Gigamon Visibility Platform for AWS Configuration Guide

Version 5.2

DOCUMENT REVISION – 1/30/18

Contents

About This Guide

This guide describes how to configure the components of the Gigamon Visibility Platform for Amazon Web Services (AWS). Use this document for instructions on configuring the Visibility Platform components and setting up the traffic monitoring sessions for AWS.

Audience

This guide is intended for users who have basic understanding of AWS. This document expects users to be familiar with Elastic Compute Cloud (EC2) Instances, subnets, Elastic Network Interfaces (ENIs), Identity and Access Management (IAM) and Access Keys (basic credentials), Security Group, and Key Pairs.

For information on AWS terminologies used in this document, refer to *Glossary* on page 119.

Licensing Information

The Gigamon Visibility Platform is available in both the public AWS cloud and in AWS GovCloud, and supports the Bring Your Own License (BYOL) model and the hourly Pay-As-You-Go (PAYG) model.

Bring Your Own License (BYOL)

The AMI for the BYOL option can be purchased based on the number of TAP points and the term of the license. Gigamon offers the following options for purchasing the license:

- Traffic visibility for up to 100 virtual TAP points (ENIs)
- Traffic visibility for up to 1000 virtual TAP points (ENIs)

NOTE: Make sure you purchase a licensing option that can provide traffic visibility to all the TAP points in the VPC. If the licensing option you have purchased cannot support all the TAP points, the ENIs are selected randomly for monitoring the traffic.

The minimum term for the license is 3 months and the maximum term is up to 12 months.

A free trial is made available in the AWS Marketplace and in the Community AMIs. The trial version provides traffic visibility for up to 10 virtual TAP points for 30 days. When a new license is purchased, the 10 virtual TAP points are replaced with how many ever TAP points the licensing option supports.

For purchasing licenses with the BYOL option, contact our Gigamon Sales. Refer to *Contacting Sales* on page 124.

Pay-As-You-Go (PAYG)

Starting in GigaVUE-FM 5.2, the AMI for the Pay-As-You-Go (PAYG) option is available in the AWS Marketplace. The hourly PAYG option charges the users for the AWS services availed on an hourly basis. For example, AWS charges the users for the period the GigaVUE-FM instance is running in the EC2 instances. When the instance stops, AWS stops charging the users. The PAYG model has no term contract.

It is a perpetual license that supports up to 100 TAP points. To support additional TAP points, a new license must be purchased from Gigamon.

NOTE: While upgrading GigaVUE-FM, make sure you choose the AMI with the same licensing option as the current AMI. For example, assume that a user has purchased GFM-AWS-100 license with hourly pricing. While upgrading GigaVUE-FM, the user must select the AMI with the same GFM-AWS-100 license associated. Else, there could be discrepancy in the number of instances monitored.

For purchasing licenses with the PAYG option, contact the Gigamon Sales. Refer to *Contacting Sales* on page 124.

Applying Licensing

After obtaining the license, use the information sent to you by Gigamon to generate the license keys.

To generate the license keys:

1. In the Email received from Gigamon, copy one or more Gigamon Installation Keys (**GIK**).
2. Locate the MAC address of the virtual network adapter. The license is only valid with the corresponding MAC address.
3. Go to https://licensing.gigamon.com to generate GIK.

4. In the Generate License page, enter the appropriate information. Multiple GIKs can be entered by clicking the + button.

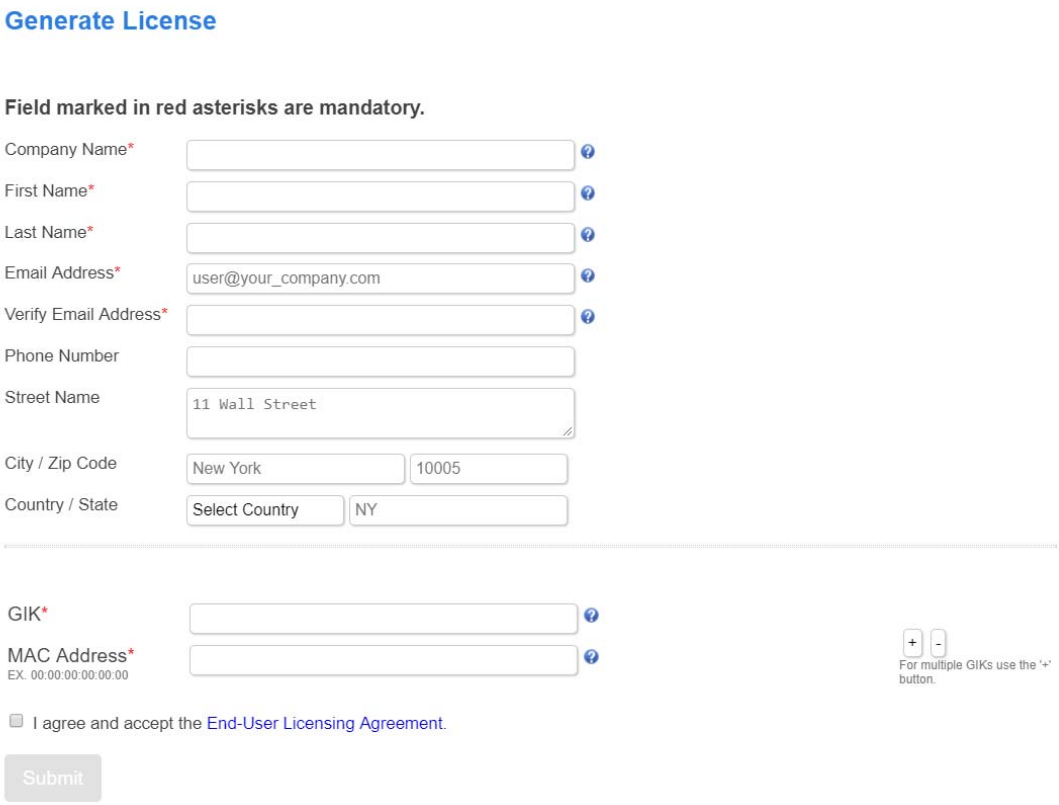

5. Select the **I agree and accept the End-User Licensing Agreement** check box and click **Submit**. The license keys are generated.
6. Copy the license keys into a Notepad.
7. Launch the GigaVUE-FM instance. For information, refer to *Launching the GigaVUE-FM Instances* on page 24.
8. After launching the GigaVUE-FM instance, log in to GigaVUE-FM.
9. In the GigaVUE-FM instance, go to **Administration > System > License** page.

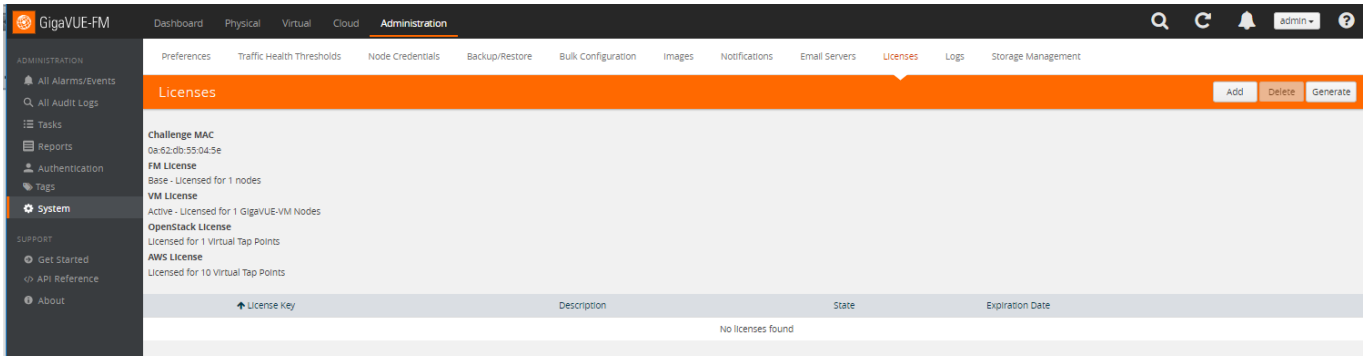

10. Click **Add** and enter the license key or keys copied in step 6 into the Add License box, and then click **Save** to apply the license.

Installing and Upgrading GigaVUE Fabric Manager

You can install and upgrade the GigaVUE® Fabric Manager (GigaVUE-FM) on cloud or on-premise.

- Cloud—To install and upgrade GigaVUE-FM inside your AWS environment, you can simply launch the GigaVUE-FM instance in your VPC. For installing the GigaVUE-FM instance, refer to *Configuring the Components in AWS* on page 17. For upgrading the GigaVUE-FM instance, refer to *Upgrading the GigaVUE-FM Instance* on page 107.
- On-premise—To install and upgrade GigaVUE-FM in your enterprise data center, refer to *GigaVUE-FM and GigaVUE-VM User's Guide* available in the Customer Portal.

Chapter 1

Overview

This chapter introduces the components of Gigamon Visibility Platform for AWS and the supported architecture. Refer to the following sections for details:

Introduction to Gigamon Visibility Platform for AWS

The Gigamon Visibility Platform for AWS provides consistent visibility into data in motion across the entire enterprise: on-premise, remote sites, private, hybrid, and public clouds.

This platform integrates with the AWS API endpoints, deploys its components, mirrors the application traffic, and replicates the customized traffic to network and security tools that reside on cloud or on-premise.

The Visibility Platform for AWS offers the following benefits:

- Provides elastic and deep visibility into data traversing through AWS
- Provides pervasive visibility into EC2 instances within hybrid and public clouds
- Enables a consistent way to access, categorize, and consolidate the delivery of network traffic to tools
- Delivers intelligent traffic to network and security tools on cloud or on-premise using Flow Mapping™, slicing, masking, and sampling
- Provides open APIs for integration, orchestration, and automation

Gigamon Visibility Platform Components

Figure 1-1 on page 12 shows how the Gigamon Visibility Platform components for AWS are configured within a VPC.

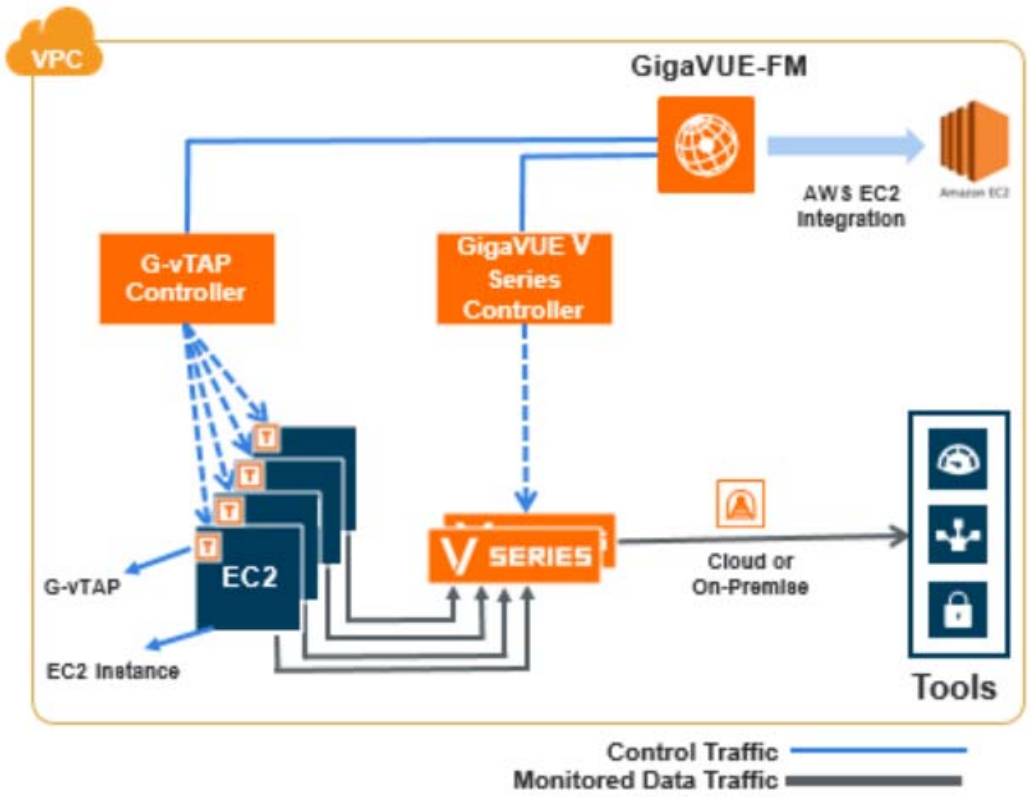

Figure 1-1: Gigamon Visibility Solution Components in AWS

The Visibility Platform for AWS includes the following components:

- **GigaVUE® Fabric Manager (GigaVUE-FM)** is a web-based fabric management interface that provides a single pane of glass visibility and management of both the physical and virtual traffic that form the Gigamon Visibility Platform.

  GigaVUE-FM can be installed on-premise or launched as an Amazon Machine Image (AMI) in AWS. GigaVUE-FM manages the configuration of the following visibility components in your Amazon Virtual Private Clouds (VPC):

  - GigaVUE V Series nodes
  - G-vTAP Controllers
  - GigaVUE V Series Controllers

  To launch the AMI in AWS, refer to *Obtaining the AMI* on page 24 and *Launching the GigaVUE-FM Instances* on page 24.

  To install GigaVUE-FM on premise, refer to *GigaVUE-FM and GigaVUE-VM User's Guide* available in the Customer Portal.

- **G-vTAP agent** is an agent that is deployed in the Elastic Compute Cloud (EC2) instance. This agent mirrors the selected traffic from the instances (virtual machines) to the GigaVUE® V Series node.

The G-vTAP agent is offered as a Debian or Redhat Package Manager (RPM) package. You can download the G-vTAP agent from the Customer Portal > **Software Download** page and install it on your instance (virtual machine). Refer to *Installing the G-vTAP Agents* on page 35.

- **GigaVUE® V Series node** is a visibility node that aggregates mirrored traffic from multiple G-vTAP agents. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to Gigamon Visibility Platform using the standard IP GRE tunnels.
- **G-vTAP Controller** manages multiple G-vTAP agents and orchestrates the flow of mirrored traffic to GigaVUE V Series nodes. GigaVUE-FM uses one or more G-vTAP Controllers to communicate with the G-vTAP agents.
- **GigaVUE V Series Controller** manages multiple GigaVUE V Series nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Controllers to communicate with the GigaVUE V Series nodes.

Supported Architecture

Gigamon's Visibility Platform supports the following cloud deployment models:

- *Public Cloud* on page 13
- *Hybrid Cloud* on page 15

Public Cloud

The Visibility Platform for AWS provides a consistent way to access network traffic within and across VPCs. It effectively distributes traffic to multiple tools, customizes network traffic to specific tools using policies, and delivers elastic, on-demand visibility as instances scale-out.

This model represents two types of deployment:

- **Single VPC** — You can have the G-vTAP Controllers, GigaVUE V Series Controllers, GigaVUE V Series nodes and GigaVUE-FM configured in the same VPC as the monitoring tools. The aggregated and optimized network traffic is sent to the monitoring tools residing in the same VPC through the management subnet. Refer to Figure 1-2 on page 14.

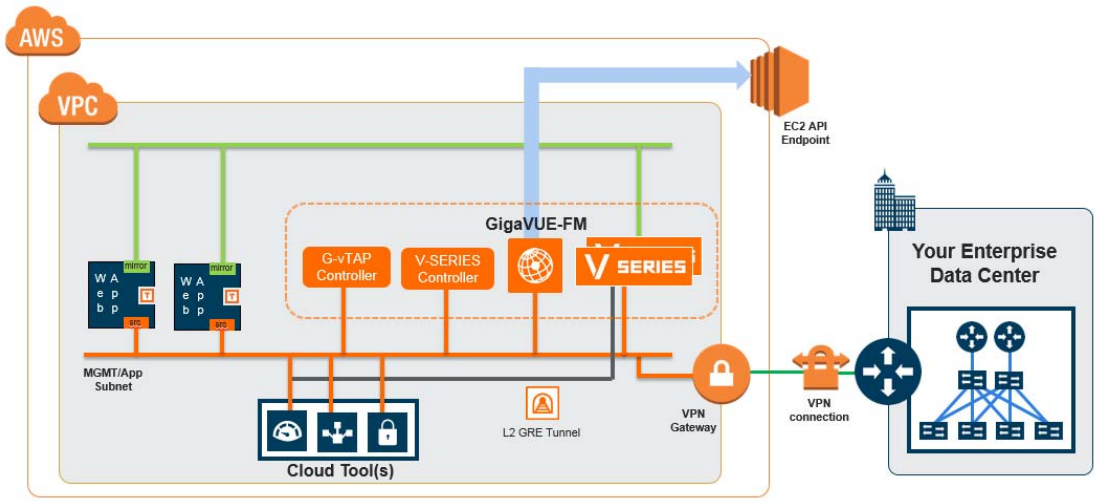

Figure 1-2: Tools in the Same VPC

- **Multiple VPCs—**You can have G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE V Series nodes configured within each availability zone, and in multiple VPCs. The GigaVUE V Series node in these VPCs optimizes the traffic and sends the traffic to the monitoring tools residing in a different VPC.

  GigaVUE-FM can be configured in any one of the VPCs. It can centrally manage the components residing in different VPCs. In this hub-and-spoke topology, the tools VPC acts as a central tool and management VPC. The traffic from the VPCs is sent to the central tool and management VPC through a tunnel subnet. Refer to Figure 1-3 on page 14.

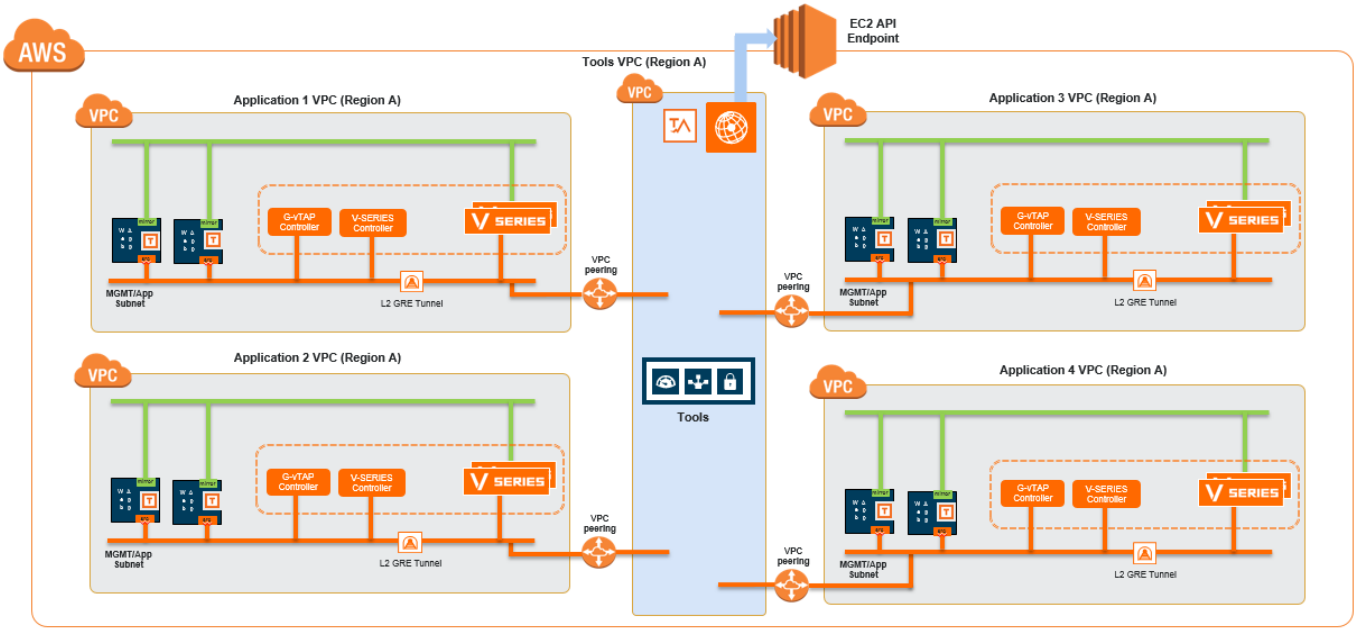

Figure 1-3: Tools in Different VPCs

Hybrid Cloud

The Gigamon Visibility Platform for AWS distributes the aggregated and optimized traffic from VPCs to the on-premise monitoring tools through a tunnel over the existing VPN connections.

GigaVUE-FM connects to the VPC through the EC2 API Endpoint and launches the G-vTAP Controller, GigaVUE V Series Controller, and GigaVUE V Series node. The mirrored traffic from the G-vTAP agents is sent to the GigaVUE V Series node through a data subnet. The customized traffic is then sent from the GigaVUE V Series node to the destination tunnel endpoints through a tunnel subnet. Refer to Figure 1-4 on page 15.

