. Refer to Shutdown or Restart of OVS traffic to know more about the manual or automatic shut down and restart.
Use UCT-V conf file	Enable this option to allow interface mirroring to follow the configuration defined in the file. Disable it to mirror traffic from all physical interfaces. **Note**: • When changing the UCT-V conf file option from **enabled to disabled**, ensure to undeploy the Monitoring Session and delete the Monitoring Domain. Once changed, you should create a new Monitoring Domain and configure the Monitoring Session. • When changing the UCT-V conf file option from **disabled to enabled**, do the following: 1. Edit the uctv.conf file a. **Windows**: C:\ProgramData\Uctv\uctv.conf b. **Linux**: /etc/uctv/uctv.conf 2. Delete the skipConf file from the backup folder a. **Windows**: C:\ProgramData\Uctv\bak\skipConf b. **Linux**: /var/lib/uctv/bak/skipConf 3. Restart the UCT-V a. **Windows**: Restart from the Task Manager b. **Linux**: sudo service uctv restart
Fabric Node Unhealthy Timeout (secs)	Specifies the duration GigaVUE-FM waits before marking a fabric node as unhealthy when heartbeats are not received from it, before declaring the un-registration timeout. The supported range is 30 to 300 seconds, the value must be a multiple of the 30-second heartbeat interval, and the default is 90 seconds.
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 down state. GigaVUE-FM retries every 60 seconds.

- UCT-V OVS module supports a maximum of 255 source interfaces per OpenStack node.
- A maximum of 100 OpenStack connections are allowed for an OpenStack module.

Shutdown or Restart of OVS traffic

GigaVUE-FM allows you to stop or restart the traffic through OVS Mirroring based on the availability of V Series node.

GigaVUE-FM helps you to stop the traffic when the V Series node is unreachable or unrecoverable, and restart it when the GigaVUE V Series Node is reachable again in the following ways:

- Manual shutdown or restart of OVS traffic
- Automatic shutdown or restart of OVS traffic

Manual shutdown or restart of OVS traffic

The traffic sent from the OVS Mirroring Module can be manually stopped and started.

To shut down or restart the OVS traffic manually, follow these steps:

1. Go to **Inventory > VIRTUAL > OpenStack,** and then click **Settings**
2. Click the **Settings** drop-down, and then select **Advanced Settings**.
3. Enable the check box **OVS Agent Traffic when V Series is down**.
4. Click the **Fabric** tab.
5. Select the V Series node.
6. Click the **Actions** drop-down list and select **Shut down OVS Traffic** or **Restart OVS Traffic** as required.

NOTE: You can view the **Shut down OVS Traffic** or **Restart OVS Traffic** options only when you enable the check box **OVS Agent Traffic when V Series is down** in the Advanced Settings.

Automatic shutdown or restart of OVS traffic

When the GigaVUE Cloud Suite V Series node is deleted or changed to an unrecoverable state in the OpenStack platform, GigavUE-FM performs the action as explained in the following table:

V Series node Status in OpenStack	Action in GigaVUE-FM	Action in GigaVUE-FM when you enable the option
When a V Series node is deleted from the OpenStack platform	GigaVUE-FM automatically sets the status of that V Series node as terminated.	In the **Advanced Settings**, when the **OVS Agent Traffic when V Series is down** checkbox is enabled, GigaVUE-FM removes the source interfaces of OVS Mirroring module
When a V Series Node is changed to a stopped or shutoff state in Openstack	GigaVUE-FM Health monitoring module tries to start the V Series node. If unsuccessful after two attempts, GigaVUE-FM considers the V Series node to be unrecoverable.	In the **Advanced Settings**, when the **OVS Agent Traffic when V Series is down** checkbox is enabled, GigaVUE-FM removes the source interfaces of OVS Mirroring module.
When a V Series node is in an active state and its connection to GigaVUE-FM is restored	GigaVUE-FM Health Monitoring module determines V Series Node is in a healthy state.	In the **Advanced Settings**, when the **OVS Agent Traffic when V Series is down** checkbox is enabled, GigaVUE-FM adds the source interfaces of the OVS Mirroring module and restarts the OVS traffic.

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 Fabric Deployment • Cloud Configurations • Sys Dump • Syslog • Cloud licenses • Cloud Inventory	• Configure GigaVUE Cloud Components • Create Monitoring Domain and Launch Visibility Fabric
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. Refer to Configure Role-Based Access for Third Party Orchestration for more details on how to create users, roles, and user groups for Third Party Orchestration.

NOTE: Cloud APIs are also RBAC enabled.

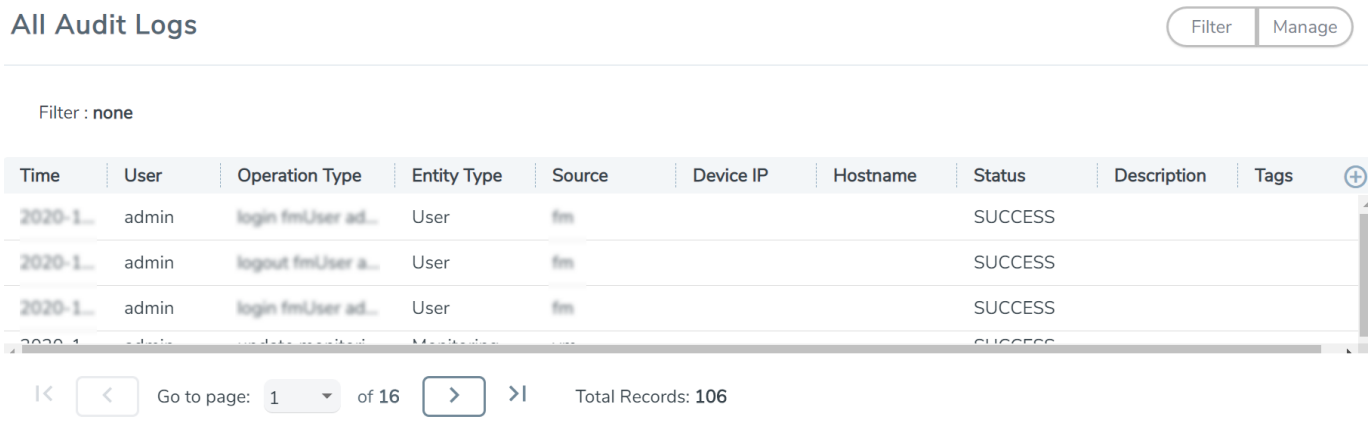
Refer to the *GigaVUE Administration Guide* for detailed information about Roles, Tags, User Groups.

About 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.



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.

