



GigaVUE-FM Copilot User and Deployment Guide

Product Version: 6.14.00

Document Version: 1.0

(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.

Document Version	Date Updated	Change Notes
1.0	07/10/2026	The original release of this document with 6.14 LA.

Contents

GigaVUE-FM Copilot User and Deployment Guide	1
Change Notes	3
Contents	4
GV-FM-Copilot	6
GigaVUE-FM Copilot	6
What GigaVUE-FM Copilot Can Do	6
How GigaVUE-FM Copilot Works	7
Supportability	7
Supported Versions	7
Supported Platforms	7
VMware ESXi Hardware Requirements	7
Open Ports Access Requirements	8
Unsupported Features	9
GigaVUE-FM Copilot with LLM Connectivity Options	9
GigaVUE-FM Copilot with Public LLM Connectivity (Amazon Bedrock over Internet)	10
How it works	10
What You Manage	10
When to Choose this Option	10
GigaVUE-FM Copilot with Private LLM Connectivity (AWS Bedrock via PrivateLink)	11
How it works	11
What You Manage	11
When to Choose this Option	11
GigaVUE-FM Copilot with Local LLM Connectivity (AWS Native Deployment with PrivateLink)	12
How it works	12
What You Manage	12
When to Choose this Option	12
Deploy GigaVUE-FM Copilot on VMware vCenter	13
Deploy GigaVUE-FM Copilot on AWS	16
Configure GigaVUE-FM Copilot in GigaVUE-FM	17
Generate the API key from the GigaVUE-FM Copilot VM (VMware ESXi)	18
Generate the API key from the GigaVUE-FM Copilot VM (AWS)	18

Acquire Amazon Bedrock Credentials	18
Create IAM Policy	18
Create IAM User	20
Enable GigaVUE-FM Copilot in GigaVUE-FM	21
Get Started with GigaVUE-FM Copilot	22
Launch GigaVUE-FM Copilot	22
How to Use Prompts Successfully	23
Be clear and specific	23
Start simple, and then improve	23
Avoid ambiguity	23
Alias Handling with Numeric/Common Terms	24
Sample Prompts	24
Debuggability and Troubleshooting	26
Generate GigaVUE-FM Copilot Sysdump in GigaVUE-FM	26
Download GigaVUE-FM Copilot Sysdump file	26
Delete the GigaVUE-FM Copilot Sysdump file	27
Troubleshoot GigaVUE-FM Copilot	27
Resolve UI errors, timeouts, or unresponsiveness	27
GigaVUE-FM Copilot returns unexpected or incorrect answers	27
GigaVUE-FM Copilot VM does not initialize correctly	28
AI Telemetry Usage	28

GV-FM-Copilot

GigaVUE-FM Copilot

NOTE: GigaVUE-FM Copilot is released as a Limited Availability (LA) feature, providing you with the opportunity to evaluate its capabilities before its General Availability (GA) release. For access to LA software please contact your Gigamon account team or file a support case and ensure you have reviewed the LA Terms and Conditions. You can file support cases with Gigamon support during the LA period and provide feedback to Gigamon on your experience with GigaVUE-FM Copilot.

GigaVUE-FM Copilot is an AI-powered conversational assistant integrated directly into GigaVUE-FM. With GigaVUE-FM Copilot, you can:

- Get accurate answers directly sourced from Gigamon product documentation and GigaVUE-FM.
- Streamline deployment and daily operations with step-by-step configuration and troubleshooting guidance.
- Access both CLI and GUI instructions tailored to your needs.
- Receive secure, role-based responses—administrators get configuration-aware information; others receive documentation-only answers.
- Summarize system health and status to highlight the most relevant details for your GigaVUE-FM managed deployment.

What GigaVUE-FM Copilot Can Do

- **Find the right documentation and get answers:** Jump straight to the most relevant content in Gigamon documentation, with links and citations to exact answers. Receive detailed, accurate responses about GigaVUE-FM features and workflows, GigaVUE-FM physical and virtual visibility nodes and integrations, and networking concepts, all while maintaining conversation context for natural follow-up questions.
- **Assist with configuration:** Answer queries and provide clear step-by-step instructions for configuring the GigaVUE-FM system, physical platforms, and cloud suite. The system retrieves information from the documentation and from GigaVUE-FM.
- **System Status:** Answer questions and provide summaries about current system status, health, usage, and configuration.
- **Troubleshoot faster:** Walk through open issues (alarms), errors and provide insights with specific root causes and further steps necessary to understand and fix based on GigaVUE-FM data and Gigamon documentation.

- **AI Usage Telemetry:** Continuously improve recommendations over time using anonymized AI usage telemetry and user feedback, while protecting sensitive information through automatic redaction of common identifiers.

How GigaVUE-FM Copilot Works

The GigaVUE-FM Copilot applies Gen AI capabilities (Large Language Models enhanced with Retrieval Augmented Generation) to Gigamon documentation along with configuration and system data on your GigaVUE-FM instance. GigaVUE-FM Copilot uses an external LLM provider to process your natural language queries and generate contextual responses tailored to your Gigamon environment.