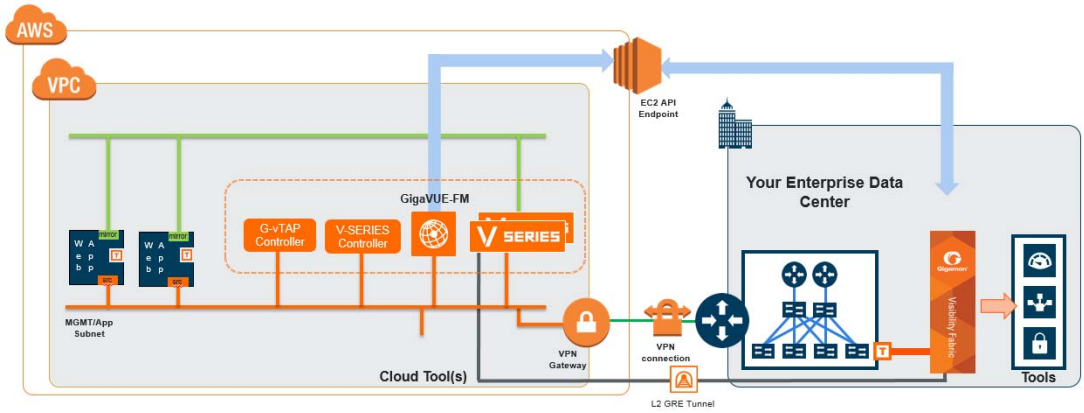

Figure 1-4: Hybrid Cloud Deployment

Chapter 2

Configuring the Components in AWS

This chapter describes how to launch a GigaVUE-FM instance and how to configure G-vTap Controllers, GigaVUE V Series nodes, and GigaVUE V Series Controllers in your VPC.

Refer to the following sections for details:

Before You Begin

This section describes the requirements for configuring the Visibility Platform for AWS. Refer to the following section for details.

AWS Permissions and Policies

Before you begin configuring the components, you must have the following permissions and policies assigned to your account:

- Full EC2 Instance access
- Read-only permission for IAM role
- EC2 pass role permission
- GigaVUE-FM Instance Role Policy

In addition, you must associate the following policies to your account before launching the GigaVUE-FM instance from the AWS Marketplace:

---EC2 Permissions

"ec2:Describe*",

"ec2:RebootInstances",

"ec2:RunInstances",

"ec2:StartInstances",

"ec2:StopInstances",

"ec2:TerminateInstances",

"ec2:ReportInstanceStatus",

"ec2:Disassociate*",

"ec2:CreateTags",

"ec2:AttachVolume",

"ec2:AttachNetworkInterface",

"ec2:Associate*",

"ec2:Allocate*",

"ec2:DeleteTags",

"ec2:DeleteVolume",

"ec2:DeleteNetworkInterface",

"ec2:ModifyInstanceAttribute",

"ec2:ModifyNetworkInterfaceAttribute",

"ec2:ModifyVolumeAttribute",

"ec2:ReleaseAddress",

"elasticloadbalancing:Describe*",

"autoscaling:Describe*",

"cloudwatch:*",

"logs:*",

"sns:*",

"sqs:*",

"events:*"

---S3 Permissions

"s3:CreateBucket",

"s3:DeleteBucket",

"s3:DeleteObject",

"s3:DeleteObjectVersion",

"s3:Get*",

"s3:ListAllMyBuckets",

"s3:PutBucketNotification",

"s3:PutBucketTagging",

"s3:PutBucketVersioning",

"s3:PutObject",

"s3:PutObjectTagging",

"s3:ReplicateDelete",

"s3:ReplicateObject",

"s3:RestoreObject"

---IAM Permissions

"iam:PassRole"

AWS Security Credentials

When you first connect GigaVUE-FM with AWS, you need the security credentials for AWS to verify your identity and check if you have permission to access the resources that you are requesting. AWS uses the security credentials to authenticate and authorize your requests.

You need one of the following security credentials:

- **Identity and Access Management (IAM) role**—If GigaVUE-FM is configured in a VPC, it is highly recommended to use an IAM role because it can securely make API requests from the instances.

  Create an IAM role and ensure that the permissions and policies listed in *AWS Permissions and Policies* on page 17 are associated to the role.

  For detailed instructions on creating an IAM role, refer to the AWS documentation on IAM Roles for Amazon EC2.

- **Access Keys**—If GigaVUE-FM is configured in the enterprise data center, then you need to use the access keys or basic credentials to connect to the VPC. Basic credentials allow full access to all the resources in your AWS account.

An access key consists of an access key ID and a secret access key.

For detailed instructions on creating access keys, refer to the AWS documentation on Managing Access Keys for Your AWS Account.

To obtain the IAM role or access keys, contact your AWS administrator.

NOTE: You cannot launch the GigaVUE-FM instance from AWS without having one of these security credentials. If you are configuring the GigaVUE-FM instance from the AWS Marketplace, you need to have only the IAM roles.

Network Requirements

To enable the flow of traffic between the components and the monitoring tools, your VPCs and instances should meet the following requirements:

- *Subnets for VPC*
- *Elastic Network Interfaces (ENIs) for Instances*

Subnets for VPC

Table 2-1 on page 20 lists the three recommended subnets that your VPC must have to configure the Visibility Platform components in AWS.

Table 2-1: Types of Subnets

Subnet	Description
Management Subnet	Subnet that the GigaVUE-FM uses to communicate with the GigaVUE V Series nodes and controllers.
Tunnel Subnet	Subnet that the GigaVUE V Series node uses to communicate with the monitoring tools that reside inside or outside of AWS, or GigaVUE H Series node that resides in your enterprise data center. The tunnel subnet can be the same as the management subnet.
Data Subnet	Subnet that receives the mirrored GRE tunnel traffic from the G-vTAP agents to the GigaVUE V Series node.

Elastic Network Interfaces (ENIs) for Instances

One or more Elastic Network Interfaces (ENIs) can be configured on the EC2 instances. If there is only one interface configured, the G-vTAP agent sends the monitored egress and ingress traffic out using the same interface. If there are two interfaces configured, the G-vTAP agent monitors the egress and ingress traffic on one interface and sends the mirrored traffic out using the second interface.

Security Group

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 Controllers, GigaVUE V Series nodes, and G-vTAP Controllers in your VPC, you add rules that control the inbound traffic to instances, and a separate set of rules that control the outbound traffic.

Table 2-2 on page 21 lists the rules and port numbers for each component.

Table 2-2: Security Group Rules

Direction	Type	Protocol	Port Range	Source CIDR, IP, or Security Group	Purpose
GigaVUE-FM Inside AWS					
Inbound	HTTPS	TCP(6)	443	Anywhere Any IP	Allows G-vTAP controllers and GigaVUE V Series controllers to communicate with GigaVUE-FM
G-vTAP Controller					
Inbound	Custom TCP Rule	TCP(6)	9900	Custom GigaVUE-FM IP	Allows G-vTAP controllers to communicate with GigaVUE-FM
G-vTAP Agent					
Inbound	Custom TCP Rule	TCP(6)	9901	Custom G-vTAP Controller IP	Allows G-vTAP controllers to communicate with G-vTAP agents
GigaVUE V Series Controller					
Inbound	Custom TCP Rule	TCP(6)	9902	Custom GigaVUE-FM IP	Allows GigaVUE-FM to communicate with GigaVUE V Series Controllers
GigaVUE V Series node					
Inbound	Custom TCP Rule	TCP(6)	9903	Custom GigaVUE V Series Controller IP	Allows GigaVUE V Series Controllers to communicate with GigaVUE V Series nodes
GRE Traffic					
Inbound	Custom Protocol Rule	GRE (47)	ALL	Anywhere Any IP	Allows mirrored traffic from G-vTAP agents to be sent to GigaVUE V Series nodes using L2 GRE or VXLAN tunnel Allows monitored traffic to be sent from GigaVUE V Series nodes to the tools using L2 GRE or VXLAN tunnel

It is recommended to create a separate security group for each component using the rules and port numbers listed in Table 2-2 on page 21.

Creating a Security Group

To create an inbound security group:

1. In the Amazon EC2 dashboard, click **Security Groups** in the navigation pane.
2. Click **Create Security Group**.
3. In the Security group name, enter a name.
4. In **Description**, specify the purpose for creating the security group.

5. In **VPC**, select the VPC ID.
6. Click **Add Rule** and enter the appropriate details. Refer to Table 2-2 on page 21.

   NOTE: The Source and the CIDR must be entered according to your requirement.
7. Click **Create**.

Name	Group ID	Group Name	VPC ID	Description
GigaVUE-FM	sg-f624dc8a	sg_gigavue-fm	vpc-48b0ac2c	For G-vTAP controllers and GigaVUE V Series controllers to comm
GigaVUE V Series Controller	sg-731ce40f	sg_gigavue-vseries-controller	vpc-48b0ac2c	For GigaVUE-FM to communicate with GigaVUE V Series Controlle
GigaVUE V Series Node	sg-a526ded9	sg_gigavue-vseries-node	vpc-48b0ac2c	For GigaVUE V Series Controllers to communicate with GigaVUE V
GRE	sg-cd1fe7b1	sg_gre-traffic	vpc-48b0ac2c	For traffic to be sent to L2 GRE or VXLAN tunnel
G-vTAP agent	sg-d519e1a9	sg_gvtap-agent	vpc-48b0ac2c	For G-vTAP controller to communicate with G-vTAP agents
G-vTAP Controller	sg-b322dacf	sg_gvtap-controller	vpc-48b0ac2c	For G-vTAP controller to communicate with GigaVUE-FM

Figure 2-1: Security Groups for Gigamon Visibility Platform

Key Pairs

A key pair consists of a public key and a private key. You must create a key pair and specify the name of this key pair when you launch the G-vTAP Controllers, GigaVUE V Series nodes, and GigaVUE V Series Controllers in your VPC. Then, you must provide the private key to connect to these instances.

For information about creating a key pair, refer to creating a key pair in the AWS documentation.

VPN Connectivity

GigaVUE-FM requires Internet access to integrate with the AWS API endpoints and deploy its visibility components. For more information about the VPN connectivity options, refer to Amazon Virtual Private Cloud Connectivity Options.

If there is no Internet access to the VPC, refer to *Configuring the Proxy Server* on page 99.

At a Glance

Refer to Figure 2-2 on page 23 for the steps to configure Gigamon Visibility Platform for AWS.

Figure 2-2: Steps for Configuring the Gigamon Visibility Platform for AWS

Obtaining the AMI

The AMI for the Gigamon Visibility Platform is available in both the AWS Public Cloud and in AWS GovCloud.

Gigamon Visibility Platform in AWS Public Cloud

The AMI for the Gigamon Visibility Platform is available in the AWS Marketplace for both the Bring Your Own License (BYOL) and the Pay-As-You-Go (PAYG) options. Figure 2-3 on page 24 shows both the licensing models in the AWS Marketplace.

Figure 2-3: AMI in the AWS Public Cloud

For purchasing licensing with the BYOL option, contact the Gigamon Sales. Refer to *Contacting Sales* on page 124.

Gigamon Visibility Platform in AWS GovCloud

AWS GovCloud is an isolated AWS region that contains specific regulatory and compliance requirements of the US government agencies. The AWS GovCloud (US) Region adheres to U.S. International Traffic in Arms Regulations (ITAR) requirements.

To monitor the instances that contain all categories of Controlled Unclassified Information (CUI) data and sensitive government data in the AWS GovCloud (US) Region, the AWS GovCloud AMI provides the same robust features in the AWS GovCloud as in the AWS public cloud.

Launching the GigaVUE-FM Instances

The AMI for the GigaVUE-FM instance can be launched from the AWS Marketplace or AWS EC2 dashboard. Refer to the following sections:

- *Launching the GigaVUE-FM Instance from the AWS Marketplace* on page 25
- *Launching the GigaVUE-FM Instance from the AWS EC2 Dashboard* on page 30

Launching the GigaVUE-FM Instance from the AWS Marketplace

To launch the GigaVUE-FM instance:

1. Login to the AWS account.
2. Go to https://aws.amazon.com/marketplace/.
3. In the **Search** field, type Gigamon and press **Enter**. Refer to Figure 2-4.

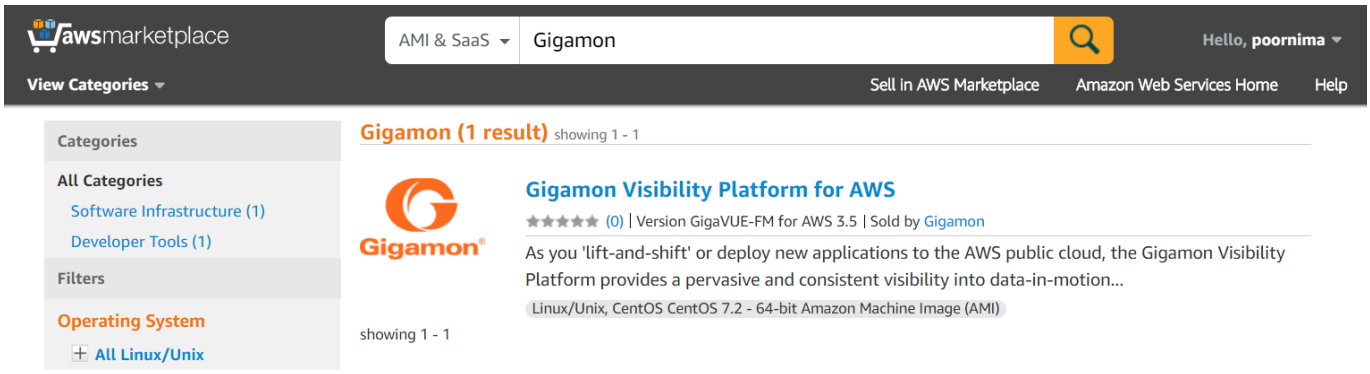

Figure 2-4: Searching for Gigamon on AWS Marketplace

4. Click the **Gigamon Visibility Platform for AWS** link to view the complete details about the product. Refer to Figure 2-5.

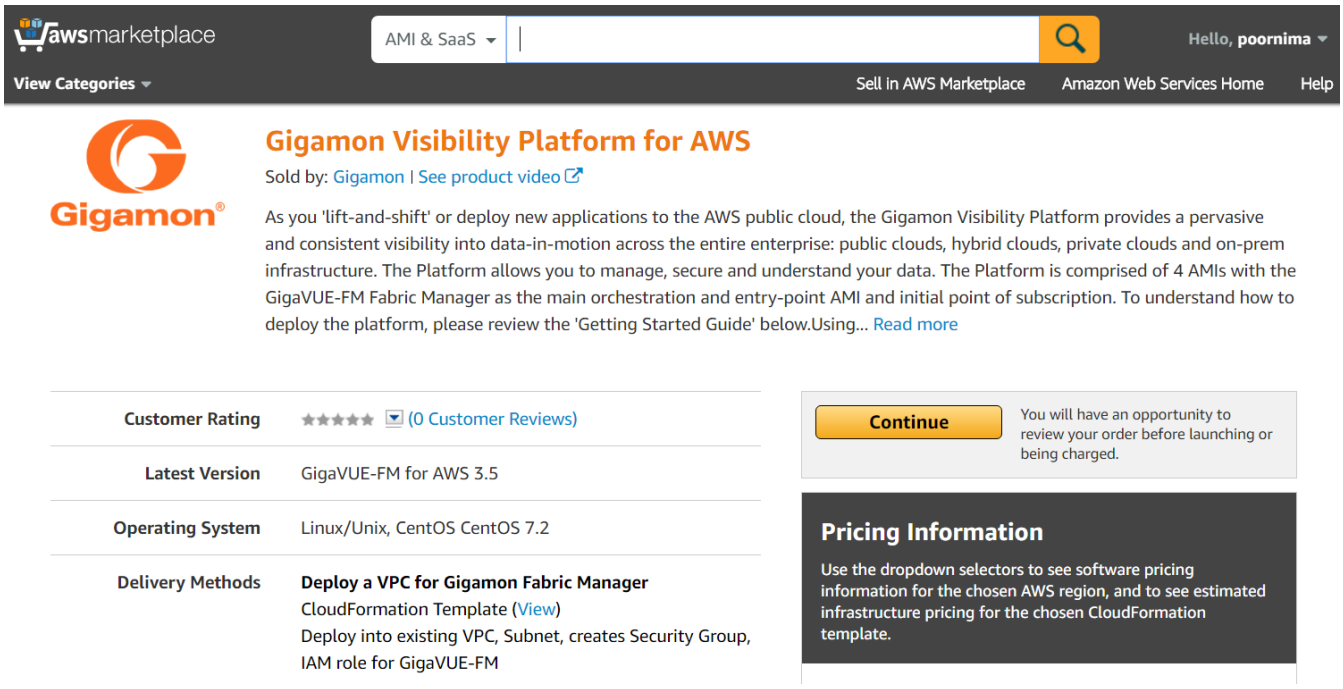

Figure 2-5: Gigamon Visibility Platform for AWS page in AWS Marketplace

5. Click **Continue**. The Launch page is displayed. Refer to Figure 2-6.

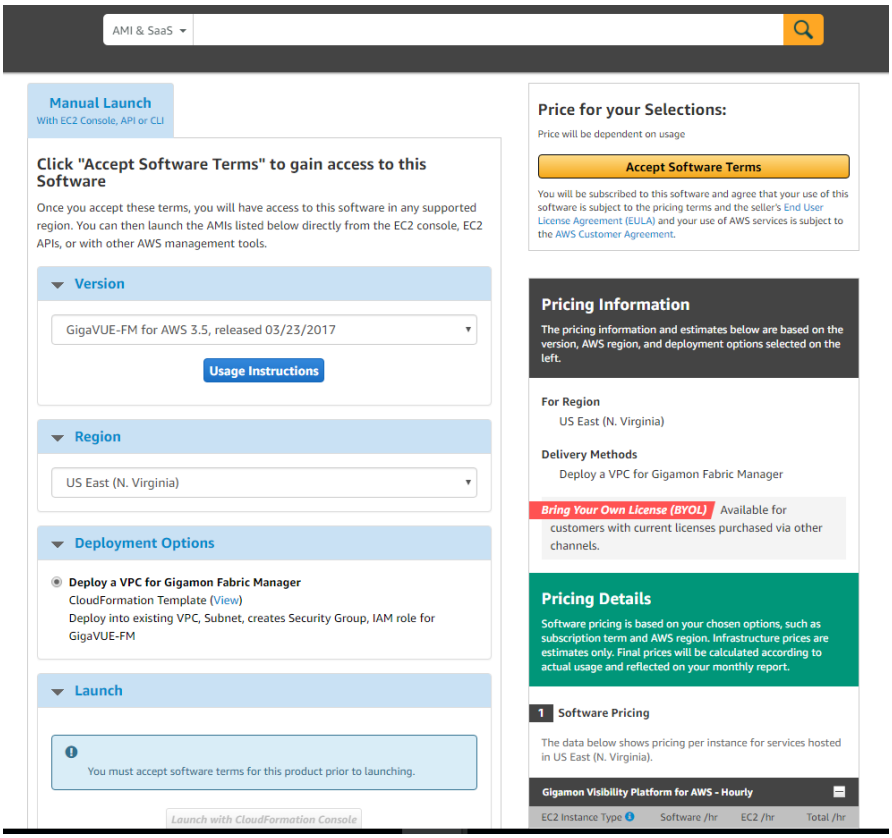

Figure 2-6: Launch on EC2 Page

6. In the Launch on EC2 page, select the following:
   a. From the **Version** drop-down list, select the latest version.
   b. From the **Region** drop-down list, select the appropriate region.
   c. By default, the **Deploy a VPC for Gigamon Fabric Manager** option is selected.

d. Click the **Accept Software Terms** button to subscribe to the Gigamon Visibility Platform for AWS software. A message is displayed to confirm the subscription. Refer to Figure 2-7. Click **Return to Launch Page**.

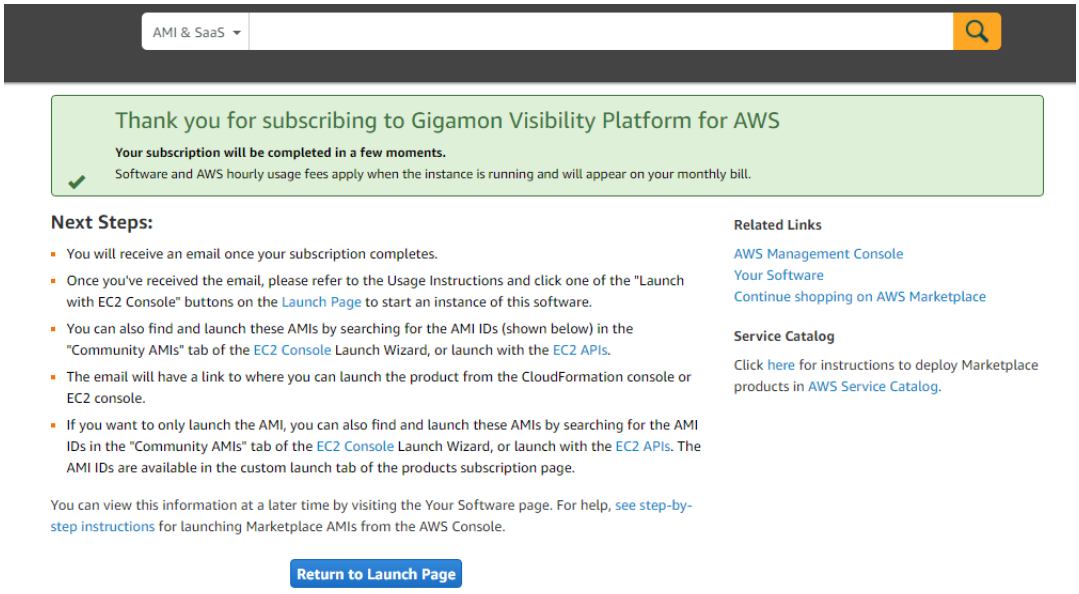

Figure 2-7: Subscription Confirmation Page

e. In the Launch on EC2 page, the **Launch with CloudFormation Console** button is enabled. Click this button. The Select Template page is displayed. Refer to Figure 2-8.