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:

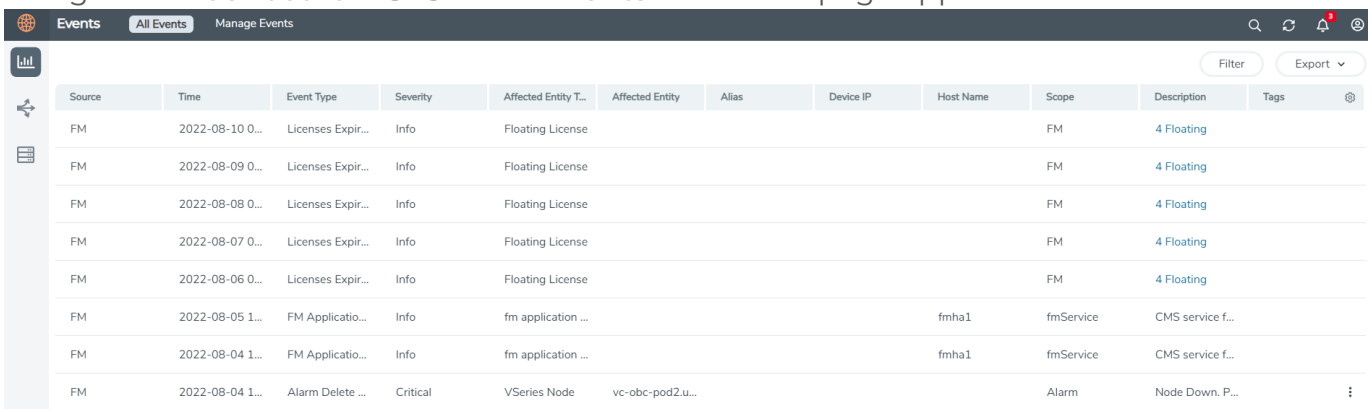
- Cloud provider License Expiry
- UCT-V Inventory Update Completed
- Cloud provider Connection Status Changed

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.



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 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.

Controls/ Parameters	Description
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.

High Availability for Cloud Infrastructure

This section provides details about the GigaVUE-FM High Availability (HA) feature in Cloud environment and describes how to configure, upgrade, and troubleshoot the feature.

Refer to the following topics for details:

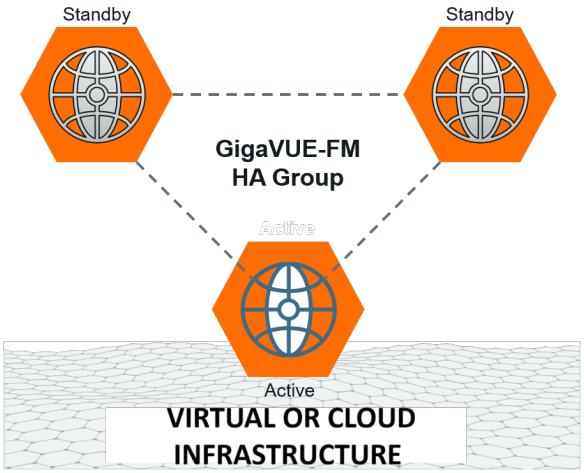
- About GigaVUE-FM High Availability
- Supportability Information
- Recommendations, Rules, Notes, and Limitations
- Licensing Information
- Hostname Setups
- Configure GigaVUE-FM HA

- Configure Secure Communication between Fabric Components in FMHA
- GigaVUE-FM HA Landing Page
- Remove Standby GigaVUE-FM Instance
- Disassemble GigaVUE-FM High Availability Group
- GigaVUE-FM High Availability States
- Fail over Mechanism
- GigaVUE-FM High Availability Scenarios
- Troubleshoot GigaVUE-FM High Availability Issues

About GigaVUE-FM High Availability

The GigaVUE-FM High Availability (HA) feature supports a highly available fabric management environment with minimal interruption. The GigaVUE-FM HA architecture consists of a minimum of three GigaVUE-FM instances that run together as a highly available group. The highly available group provides protection from failure of any one of the members in the group. The communication between the GigaVUE-FM instances in the highly available group can be encrypted using Pre-shared Secret Key or Certificate based Authentication.

The following image shows the high-level architecture of the GigaVUE-FM HA feature.



GigaVUE-FM allows you to add more than three GigaVUE-FM instances to the HA group (up to 7). The hardware requirements and resource allocation process for the additional instances are the same as that of the active and standby instances. The additional instances can be configured either as supporting instances or as active-eligible instances:

Supporting instances

Supporting instances are GigaVUE-FMs that run minimal services, thereby providing more resources to OpenSearch services. These instances:

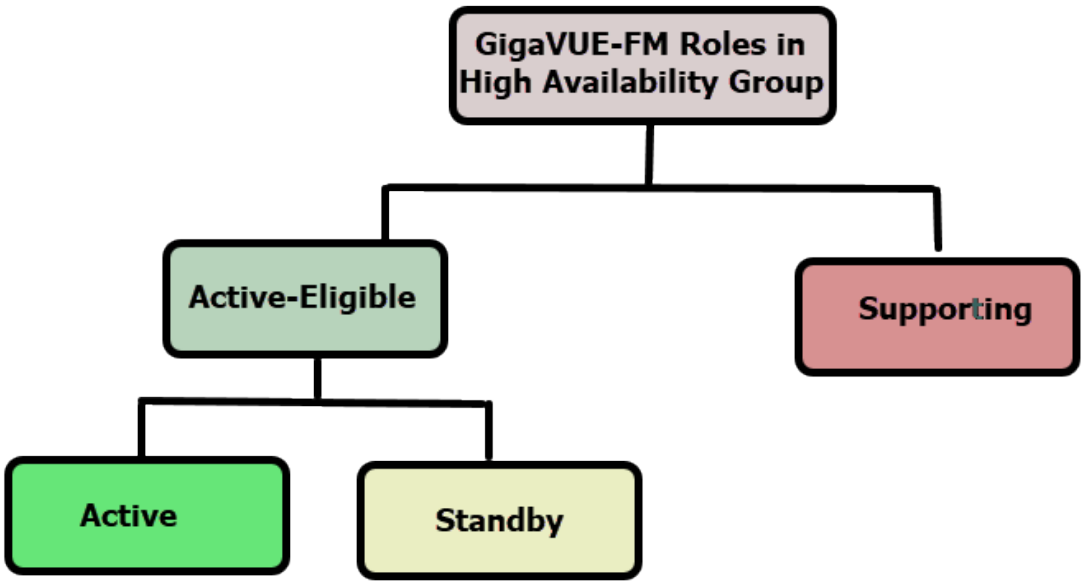
- Do not distribute application requests, as all application services are disabled by default.
- Do not participate in distributed services
- Provide enhanced scalability.
- Are excluded from active instance election

Active-Eligible instances

Active-Eligible instances are GigaVUE-FMs that distribute application requests and run distributed services. These instances:

- Enhance fault-tolerant and high available system.
- Are eligible for active instance election.

Refer to the following diagram for the roles in the HA group.



NOTE: When you add a GigaVUE-FM instance as either a supporting instance or a standby instance, the current configuration of that GigaVUE-FM instance will be removed once it is added to the existing GigaVUE-FM HA cluster.