Key capabilities include:

- Delivering guided steps based on your environment
- Providing direct references to relevant sections of the Gigamon documentation set
- Using your existing GigaVUE-FM resources and Amazon Bedrock environment.
- Augmenting responses by combining Gigamon documentation with your internal GigaVUE-FM data.

Supportability

This section details the supported versions and hardware requirements for deployment.

Supported Versions

GigaVUE-FM Copilot is available in version 6.14 of GigaVUE-FM and subsequent releases. It references documentation from versions 6.12 and above.

Supported Platforms

GigaVUE-FM Copilot supports the following platforms:

- VMware ESXi
- AWS

VMware ESXi Hardware Requirements

The following table describes the hardware requirements for VMware ESXi to run GigaVUE-FM Copilot.

Table 1: Hardware Requirements for VMware Hypervisor

Hardware Requirements	
VMware Hypervisor	vSphere ESXi: v8.xx and above. Refer to Supported Hypervisors for VMware for more detailed information.
CPU	4 vCPUs
RAM	16GB
Disk Space	204GB
Network	At least one 1Gb NIC

Open Ports Access Requirements

NOTE: To ensure the GigaVUE-FM Copilot feature functions correctly, you must establish network connectivity between GigaVUE-FM and GigaVUE-FM Copilot by opening the required firewall ports listed below.

The following table describes the open ports access requirements.

Table 2: Open Ports Access Requirements

Direction	Protocol	Port Number	Service	Source CIDR	Destination	Purpose
Inbound	TCP	22	SSH	Administrat or Subnet	GigaVUE-FM Copilot	Allows CLI access to user-initiated management and diagnostics.
Inbound	TCP	443	HTTPS	GigaVUE-FM	GigaVUE-FM Copilot	Allows GigaVUE-FM to reach GigaVUE-FM Copilot for configuration and health check.
Outbound	TCP	443	HTTPS	GigaVUE-FM Copilot	GigaVUE-FM	Allows GigaVUE-FM Copilot to reach GigaVUE-FM only to provide

Direction	Protocol	Port Number	Service	Source CIDR	Destination	Purpose
						contextual information for the queries.
Outbound	TCP	443	HTTPS	GigaVUE-FM Copilot	LLM Service Providers (Amazon Bedrock): https://docs.aws.amazon.com/general/latest/gr/bedrock.html	Allows GigaVUE-FM Copilot to reach Amazon Bedrock for LLM access.
Outbound	TCP	443	HTTPS	GigaVUE-FM Copilot	Azure: - login.microsoftonline.com - scpladsl.blob.core.windows.net	Allows GigaVUE-FM Copilot to export the telemetry data.

Unsupported Features

GigaVUE-FM Copilot does not currently support the following features:

- Configuration creation/modifications to GigaVUE-FM or managed devices
- Automated validation of GigaVUE-FM configuration against documentation-based insights derived from a user query
- Use of operational data sources such as audit logs, licensing data, GigaVUE-FM system logs, device logs, and selected GigaVUE-FM system or admin settings
- Subscriber Intelligence, Flex Inline, Tunnel Unification, Map templates, Map Rules, UCT-C configurations, Fabric Health Analytics (FHA), Fabric Maps, Certificates functionalities or features are not supported.
- For Application Intelligence traffic related questions, only configuration and health queries for Application Intelligence are supported.

GigaVUE-FM Copilot with LLM Connectivity Options

GigaVUE-FM Copilot supports an external LLM deployment model with the following characteristics:

- GigaVUE-FM Copilot supports only the gpt-oss-20b model from Amazon Bedrock.
- Requires customers to provision Amazon Bedrock in their own AWS account.

- Requires customers to provide the corresponding Amazon Bedrock API keys to GigaVUE-FM Copilot during configuration.

Each option uses the same external LLM (gpt-oss-20b) but differs in where you deploy GigaVUE-FM Copilot and how you connect to Amazon Bedrock.

The following deployment options are supported for connecting to Amazon Bedrock:

- [GigaVUE-FM Copilot with Public LLM Connectivity \(Amazon Bedrock over Internet\)](#)
- [GigaVUE-FM Copilot with Private LLM Connectivity \(AWS Bedrock via PrivateLink\)](#)
- [GigaVUE-FM Copilot with Local LLM Connectivity \(AWS Native Deployment with PrivateLink\)](#)

GigaVUE-FM Copilot with Public LLM Connectivity (Amazon Bedrock over Internet)

In this option, you deploy GigaVUE-FM Copilot on-premises as a VM on VMware ESXi. The GigaVUE-FM Copilot connects to Amazon Bedrock over the public internet using secure HTTPS to access Large Language Model (LLM) features.

How it works

GigaVUE-FM Copilot starts an outbound HTTPS session to Amazon Bedrock for all LLM requests. You do not need inbound access from the internet. The setup does not use a private link or a cloud-hosted GigaVUE-FM instance.