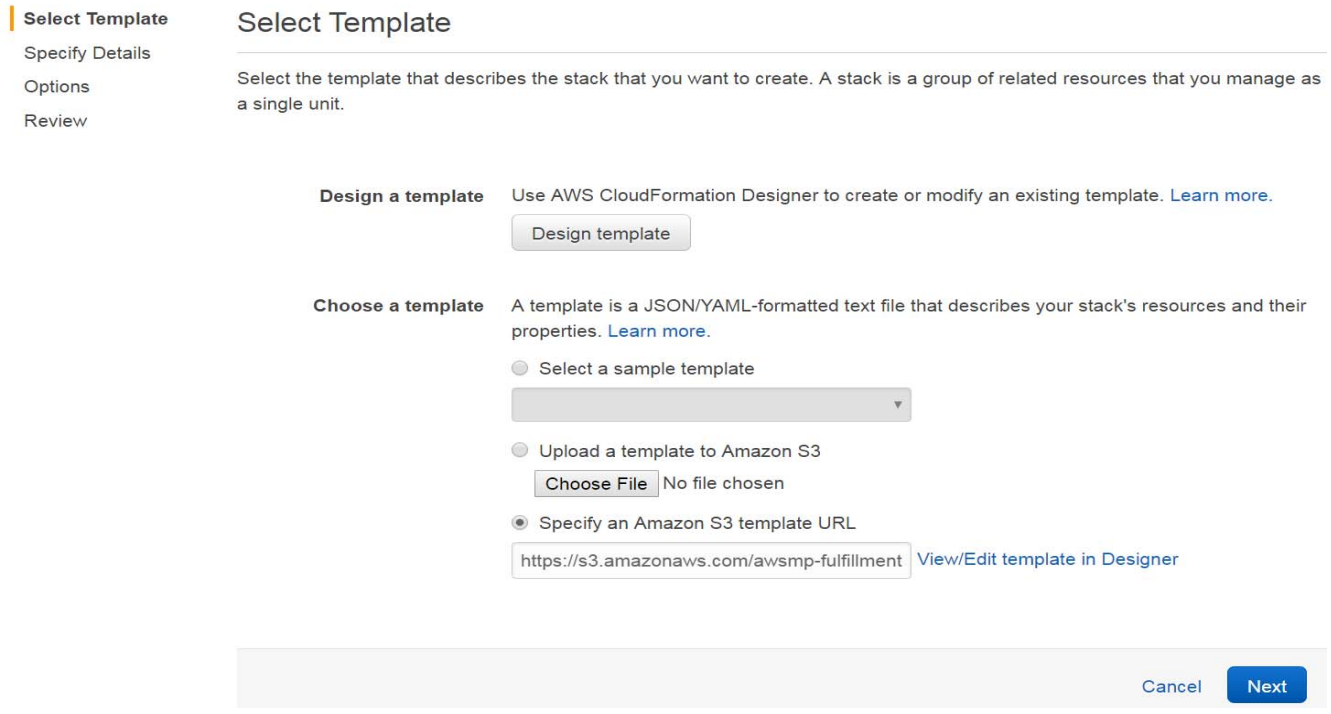

Figure 2-8: Select Template Page

7. In the Select Template page, the Gigamon Fabric Manager CloudFormation template is selected by default. Click **Next**. The Specify Details page is displayed. Refer to Figure 2-9.

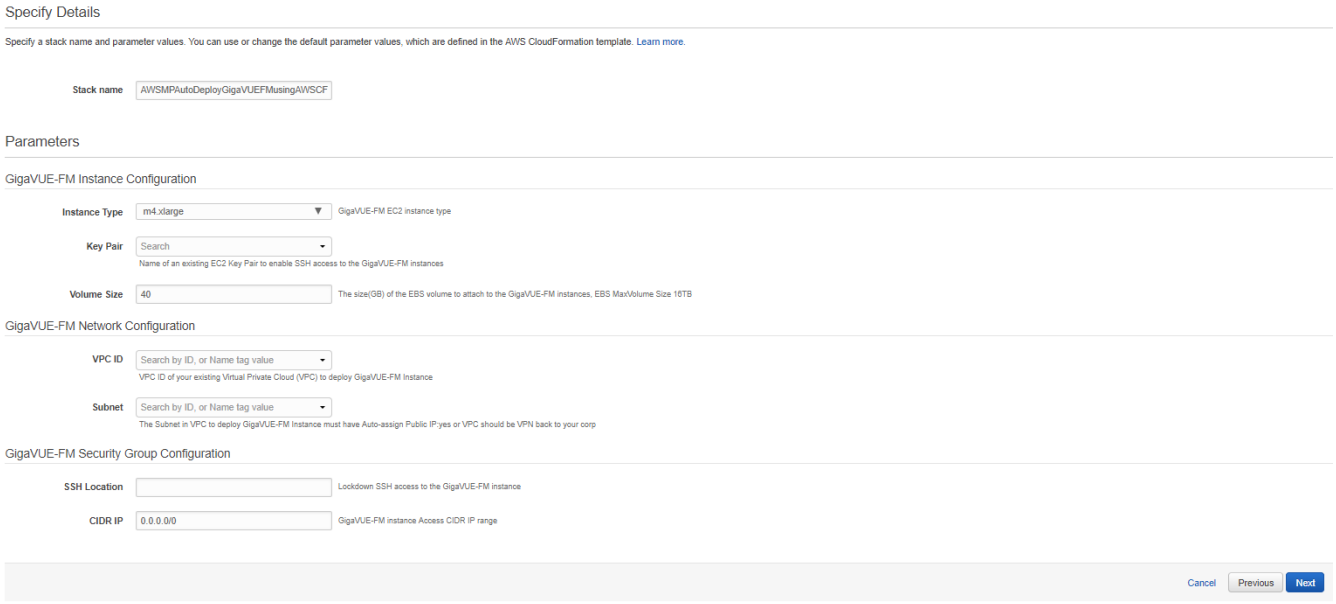

Figure 2-9: Specify Details Page

8. In the Specify Details page, enter the following:
   a. In the Stack name field, enter a stack name.
   b. From the Instance Type drop-down list, select m4.xlarge as the minimum instance type for GigaVUE-FM.

   NOTE: The t2 instance types are not supported.

   c. From the Key Pair drop-down list, select the name of an existing EC2 key pair.
   d. In the Volume Size field, by default 40 is selected. Change the volume size based on your requirement.
   e. From the VPC ID drop-down list, select the appropriate VPC ID.
   f. From the My Subnet drop-down list, select the appropriate public subnet ID.
   g. In the SSH Location field, enter the SSH to lock down the SSH access to the Gigamon FM instance.
   h. In the CIDR IP, enter a CIDR block to associate with the instance and click **Next**.
9. In the Review page, review the complete details and then select the check box to acknowledge that AWS CloudFormation might create IAM resources.
10. Click **Create**.
11. It will take several minutes for the instance to initialize. After the initialization is completed, you can verify the instance through the Web interface as follows:
    a. Find your instance and expand the page in the **Descriptions** tab to view the instance information, if necessary.

b. Copy the Public DNS value and paste the value into a new browser window or tab.

c. Copy the Instance ID.

   If GigaVUE-FM is deployed inside AWS, use the **Instance ID** as the password for the admin user to login to GigaVUE-FM as shown in Figure 2-18 on page 34, for example, admin/i-033f5644b3e265d27. You can change the password after logging in to GigaVUE-FM.

   If GigaVUE-FM is deployed outside AWS, use admin123A! as the default admin password.

Figure 2-10: GigaVUE-FM Login Screen

Launching the GigaVUE-FM Instance from the AWS EC2 Dashboard

This section describes how to launch the GigaVUE-FM instance in your VPC.

To launch the GigaVUE-FM instance:

1. Login to the AWS account and select **Services > EC2**. Refer to Figure 2-11 on page 30.

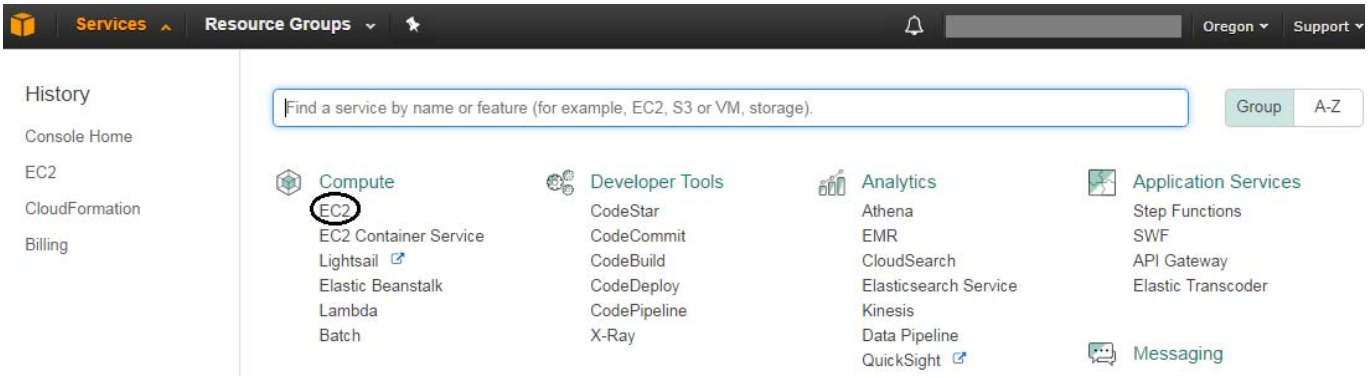

Figure 2-11: Selecting EC2

2. Click **Launch Instance** and go to **AWS Marketplace** or **Community AMIs**.

NOTE: If you are logged into AWS GovCloud (US) account, the GigaVUE-FM instance will be available in the Community AMIs.

3. Search for **Gigamon**, locate the latest version of the GigaVUE-FM AMI, and click **Select**. Refer to Figure 2-12 on page 30.

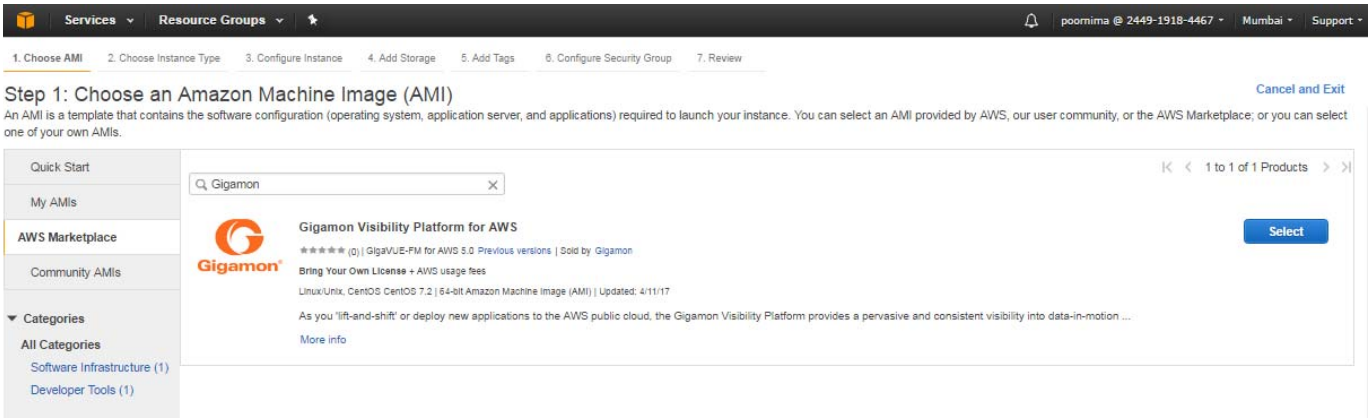

Figure 2-12: Searching for Gigamon in EC2 Console

4. Choose the Instance Type. The recommended instance type is **m4.xlarge**. Figure 2-13 on page 31.

NOTE: The t2 instance types are not supported.

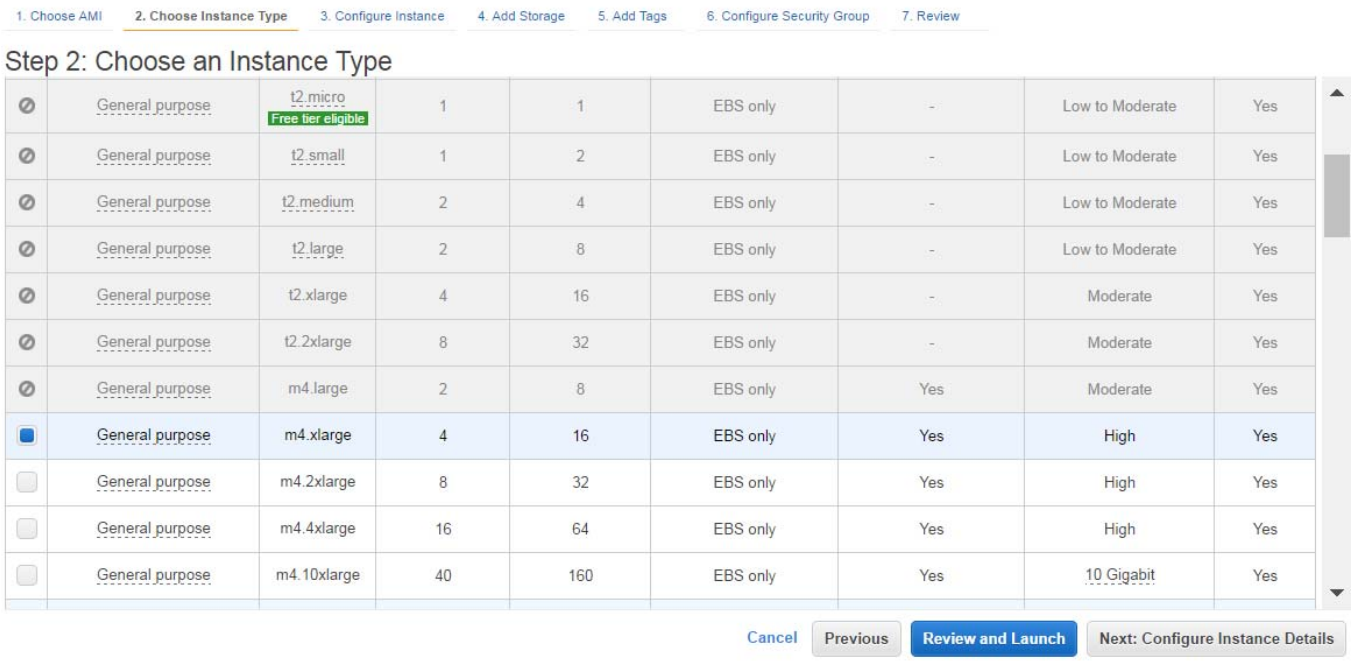

Figure 2-13: Selecting the Instance Type

5. Click **Next: Configure Instance Details** and perform the following as shown in Figure 2-14 on page 32:
   - **Network**— Select the VPC in which you want to launch the instance.
   - **Subnet**— Select the management subnet that the instance will use after launch.
   - **Auto-assign Public IP**— Select **Enable**.
   - **IAM role**—Select an existing IAM role or create a new IAM role to associate with the instance. For creating an IAM role, refer to *AWS Security Credentials* on page 19. Ensure that the policies listed in *AWS Permissions and Policies* on page 17 are attached to the IAM role.

NOTE: If you do not have an IAM role, contact your AWS administrator.

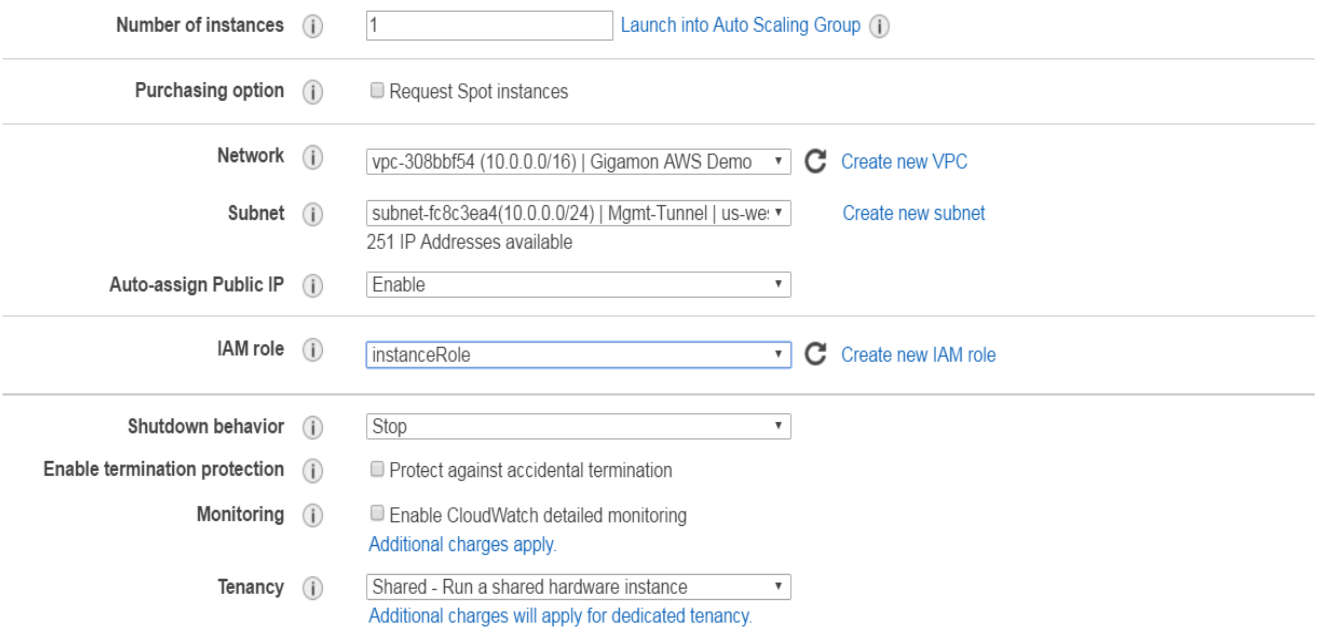

Figure 2-14: Configuring an Instance

6. Click **Next: Add Storage**, click **Add New Volume**, and then select the following storage device settings as shown in Figure 2-15 on page 32:
   - **Size (GiB)**— Enter a minimum of 40Gb of storage.
   - **Volume Type**— Select a volume type. The recommended volume is General Purpose SSD (GP2).
   - **Delete on Termination**— Select this check box to make sure the volumes are cleaned up when the GigaVUE-FM instance is removed.

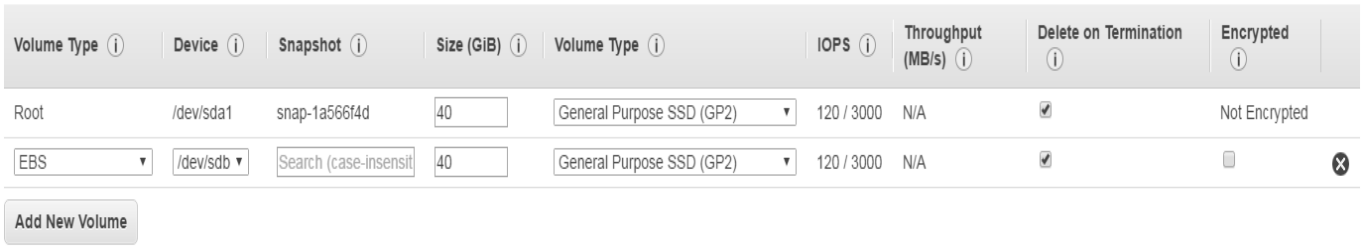

Figure 2-15: Adding a Storage

7. Click **Next: Add Tags,** and then click the **Add Tag** button to add a key-value pair to identify the instance. (Refer to Figure 2-16 on page 33.)