Gigamon cloud solution deployments in GigaVUE-FM HA are supported on both public and private platforms. These deployments can operate in one of the following modes:

- GigaVUE-FM Orchestrated
- Third-party Orchestrated (Generic / Integrated)

In either scenario, you need to first launch GigaVUE-FM and run.

GigaVUE-FM Orchestration

Using GigaVUE-FM orchestration, you can deploy GigaVUE Fabric Components through the GigaVUE-FM user interface. After launching GigaVUE-FM and applying the required license, you can deploy components such as the GigaVUE V Series Node, GigaVUE V Series Proxy, and UCT-V Controller.

Once the deployment is complete, GigaVUE-FM continuously monitors the health and status of the deployed components. If any error condition occurs, GigaVUE-FM automatically relaunches or replaces the affected component to ensure system reliability.

Third-party Orchestration

With this method, you deploy GigaVUE Fabric Components using external orchestration solutions. GigaVUE-FM does not manage the deployment process directly, but you can still use it to monitor the health of the deployed components.

To enable monitoring, you need to provide an additional configuration file during deployment. This file allows components like the GigaVUE V Series Node, GigaVUE V Series Proxy, UCT-V Controller, UCT-V, UCT-C, and GCB to register with the Monitoring Domain in GigaVUE-FM . Refer to the following sections for more details:

- Deploy UCT-C Solution in Kubernetes
- Deploy GCB Controller Service and Pods

Supportability Information

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

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

Orchestration Method	Platforms/components Supported
	• GCB **Integrated Mode** • AWS • Azure • OpenStack

Recommendations, Rules, Notes, and Limitations

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

Recommendations

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

Rules and Notes

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

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

- You can deploy the GigaVUE-FM HA virtual machines on a WAN link with a maximum latency of 200 ms.
- You cannot add a GigaVUE-FM Hardware Appliance and a GigaVUE-FM virtual machine in the same HA group.
- The `Reload All` or `Reload any FM` commands will work only if there is an active GigaVUE-FM instance in the HA group.

- You must ensure that NTP is configured across all FM HA nodes.
- The GigaVUE-FM HA Cluster functions like standalone GigaVUE-FM's when a fail-over occurs during long-running operations, such as Fabric Launch, Monitoring Session Deployment, or Fabric Upgrade.
- When FMHA is configured with a reverse proxy, access to GigaVUE-FM requires non-SAML authentication, such as local authentication. Use the following sample URL to access GigaVUE-FM:

  `https://<proxy_ip>/admin`
- If GigaVUE-FM has multiple IPv6 addresses, ensure FMHA uses the permanent IPv6 as the source IP. GigaVUE-FM excludes temporary IPv6 addresses from the IP set, which causes connection failures and FMHA formation issues.

Shrinking the size of the HA group:

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

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

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

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

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

Supporting GigaVUE-FM Instances:

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

Notes:

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

Licensing Information

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

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

Hostname Setups

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

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

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

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

Configure GigaVUE-FM HA

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

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

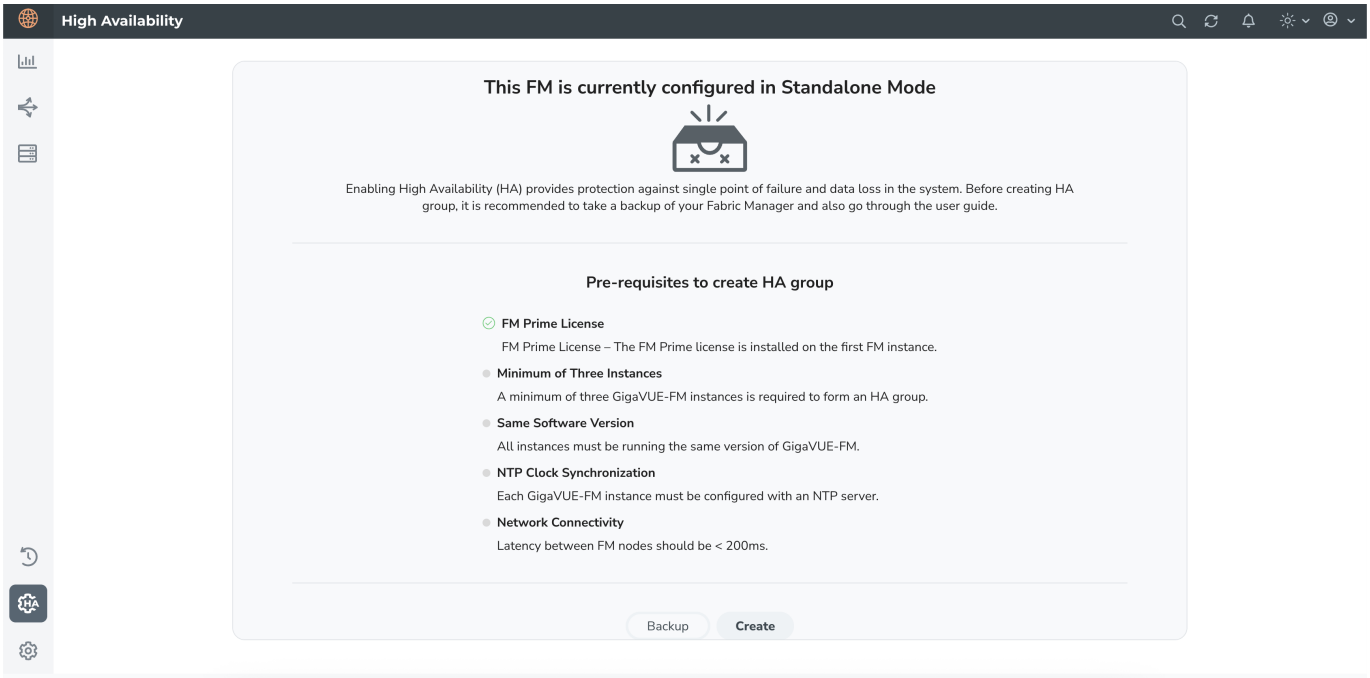
Refer to the following sections:

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

Create and Add GigaVUE-FM Instances to HA Group

To add instances and to configure the HA group:

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



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

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

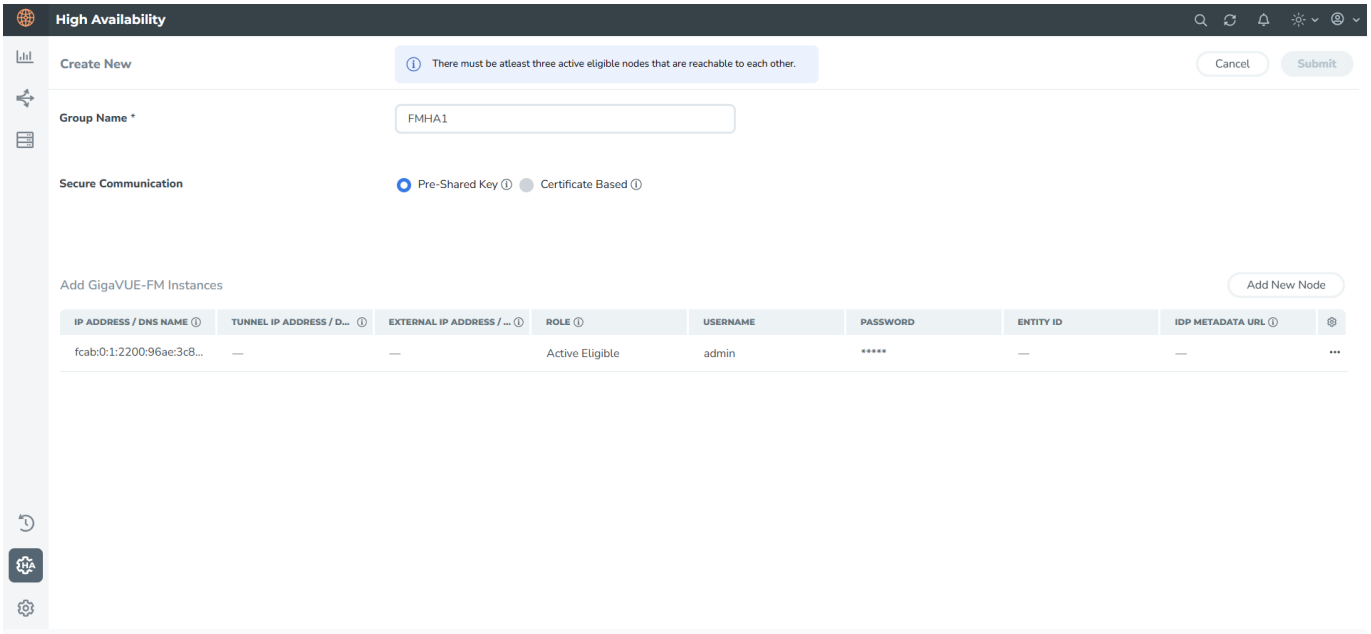
3. Choose any one of the following options to encrypt the communication between nodes:
   - **Pre-shared Secret Key** – Encrypts communication between GigaVUE-FM nodes using an application-managed shared secret key. Refer to Encryption of Communication between FMHA Nodes.
   - **Certificate Based** –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

Edit the Primary Node

1. On the first instance, click the ellipses icon and select **Edit Node**.
2. Enter the following details: (These fields are optional)

- Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
  - **Tunnel IP Address/DNS Name**: IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.
  - **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name**: If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
- Enter the Entity ID and the IdP Metadata URL.

3. Click the **save** icon to save the configuration. Refer to Additional Fields for details about other fields displayed on the HA page.



Add New Node

1. Click **Add New Node** to add the GigaVUE-FM instances. Enter the following details.:
   - **IP Address/DNS Name** - IP address/DNS name of the GigaVUE-FM instances (for communication between the GigaVUE-FMs and to configure the HA group). It is always recommended to use the DNS name while creating the GigaVUE-FM High Availability group.
   - (Optional) Select the **Use Tunnel IP** check box to enable tunnel configuration. Selecting this check box reveals two additional fields:
     - **Tunnel IP Address/DNS Name**: IP address used for establishing the secure communication tunnel between HA instances. This is required when the instances are reachable only through public or external IP addresses.

- **External IP Address / DNS Name is same as Tunnel IP Address / DNS Name**: If you check this box, GigaVUE-FM automatically copies the value from the Tunnel IP field to the External IP field and disables manual editing. If you uncheck it, the External IP field becomes editable again, retaining the previously entered value.
- When you add a new node with a Tunnel IP address to an existing HA group, the FMHA table updates to reflect the information in the newly added column.

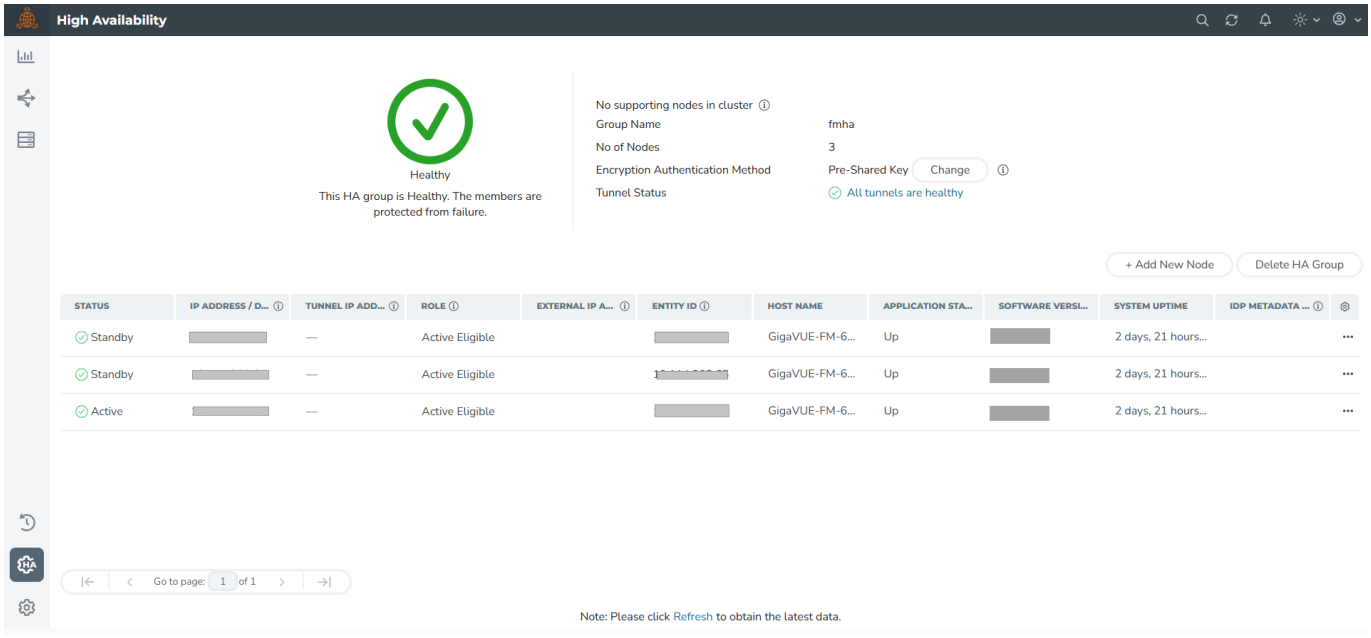
Notes:

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

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

2. Click **Save** to validate the details entered for the instances.

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



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 1: Additional Fields

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

Field	Description
	• **Down**: Indicates service is not ready to handle the user requests with one or more underlying services being down. • **In Progress**: A state that appears between up and down indicating that the services are initializing.
Software Version	Software version of the GigaVUE-FM instance.
System Uptime	Time since the instance is up and active.
Third Party Authentication URL	Third party authentication URL required for SSO configuration.
Upgrade Status	Upgrade status of GigaVUE-FM. This is displayed only when the upgrade is triggered. Otherwise, this field is empty.