What You Manage


You are responsible for:

- [Provisioning Amazon Bedrock in your AWS account.](#)
- [Creating API keys and entering them during GigaVUE-FM setup.](#)
- Allowing outbound HTTPS to the Amazon Bedrock API endpoint in your selected region.

When to Choose this Option

Choose this option to:

- Reduce cost (no GPU or special hardware).
- Provide predictable behavior.
- Simplify networking with only outbound internet access.



GigaVUE-FM Copilot with Private LLM Connectivity (AWS Bedrock via PrivateLink)

In this option, you deploy GigaVUE-FM Copilot on-premises as a VM on VMware ESXi without GPU requirements. The GigaVUE-FM Copilot connects to AWS Bedrock through a private network path using your existing VPN or AWS Direct Connect link to your AWS VPC. This avoids public internet endpoints.

How it works

The GigaVUE-FM Copilot sends all LLM requests to Amazon Bedrock through your private network. The connection uses your VPN or Direct Connect link to your AWS environment. The GigaVUE-FM Copilot does not require public internet access for LLM features.

What You Manage


You are responsible for:

- [Provisioning Amazon Bedrock in your AWS account.](#)
- [Creating API keys and entering them during GigaVUE-FM setup.](#)
- Ensuring private connectivity to the Amazon Bedrock API endpoint in your selected region.

When to Choose this Option

Choose this option when:

- You already have AWS workloads running in your VPC.
- Your security policy restricts the use of public API endpoints.
- You prefer all LLM traffic to stay within a controlled, private path.



GigaVUE-FM Copilot with Local LLM Connectivity (AWS Native Deployment with PrivateLink)

In this option, you deploy GigaVUE-FM Copilot directly in AWS as a cloud-native appliance. The VM does not require a GPU. GigaVUE-FM Copilot connects Amazon AWS Bedrock through an AWS PrivateLink endpoint, ensuring that all LLM traffic stays within the AWS network and never uses public internet paths.

How it works

You deploy the GigaVUE-FM Copilot in your AWS account. GigaVUE-FM Copilot sends all LLM requests to Amazon Bedrock through PrivateLink, keeping traffic on the AWS backbone. The deployment uses AWS-native networking and does not require on-prem infrastructure.

What You Manage

You are responsible for:


- [Provisioning Amazon Bedrock in your AWS account.](#)
- [Creating API keys and entering them during GigaVUE-FM setup.](#)
- Ensuring private connectivity to the AWS Bedrock API endpoint in your selected region.

When to Choose this Option

Choose this option when you:

- Prefer a cloud-first deployment strategy.

- Have challenges or restrictions that prevent deploying VMs on-premises.
- Want a simple, fully cloud-based solution without local hardware.



Deploy GigaVUE-FM Copilot on VMware vCenter

Before you begin:

- Ensure that you have installed VMware vSphere Standard, Enterprise, or Enterprise Plus on compatible hardware. Refer to [Hardware Requirements](#) for the minimum hardware requirements and [Open Ports Access Requirements](#) for port numbers.
- Make sure the VMware ESXi host and the GigaVUE-FM Copilot VM are time-synchronized. If the ESXi host time is incorrect, the GigaVUE-FM Copilot VM inherits it, which can cause TLS or signature errors when GigaVUE-FM Copilot connects to services such as Amazon Bedrock and prevent successful registration in GigaVUE-FM.

The GigaVUE-FM Copilot software package is distributed as an OVA file. You can download the "gigamon-gigavue-copilot-6.12.03.ova" file from the [VUE Community](#).

Use the vSphere Client to install the GigaVUE-FM Copilot OVA file. You cannot deploy GigaVUE-FM Copilot directly from the ESXi host. You must login to the vCenter on the vSphere client to deploy a GigaVUE-FM Copilot instance.

Deploy GigaVUE-FM Copilot on VMware vCenter

GigaVUE-FM Copilot with Local LLM Connectivity (AWS Native Deployment with PrivateLink)

IMPORTANT NOTE:



The OVA file must be stored in a location that is accessible to the vSphere Client. It cannot be stored on a datastore that is accessible to the ESXi host designated for the deployment.

To deploy a GigaVUE-FM Copilot using OVA file:

1. Log in to the VMware vCenter web interface.
2. In the vSphere Client, select an inventory object that is a valid parent object of a virtual machine, such as a data center, cluster, or ESXi host.
3. Right-click the ESXi Host, Cluster, or data center on which you want to deploy GigaVUE-FM Copilot, and then select **Deploy OVF Template**.

The **Deploy OVF Template** wizard opens.