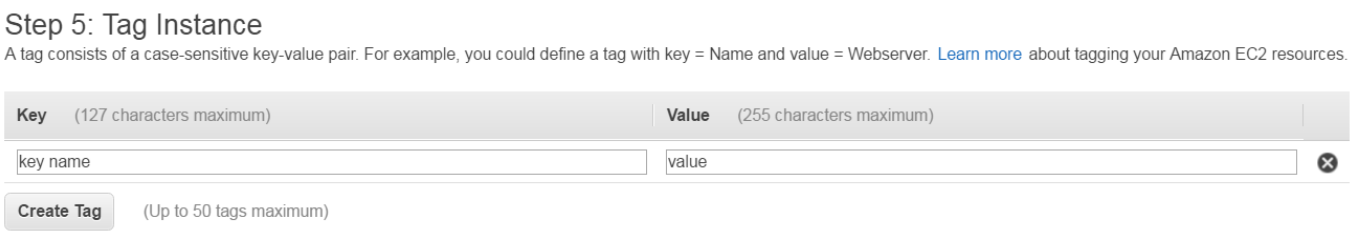

Figure 2-16: Adding a Tag to an Instance

8. Click **Next: Add Security Group**. Click the **Select an existing security group** check box if you have a security group already created. Otherwise, select the **Create a new security group** check box and click **Add Rule**. For more information on creating a security group, refer to *Security Group* on page 20.
9. Click **Review and Launch**. Review the instance launch details and click **Launch**.
10. Select the SSH key pair, check the acknowledgment check box, and click **Launch Instances** as shown in Figure 2-17 on page 33.

    NOTE: If you are unable to create a key pair, check the Administrator Privileges to make sure you have the permission to create a key pair.

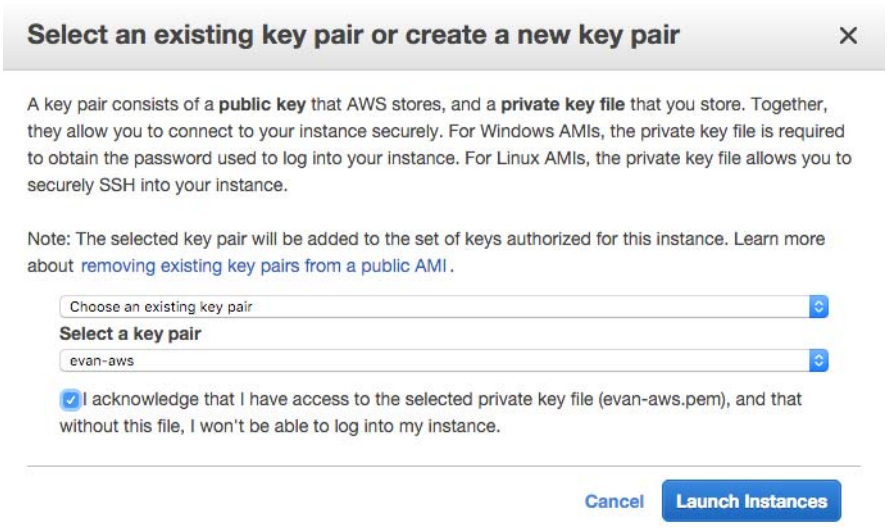

Figure 2-17: Selecting an SSH Key Pair

11. It will take several minutes for the instance to initialize. After the initialization is completed, you can verify the instance through the Web interface as follows:
    a. Find your instance and expand the page in the **Descriptions** tab to view the instance information, if necessary.
    b. Copy the Public DNS value and paste the value into a new browser window or tab.
    c. Copy the Instance ID.

If GigaVUE-FM is deployed inside AWS, use the **Instance ID** as the password for the admin user to login to GigaVUE-FM as shown in Figure 2-18 on page 34, for example, admin/i-033f5644b3e265d27. You can change the password after logging in to GigaVUE-FM.

If GigaVUE-FM is deployed outside AWS, use admin123A! as the default admin password.

Figure 2-18: GigaVUE-FM Login Screen

G-vTAP Agents

A **G-vTAP agent** is an agent that is deployed in the EC2 instances. This agent mirrors the selected traffic from the instances (virtual machines), encapsulates it using GRE or VXLAN tunneling, and forwards it to the GigaVUE® V Series node.

A G-vTAP agent consists of a source interface and a destination interface. The network packets collected from the source interface are sent to the destination interface. From the destination interface, the packets traverse through the L2 GRE or VXLAN tunnel interface to the GigaVUE V Series node.

A source interface can be configured with one or more ENIs. While configuring a source interface, you can specify the direction of the traffic to be monitored in the instance. The direction of the traffic can be egress or ingress or both.

Linux Agent Installation

Refer to the following sections for the Linux agent installation:

- *Single ENI Configuration* on page 35
- *Dual ENI Configuration* on page 35
- *Installing the G-vTAP Agents* on page 35
- *Installing from an Ubuntu/Debian Package* on page 36
- *Installing from an RPM package* on page 36

Then refer to *Creating Images with Agent Installed* on page 38.

Single ENI Configuration

A single ENI acts both as the source and the destination interface. A G-vTAP agent with a single ENI configuration lets you monitor the ingress or egress traffic from the ENI. The monitored traffic is sent out using the same ENI.

For example, assume that there is only one interface eth0 in the monitoring instance. In the G-vTAP configuration, you can configure eth0 as the source and the destination interface, and specify both egress and ingress traffic to be selected for monitoring purpose. The egress and ingress traffic from eth0 is mirrored and sent out using the same interface.

Using a single ENI as the source and the destination interface can sometimes cause increased latency in sending the traffic out from the instance.

Dual ENI Configuration

A G-vTAP agent lets you configure two ENIs. One ENI can be configured as the source interface and another ENI can be configured as the destination interface.

For example, assume that there is eth0 and eth1 in the monitoring instance. In the G-vTAP agent configuration, eth0 can be configured as the source interface and egress traffic can be selected for monitoring purpose. The eth1 interface can be configured as the destination interface. So, the mirrored traffic from eth0 is sent to eth1. From eth1, the traffic is sent to the GigaVUE V Series node.

Installing the G-vTAP Agents

You must have sudo/root access to edit the G-vTAP agent configuration file.

For dual or multiple ENI configuration, you may need to modify the network configuration files to make sure that the extra ENI will initialize at boot time.

You can install the G-vTAP agents either from Debian or RPM packages as follows:

- *Installing from an Ubuntu/Debian Package*
- *Installing from an RPM package*

Installing from an Ubuntu/Debian Package

To install from a Debian package:

1. Download the G-vTAP Agent Debian (.deb) package from the following location:

   https://s3.amazonaws.com/gvtap-agent/1.2-1/gvtap-agent_1.2-1_amd64.deb

2. Copy this package to your instance. Install the package with root privileges, for example:

```
ubuntu@ip-10-0-0-246:~$ ls gvtap-agent_1.2-1_amd64.deb
ubuntu@ip-10-0-0-246:~$ sudo dpkg -i
   gvtap-agent_1.2-1_amd64.deb
```

3. Once the G-vTAP package is installed, modify the file /etc/gvtap-agent/gvtap-agent.conf to configure and register the source and destination interfaces.

   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
```

4. Save the file.

5. Reboot the instance.

   The G-vTAP agent status will be displayed as running. Check the status using the following command:

```
ubuntu@ip-10-0-0-246:~$ sudo service gvtap-agent status
G-vTAP Agent is running
```

Installing from an RPM package

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

1. Download the G-vTAP Agent RPM (.rpm) package from the following location:

   https://s3.amazonaws.com/gvtap-agent/1.2-1/gvtap-agent_1.2-1_x86_64.rpm

2. Copy this package to your instance. Install the package with root privileges, for example:

```
[ec2-user@ip-10-0-0-214 ~]$ ls
gvtap-agent_1.2-1_x86_64.rpm
[ec2-user@ip-10-0-0-214 ~]$ sudo rpm -i
   gvtap-agent_1.2-1_x86_64.rpm
```

3. Modify the file /etc/gvtap-agent/gvtap-agent.conf to configure and register the source and destination interfaces.

   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
   ```

4. Save the file.
5. Reboot the instance.

   Check the status with the following command:

   ```
   [ec2-user@ip-10-0-0-214 ~]$ sudo service gvtap-agent status
   G-vTAP Agent is running
   ```

Windows Agent Installation

To install the Windows agent:

1. Download the Windows agent package from the following location:

   Windows G-vTAP Agent package (gvtap-agent_1.3-2.zip)

2. Extract the contents of the .zip file into a convenient location.
3. Run ‘WinPcap_4_1_3.exe’ (located in the ‘winpcap’ folder) as **Administrator**.
4. Run ‘install.bat’ as **Administrator**.
5. If you want to build an Azure image, this would be the point to go about creating the image.
6. If you want to start the Windows G-vTAP agent, you may do one of the following:
   - Reboot the VM.
   - Run ‘sc start gvtap’ from the command prompt.
   - Start the G-vTAP Agent from the Task Manager.

Next, refer to *Creating Images with Agent Installed* on page 38.

Creating Images with Agent Installed

To avoid downloading and installing the G-vTAP agents multiple times, save the G-vTAP agent running on an instance as a private AMI. Launch the G-vTAP agent AMI when ever a new instance needs to be monitored. When there is a change in the number of instances being monitored, GigaVUE-FM automatically updates the number in the monitoring session.

To save the G-vTAP agent as an AMI:

1. From the EC2 console, right click the instance.
2. Click **Image** > **Create Image**.

To launch the G-vTAP agent:

1. Follow steps 1 to 11 as described in *Launching the GigaVUE-FM Instance from the AWS EC2 Dashboard* on page 30 to launch the G-vTAP agent.
2. In Step 4, choose **t2 medium** as the instance type.
3. For configuration examples 2 and 3, in Step 5, click **Add Device** and add another ENI which acts as a mirror subnet.

Configuring the Visibility Platform Components in AWS

First, you must establish a connection between GigaVUE-FM and your AWS environment. Then, GigaVUE-FM lets you launch the G-vTAP Controllers, GigaVUE V Series Controllers, and GigaVUE V Series nodes in the specified VPC.

Pre-Configuration Checklist

Table 2-3 on page 38 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 Visibility Platform for AWS:

Table 2-3: Pre-configuration Checklist

	Required Information
☐	VPC ID
☐	Instance ID of the GigaVUE-FM
☐	Public or Private IP of the GigaVUE-FM
☐	Elastic IP **NOTE:** If GigaVUE-FM is installed in the enterprise data center, an Elastic IP is required for G-vTAP controllers and GigaVUE V Series controllers to communicate with GigaVUE-FM.
☐	Region name for the VPC
☐	Availability zone of the VPC
☐	IAM role name OR Access key ID and Secret Access key
☐	SSH Key Pair

	Required Information
☐	Subnets
☐	Security groups

Logging in to GigaVUE-FM

To login to GigaVUE-FM, do the following:

1. Copy the Public or Private IP of GigaVUE-FM into a browser. The GigaVUE-FM login page is displayed.
2. Enter admin as the user name and Instance ID of GigaVUE-FM as the password.
3. Click **Login**.

If you want to enable CloudWatch Events to track instance state changes in GigaVUE-FM, you must enable it before connecting to AWS.

To enable **AWS Cloud Watch Event Based Inventory Refresh** for the connection, go to **AWS > Configuration > Settings** and select the **AWS CloudWatch event-based inventory refresh** check box. AWS Cloud Watch Event Based Inventory Refresh lets you create a CloudWatch Event Rule and SQS queue that sends and receives the instance state change events. The GigaVUE-FM will poll the SQS queue for instance state change events.

Connecting to AWS

GigaVUE-FM connects to the VPC through the EC2 API endpoint. The default protocol GigaVUE-FM uses to communicate with the EC2 API endpoint is HTTPS. For more information about the endpoint and the protocol used, refer to http://docs.aws.amazon.com/general/latest/gr/rande.html#ec2_region. Once the connection is established, GigaVUE-FM launches the G-vTAP Controller, GigaVUE V Series Controller, and GigaVUE V Series node.

GigaVUE-FM provides you the flexibility to connect to multiple VPCs. You can choose the VPC ID and launch the Visibility Platform components in the desired VPCs.

To connect to AWS using GigaVUE-FM:

1. Click **Cloud** in the top navigation link.

2. Under AWS, select **Configuration > Connections**, and then click **New**. The AWS Connection page is displayed as shown in Figure 2-19 on page 40.

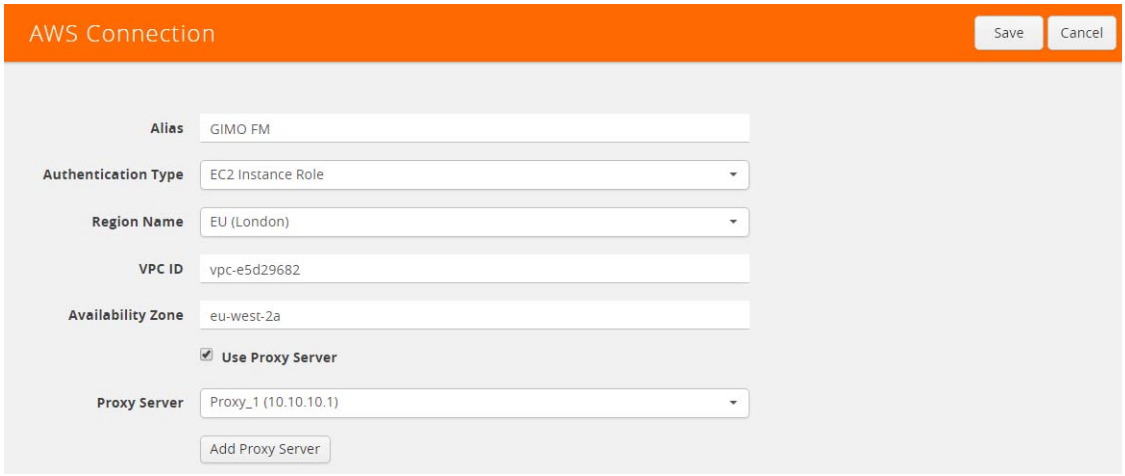

Figure 2-19: Connecting to AWS

3. Enter or select the appropriate information as shown in Table 2-4 on page 41.

Table 2-4: AWS Connection

Field	Description
Alias	An alias used to identify the connection to AWS. For example, vpcs-48b0ac2c-Oregon.
Authentication Type	The authentication type for the connection. For more information, refer to *AWS Security Credentials* on page 19.
Region Name	The AWS region for the connection. For example, EU (London). **NOTE:** If the region you want to choose is not available in the Region Name list, you can add a custom region. **Adding a Custom Region** To add a custom region: **a.** In the Region Name drop-down list, select **Custom Region**. Refer to Figure 2-20. 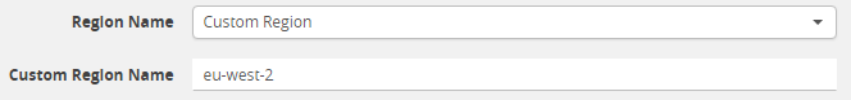 *Figure 2-20: Adding a Custom Region* **b.** In the Custom Region Name field, enter the name of the region that is not available in the list.
VPC ID	The ID of the target VPC for establishing the connection.
Availability Zone	The availability zone of the VPC. For example, US-West-2c.
Use Proxy Server	The check box to add a proxy server. Proxy server enables communication from GigaVUE-FM to the Internet, if there is no Internet access to the VPC.
Proxy Server	The list of proxy servers already configured in GigaVUE-FM. For more information on adding the proxy servers before configuring the AWS connection, refer to *Configuring the Proxy Server* on page 99
Add Proxy Server	The proxy sever can be configured from the AWS Connection page. Click **Add Proxy Server**. For more information, refer to *Configuring the Proxy Server* on page 99.

4. Click **Save**.

   If the connection is established, the status is displayed as **Connected** in the Connections page. GigaVUE-FM discovers the inventory of the VPC in the background.

   If the connection fails, a **Connection Failed** error message is displayed when **Save** is clicked.

The connection status is also displayed in **Cloud** > **Audit Logs**. Refer to Figure 2-21 on page 42.

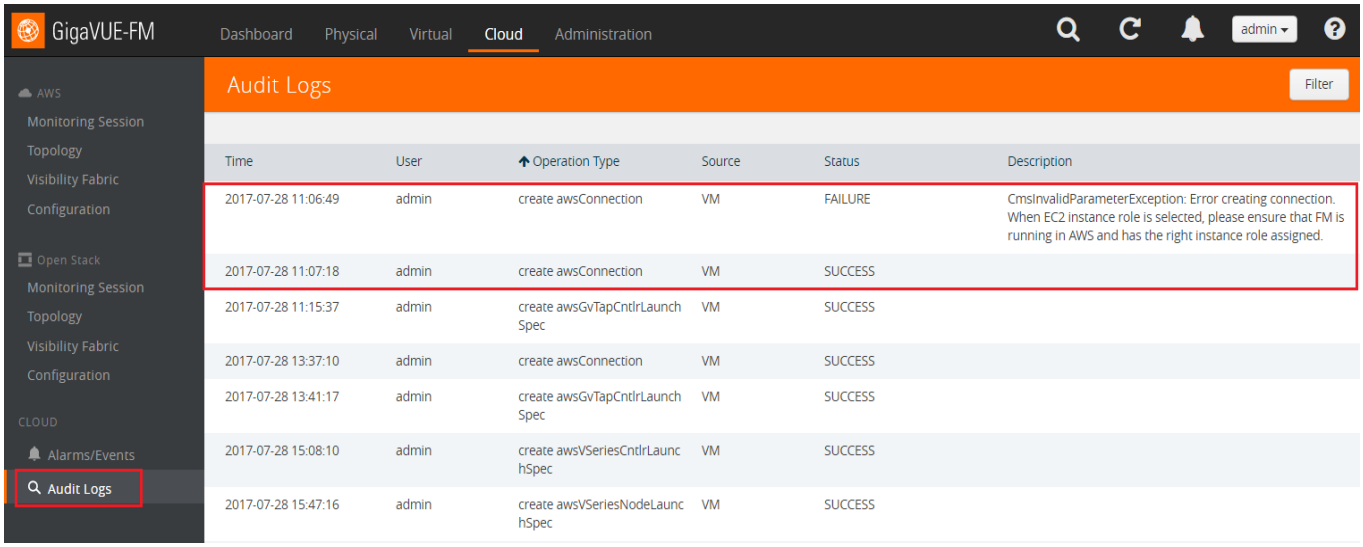

Figure 2-21: Audit Logs

Configuring the G-vTAP Controllers

A G-vTAP Controller manages multiple G-vTAP agents and orchestrates the flow of mirrored traffic to GigaVUE V Series nodes. A single G-vTAP Controller can manage up to 100 G-vTAP agents deployed in the EC2 instances.

A G-vTAP Controller can only manage G-vTAP agents with the same version. For example, the G-vTAP Controller v1.3 can only manage G-vTAP agents v1.3. So, if you have G-vTAP agents v1.2 still deployed in the EC2 instances, you must configure both G-vTAP Controller v1.2 and v1.3.

While configuring the G-vTAP Controllers, you can also specify the tunnel type to be used to carry the mirrored traffic from the G-vTAP agents to the GigaVUE V Series nodes. The tunnel type can be L2GRE or VXLAN.

To configure the G-vTAP Controllers:

1. Click **Cloud** in the top navigation link.
2. Under AWS, click **Configuration > G-vTAP Controllers**.

3. Click **New**. The G-vTAP Configuration page is displayed as shown in Figure 2-22 on page 43.

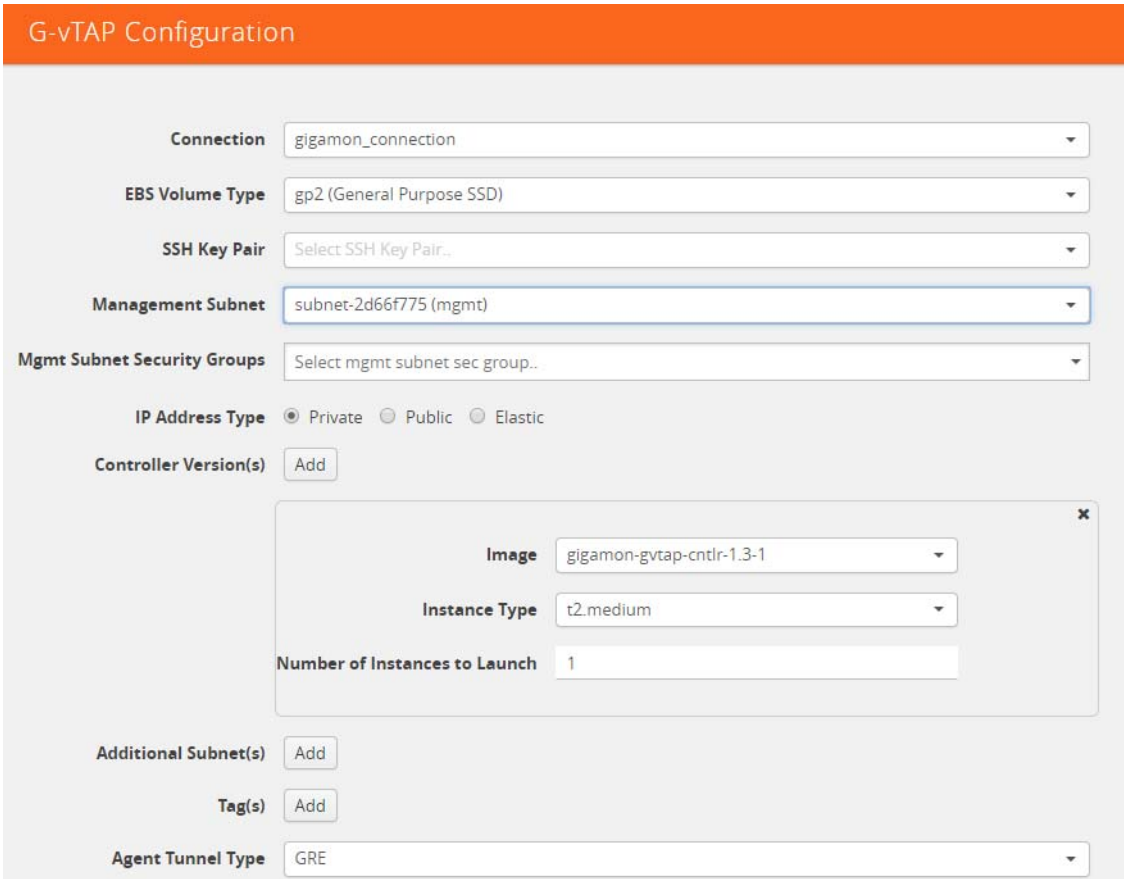

Figure 2-22: Configuring the G-vTAP Controller

4. Enter or select the appropriate information as shown in Table 2-5 on page 43.

Table 2-5: Fields for G-vTAP Configuration

Fields	Description
Connection	The name of the AWS connection.
EBS Volume Type	The Elastic Block Store (EBS) volume that you can attach to a single G-vTAP Controller instance. The available options are gp2 (General Purpose SSD), io1 (Provisioned IOPS SSD), and standard (Magnetic).
SSH KeyPair	The SSH key pair for the G-vTAP Controller. For more information about SSH key pair, refer to *Key Pairs* on page 22.
Management Subnet	The public subnet that is used for communication between the G-vTAP Controllers and the G-vTAP agents.
Mgmt Subnet Security Groups	The security group created for the management subnet. For more information, refer to *Security Group* on page 20.