You can also remove the instances from the HA group and reduce the size of the HA group as required. However, while removing the instances ensure to retain the quorum[1] number of instances.

NOTE: To scale components such as Fabric Health Analytics, topology visualization and to improve the performance of GigaVUE-FM by fully utilizing the resource available, it is recommended to configure three active eligible instances and additional supporting instances.

The following are the behavioral changes observed in the HA group with active, standby and supporting instances:

Task	Scenario	Behavior
Upgrade Operation	Few instances not being up (either active-eligible or supporting instances not up)	Upgrade will not be initiated.
	Few supporting instances are removed from the HA group after formation of the group	Upgrade will be initiated, and the supporting instances present in HA group will be upgraded first followed by the active-eligible instances.
	All instances are up and running	Upgrade will be initiated in an active instance. Upgrade will be done in a rolling fashion starting with supporting instances and ending with the active-eligible instances.
System Configurations	All instances are active-eligible	All GigaVUE-FM's GUI will be accessible, and user

[1]Quorum number of instances is the minimum number of instances that must be active-eligible in a HA group. In a HA group with five instances, the quorum number is 3 (based on the formula n/2+1)).

		can configure the system through the same
	HA with active-eligible and few Supporting instances	For supporting instances, GUI will not be available (as application service is shut down). An option is provided in the GUI to turn on the application service and its dependent service(s) and the status can be seen in HA page. You can then configure the system related configuration from the logged in GigaVUE-FM instance once the application status of the supporting instance is up.
Background processes		
-FM Service Distribution	All GigaVUE-FMs are active-eligible instances.	Services will be equally distributed/redistributed among the instances based on the availability of the instances.
	HA group with active-eligible instances and supporting instances.	Services will be distributed/re-distributed only to active eligible instances
-Load Balancer	All GigaVUE-FMs are active-eligible instances.	All GigaVUE-FM API requests will be distributed to all instances in the HA group.
	With supporting instances.	All GigaVUE-FM API requests will be distributed only to active-eligible instances.

Connect GigaVUE-FM High Availability to a Monitoring Domain using GigaVUE-FM Orchestration

When fabric nodes need to communicate with GigaVUE-FM using public IP addresses, you should enable the **Use Public IP** option during Monitoring Domain creation. When **Use Public IP** option is enabled, GigaVUE-FM pushes its public IP address to Fabric Nodes. Fabric Nodes use this IP to communicate back to FM. FM does not initiate communication using its private IP. This setting is essential in cloud or hybrid environments where private

IPs are not reachable, such as across cloud platforms. It ensures that GigaVUE-FM can send notifications and accept registrations using an externally accessible IP. Refer to Supportability Information for the list supported platforms.

Connect GigaVUE-FM High Availability to a Monitoring Domain using Third Party Orchestration

Fabric Nodes can be configured with either the IP address (remoteIP) or the Fully Qualified Domain Name (FQDN) (remoteAddress) of GigaVUE-FM to establish communication. The node sends registration requests to GigaVUE-FM until it successfully registers. In High Availability mode, if a node fails to register with one GigaVUE-FM instance, it automatically attempts registration with the next available GigaVUE-FM in its list until successful. Refer to Supportability Information for the list supported platforms.

Use the following data to manually configure **/etc/gigamon-cloud.conf** with registration configurations to register the fabric components with GigaVUE-FM.

```
Registration:
    groupName: <Monitoring Domain Name>
    subGroupName: <Connection Name/Alias>
    token: <Token>
    remoteAddress: <remoteIPs/FQDN of the GigaVUE-FM>
    remotePort: 443
    sourceIP: <IP address of the fabric node>
    nameServer: <DNS server's IP address>
    usePublicIP: <true/false>
    preferV6: <true/false><If set to true IPv6 will be preferred over IPv4, used only for DNS
resolution>
```

Notes:

- Configure either remoteIP or remoteAddress—not both. Using both prevents the node from registering with GigaVUE-FM.
- Remote IP will continue to work as expected. However, if you plan to use FQDN, update the configuration to use the Remote Address field instead.
- When configuring a remote address, you can specify multiple IP addresses. However, if you use a FQDN, ensure that you configure only one address
- FQDN configuration requires a valid nameserver.
- Ensure that the FQDN configured for Third Party Orchestration resolves to all IP addresses in the High Availability cluster.

Integrate GigaVUE-FM High Availability with UCT-C

UCT-C supports integration with GigaVUE-FM in high availability mode. You can configure the controller with either a Fully Qualified Domain Name (FQDN) or a comma-separated list of GigaVUE-FM IP addresses in the Helm configuration. The controller iterates through the provided FQDN or IP list and attempts to register with the master GigaVUE-FM node. If the

GigaVUE-FM IP changes, you must update the configuration again. This enhancement enables UCT-C to maintain connectivity and recover from GigaVUE-FM failover events in HA deployments.

Below are the communication parameters:

- **fm_fqdn** – Specifies the GigaVUE-FM HA FQDN value used by the UCT-C Controller for communication with the GigaVUE-FM HA group. When this parameter is set, the controller uses the IP addresses resolved from the FQDN.
- **fmha_ip_list** – Specifies a comma-separated list of GigaVUE-FM HA IP addresses used by the UCT-C Controller for communication with the GigaVUE-FM HA group. If fm_fqdn is not set, the controller uses the IP addresses provided in fmha_ip_list by default.

Integrate GigaVUE-FM High Availability with Gigamon Containerized Broker

GigaVUE-FM High Availability group can be configured in an Gigamon Containerized Broker through a DNS Server in which the GigaVUE-FM instances are launched and deployed in the Gigamon Containerized Broker environment.

For GCB to make use of high availability feature of GigaVUE-FM, you must configure the FQDN (Fully Qualified Domain Name) of the GigaVUE-FM.

In a standalone GigaVUE-FM, the GCB uses the FQDN name of the GigaVUE-FM (if configured). It uses the legacy method of using the configured IP address when:

- The FQDN is not configured or
- The GCB fails to resolve to an IP address.

For example:

GCB Controller YAML file:

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

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

For more details on the GigaVUE-FM High Availability configuration, refer to the *GigaVUE Administration Guide guide*.

Configure Secure Communication between Fabric Components in FMHA

IMPORTANT: Before upgrading the Fabric Components to version 6.10 or later, complete the following steps after upgrading GigaVUE-FM to version 6.10 or later.

Follow these steps:

1. Access the active GigaVUE-FM via CLI.
2. Archive the stepCA directory using the following commands:
   ```
   sudo su
   cd /var/lib
   tar -cvf /home/admin/stepca.tar stepca
   ```
3. Set the permissions of the tar file using the following commands:
   ```
   chmod 666 /home/admin/stepca.tar
   ```
4. Copy the tar file to all standby instances in the **/home/admin/ directory** using scp:
   ```
   scp /home/admin/stepca.tar <standby-node>:/home/admin/
   ```
5. Download the **runstepca_fmha** script from the Community Portal.
6. Log in to the standby instance using CLI.
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 1Remove Standby GigaVUE-FM Instance.