4. In the **Select OVF Template**, choose one of the following options:
 - **URL**—Enter the URL from where you want to download and install the OVF package.
 - **Local file**—Click **Browse** to navigate to the OVA file available on your local machine, and then select the OVA file.
5. Click **Next**. The **Select a name and folder** page of the **Deploy OVF Template** wizard appears. Specify a unique name for the GigaVUE-FM Copilot instance, and then select a location and host to which you want to deploy the GigaVUE-FM Copilot instance.
6. Click **Next**. The **Select a compute resource** page of the **Deploy OVF Template** wizard appears. Select a destination compute host for the OVF deployment. The Deploy OVF Template wizard performs validation to ensure that the selected host has all the resources required for the GigaVUE-FM Copilot deployment.
7. Click **Next**. The **Review details** page of the **Deploy OVF Template** wizard appears. Verify the OVF template details.
8. Click **Next**. The **Select storage** page of the **Deploy OVF Template** wizard appears. In this page:
 - a. Define where and how to store the files for the deployed OVA template.
 - b. Select the datastore where the virtual machine files will be stored.
9. Click **Next**. The **Select networks** page of the **Deploy OVF Template** wizard appears.
 - a. Select a network that provides the necessary connectivity for the GigaVUE-FM Copilot VM, including management access and any required data-plane communication.
 - b. Ensure the IP Protocol drop-down remains at its default value. Altering this setting may result in vCenter-related issues.

10. Click **Next**. The **Customize template** page of the **Deploy OVF Template** wizard appears. Customize the deployment properties:
 - a. Configure the hostname and, if required, the domain name for the appliance. These settings specify how the system is identified on the network.
 - b. **IPv4 and IPv6 Networking Configuration**: Use this section to configure the IP settings for the management interface.
 - c. From the Management Port Connection drop-down menu, select one of the following options:
 - **DHCP**: The IP address is automatically assigned by the DHCP server.
 - **Static**: Enter the **IP Address**, **Subnet Mask**, **Gateway** and **DNS Server**.
 - **Disabled**: Disables the selected IP protocol (IPv4 or IPv6).
11. Click **Next**. The **Ready to Complete** page of the **Deploy OVF Template** wizard appears.
 - a. Verify that all of the settings are correct, and then click **Finish**.
 - b. The **Recent Tasks** pane of the vSphere Client Home page shows the progress of the deployment operation. When the operation is complete, you will successfully deploy a GigaVUE-FM Copilot instance.

Deploy OVF Template

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

Ready to complete

▼ Select storage

Size on disk	254.0 GB
Storage mapping	1
All disks	Datastore: datastore1 (6); Format: Thick provision lazy zeroed

▼ Select networks

Network mapping	1
management	VM Network

IP allocation settings

IP protocol	IPv6
IP allocation	Static - Manual

▼ Customize template

Properties	01. System Hostname = FM-COPILOT-61201 02. System Domain Name = 01. Management Port IPv4 Connection = Static 02. Management Port IPv4 Address = 10.115.203.27 03. Management Port IPv4 Netmask = 255.255.248.0 04. Management Port IPv4 Gateway = 10.115.200.1 05. IPv4 Domain Name Server = 10.10.1.20 01. Management Port IPv6 Connection = Disabled 02. Management Port IPv6 Address = ::1 03. Management Port IPv6 Prefix = 128 04. Management Port IPv6 Gateway = ::1 05. IPv6 Domain Name Server = ::1
------------	---

CANCEL
BACK
FINISH

12. Select the newly created VM and click **Power on**.

13. Open the VM console to monitor the boot process. It may take approximately 15-20 minutes for the VM to boot.



Note:

When you SSH into the appliance for the first time, use the following default credentials:

- Username: admin
- Password: gigamon123A!!

For security reasons, you are prompted to change the password immediately after your first login. We strongly recommend that you change the default password to secure the appliance.

Deploy GigaVUE-FM Copilot on AWS

To subscribe to the GigaVUE-FM Copilot, perform the following steps:

1. Log in to your AWS account.
2. Go to AWS Marketplace: <https://aws.amazon.com/marketplace>.
3. In the **Search** field, type "GigaVUE Fabric Manager Copilot" and select **Search**.
4. Select **View Purchase Options**. The terms and conditions page is displayed.
5. Review the Terms and Conditions and select **Accept Terms**.
6. Review the summary and then select **Continue to Configuration**.
7. In the **Configure this software** page, enter the following details for your deployment:
 - a. Set **Fulfillment Option** to the default value.
 - b. Select the latest version in the **Software Version** field.
 - c. Choose your deployment **Region**.
 - d. Select **Continue to Launch**.
8. In the **Configure this software** page, perform the following:
 - a. In the **Choose Action** field, select the **Launch from Website** option.
 - b. From the **EC2 Instance Type** drop-down list, select the **m5.xlarge** instance type. Refer to the [Recommended and Supported Instance Types for AWS](#).

- c. Under **Key pair**, choose **Create new key pair** or select an existing one. Download the .pem or .ppk file if creating a new key pair.

NOTE: You cannot download the key again later.