Table 2-5: Fields for G-vTAP Configuration

Fields	Description
IP Address Type	The IP address type. Select one of the following: • Select **Private** if you want to assign an IP address that is not reachable over Internet. You can use private IP address for communication between the G-vTAP Controller instances and GigaVUE-FM instances in the same network. • Select **Public** if you want the IP address to be assigned from Amazon's pool of public IP address. The public IP address gets changed every time the instance is stopped and restarted. • Select **Elastic** if you want a static IP address for your instance. The Elastic IPs drop-down list is displayed under Controller Version(s). • The elastic IP address does not change after you stop or start the instance.
Controller Version(s)	The G-vTAP Controller version. The G-vTAP Controller version you configure must always be the same as the G-vTAP agents' version number deployed in the EC2 instances. This is because the G-vTAP Controller v1.2 can only manage G-vTAP agents v1.2. Similarly, the G-vTAP Controller v1.3 can only manage G-vTAP agents v1.3. If there are multiple versions of G-vTAP agents deployed in the EC2 instances, then you must configure multiple versions of G-vTAP Controllers that matches the version numbers of the G-vTAP agents. **NOTE:** If there is a version mismatch between G-vTAP controllers and G-vTAP agents, GigaVUE-FM cannot detect the agents in the instances. To add multiple versions of G-vTAP Controllers: **a.** Under **Controller Versions**, click **Add**. **b.** From the **Image** drop-down list, select a G-vTAP Controller image that matches with the version number of G-vTAP agents installed in the instances. **c.** From the **Instance Type** down-down list, select an instance type for the G-vTAP Controller. The recommended instance type is t2.medium. **NOTE:** The instance type t2.nano is not supported. **d.** In **Number of Instances to Launch**, specify the number of G-vTAP Controllers to launch. The minimum number you can specify is 1. **e.** The Elastic IPs drop-down list appears only if the **Elastic** option is selected in the IP Address Type. From the **Elastic IPs** drop-down list, select an IP.

Table 2-5: Fields for G-vTAP Configuration

Fields	Description
Controller Version(s) (continued)	An older version of G-vTAP Controller can be deleted once all the G-vTAP agents are upgraded to the latest version. To delete a specific version of G-vTAP Controller, click **x** (delete) next to its G-vTAP Controller image. 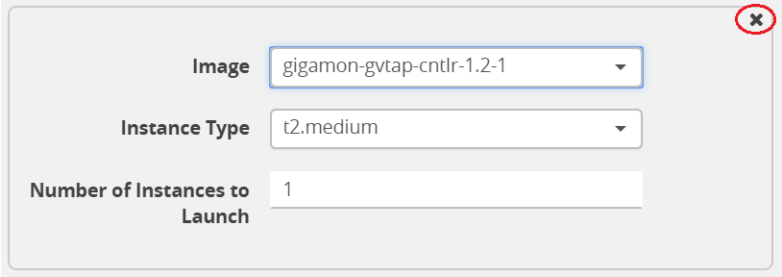 *Figure 2-23: Delete a G-vTAP Controller Version* Once you delete a G-vTAP Controller image from the G-vTAP Configuration page, all the G-vTAP Controller instances of that version are deleted from AWS.
Additional Subnet(s)	(Optional) If there are G-vTAP agents on subnets that are not IP routable from the management subnet, additional subnets must be specified so that the G-vTAP Controller can communicate with all the G-vTAP agents. Click **Add** to specify additional subnets, if needed. Also, make sure that you specify a list of security groups for each additional subnet.
Tag(s)	(Optional) The key name and value that helps to identify the G-vTAP Controller instances in your AWS environment. For example, you might have G-vTAP Controllers deployed in many regions. To distinguish these G-vTAP Controllers based on the regions, you can provide a name that is easy to identify such as us-west-2-gvtap-controllers. 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-gvtap-controllers. When the G-vTAP Controllers are launched in the VPC, they appear as shown in Figure 2-24 on page 45: 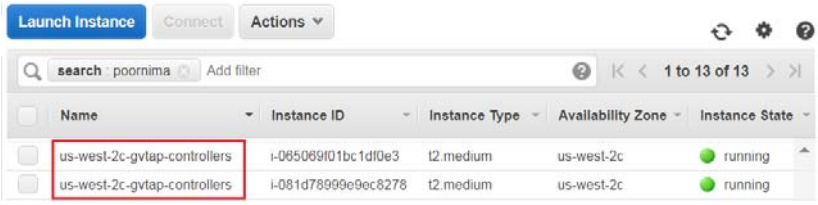 *Figure 2-24: G-vTAP Controllers with Custom Tag Name*

Table 2-5: Fields for G-vTAP Configuration

Fields	Description
Agent Tunnel Type	The type of tunnel used for sending the traffic from G-vTAP agents to GigaVUE V Series nodes. The options are GRE or VXLAN tunnels.
G-vTAP Agent MTU (Maximum Transmission Unit)	The Maximum Transmission Unit (MTU) is the maximum size of each packet that the tunnel endpoint can carry from the G-vTAP agent to the GigaVUE V Series node. For GRE, the default value is 9001. For VXLAN, the default value is 8951. However, the G-vTAP agent tunnel MTU should be 50 bytes less than the agent's destination interface MTU size.

5. Click **Save**.
6. To view the G-vTAP Controllers connection status, click **Visibility Fabric > G-vTAP Controllers**.

   The G-vTAP Controller instance takes a few minutes to fully initialize. After the initialization is complete, the connection status is displayed as **OK**. Refer to Figure 2-25 on page 46.

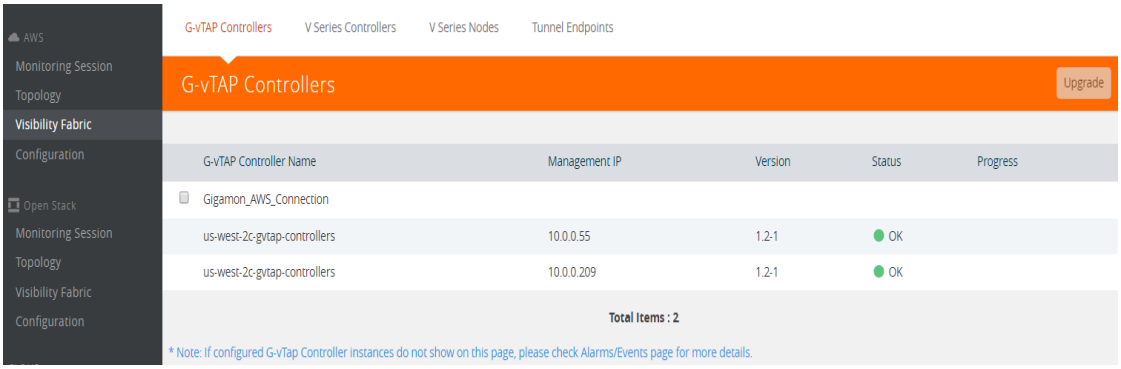

Figure 2-25: G-vTAP Controllers Connection Status

The G-vTAP Controller launch is displayed as an event in the **Cloud** > **Alarms/Events** page.

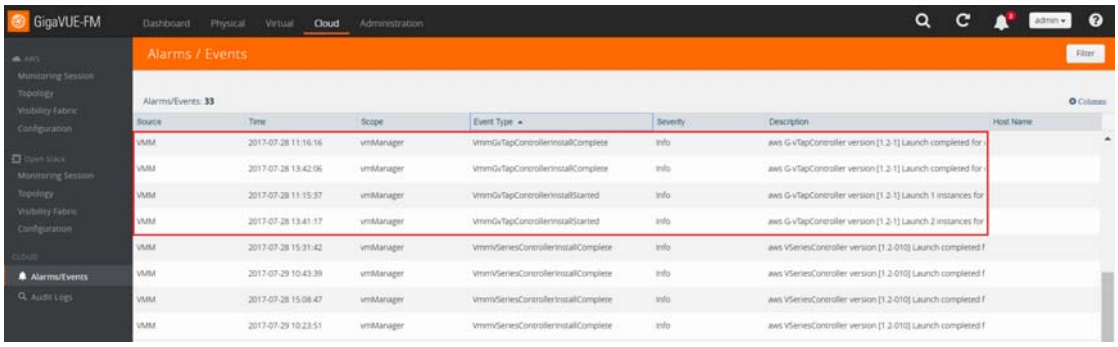

Figure 2-26: G-vTAP Controllers Events in Alarms/Events Page

To view the G-vTAP Controllers launched in your VPC:

1. Login to the AWS account and select **Services > EC2**.

2. In the left navigation pane, click **Instances**. The G-vTAP Controllers launched in your VPC can be seen as shown Figure 2-27 on page 47.

Name	Instance ID	Instance Type	Availability Zone	Instance State
us-west-2c-gvtap-controllers	i-081d78999e9ec8278	t2.medium	us-west-2c	running
us-west-2c-gvtap-controllers	i-065069f01bc1df0e3	t2.medium	us-west-2c	running

Figure 2-27: G-vTAP Controllers Configured in AWS

Configuring the GigaVUE V Series Controllers

GigaVUE V Series Controller manages multiple GigaVUE V Series nodes and orchestrates the flow of traffic from GigaVUE V Series nodes to the monitoring tools. GigaVUE-FM uses one or more GigaVUE V Series Controllers to communicate with the GigaVUE V Series nodes.

A single GigaVUE V Series Controller can manage up to 100 GigaVUE V Series nodes.

To configure the GigaVUE V Series Controller, do the following:

1. Select **AWS > Configuration > V Series Controllers**.
2. Click **New**. The V Series Controller Configuration page opens.

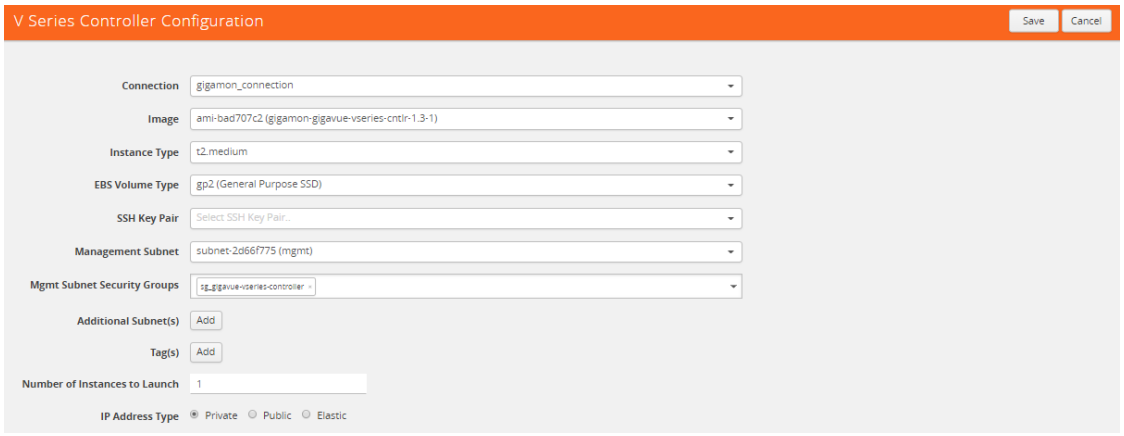

Figure 2-28: Configuring the GigaVUE V Series Controller

3. Follow the Step 4, Step 5, and Step 6 as described in *Configuring the G-vTAP Controllers* on page 42 and select the appropriate information for GigaVUE V Series Controllers.

To view the *GigaVUE V Series Controller* configured in your VPC:

1. Login to the AWS account and select **Services > EC2**.

2. In the left navigation pane, click **Instances**. The *GigaVUE V Series Controller* configured in your VPC can be seen as shown Figure 2-27 on page 47.

Figure 2-29: GigaVUE V Series Controllers Configured in AWS

Configuring the GigaVUE V Series Nodes

GigaVUE® V Series node is a visibility node that aggregates mirrored traffic from multiple G-vTAP agents. It applies filters, manipulates the packets using GigaSMART applications, and distributes the optimized traffic to cloud-based tools or backhaul to Gigamon Visibility Platform using the standard IP GRE tunnels.

GigaVUE V Series nodes can be successfully launched only after GigaVUE V Series Controller is fully initialized and the status is displayed as OK.

To launch a GigaVUE V Series node, do the following:

1. Select **AWS > Configuration > V Series Nodes**.
2. Click **New**. The V Series Node Configuration page is displayed as shown in Figure 2-30 on page 48.

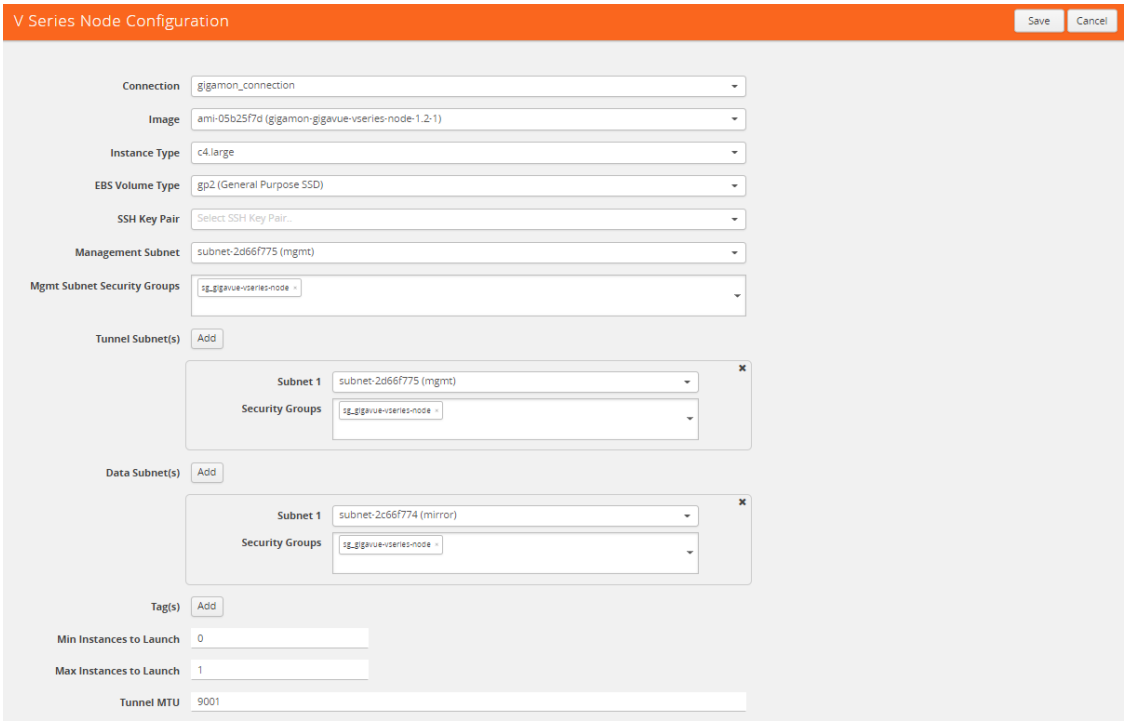

Figure 2-30: Configuring the GigaVUE V Series Node

NOTE: Make sure the GigaVUE V Series node version matches with the GigaVUE V Series Controller version that is already configured.

3. Enter or select the appropriate information as shown in Table 2-5 on page 43.

Table 2-6: Fields for GigaVUE V Series Configuration

Fields	Description
Connection	The name of the AWS connection.
Image	The GigaVUE V Series node image. **NOTE:** For GigaVUE-FM 5.2 and above, only the GigaVUE V Series node v1.3 is supported. The version number of GigaVUE V Series node must match with the version number of the GigaVUE V Series Controller.
Instance Type	The instance type for the GigaVUE V Series node. The recommended minimum instance type is c4. large.
EBS Volume Type	The Elastic Block Store (EBS) volume that you can attach to a single G-vTAP Controller instance. The available options are gp2 (General Purpose SSD), io1 (Provisioned IOPS SSD), and standard (Magnetic).
SSH KeyPair	The SSH key pair for the GigaVUE V Series node. For more information about SSH key pair, refer to *Key Pairs* on page 22.
Management Subnet	The public subnet that is used for communication between the GigaVUE V Series Controller and the GigaVUE V Series node.
Mgmt Subnet Security Groups	The security group created for the management subnet. For more information, refer to *Security Group* on page 20.
Tunnel Subnet(s)	The subnet that the GigaVUE V Series node uses to communicate with the monitoring tools or GigaVUE H Series node. The tunnel subnet can be same as the management subnet.
Data Subnet(s)	The subnet that receives the mirrored GRE tunnel traffic from the G-vTAP agents.
Tag(s)	(Optional) The key name and value that helps to identify the GigaVUE V Series node instances in your AWS environment. For example, you might have GigaVUE V Series node deployed in many regions. To distinguish these GigaVUE V Series node based on the regions, you can provide a name that is easy to identify such as us-west-2-vseries. 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-vseries.
Min Instances to Launch	The minimum number of GigaVUE V Series nodes to be launched in the AWS connection. The minimum number of instances that can be entered is 0. When 0 is entered, no GigaVUE V Series nodes are launched.
Max Instances to Launch	The maximum number of GigaVUE V Series nodes that can be launched in the AWS connection. When the number of instances per V Series node exceeds the max instances specified in this field, increase the number in the Max Instances to Launch. When additional V Series nodes are launched, GigaVUE-FM rebalances the instances assigned to the nodes. This can result in a brief interruption of traffic.
Tunnel MTU (Maximum Transmission Unit)	The Maximum Transmission Unit (MTU) on the outgoing tunnel endpoints of the GigaVUE V Series node when a monitoring session is deployed. The default value is 9001.

To view the *GigaVUE V Series nodes* launched in your VPC:

1. Login to the AWS account and select **Services > EC2**.
2. In the left navigation pane, click **Instances**. The *GigaVUE V Series nodes* launched in your VPC can be seen as shown Figure 2-31 on page 50.

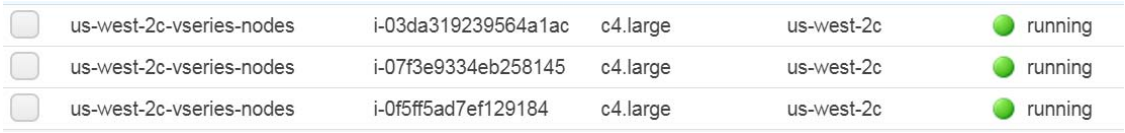

Figure 2-31: GigaVUE V Series Nodes Configured in AWS

NOTE:

- The recommended minimum instance type for the GigaVUE V Series node is c4.large.
- Certain availability zones may sometimes throw an insufficient instance capacity error. This is because AWS does not currently have enough capacity to service your request. When this error is displayed, you can launch the instance using a different instance type and resize at a later stage. Refer to the following link to select another instance type:

  http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ebs-ec2-config.html
- The insufficient instance capacity error can be viewed only on Alarms/Events page. Refer to *Alarms and Events* on page 101.
- To change the instance type at a later stage, the active monitoring sessions must be undeployed and the GigaVUE V Series nodes must be relaunched with the new configuration settings.