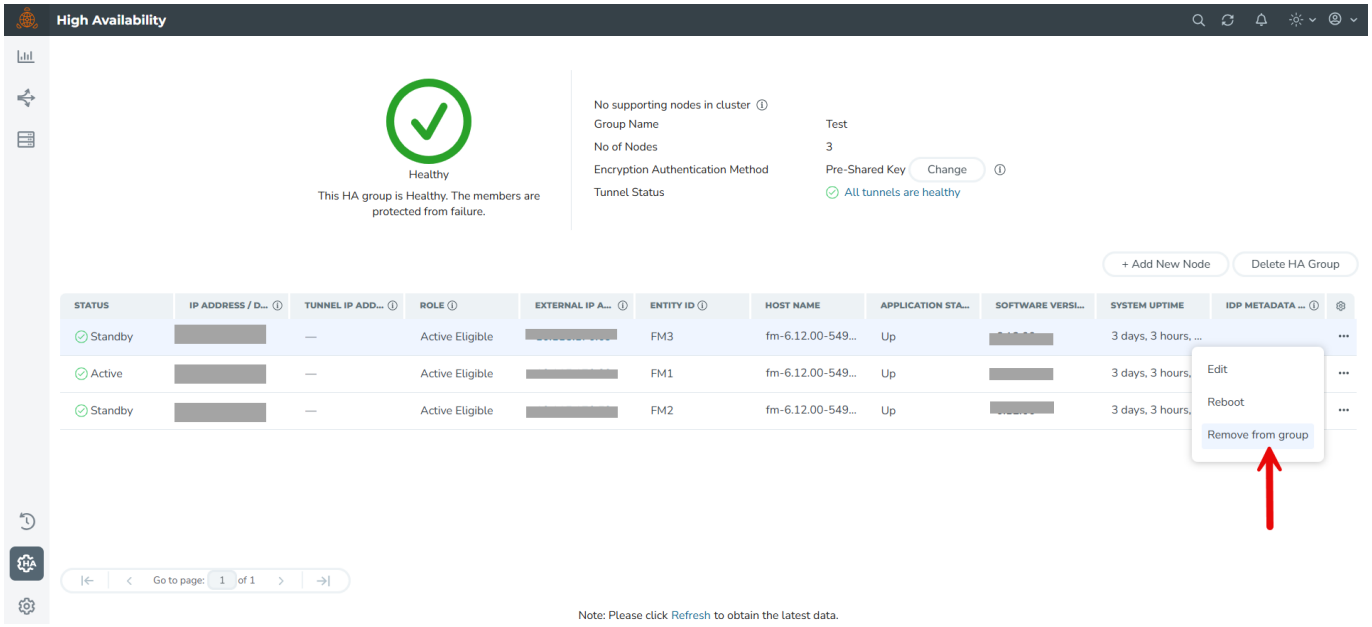


Figure 1 *Remove Standby GigaVUE-FM Instance*

3. Select the **Remove from group** option and confirm your chnages. 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.

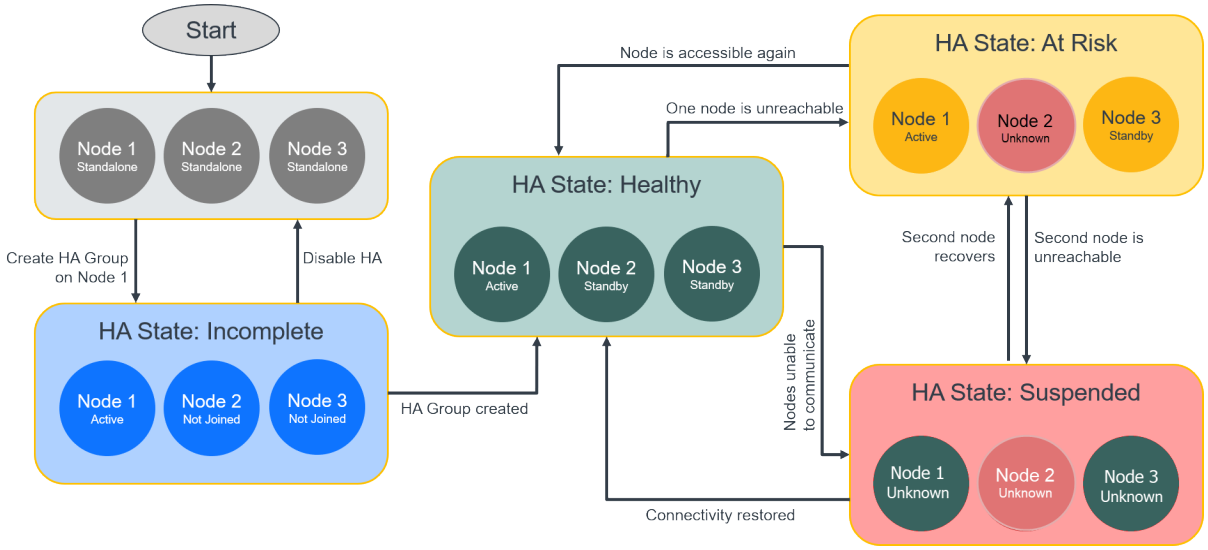
GigaVUE-FM High Availability States

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 2: High Availability States

State	Number of GigaVUE-FM Instances	Description
Healthy	Three GigaVUE-FM instances are up and running	One GigaVUE-FM instance is in active state. Other two instances are in standby state.
At Risk	Two GigaVUE-FM instances are up and running	One GigaVUE-FM instance is in active state. Another GigaVUE-FM instance is in standby state. The third GigaVUE-FM instance has either not joined the HA group or has left the HA group. **NOTE:** You must recover the standby GigaVUE-FM instance with in 3 days (72 hours). Failure to do so will cause the instance to move out of the High Availability group.

State	Number of GigaVUE-FM Instances	Description
Incomplete	One GigaVUE-FM instance is up and running	Only one GigaVUE-FM instance is in active state. The other two GigaVUE-FM instances have either not joined the HA group or have left the HA group.
Standalone	One GigaVUE-FM instance is up and running	HA is not configured on the GigaVUE-FM instance.
Suspended	Two or more GigaVUE-FM instances are up and running	HA is configured, but an active GigaVUE-FM instance is yet to be elected. **NOTE:** The GigaVUE-FM high availability group can move to a **Suspended** state from the **At Risk** state when the HA group has only one active-eligible and reachable GigaVUE-FM instance. Events will not be captured for the status of the GigaVUE-FM instance that went down (which causes the HA group to move into a suspended state) as there is no active/leader instance in the High Availability group.



In case of FMHA going to suspended mode from At-Risk mode, the reachability status of node which went down (due to which HA went into suspended mode) won't be audited as writes are not accepted if there is no active/leader instance in the DB cluster

Figure 2 *High Availability States*

Fail over Mechanism

The active GigaVUE-FM instance in the high availability group may fail at times resulting in one of the standby instances to take over and become the active instance. This process is called failover.