- d. From the **VPC Settings** drop-down list, select the VPC for deploying GigaVUE-FM Copilot.
- e. In the **Subnet Settings**, select your desired Subnet.
- f. Disable the Auto-assign Public IP. We recommend using a private IP address for this instance.
- g. In the **Security Group Settings**, configure the security group to match your access and permissions needs. For details, refer to [Security Group](#). Enable **Allow SSH traffic from** (required to generate the API key) and **Allow HTTPS traffic from the internet** (used for communication between GigaVUE-FM and GigaVUE-FM Copilot).
- h. In the **Advanced details** section, from the **Metadata version** dropdown list, select **V2 only (token required)**.



Important:

When you SSH into the appliance for the first time, use the following default credentials:

- Username: admin
- Password: gigamon123A!!

For security reasons, you are prompted to change the password immediately after your first login. We strongly recommend that you change the default password to secure the appliance.

Configure GigaVUE-FM Copilot in GigaVUE-FM

After you deploy the GigaVUE-FM Copilot VM, configure GigaVUE-FM Copilot in GigaVUE-FM. Before you begin, make sure you have acquired:

- [Acquire Amazon Bedrock Credentials](#)
- Generate the GigaVUE-FM Copilot API key for your deployment type:
 - [Generate the API key from the GigaVUE-FM Copilot VM \(VMware ESXi\)](#)
 - [Generate the API key from the GigaVUE-FM Copilot VM \(AWS\)](#)

Generate the API key from the GigaVUE-FM Copilot VM (VMware ESXi)

After the GigaVUE-FM Copilot VM installation, generate the GigaVUE-FM Copilot API key from the GigaVUE-FM Copilot VM by following the steps given:

1. SSH into the GigaVUE-FM Copilot VM:

```
ssh <username>@<GigaVUE-FM Copilot-ip-or-hostname>
```
2. When prompted, enter the default password: gigamon123A!!
3. You are prompted to change the password immediately after your first login. Enter and confirm a new, secure password.
4. After updating your password, run the following command to retrieve the API key:

```
copilotctl secrets get copilot --api-key
```
5. Make a note of the API key. Use this API key to enable GigaVUE-FM Copilot in GigaVUE-FM.

Generate the API key from the GigaVUE-FM Copilot VM (AWS)

After the GigaVUE-FM Copilot installation, generate the GigaVUE-FM Copilot API key from the GigaVUE-FM Copilot VM by following the steps given:

1. SSH into the GigaVUE-FM Copilot VM:

```
ssh -i ~<path to copilot-aws-ssh-key.pem><username>@<GigaVUE-FM Copilot-ip-or-hostname>
```
2. When the terminal prompts, type "yes" to continue connecting.
3. Run the following command to retrieve the API key:

```
copilotctl secrets get copilot --api-key
```
4. Make a note of the API key. Use this API key to enable GigaVUE-FM Copilot in GigaVUE-FM.

Acquire Amazon Bedrock Credentials

To acquire Amazon Bedrock credentials, perform the following steps:

Create IAM Policy

1. Log in to <https://aws.amazon.com/console>.
2. Navigate to IAM > Access Management > Policies and click **Create policy**.

3. Select an AWS service: **Bedrock**.
4. Filter Actions allowed: **Invoke**.
5. Select **InvokeModel** and **InvokeModelWithResponseStream**.
6. Under Resources, add ARNs to grant access to both the foundation model and the inference profile.
7. Click Add ARNs and add the foundation-model ARN:
 - Resource Region: Select Region (or any)
 - Resource resource: **openai.gpt-oss-20b-1:0**
8. To grant inference profile access, click **Add permissions**.
9. Select **Bedrock** as the AWS service.
10. Filter Actions by **Invoke**, and select: **InvokeModel** and **InvokeModelWithResponseStream**
 - a. Under Resources, add the inference-profile ARN.
 - Resource Region: Select Region (or any)
 - Resource resource: **openai.gpt-oss-20b-1:0**
 - b. Go to Request conditions > **Add another condition** for the inference-profile ARN.
 - Condition key: **aws:ResourceAccount**
 - Operator: **StringEquals**
 - Value: **\${aws:PrincipalAccount}**
 - c. Then click **Add Condition**, followed by **Next**
11. Under Policy details, provide a Policy name: fm-copilot-llm-access (example).
12. Review all permissions and click **Create Policy** to save your new policy.

Your IAM policy is created and ready to be attached to the IAM user.

The following is an example JSON template after the required permissions are configured.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "AllowInferenceProfileAccess",
      "Effect": "Allow",
      "Action": [
        "bedrock:InvokeModel",
        "bedrock:InvokeModelWithResponseStream"
      ],
      "Resource": [
        "arn:aws:bedrock:*:*:inference-profile/openai.gpt-oss-20b-1:0"
      ]
    }
  ]
}
```