Chapter 3

Configuring Monitoring Sessions in AWS

This chapter describes how to setup the tunnel endpoints to receive and send traffic from the GigaVUE V Series node, and how to filter, manipulate, and send the traffic from the GigaVUE V Series node to the monitoring tools or GigaVUE H Series node.

Refer to the following sections for details:

Overview of Visibility Components

The GigaVUE V Series node aggregates the traffic from multiple G-vTAP agents and filters them using maps. It applies intelligence and optimization to the aggregated traffic using GigaSMART applications such as Flow Mapping™, sampling, slicing, and masking, and distributes them to the tunnel endpoints.

Table 3-1 on page 52 lists the components of the monitoring session:

Table 3-1: Components of Traffic Visibility Sessions

Parameter	Description
Map	A map (M) is used to filter the traffic flowing through the GigaVUE V Series node. It is a collection of one or more rules (R). The traffic passing through a map can match one or more rules defined in the map.
Rule	A rule (R) contains specific filtering criteria that the packets must match. The filtering criteria lets you determine the target instances and the (egress or ingress) direction of tapping the network traffic. The rules must contain the appropriate Layer 2 (L2) to Layer 4 (L4) filters defined in them. For example, if you want to filter the traffic for HTTP Port 80, you must select the following criteria: • Layer 2—Ethertype IPv4 or IPv6 • Layer 3—Protocol TCP • Layer 4—Port Destination 80 By default, a rule always displays conditions based on the attributes of L2. Refer to Figure 3-1 on page 52. 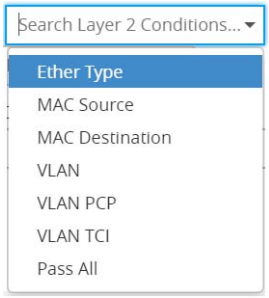 *Figure 3-1: Layer 2 Rule Conditions* A rule is also associated with priority and action set.
Priority	A priority determines the order in which the rules are executed. The greater the value, the higher the priority. The priority value can range from 0 to 99.

Table 3-1: Components of Traffic Visibility Sessions

Parameter	Description
Action Set	An action set is an exit point in a map that you can drag and create links to the other maps, applications, and the monitoring tools. A single map can have multiple action sets. A single action set can have multiple links connecting to maps and applications. In the following example (refer to Figure 3-2 on page 53), the packets that match the rules in Action Set 0 are forwarded to a tunnel endpoint. The packets that match the rules in Action Set 1 are forwarded to another map. 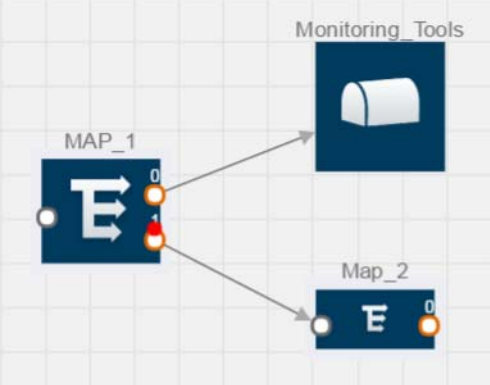 *Figure 3-2: Action Set* A single action set can have up to 8 links connecting the same destination point. The same packets from the map are replicated in 8 different links. Refer to Figure 3-3 on page 53. 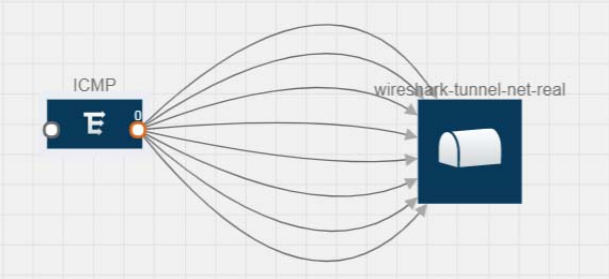 *Figure 3-3: Action Set with Multiple Links*
Link	A link directs the packets to flow from a map to the destination. The destination could be the other maps, applications, and the monitoring tools. In Figure 3-2 on page 53, the link originating from action set 0 is moving the traffic from MAP_1 to Monitoring_Tools. A link lets you add header transformation to the packets passing through it before they are sent to the destination. This is called Header Transformation. This transformation is supported only with GigaVUE V Series node v1.2-1 and above. For more information about Header Transformation, refer to *Adding Header Transformations* on page 91.
Group	A group is a collection of maps that are pre-defined and saved in the map library for reuse.
Application	An application performs operations such as sampling, slicing, and masking on the traffic.
Inclusion Map	An inclusion map determines the instances or ENIs to be included for monitoring. This map is used only for target selection.
Exclusion Map	An exclusion map determines the instances or ENIs to be excluded from monitoring. This map is used only for target selection.

Table 3-1: Components of Traffic Visibility Sessions

Parameter	Description
Target	A target determines the instances that are to be monitored. Targets are determined based on the following formula: Target = (Maps ∩ Inclusion map) – Exclusion map
Automatic Target Selection (ATS)	A built-in feature that automatically selects the EC2 instances and ENIs based on the rules defined in the maps, inclusion maps, and exclusion maps in the monitoring session. For example, if you create a rule determining the MAC source address in a map and a subnet in the inclusion map, the egress traffic from all instances or ENIs matching the MAC address in the specified subnet is selected for tapping the traffic.
Tunnel	A tunnel lists the monitoring tools to which the traffic matching the filtered criteria is routed.

Creating Tunnel Endpoints

The customized traffic from the GigaVUE V Series node is distributed to the tunnel endpoints using a standard L2 Generic Routing Encapsulation (GRE) or Virtual Extensible LAN (VXLAN) tunnel.

To create the tunnel endpoints:

1. Select **AWS > Configuration > Tunnel Library**.
2. Click **New**. The Add Tunnel page is displayed as shown in Figure 3-4 on page 54.

Figure 3-4: Adding a Tunnel Endpoint

3. Select or enter the appropriate information as shown in Table 3-2 on page 54.

Table 3-2: Fields for Tunnel Endpoint

Field	Description
Alias	The name of the tunnel endpoint. **NOTE:** Do not enter spaces in the alias name.
Description	The description of the tunnel endpoint.
Type	The type of the tunnel. Select L2GRE or VXLAN to create a tunnel. If you choose VXLAN, you must enter the remote tunnel port.

Table 3-2: Fields for Tunnel Endpoint

Field	Description
Traffic Direction	The direction of the traffic flowing through the GigaVUE V Series node. Choose **Out** for creating a tunnel from the GigaVUE V Series node to the destination endpoint. **NOTE:** Traffic Direction **In** is not supported in the current release.
Remote Tunnel IP	The IP address of the tunnel destination endpoint. **NOTE:** You cannot create two tunnels from a GigaVUE V Series node to the same IP address.
Remote Tunnel Port	(Optional) The port number of the tunnel destination endpoint. **NOTE:** This option is displayed only if you select VXLAN as the tunnel type.
Source Subnet	(Optional) The source subnet CIDR that GigaVUE V Series nodes must use to send the traffic to the remote tunnel endpoint. For example, if a GigaVUE V Series node has two tunnel subnets (subnet A and B) and the V series node has to use subnet B to send traffic to the remote tunnel endpoint, then the subnet CIDR of network B must be specified in the Source Subnet text box.

4. Click **Save**. The tunnel endpoints are added successfully. Refer to

Figure 3-5: Tunnel Endpoints Created

Creating a Monitoring Session

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

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

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

- *Creating a New Session* on page 56
- *Creating a Map* on page 57
- *Adding Applications to the Monitoring Session* on page 62
- *Deploying the Monitoring Session* on page 88

- *Adding Header Transformations* on page 91
- *Viewing the Statistics* on page 94
- *Viewing the Topology* on page 95

Creating a New Session

You can create multiple monitoring sessions within a single VPC connection.

To create a new session:

1. Select **AWS > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Enter the appropriate information in the MONITORING SESSION INFO as shown in the Table 3-3 on page 56.

Table 3-3: Fields for Session Info

Field	Description
Name	The name of the monitoring session.
Connection	The alias name of the AWS connection.

Creating a Map

Each map can have up to 32 rules associated with it. Table 3-4 on page 57 lists the various rule conditions that you can select for creating a map, inclusion map, and exclusion map.

Table 3-4: Conditions for the Rules

Conditions	Description
L2, L3, and L4 Filters	
Ether Type	The packets are filtered based on the selected ethertype. The following conditions are displayed: • IPv4 • IPv6 • ARP • RARP • Other **L3 Filters** If you choose IPv4 or IPv6, the following L3 filter conditions are displayed: • Protocol • IP Fragmentation • IP Time to live (TTL) • IP Type of Service (TOS) • IP Explicit Congestion Notification (ECN) • IP Source • IP Destination **L4 Filters** If you select TCP or UDP protocol, the following L4 filter conditions are displayed: • Port Source • Port Destination
MAC Source	The egress traffic from the instances or ENIs matching the specified source MAC address is selected.
MAC Destination	The ingress traffic from the instances or ENIs matching the specified destination MAC address is selected.
VLAN	All the traffic matching the specified IEEE 802.1q Virtual LAN tag is filtered. Specify a number from 0 to 4094.
VLAN Priority Code Point (PCP)	All the traffic matching the specified IEEE 802.1q Priority Code Point (PCP) is filtered. Specify a value between 0 to 7.
VLAN Tag Control Information (TCI)	All the traffic matching the specified VLAN TCI value is filtered. Specify the exact TCI value.
Pass All	All the packets coming from the monitored instances are passed through the filter. When Pass All is selected, the L3 and L4 filters are disabled.

When you select a condition without source or destination specified, then both egress and ingress traffic is selected for monitoring the traffic. For example, if you select IPv4

as the Ether Type, TCP as the protocol, and do not specify IP source or destination, then both egress and ingress traffic is selected for monitoring purpose.

When you select a condition with either source or destination specified, it determines the direction based on the selection. For example, if only IP source is selected as shown in Figure 3-6 on page 58, the egress traffic from the instances in the subnet 10.0.1.0/24 is selected for monitoring the traffic.

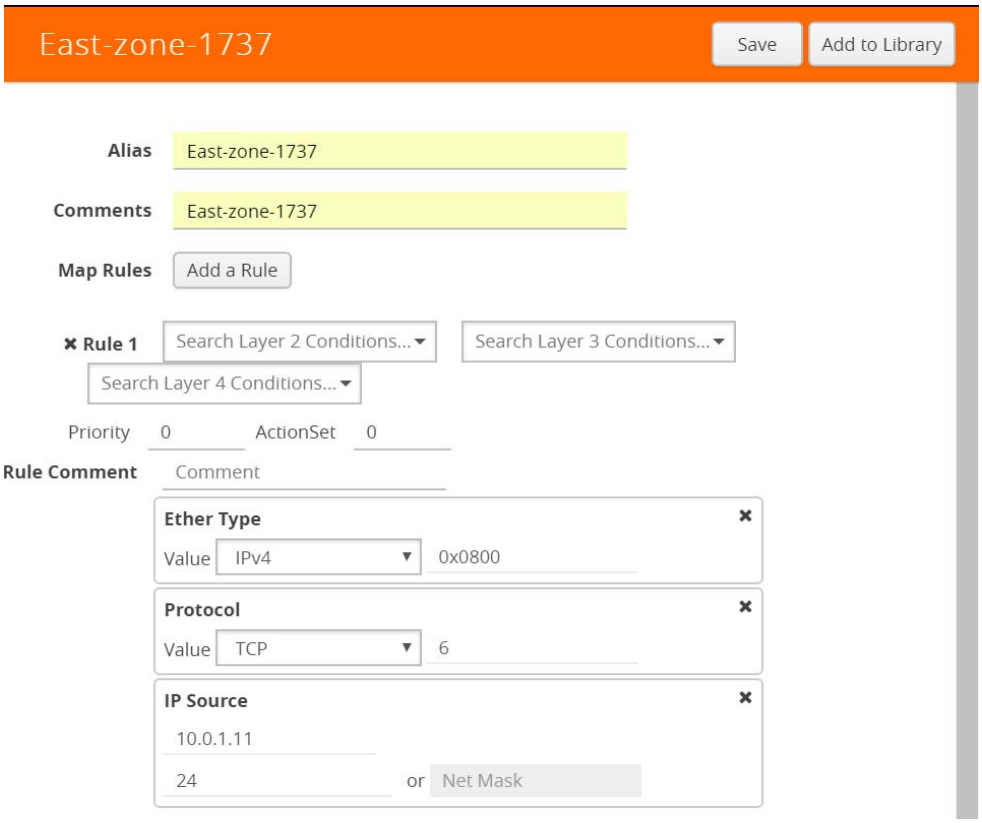

Figure 3-6: Creating a Map for Tapping Egress Traffic

NOTE: You can create Inclusion and Exclusion Maps using all default conditions except Ether Type and Pass All.

To create a new map:

1. Select **AWS > Monitoring Session**.
2. Click **New**. The Monitoring Sessions page is displayed.
3. Create a new session. Refer to *Creating a New Session* on page 56.
4. From **Maps**, drag and drop a new map template to the workspace. If you are creating an exclusion or inclusion map, drag and drop a new map template to their respective section at the bottom of the workspace.

The new map page is displayed as shown in Figure 3-7 on page 59.

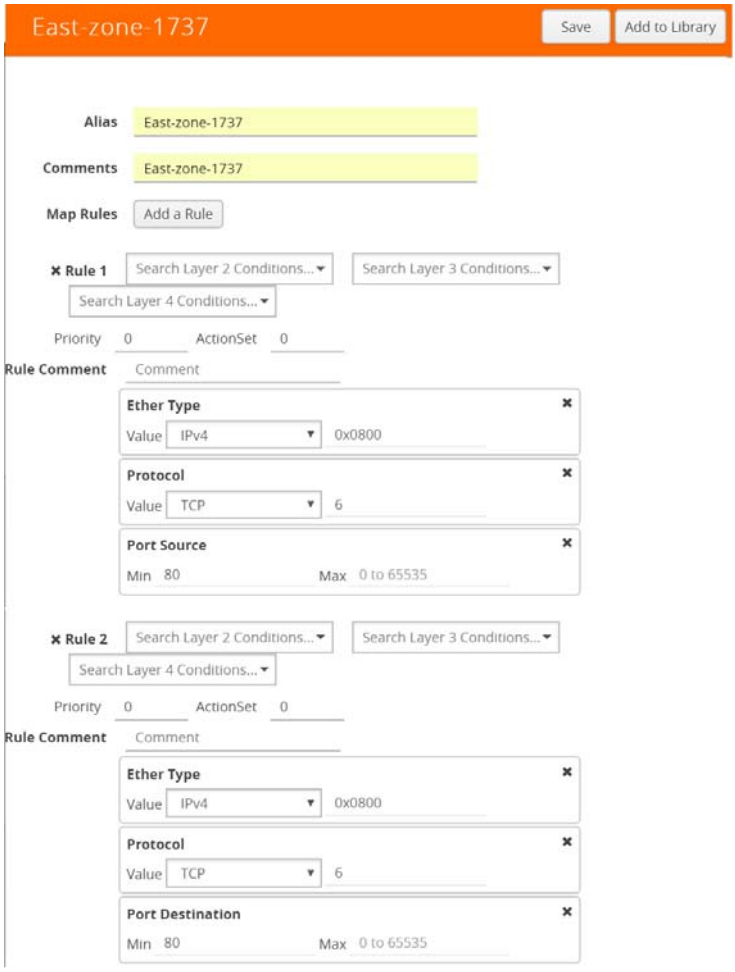

Figure 3-7: Creating a New Map

5. Enter the appropriate information for creating a new map as shown in Table 3-5 on page 59.

Table 3-5: Fields for Creating a New Map

Parameter	Description
Alias	The name of the new map. **NOTE:** The name can contain alphanumeric characters with no spaces.
Comments	The description of the map.

Table 3-5: Fields for Creating a New Map

Parameter	Description
Map Rules	The rules for filtering the traffic in the map. To add a map rule: **a.** Click **Add a Rule**. **b.** Select a condition from the **Search L2 Conditions** drop-down list and specify a value. Based on this selection, the Search L3 Conditions drop-down list is automatically updated. Refer to Figure 3-8 on page 60.. 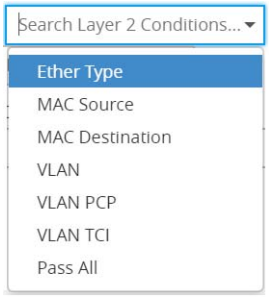 *Figure 3-8: L2 Conditions* **c.** Select a condition from the **Search L3 Conditions** drop-down list and specify a value. Refer to Figure 3-9 on page 60. 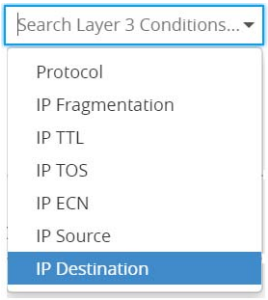 *Figure 3-9: L3 Conditions* **d.** (Optional) If you have selected TCP or UDP as the protocol in the L3 conditions, then select Port Source or Port Destination from the Search L4 Conditions drop-down list and specify a value. If you have selected conditions other than TCP or UDP, then the Search L4 Conditions drop-down list is disabled. Refer to Figure 3-10 on page 60. 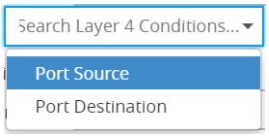 *Figure 3-10: L4 Conditions*

Table 3-5: Fields for Creating a New Map

Parameter	Description
Map Rules	**e.** (Optional) In the Priority and Action Set box, assign a priority and action set. **f.** (Optional) In the Rule Comment box, enter a comment for the rule. **NOTE:** Repeat steps **b** through **f** to add more conditions. **NOTE:** Repeat steps **a** through **f** to add nested rules.

NOTE: Do not create duplicate map rules with the same priority.

6. Click **Add to Library** and save the map for reuse using one of the following ways:
   - Select an existing group from the **Select Group** list and click **Save**.
   - Enter a name for the new group in the **New Group** field and click **Save.**

   NOTE: The maps saved in the Map Library can be reused in all the monitoring sessions in the VPC.
7. Click **Save**.

To edit or delete a map, click a map and select **Details** to edit the map or **Delete** to delete the map as shown in Figure 3-11 on page 61.

Figure 3-11: Editing or Deleting a Map

Click the **Show Targets** button to view the monitoring targets highlighted in blue. Refer to Figure 3-12 on page 61.

Figure 3-12: Viewing the Topology

Click on ⤢ to expand the **Targets** dialog box. Click on ≡ to change the view from topology to viewing the instance names. To view more details about the instance tag

name, direction of tapping, and so on, click the arrow next to the instance name. Refer to Figure 3-13 on page 62.

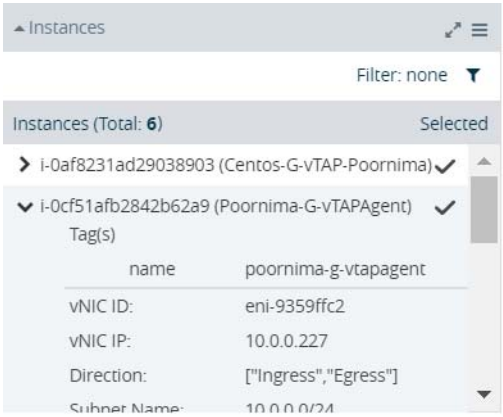

Figure 3-13: Viewing the instance Details

Filter the instances based on the Instance Name Prefix, IP address, or the MAC address. Refer to Figure 3-14 on page 62.

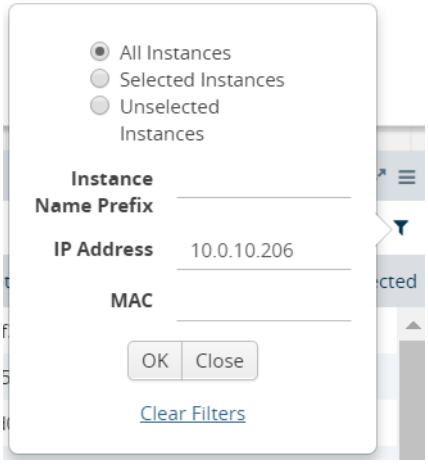

Figure 3-14: Filtering the instances

Adding Applications to the Monitoring Session

Gigamon supports the following GigaSMART applications with the Visibility Platform for AWS:

- *Sampling* on page 62
- *Slicing* on page 64
- *Masking* on page 66
- *NetFlow* on page 67

Sampling

Sampling lets you sample the packets randomly based on the configured sampling rate and then forwards the sampled packets to the monitoring tools.