The following table provides the reasons for failover:

Reason for Failover	Description
Reloading the active GigaVUE-FM instance	An active GigaVUE-FM instance is reloaded (using the Reboot option) to bring back the HA group to healthy state again.
Planned downtime of the active GigaVUE-FM instance	An active GigaVUE-FM instance is brought down due to various reasons, for example to upgrade to a newer software version.
Monitoring session deployment during HA leader switchover	The new leader automatically redeploys Monitoring Sessions that failed during the switchover.
Node registration during HA leader switchover	Registration requests may temporarily fail or be delayed. Fabric nodes automatically retry registration. The new leader's IP address is configured as the RMQ notification target on all fabric nodes.
Certificate handling or fabric deployment/upgrade during HA leader switchover	The newly elected leader resumes certificate handling and continues any in-progress fabric deployment or upgrade from the last known state. No manual intervention is required.

GigaVUE-FM High Availability Scenarios

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

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

Scenario	Changes in GUI
What happens to the High Availability page immediately after a failover?	The High Availability page may not update immediately or may not show all the GigaVUE-FM instances. **NOTE:** Refresh the page after a minute to view the new active GigaVUE-FM instance. However, the page will get updated automatically after 5 minutes.

Scenario	Changes in GUI
What happens to an active GigaVUE-FM instance when it fails (either by itself or if failover is triggered manually)?	• One of the standby GigaVUE-FM instances changes to the active state. The GigaVUE-FM instance that was initially in the active state changes to the standby state. It takes a few seconds for this transition. • The GigaVUE-FM GUI of the new active instance will have all the menus and dashboards.
What happens to the GigaVUE-FM instances that were previously in standby state?	• One of the standby GigaVUE-FM instances changes to the active state • The other standby GigaVUE-FM instance remains in the standby state.
What happens to the embedded devices when a new Active instance takes over?	There are no changes to the devices except that they are being managed by the new active GigaVUE-FM instance.
How do you trigger a failover?	Click on the 'Reboot' option on the current active GigaVUE-FM instance to trigger a failover.

Troubleshoot GigaVUE-FM High Availability Issues

Use the following table to troubleshoot issues that you might encounter while working with the HA feature.

Problem	Solution
Unable to add a license to GigaVUE-FM HA group after a failover **Reason:** Using the MAC Address in the **About** page to generate the license	Always use the **Challenge MAC Address** in the **Licenses** page of the active GigaVUE-FM instance to generate licenses. Add the licenses to the HA group.
Unable to join the GigaVUE-FM HA group after changing the password **Reason:** Not logging out of GigaVUE-FM after changing the password and before joining the HA group	Always logout of GigaVUE-FM after changing the password and login again with the new password before joining the HA group.

(missing or bad snippet)(missing or bad snippet)

Analytics for Virtual Resources

Analytics in GigaVUE-FM is a standalone service that provides data visualization capabilities. Using Analytics [1], you can create visual elements such as charts that are embedded as visualizations. The visualizations are grouped together in dashboards.

You can also create search objects using Analytics. Dashboards, Visualizations and Search Objects are called Analytics objects. For details, refer to Analytics.

Rules and Notes:

- You cannot edit or delete these default dashboards. However, you can clone the dashboards and visualizations.
  Refer to the Clone Dashboard section in GigaVUE-FM Installation and Upgrade Guidefor more details.
- Use the **Time Filter** option to select the required time interval for which you need to view the visualization.

Virtual Inventory Statistics and Cloud Applications Dashboard

Analytics dashboards allow users to monitor the physical and virtual environment and detect anomalous behavior and plan accordingly.

For details, refer to the Analytics section in *GigaVUE Fabric Management Guide* for details on how to create a new dashboard, clone a dashboard, create a new visualization, and other information about the Discover page and Reports page.

How to access the dashboards

1. Go to  -> **Analytics -> Dashboards.**
2. Select the required dashboard to view the visualizations.

The following table lists the various virtual dashboards:

[1]Analytics 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	*V Series Node Status by Platform*	Number of healthy and unhealthy V Series Nodes for each of the supported cloud platforms.
		Monitoring Session Status by Platform	Number of healthy and unhealthy monitoring sessions for each of the supported cloud platforms
		Connection Status by Platform	Number of healthy and unhealthy connections for each of the supported cloud platforms
		GCB Node Status by Platform	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	*V Series Node Maximum CPU Usage Trend*	Line chart that displays maximum CPU usage trend of the V Series node in 5 minutes interval, for the past one hour. **NOTE:** The maximum CPU Usage trend refers to the CPU usage for service cores only. Small form factor V Series nodes do not have service cores, therefore the CPU usage is reported as 0.
		V Series Node with Most CPU Usage For Past 5 minutes	Line chart that displays Maximum CPU usage of the V Series node for the past 5 minutes. **NOTE:** You cannot use the time based filter options to filter and visualize the data.
		V Series Node Rx	Receiving trend of the V

Dashboard	Displays	Visualizations	Displays
		Trend	Series node in 5 minutes interval, for the past one hour.
		V Series Network Interfaces with Most Rx for Past 5 mins	Total packets received by each of the V Series network interface for the past 5 minutes. **NOTE:** You cannot use the time based filter options to filter and visualize the data.
		V Series Node Tunnel Rx Packets/Errors	Displays the reception of packet at the Tunnel RX. This is the input to V Series Node, Grouping by tunnel identifier comprising {monDomain, conn, VSN, tunnelName}, before aggregation.
		V Series Node Tunnel Tx Packets/Errors	TX is for output tunnels from VSN. V Series Node Tunnel Tx Packets/Errors
Dedup	Displays visualizations related to Dedup application. You can filter the visualizations based on the following control filters: • Platform • Connection • V Series Node	*Dedup Packets Detected/Dedup Packets Overload*	Statistics of the total de-duplicated packets received (ipV4Dup, ipV6Dup and nonIPDup) against the de-duplication application overload.
		Dedup Packets Detected/Dedup Packets Overload Percentage	Percentage of the de-duplicated packets received against the de-duplication application overload.
		Total Traffic In/Out Dedup	Total incoming traffic against total outgoing traffic
Tunnel (Virtual)	Displays visualizations related to the tunneled traffic in both bytes as well as the number of packets. You can select the following control filters, based on which the visualizations	*Tunnel Bytes*	Displays received tunnel traffic vs transmitted tunnel traffic, in bytes. • For input tunnel, transmitted traffic is

Dashboard	Displays	Visualizations	Displays
	will get updated: • **Monitoring session**: Select the required monitoring session. The cloud platform, monitoring domain and connection within the monitoring domain that is used by the V Series node are shown in square brackets, comma-separated, after the name, to distinguish the whole path to it. • **V Series node**: Management IP of the V Series node. Choose the required V Series node from the drop-down. • **Tunnel:** Select any of the tunnels shown in the Tunnel drop-down. The direction for each tunnel is shown with the prefix in or out. The following statistics are displayed for the tunnel: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets		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: • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets	*App Bytes*	Displays received traffic vs transmitted traffic, in Bytes.