```

    ],
    "Condition": {
      "StringEquals": {
        "aws:ResourceAccount": "${aws:PrincipalAccount}"
      }
    }
  },
  {
    "Sid": "AllowFoundationModelAccess",
    "Effect": "Allow",
    "Action": [
      "bedrock:InvokeModel",
      "bedrock:InvokeModelWithResponseStream"
    ],
    "Resource": [
      "arn:aws:bedrock:*::foundation-model/openai.gpt-oss-20b-1:0"
    ]
  }
]
}

```

Create IAM User

1. In the search bar of the AWS Management Console, type IAM, and then select IAM (Identity and Access Management).
2. In the left navigation pane, under Access Management, click **Create Users**.
3. Enter the following information:
 - a. User name: fm-copilot-llm-user (example)
 - b. Under Permission options, select **Attach policies directly**. In the search bar, type the name of your customer-managed policy (the one you created earlier).
 - c. Select the policy check box.
 - d. Select permissions: fm-copilot-llm-access.
4. Click **Next**, then **Create User**. Search for user and click on the returned name.
 - a. Under Summary, click on **Create access key**.
 - b. For Use case select: **Other**.
 - c. Click **Next**. Provide a Description tag value: fm-copilot-key.
 - d. Click **Create access key**, your access key and secret access key is generated. Make a note of these access keys to configure them to enable this LLM for Gigamon Insights.

Enable GigaVUE-FM Copilot in GigaVUE-FM


Follow the steps below to enable GigaVUE-FM Copilot in GigaVUE-FM:

1. In the left navigation pane, go to **Settings > System > GigaVUE-FM Copilot**.
2. On the Configuration tab, click **Enabled** to enable the GigaVUE-FM Copilot.
3. Confirm your terms with the Gigamon AI Usage Telemetry Data Collection & Usage Notice.

NOTE: In this LA release, accepting AI usage telemetry is required to use GigaVUE-FM Copilot. If you do not accept the terms of use, you cannot enable this feature.

4. [Optional] Enter your Customer ID when you accept the terms and conditions.
5. In the **Server URL** field, enter the GigaVUE-FM Copilot URL:
`https://<copilot-ip-or-hostname>.`
6. In the **API Key** field, enter the API key you generated earlier.
7. You can edit the **Server URL** and **API Key** fields at any time.
8. In the External LLM Provider section, enter the following details:
 - Region - Select the LLM provider region that is geographically closest to the GigaVUE-FM Copilot appliance to minimize network latency and improve response time
 - Access Key Id - Provide the access key acquired from the Amazon Console.
 - Secret Access Key - Provide the Secret access key acquired from the Amazon Console.
9. Click **Apply** to enable the GigaVUE-FM Copilot.

To access GigaVUE-FM Copilot, ensure that the AI Assistant icon appears in the left pane, below the **Settings** icon. Ensure the configuration is correct by clicking on the GigaVUE-FM Copilot icon, asking a question, and getting a response. That's it, configuration complete!



Get Started with GigaVUE-FM Copilot

The following section provides instructions on how to start using GigaVUE-FM Copilot and using prompts effectively.

Launch GigaVUE-FM Copilot

To start using the GigaVUE-FM Copilot:

1. Click the  GigaVUE-FM Copilot icon at the bottom left of any page.
2. Confirm your agreement to the usage terms, and activate the feature.
3. Start typing a question or request and GigaVUE-FM Copilot provides:
 - Step-by-step instructions based on your question. For example, where to click, what to select.
 - If you request a CLI, it shares the available commands.
 - Links to exact sections in the documentation.
4. Ask follow-up questions to narrow the results. For example, Show only prerequisites and open ports for VMware.

5. Select the thumbs-up or thumbs-down button to share feedback on responses to help improve the accuracy and relevance of GigaVUE-FM Copilot.

How to Use Prompts Successfully

Follow these guidelines to get clear answers, thorough instructions, and helpful tips tailored to your specific context.

Be clear and specific

Reduces ambiguity and improves answer quality. To achieve the most reliable results, specify the scope of your query at the beginning. We recommend that you include aliases of ports, devices, maps and so on in your prompts for better results.

Do	Don't
"Show me total system memory metrics of this GigaVUE-FM"	"Tell me total memory"
"How many devices are managed by this GigaVUE-FM?"	"How many devices are there?"
"What is the device <device alias> upgrade status"	"What is my upgrade status?"
"List the documented transceiver support on HC and TA devices"	"List the transceiver support on HC and TA devices"

NOTE: To achieve the most reliable results, specify the scope of your query at the beginning.

Start simple, and then improve

Start with an initial prompt, review the response, and refine by adding context or narrowing scope.

Initial Query	Result	Refined Query
"Show me alarms in this GigaVUE-FM"	Lists all alarms.	"Show me open or correlated alarms in GigaVUE-FM"
"How many devices are up?"	Lists devices that are included in cluster and fabric nodes.	"How many managed devices are unhealthy in this GigaVUE-FM?"

Avoid ambiguity

The term "node" can mean a physical device, a fabric node, or GigaVUE-FM itself. Specify the entity type: "physical device," "fabric node," or "V Series node."

Do	Don't
"How is the physical device <device alias> performing?"	"How is my node <device alias> performing?"

Alias Handling with Numeric/Common Terms

Ask targeted questions using a single alias at a time and avoid combining multiple aliases in a single query.

Do	Don't
"What are the maps affected by the port change 1/1/x3?"	"I modified the port parameters on 1/1/x3 — will this affect my maps map123 and map234?"

Sample Prompts

Use the following sample prompts categorized by specific functionalities. You can type these questions as-is or refine them to match your environment:

- **Sample Prompts for General GigaVUE-FM status**
 - Tell me about the GigaVUE-FM.
 - Is the NTP server in this GigaVUE-FM in sync?