To add a sampling application:

1. Drag and drop **Sample** from **APPLICATIONS** to the graphical workspace.

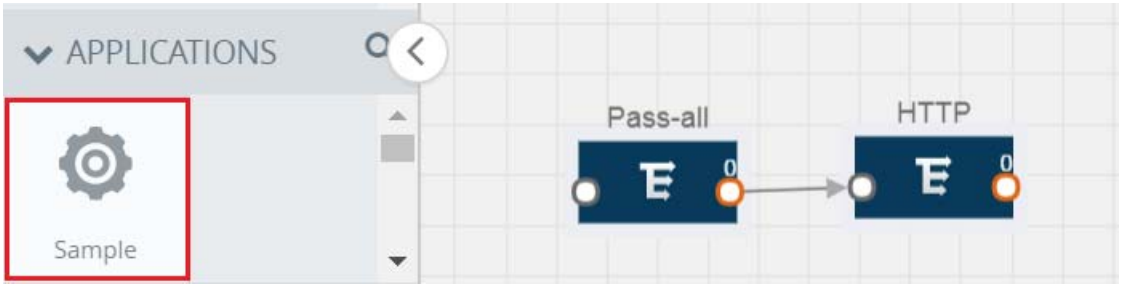

Figure 3-15: Dragging the Sample Application

2. Click **Sample** and select **Details**.

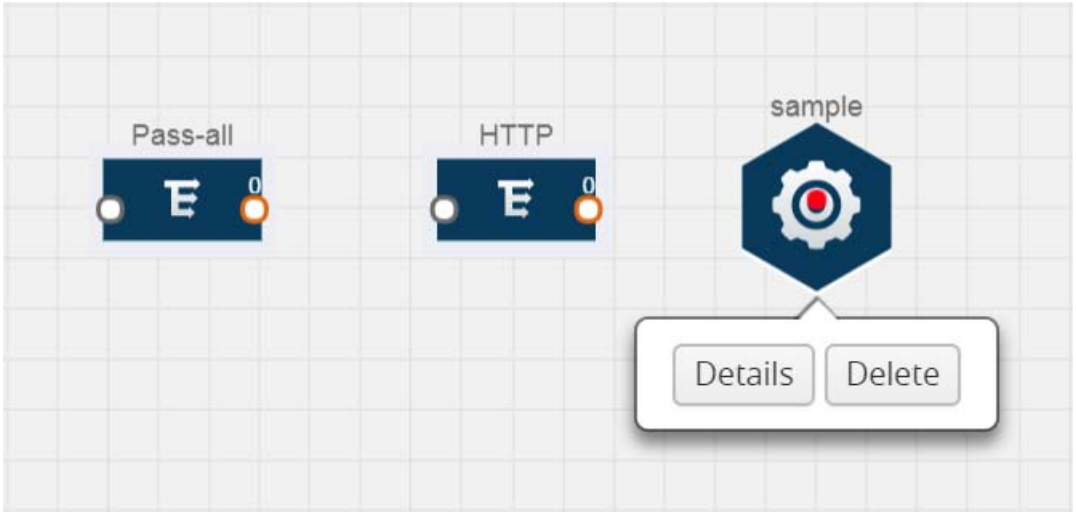

Figure 3-16: Selecting Details

3. In the **Alias** field, enter a name for the sample.

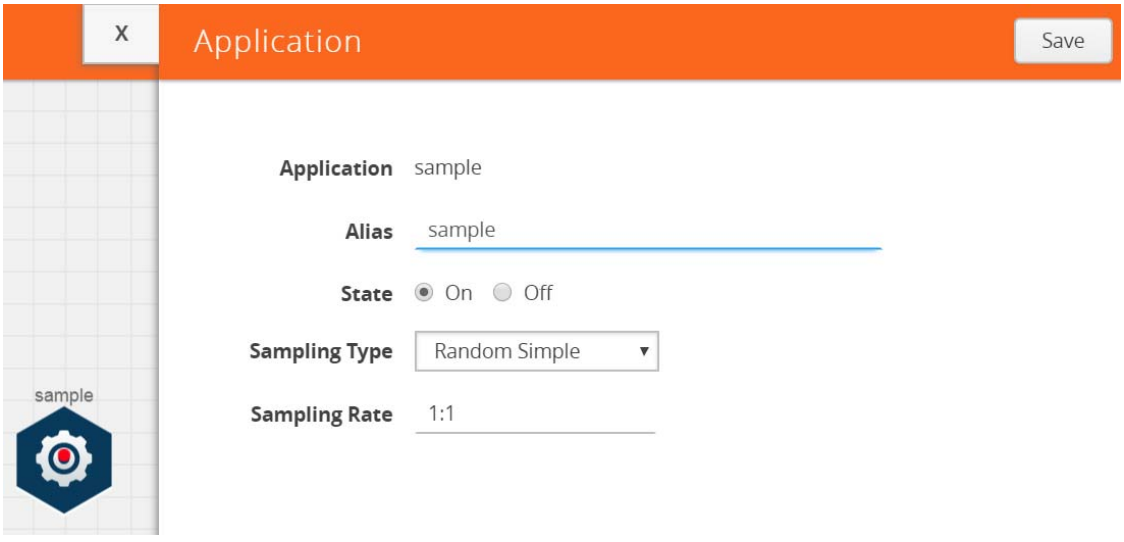

Figure 3-17: Viewing Sample Application Quick View

4. For State, select the **On** check box to determine that the application is sampling packets randomly. Select the **Off** check box to determine that the application is not currently sampling the packets. The state can be changed at anytime whenever required.
5. From the Sampling Type drop-down list, select the type of sampling:
   - **Random Simple** — The first packet is selected randomly. The subsequent packets are also selected randomly based on the rate specified in the **Sampling Rate** field.

     For example, if the first packet selected is 5 and the sampling rate is 1:10, after the 5th packet a random 10 packets are selected for sampling.
   - **Random Systematic** —The first packet is selected randomly. Then, every nth packet is selected, where n is the value specified in the **Sampling Rate** field.

     For example, if the first packet selected is 5 and the sampling rate is 1:10, then every 10th packet is selected for sampling: 15, 25, 35, and so on.
6. In the **Sampling Rate** field, enter the ratio of packets to be selected. The default ratio is 1:1.
7. Click **Save**.

Slicing

Packet slicing lets you truncate packets after a specified header and slice length, preserving the portion of the packet required for monitoring purposes.

To add a slicing application:

1. Drag and drop **Slice** from **APPLICATIONS** to the graphical workspace.

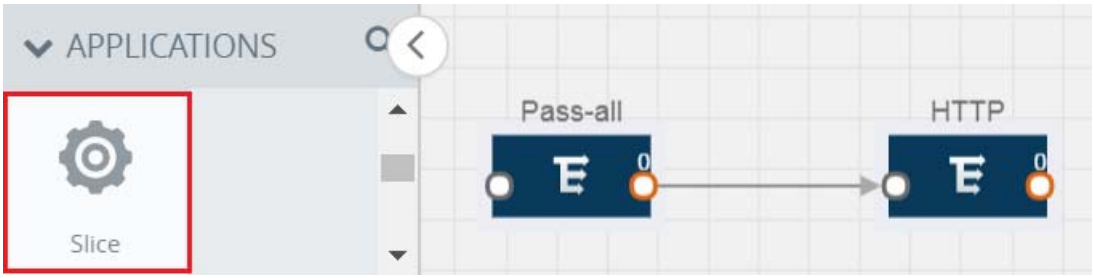

Figure 3-18: Dragging the Slice Application

2. Click the Slice application and select **Details**.

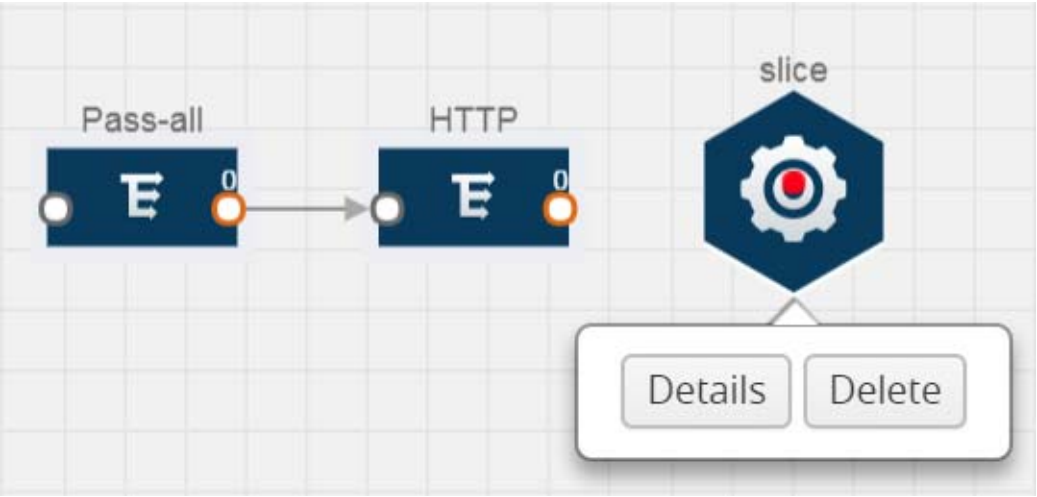

Figure 3-19: Selecting Details

3. In the **Alias** field, enter a name for the slice.

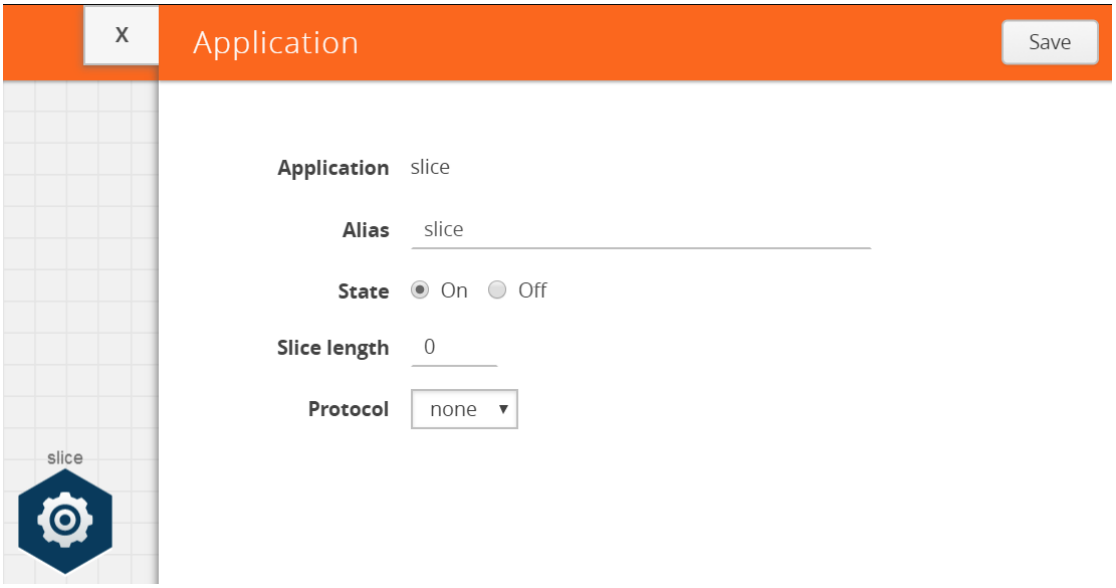

Figure 3-20: Viewing Slice Application Quick View

4. For State, select the **On** check box to determine that the application is slicing packets. Select the **Off** check box to determine that the application is not currently slicing the packets. The state can be changed at a later time whenever required.
5. In the Slice Length field, specify the length of the packet that must be sliced.
6. From the Protocol drop-down list, specify an optional parameter for slicing the specified length of the protocol. The options are as follows:
   - None
   - IPv4
   - IPv6
   - UDP

- TCP

7. Click **Save**.

Masking

Masking lets you overwrite specific packet fields with a specified pattern so that sensitive information is protected during network analysis.

To add a masking application:

1. Drag and drop **Mask** from **APPLICATIONS** to the graphical workspace.

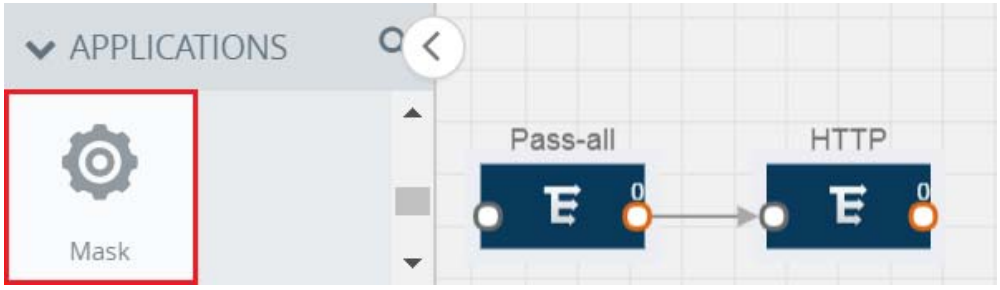

Figure 3-21: Dragging the Mask Application

2. Click the Mask application and select **Details**.

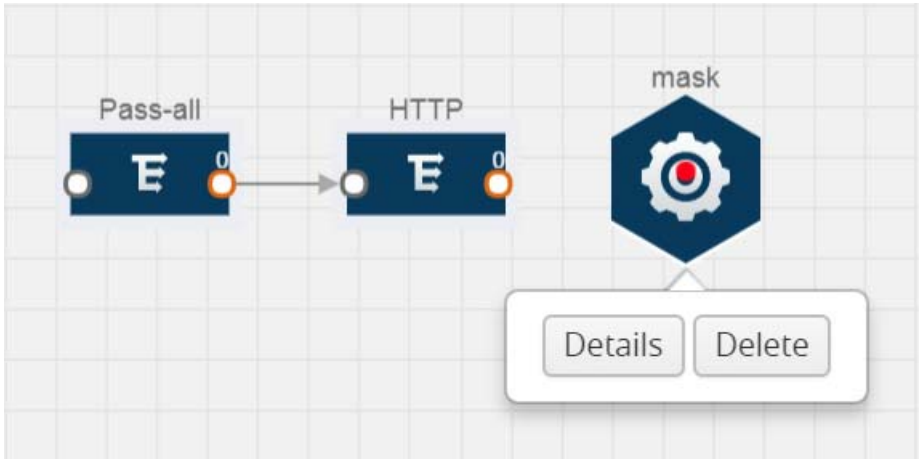

Figure 3-22: Selecting Details

3. In the **Alias** field, enter a name for the mask.

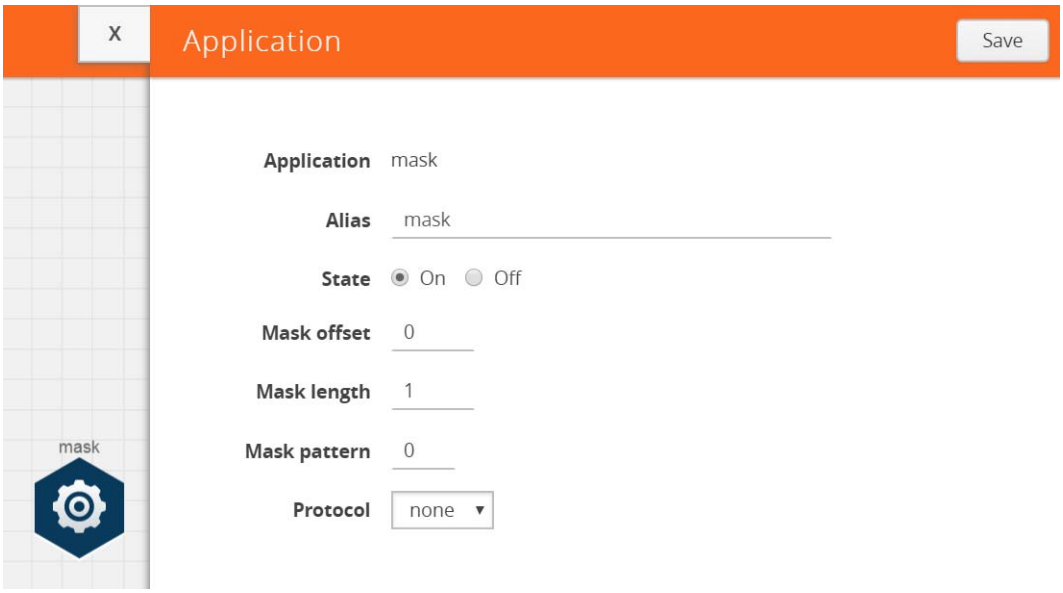

Figure 3-23: Viewing Mask Application Quick View

4. For State, select the **On** check box to determine that the application is masking packets. Select the **Off** check box to determine that the application is not currently masking the packets. The state can be changed at anytime whenever required.
5. In the Mask offset field, enter the offset from which the application should start masking data following the pattern specified in the Pattern field.

   The value can be specified in terms of either a static offset, that is, from the start of the packet or a relative offset, that is, from a particular protocol layer as specified in the Protocol field.
6. In the Mask length field, enter the length of the packet that must be masked.
7. In the Mask pattern field, enter the pattern for masking the packet. The value of the pattern is from 0 to 255.
8. From the Protocol drop-down list, specifies an optional parameter for masking packets on the data coming from the selected protocol.
9. Click **Save**.

NetFlow

NetFlow collects IP network traffic on all interfaces where NetFlow monitoring is enabled. It gathers information about the traffic flows and exports the NetFlow records, which includes data and templates, to at least one NetFlow collector. The application that serves as a NetFlow collector receives the NetFlow data sent from exporters, processes the information, and provides data visualization and security analytics.

Key benefits of NetFlow application:

- Compresses network information into a single flow record.
- Facilitates up to 99% reduction in data transferred.
- Accelerates the migration of mission-critical workloads to AWS.

- Provides summarized information on traffic source and destination, congestion, and class of service.
- Identifies and classifies DDOS attacks, viruses, and worms in real-time.
- Secures network against internal and external threats.
- Identifies top consumers and analyzes their statistics.
- Reduces the cost of security monitoring.
- Analyzes the network flows based on algorithms and behavior rather than signature matching.
- Analyzes east-west traffic between flows within and across VPCs.

The NetFlow application contains key elements that specify what to match in the flow, such as all packets with the same source and destination port, or the packets that come in on a particular interface. For information about Match/Key fields, refer to *Match/Key Fields* on page 69. A NetFlow record is the output generated by NetFlow. A flow record contains non-key elements that specify what information to collect for the flow, such as when the flow started or the number of bytes in the flow. For information about Match/Key fields, refer to *Collect/Non-Key Fields* on page 71.

Figure 3-24 on page 68 shows an example of a NetFlow application created on a GigaVUE V Series node in the monitoring session.

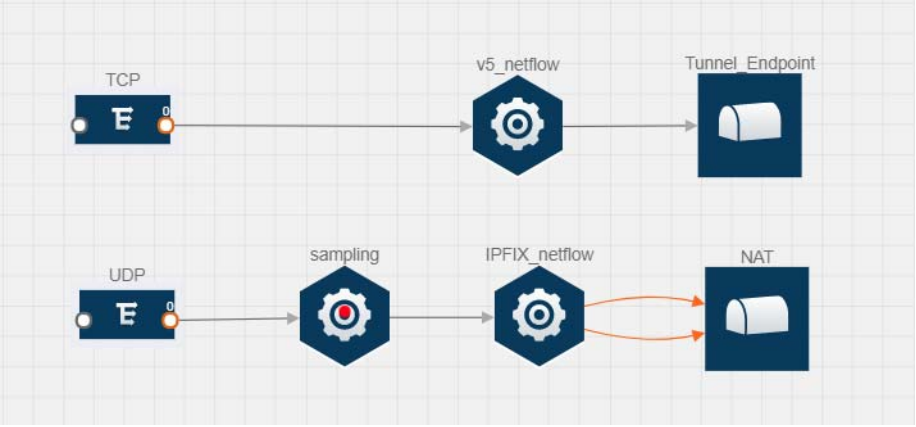

Figure 3-24: NetFlow on GigaVUE V Series Node

The NetFlow record generation is performed on GigaVUE V Series node running the NetFlow application. In Figure 3-24 on page 68, incoming packets from G-vTAP agents are sent to the GigaVUE V Series node. In the GigaVUE V Series node, one map sends the TCP packet to the version 5 NetFlow application. Another map sends the UDP packet to a sampling application. The map rules and applications such as slice, mask, and sample can only be applied prior to sending the data to NetFlow.

A NetFlow application examines the incoming packets and creates a single or multiple flows from the packet attributes. These flows are cached and exported based on the active and inactive cache timeout specified in the Netflow application configuration.

The flow records can be sent to a tunnel for full packet inspection or to a NAT device for flow inspection. NAT allows the NetFlow records to be directly transmitted to a collector

without a tunnel. For more information about NAT, refer to *Network Address Translation (NAT)* on page 77.

The Netflow application exports the flows using the following export versions:

- version 5—The fields in the NetFlow record are fixed.
- version 9—The fields are configurable, thus a template is created. The template contains information on how the fields are organized and in what order. It is sent to the collector before the flow record, so the collector knows how to decode the flow record. The template is sent periodically based on the configuration.
- IPFIX—The extended version of version 9 supports variable length fields as well as enterprise-defined fields.

Match/Key Fields

NetFlow v9 and IPFIX records allow you to configure Match/Key elements.