Dashboard	Displays	Visualizations	Displays
	• Errored Packets • Dropped Packets	*App Packets*	Displays received traffic vs transmitted traffic, as the number of packets.
End Point (Virtual)	Displays Byte and packet level statistics for the un-tunneled traffic deployed on the V Series nodes. The following statistics that are shown for Endpoint (Virtual): • Received Bytes • Transmitted Bytes • Received Packets • Transmitted Packets • Received Errored Packets • Received Dropped Packets • Transmitted Errored Packets • Transmitted Dropped Packets The endpoint drop-down shows *<V Series Node Management IP address : Network Interface>* for each endpoint. You can select the following control filters, based on which the visualizations will get updated: • **Monitoring session** • **V Series node** • **Endpoint:** Management IP of the V Series node followed by the Network Interface (NIC)	*Endpoint Bytes*	Displays received traffic vs transmitted traffic, in Bytes.
		Endpoint Packets	Displays received traffic vs transmitted traffic, as the number of packets.

NOTE: The Tunnel (Virtual), App (Virtual) and Endpoint (Virtual) dashboards do not show data from the previous releases if the *Monitoring Session [Platform : Domain : Connection]* dashboard filter is applied. This is because, this filter relies on the new attributes in the OpenSearch database, which are available only from software version 5.14.00 and beyond.

Debuggability and Troubleshooting

Use the following information to help diagnose and resolve GigaVUE V Series Nodes issues.

Sysdumps

A sysdump is a log and system data package generated when a GigaVUE V Series Node experiences a crash (such as kernel, application, or hardware failure). These files are essential for debugging.

You cannot download sysdump files if the associated fabric component is deleted or unreachable.

Sysdumps—Rules and Notes

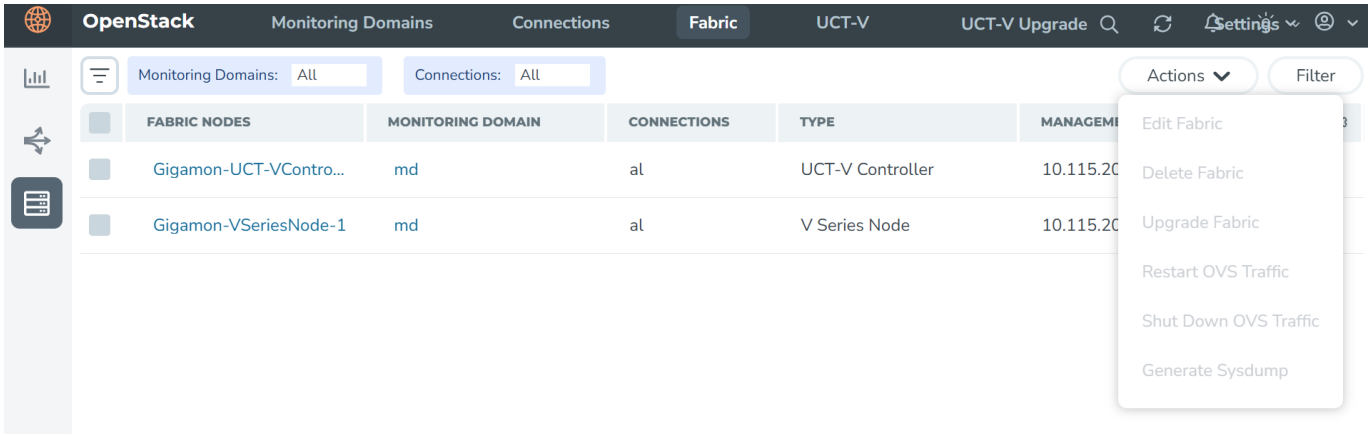
Consider the following points before you generate sysdumps:

- You can generate only one sysdump file at a time for a GigaVUE V Series Node.
- You cannot generate a sysdump file when generation of another sysdump file is in progress.
- The limit of sysdump files available per GigaVUE V Series Node is six. When you generate a seventh sysdump file, the file overwrites the first sysdump file.
- You can download only one sysdump file per GigaVUE V Series Node at a time.
- You can delete sysdump files in bulk for a GigaVUE V Series Node.
- To ensure efficient usage, the system limits the number of simultaneous sysdump generation requests to 10 GigaVUE V Series Nodes.
- GigaVUE V Series Node sysdumps are not stored in Fabric Manager but generated and stored on the GigaVUE V Series Node itself.

Generate a Sysdump File

To generate a sysdumps file:

1. Select the required node, and use one of the following options to generate a sysdump file:
   - Select **Actions > Generate Sysdump**.
   - In the lower pane, go to **Sysdump**, and select **Actions > Generate Sysdump**.
2. View the latest status, click **Refresh**.



Other Actions

- To download a sysdump file, select the file in the lower pane, and then click **Actions > Download**.

- To delete a sysdump file,
    1. Select the file in the lower pane.
    2. Select the desired sysdump file.
    3. Select **Actions > Delete**.
- To bulk delete, select all the sysdump files, and then select **Actions > Delete All.**

Troubleshooting

This section provides the information needed to troubleshoot GigaVUE-FM integration with OpenStack. Refer to

Refer to the following sections for details:

- OpenStack Connection Failed
- Handshake Alert: unrecognized_name
- GigaVUE V SeriesNode or UCT-V Controller is Unreachable

OpenStack Connection Failed

The connFailed state indicates that the OpenStack connection has failed. Check the following troubleshoot tips to restore the connection:

- Verify if GigaVUE-FM is able to reach the OpenStack cloud controller.
- Check if the OpenStack cloud controller is DNS resolvable from GigaVUE-FM.
- Verify if the region name provided while launching the instance is accurate.
- Ensure that all the security group rules required for communication between GigaVUE-FM and OpenStack cloud controller OR GigaVUE-FM and DNS server are accurately setup.
- Check if the Compute Servers that the nova API returns are reachable from GigaVUE-FM. Refer to Handshake Alert: unrecognized_name.

Handshake Alert: unrecognized_name

When setting up the OpenStack connection in GigaVUE-FM, the GigaVUE-FM logs might show a handshake alert: unrecognized_name error. This error is related to a Server Name Indication (SNI) error. Starting with Java 7, the JDK does not ignore the unrecognized name warning. To resolve this issue, perform either of the following:

- Fix the configuration on the server where the error is occurring.
- Ignore the warning on the client side (GigaVUE-FM server) by using the Java system property **--Djsse.enableSNIExtension=false** while launching GigaVUE-FM.

Contact support for information on how to use the Java system property. However, this is not recommended for security reasons.

GigaVUE V SeriesNode or UCT-V Controller is Unreachable

If GigaVUE V Series node or UCT-V Controller is unreachable, verify the following:

- The correct version of the image is uploaded.
- The network is reachable.

FAQs - Secure Communication between GigaVUE Fabric Components (OpenStack)

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 Security Group for OpenStack.

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.