 - Show me the active users logged in to my GigaVUE-FM.
 - Is GigaVUE-FM configuration backup working?
 - What is the GigaVUE-FM upgrade status?
 - What is the status of my most recent GigaVUE-FM upgrade and when was it performed?
 - When was the FM/FMHA backup taken before the last FM upgrade?
 - What is the GigaVUE-FM uptime?
 - Show me the GigaVUE-FM image name?
 - What is the GigaVUE-FM version and build information?
 - How many devices with the version 6.11.00 are managed by this GigaVUE-FM?
- **Sample Prompts for Users, groups, and roles**
 - Show me the list of user groups configured in this GigaVUE-FM?
 - Show me the users part of the user group <user group> ?
 - Are any users currently locked out of GigaVUE-FM?
 - How many users are accessing the GigaVUE-FM now?
 - List the last three active users logged into GigaVUE-FM?
- **Sample Prompts for Cluster and connectivity**
 - How are the devices connected in cluster <cluster name>?
 - What are the network ports connected to tool port 1/1/x1?
 - What are the maps that are contributing to tool port 1/1/x2?
 - How many maps are unhealthy in host <name>?

- Show me the maps in the host <name>?
- What are the maps disabled in cluster <name>?
- How many ports are available in my cluster <name>?
- **Sample Prompts for Ports and devices overview**
 - Summarize port speeds configured in my GigaVUE-FM.
 - Can you summarize the ports in my system?
 - Can you summarize devices managed in my system by model?
 - Summarize GigaSMART cards by their health.
 - Show me the port-pair configuration on host <name>?
 - What are the TA node names by this GigaVUE-FM?
 - Show me the ports which are down in host <name>?
- **Sample Prompts for Application Intelligence and monitoring**
 - Show me the Application Intelligence solutions configured in my GigaVUE-FM.
 - Why is the HCI-Solution1 Application Intelligence solution unhealthy?
 - To which tool is the above solution exporting traffic?
 - List all Application Intelligence solutions deployed in GigaVUE-FM.
 - List connections in my system/monitoring domains.
 - Summarize flow maps by health/type (first level/second level).
- **Sample Prompts for Monitoring sessions, tunnels, and applications**
 - What are the fabrics deployed in my vCenter <IP address/hostname> and what is their status?
 - What are the tunnels configured in monitoring session <MS1>?
 - What are the applications configured in monitoring session <MS3>?
 - List the monitoring sessions enabled with dedup.
 - What are the agents present in my monitoring domain <MD>?
- **Sample Prompts for Tool port utilization and drops**
 - Show me the tool ports that are used more than 50%.
 - Are there any tool ports dropping traffic?
 - Show me the port total capacity of node <node name>.
 - Show me the Rx and Tx rates of tool ports in host <name>.
- **Sample Prompts for Alarms and system health**
 - Provide the open alarms that show why port <port number> is unhealthy.
 - How is my GigaVUE-FM performing on system metrics?
 - Provide the unacknowledged alarms for cluster <cluster name>.
 - Show the critical alarms for ports <port list> or <port name>.
 - Show me the system health alarms.
 - How many alarms are currently active?
 - Show me all critical active alarms.
 - Show me the alarms.
 - Provide the issues in the power modules in cluster <cluster name>.
- **Sample Prompts for Specific issues and how to fix them**

- Why is port <port number> down?
- How do I fix the device power module alarms?
- What are the issues in the GigaStream on cluster <cluster name>?
- Why is my GigaStream <alias> not working?
- Why is my vSeries <name> down?
- Why is port <name> dropping traffic?
- Why is traffic not being sent to my tool on port 3/1/x1?
- Why did my last GigaVUE-FM backup fail?
- Why did my last GigaVUE-FM upgrade fail?
- Why user <name> cannot login to GigaVUE-FM?

Debuggability and Troubleshooting

This section lists the common issues that may occur in the GigaVUE-FM Copilot environment and provides the recommended steps to troubleshoot them.

Generate GigaVUE-FM Copilot Sysdump in GigaVUE-FM

To generate a sysdump file:

1. On the left navigation pane, go to **Settings > System > GigaVUE-FM Copilot**.
2. In the **Health** tab, click **Generate** to generate a sysdump file.

GigaVUE-FM Copilot allows you to generate up to 5 sysdump files. If you try to generate more than five:

- GigaVUE-FM prompts you to delete the oldest sysdump file.
- After you confirm, GigaVUE-FM deletes the oldest file and generates the new sysdump file.

Download GigaVUE-FM Copilot Sysdump file

To download a sysdump file:

1. Select the file that you want to download in the **GigaVUE-FM Copilot System Logs** page.
2. Click the **Actions** drop-down list and select **Download**.

Delete the GigaVUE-FM Copilot Sysdump file

To delete a sysdump file:

1. Select the file that you want to delete in the **GigaVUE-FM Copilot System Logs** page.
2. Click the **Actions** drop-down list and select **Delete**.

NOTE: To delete all the sysdump files, click the **Actions** drop-down list and select **DeleteAll**. You do not need to select the files to delete.

Troubleshoot GigaVUE-FM Copilot

Refer to the following sections for detailed procedures:

Resolve UI errors, timeouts, or unresponsiveness

Problem description: The GigaVUE-FM Copilot UI may become unresponsive, load slowly, or time out.

Corrective Action:

1. Generate sysdump files. Refer to the [Generate GigaVUE-FM Copilot Sysdump in GigaVUE-FM](#)
2. Collect the sysdump files and share them with the Gigamon support team. Refer to the instructions in [Contacting Technical Support](#) to send the logs to the support team.

GigaVUE-FM Copilot returns unexpected or incorrect answers

Problem description: GigaVUE-FM Copilot may return unexpected, incomplete, or incorrect answers, or display UI errors when you submit queries.

Corrective Action:

1. Open the VM console or connect to the GigaVUE-FM Copilot VM over SSH.
2. Log in as admin with the updated password.
3. Go to the logs directory:

```
/cp-data/fm-copilot/logs
```

4. Identify and collect log files related to the time when the issue occurred.

5. Send the logs to the Gigamon support team. Refer to the instructions in [Contacting Technical Support](#) to send the logs to the support team.

GigaVUE-FM Copilot VM does not initialize correctly

Problem description: After deployment, the GigaVUE-FM Copilot UI may not load after several minutes, or initialization may appear stuck.

Corrective Action:

1. Open a console to the GigaVUE-FM Copilot VM from your hypervisor, or connect over SSH.
2. Log in as admin with the updated password.
3. Run the following command to view initialization logs:

```
sudo journalctl -u fm-copilot-init -f
```

This command displays logs for the GigaVUE-FM Copilot Init Service.

4. Confirm that the “GigaVUE-FM Copilot Init Service” started successfully. If you see repeated errors or the service does not start, capture the log output and proceed to log collection.
5. If errors occur or initialization fails :
 - a. Capture the console log output.
 - b. Generate and download the GigaVUE-FM Copilot sysdump file.
 - c. Send the logs to the Gigamon support team. Refer to the instructions in [Contacting Technical Support](#) to send the logs to the support team.

AI Telemetry Usage

For the LA release of GigaVUE-FM Copilot, you must enable AI Usage Telemetry. This will periodically share non-sensitive AI usage data with Gigamon for the purpose of improving AI quality. Gigamon is committed to security and privacy. Any Personally Identifiable Information (PII) is sanitized and no data is shared with third-parties. All possible security measures are also implemented to protect and isolate the data collected.

Telemetry collection is designed to minimize data storage, exclude AI-generated responses, and protect user privacy by automatically removing personally identifiable information (PII).

The system collects the following categories of telemetry.

User Prompts (after PII Scrubbing)

The text you enter as a prompt is processed by an automated PII scrubbing pipeline, and any sensitive or identifiable information is removed before it is used. Only the sanitized version of the prompt is stored. PII scrubbing process detects and removes or masks commonly sensitive identifiers, including:

- IP addresses
- MAC addresses
- Email addresses
- Phone numbers
- Government identifiers (SSN)
- Credit card numbers
- Street/mailing addresses

Detected PII elements (for example, IP addresses or email addresses) are replaced with generic placeholders such as <IP>, <MAC_ADDRESS>, or <EMAIL_ADDRESS>.

For example,

Original Prompt: Show me all devices and interfaces in FM that are using IP address 192.168.100.200 and summarize their health status.

Stored Telemetry Version: Show me all devices and interfaces in FM that are using IP address <IP> and summarize their health status.

Operational Metadata

Operational metadata describes how each AI request is processed and includes:

- Message ID – Unique identifier for the AI request.
- Conversation ID – Identifier for the chat session.
- Parent Message ID – Identifier for the preceding message.
- Tenant ID/Tenant code – Unique deployment identifiers assigned during setup.
- Timestamps – When the message was created and last updated.
- Provider identifier – AI or LLM provider (for example, Bedrock).
- Model identifier – Model used for inference.
- Token counts – Number of input (prompt), output (completion), and total tokens.

Operational metadata and the scrubbed prompt text are stored as structured JSON records.

User Feedback

If you choose to rate or comment on an AI response, the following optional telemetry may be collected:

- User rating – Thumbs up or thumbs down.
- Optional comment – Free text feedback about the response.

Feedback comments also undergo PII scrubbing before storage. Feedback telemetry is used to identify areas where AI responses can be improved.