The supported Match/Key elements are outlined in the following table:

Table 3-6: Match/Key Elements

Match Type	Description	Supported NetFlow Versions
Data Link		
Destination MAC	Configures the destination MAC address as a key field.	v9 and IPFIX
Egress Dest MAC	Configures the post Source MAC address as a key field.	IPFIX
Ingress Dest MAC	Configures the IEEE 802 destination MAC address as a key field.	IPFIX
Source MAC	Configures the IEEE 802 source MAC address as a key field.	v9 and IPFIX
IPv4		
ICMP Type Code	Configures the type and code of the IPv4 ICMP message as a key field.	v9 and IPFIX
IPv4 Dest IP	Configures the IPv4 destination address in the IP packet header as a key field.	v9 and IPFIX
IPv4 ICMP Code	Configures the code of the IPv4 ICMP message as a key field.	IPFIX
IPv4 ICMP Type	Configures the type and code of the IPv4 ICMP message as a key field.	IPFIX
IPv4 Options	Configures the IPv4 options in the packets of the current flow as a key field.	IPFIX
IPv4 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9 and IPFIX
IPv4 Total Length	Configures the total length of the IPv4 packet as a key field.	IPFIX

Table 3-6: Match/Key Elements

Match Type	Description	Supported NetFlow Versions
Network		
IP CoS	Configures the IP Class Of Service (CoS) as a key field.	v9 and IPFIX
IP DSCP	Configures the value of a Differentiated Services Code Point (DSCP) encoded in the Differentiated Services field as a key field.	IPFIX
IP Header Length	Configures the length of the IP header as a key field.	IPFIX
IP Precedence	Configures the value of the IP Precedence as a key field.	IPFIX
IP Protocol	Configures the value of the protocol number in the IP packet header as a key field.	v9 and IPFIX
IP Total Length	Configures the total length of the IP packet as a key field.	IPFIX
IP TTL	For IPv4, configures the value of Time to Live (TTL) as a key field. For IPv6, configures the value of the Hop Limit field as a key field.	IPFIX
IP Version	Configures the IP version field in the IP packet header as a key field.	v9 and IPFIX
IPv6		
IPv6 Dest IP	Configures the IPv6 destination address in the IP packet header as a key field.	v9 and IPFIX
IPv6 Flow Label	Configures the value of the IPv6 flow label field in the IP packet header as a key field.	v9 and IPFIX
IPv6 ICMP Code	Configures the code of the IPv6 ICMP message as a key field.	IPFIX
IPv6 ICMP Type	Configures the type of the IPv6 ICMP message as a key field.	IPFIX
IPv6 ICMP Type Code	Configures the type and code of the IPv6 ICMP message as a key field.	IPFIX
IPv6 Payload Length	Configures the value of the payload length field in the IPv6 header as a key field.	IPFIX
IPv6 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9 and IPFIX
Transport		
L4 Dest Port	Configures the destination port identifier in the transport header as a key field.	v9 and IPFIX

Table 3-6: Match/Key Elements

Match Type	Description	Supported NetFlow Versions
L4 Src Port	Configures the source port identifier in the transport header as a key field.	v9 and IPFIX
TCP AcK Number	Configures the acknowledgment number in the TCP header as a key field.	IPFIX
TCP Dest Port	Configures the destination port identifier in the TCP header as a key field.	IPFIX
TCP Flags	Configures the TCP control bits observed for the packets of this flow as a key field.	v9 and IPFIX
TCP Header Length	Configures the length of the TCP header as a key field.	IPFIX
TCP Seq Number	Configures the sequence number in the TCP header as a key field.	IPFIX
TCP Src Port	Configures the source port identifier in the TCP header as a key field.	IPFIX
TCP Urgent	Configures the urgent pointer in the TCP header as a key field.	IPFIX
TCP Window Size	Configures the window field in the TCP header as a key field.	IPFIX
UDP Dest Port	Configures the destination port identifier in the UDP header as a key field.	IPFIX
UDP Src Port	Configures the source port identifier in the TCP header as a key field.	IPFIX

Collect/Non-Key Fields

NetFlow v9 and IPFIX records allow you to configure Collect/Non-Key elements.

The supported Collect/Non-Key elements are outlined in the following table:

Table 3-7: Collect/Non-Key Elements

Match Type	Description	Supported NetFlow Versions
Counter		
Byte Count	Configures the number of octets since the previous report in incoming packets for the current flow as a non-key field.	v9 and IPFIX
Packet Count	Configures the number of incoming packets since the previous report for this flow as a non-key field.	v9 and IPFIX
Data Link		
Destination MAC	Configures the destination MAC address as a non-key field.	v9 and IPFIX

Table 3-7: Collect/Non-Key Elements

Match Type	Description	Supported NetFlow Versions
Egress Des MAC	Configures the post source MAC address as a non-key field.	IPFIX
Ingress Des MAC	Configures the IEEE 802 destination MAC address as a non-key field.	IPFIX
Source MAC	Configures the IEEE 802 source MAC address as a non-key field.	v9 and IPFIX
Timestamp		
Flow End Millisec	Configures the absolute timestamp of the last packet of current flow in milliseconds as a non-key field.	IPFIX
Flow End Sec	Configures the flow start SysUp time as a non-key field.	IPFIX
Flow End Time	Configures the flow end SysUp time as a non-key field.	v9 and IPFIX
Flow Start Millisec	Configures the value of the IP Precedence as a non-key field.	IPFIX
Flow Start Sec	Configures the absolute timestamp of the first packet of this flow as a non-key field.	IPFIX
Flow Startup Time	Configures the flow start SysUp time as a non-key field.	v9 and IPFIX
Flow		
Flow End Reason	Configures the reason for Flow termination as a non-key field.	IPFIX
IPv4		
ICMP Type Code	Configures the type and code of the IPv4 ICMP message as a non-key field.	v9 and IPFIX
IPv4 Dest IP	Configures the IPv4 destination address in the IP packet header as a non-key field.	v9 and IPFIX
IPv4 ICMP Code	Configures the code of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 ICMP Type	Configures the type of the IPv4 ICMP message as a non-key field.	IPFIX
IPv4 Options	Configures the IPv4 options in the packets of the current flow as a non-key field.	IPFIX
IPv4 Src IP	Configures the IPv6 source address in the IP packet header as a non-key field.	v9 and IPFIX
IPv4 Total Length	Configures the total length of the IPv4 packet as a non-key field.	IPFIX
Network		

Table 3-7: Collect/Non-Key Elements

Match Type	Description	Supported NetFlow Versions
IP CoS	Configures the IP Class Of Service (CoS) as a key field.	v9
IP Protocol	Configures the value of the protocol number in the IP packet header as a key field.	v9
IP Version	Configures the IP version field in the IP packet header as a key field.	v9
IPv6		
IPv6 Dest IP	Configures the IPv6 destination address in the IP packet header as a key field.	v9
IPv6 Flow Label	Configures the value of the IPv6 flow label field in the IP packet header as a key field.	v9
IPv6 Src IP	Configures the IPv6 source address in the IP packet header as a key field.	v9
Transport		
L4 Dest Port	Configures the destination port identifier in the transport header as a non-key field.	v9 and IPFIX
L4 Src Port	Configures the source port identifier in the transport header as a non-key field.	v9 and IPFIX
TCP AcK Number	Configures the acknowledgment number in the TCP header as a non-key field.	IPFIX
TCP Dest Port	Configures the destination port identifier in the TCP header as a non-key field.	IPFIX
TCP Flags	Configures the TCP control bits observed for the packets of this flow as a non-key field.	v9 and IPFIX
TCP Header Length	Configures the length of the TCP header as a non-key field.	IPFIX
TCP Seq Number	Configures the sequence number in the TCP header as a non-key field.	IPFIX
TCP Src Port	Configures the source port identifier in the TCP header as a non-key field.	IPFIX
TCP Urgent	Configures the urgent pointer in the TCP header as a non-key field.	IPFIX
TCP Window Size	Configures the window field in the TCP header as a non-key field.	IPFIX
UDP Dest Port	Configures the destination port identifier in the UDP header as a non-key field.	IPFIX
UDP Src Port	Configures the source port identifier in the UDP header as a non-key field.	IPFIX

Adding a Version 5 NetFlow Application

To add a version 5 NetFlow application:

1. Drag and drop **NetFlow** from **APPLICATIONS** to the graphical workspace.

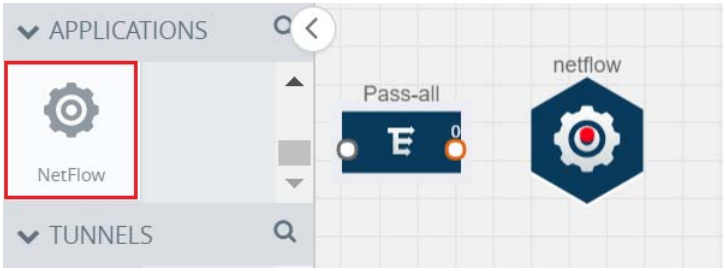

Figure 3-25: Dragging the NetFlow Application

2. Click the NetFlow application and select **Details**. A quick view is displayed for configuring the NetFlow application.

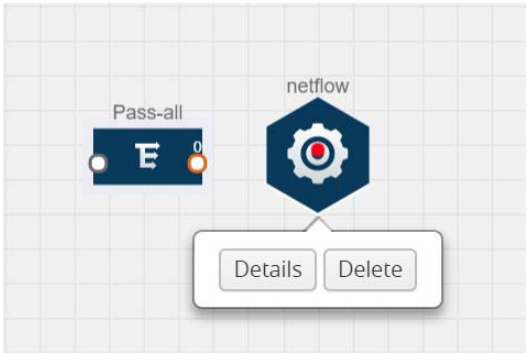

Figure 3-26: Selecting Details

3. In the **Alias** field, enter a name for the v5 NetFlow application.

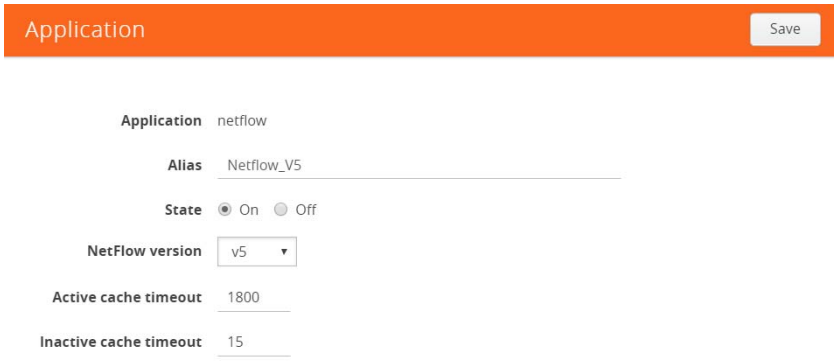

Figure 3-27: Viewing v5 NetFlow Application Quick View

4. For State, select the **On** check box to determine that the application is currently running. Select the **Off** check box to determine that the application is currently not running. The state can be changed at anytime whenever required.
5. From the **NetFlow version** drop-down list, select v5.

6. In **Active cache timeout**, enter the number of seconds that an active flow record must remain in the cache before it is exported and removed. The default value is 1800 seconds.
7. In **Inactive cache timeout**, enter the number of seconds an inactive flow record must remain in the cache before it times out. The default value is 15 seconds.
8. Click **Save**.

For some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes, refer to *NetFlow Examples* on page 80.

Adding a Version 9 and IPFIX NetFlow Application

To add a v9 and IPFIX NetFlow application:

1. Drag and drop **NetFlow** from **APPLICATIONS** to the graphical workspace.

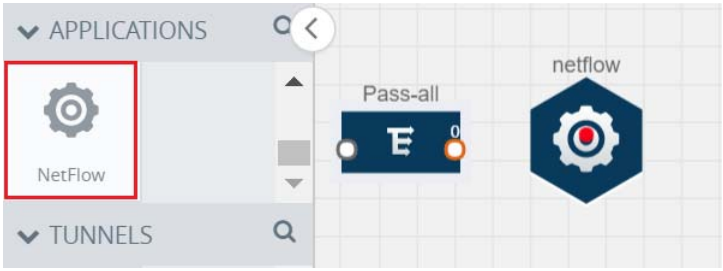

Figure 3-28: Dragging the NetFlow Application

2. Click the NetFlow application and select **Details**. A quick view is displayed for configuring the NetFlow application.

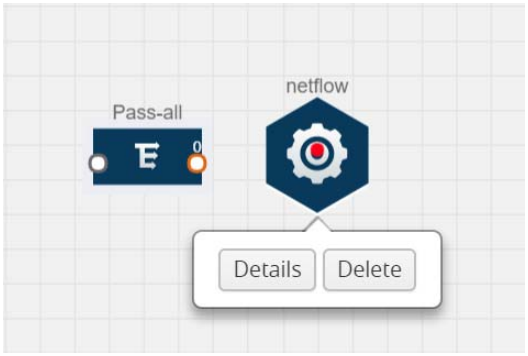

Figure 3-29: Selecting NetFlow Details

3. In the **Alias** field, enter a name for the NetFlow application.

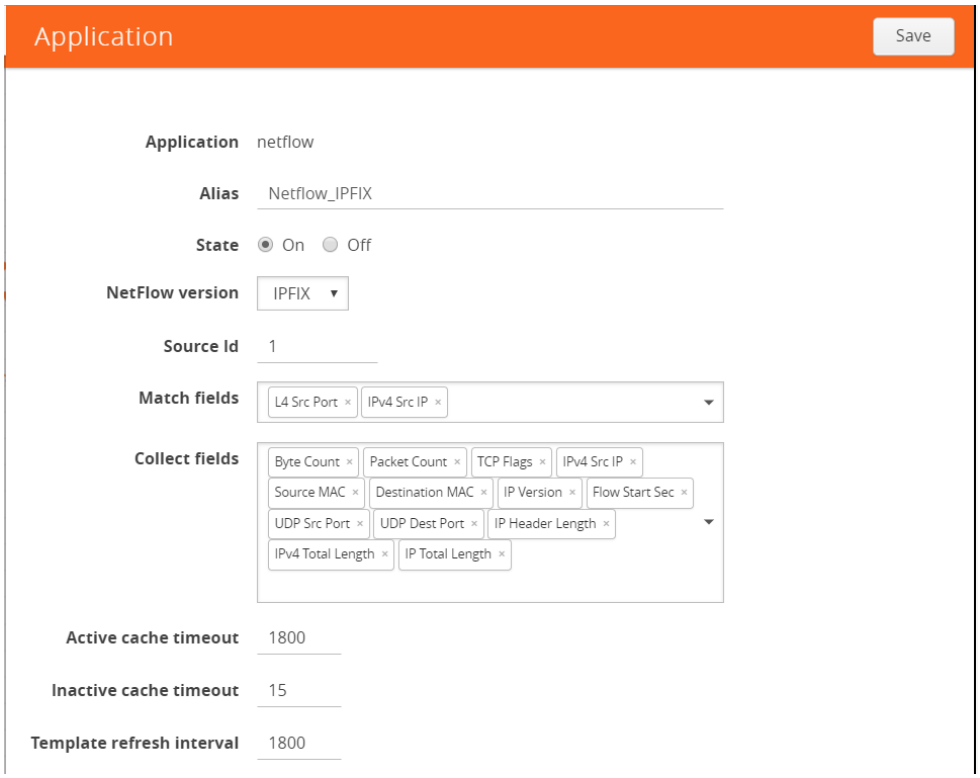

Figure 3-30: Viewing NetFlow Application Quick View

4. For State, select the **On** check box to determine that the application is generating NetFlow records from the packets coming from the G-vTAP agents. Select the **Off** check box to determine that the application is not currently generating NetFlow records. The state can be changed at anytime whenever required.
5. From the **NetFlow version** drop-down list, select the version you want to use to generate the NetFlow records. The default version selected is v5.
6. In the **Source ID** field, enter the observation domain to isolate the traffic. The NetFlow application uses source ID to segregate the records into categories. For example, you can assign source ID 1 for traffic coming over TCP. This results in generating a separate NetFlow record for TCP data. Similarly, you can assign Source ID 2 for traffic coming over UDP. This results in generating a separate NetFlow record for UDP data.
7. From the **Match fields** drop-down list, select the parameters that identify what you want to collect from the incoming packets. The Match fields displayed in the drop-down list are based on the NetFlow version selected in step 5. Refer to *Match/Key Fields* on page 69.
8. From the **Collect fields** drop-down list, select the parameters that identify what you want to collect from the NetFlow records. The Collect fields displayed in the drop-down list are based on the NetFlow version selected in step 5. Refer to *Collect/Non-Key Fields* on page 71.
9. In **Active cache timeout**, enter the number of seconds that an active flow record must remain in the cache before it is exported and removed. The default value is 1800 seconds.

10. In **Inactive cache timeout**, enter the number of seconds an inactive flow record must remain in the cache before it times out. The default value is 15 seconds.
11. In **Template refresh interval**, enter the frequency at which the template must be sent to the tool. The default value is 1800 seconds.
12. Click **Save**.

For some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes, refer to *NetFlow Examples* on page 80.

Network Address Translation (NAT)

NAT allows the NetFlow records to be directly transmitted to a collector without a tunnel. It lets you configure the destination IP of one or more collectors and the source IP of the GigaVUE V Series node interface through which the NetFlow records are sent out. The NetFlow records are exported to the collector over UDP protocol with the configurable source IP and destination IP.

NOTE: Only one NAT can be added per monitoring session.

Adding NAT

To add a NAT device:

1. Drag and drop **NAT** to the graphical workspace.

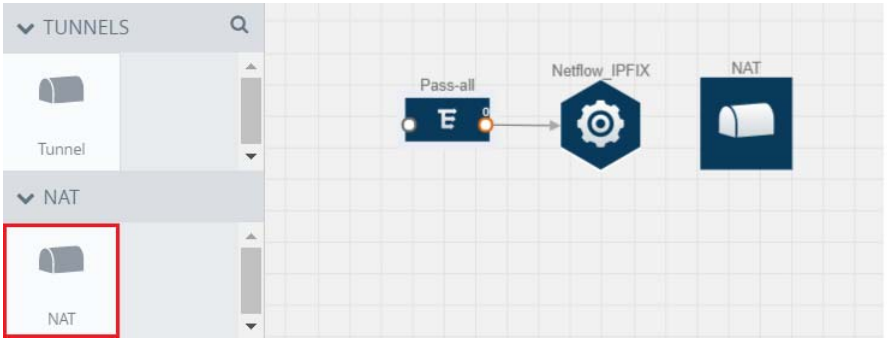

Figure 3-31: Adding NAT

2. Click **NAT** and select **Details**. A quick view is displayed for configuring a NAT device.

Figure 3-32: Selecting Details

3. In the **Alias** field, enter a name for the NAT device.

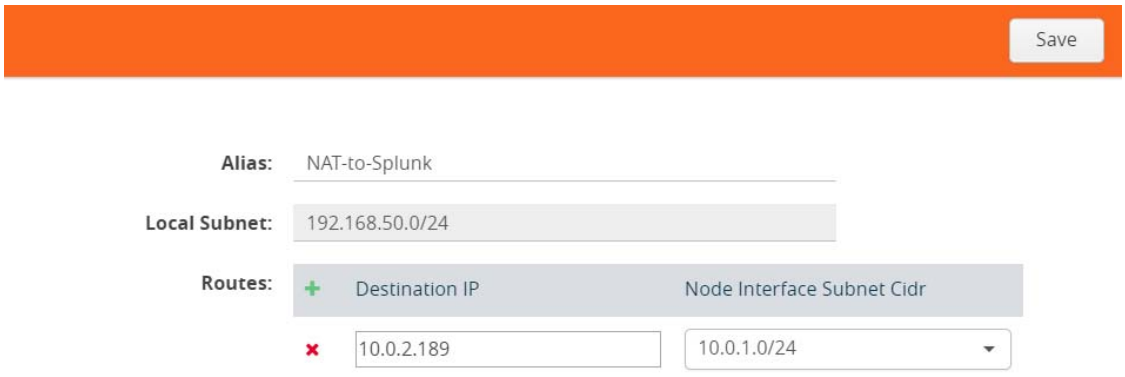

Figure 3-33: Configuring NAT

4. (Optional) In **Local Subnet**, enter a local subnet IP address that you want to assign to the NetFlow record. By default, GigaVUE V Series node auto generates a default local subnet. The subnet that you enter will override the default subnet.
5. (Optional) In **Routes**, define the routes to send the flow records to NetFlow collectors. Enter the following:
   a. In **Destination IP**, enter the IP address of the NetFlow collector. For example, if Splunk is the NetFlow collector, enter the IP address of Splunk.
   b. In **Node Interface Subnet CIDR**, enter the GigaVUE V Series node interface subnet Cidr for routing the NetFlow records out from GigaVUE V Series node.
   c. Click **+** to add more routes. Repeat steps a and b to enter the destination IP and node interface CIDR.
6. Click **Save**.

Linking a NetFlow Application to NAT

To create a link from a NetFlow application to a NAT device:

1. Drag and drop a link from the NetFlow application to a NAT device. A Link quick view is displayed. It is a header transformation operation that lets you configure the IPv4 destination IP of the NetFlow collector.

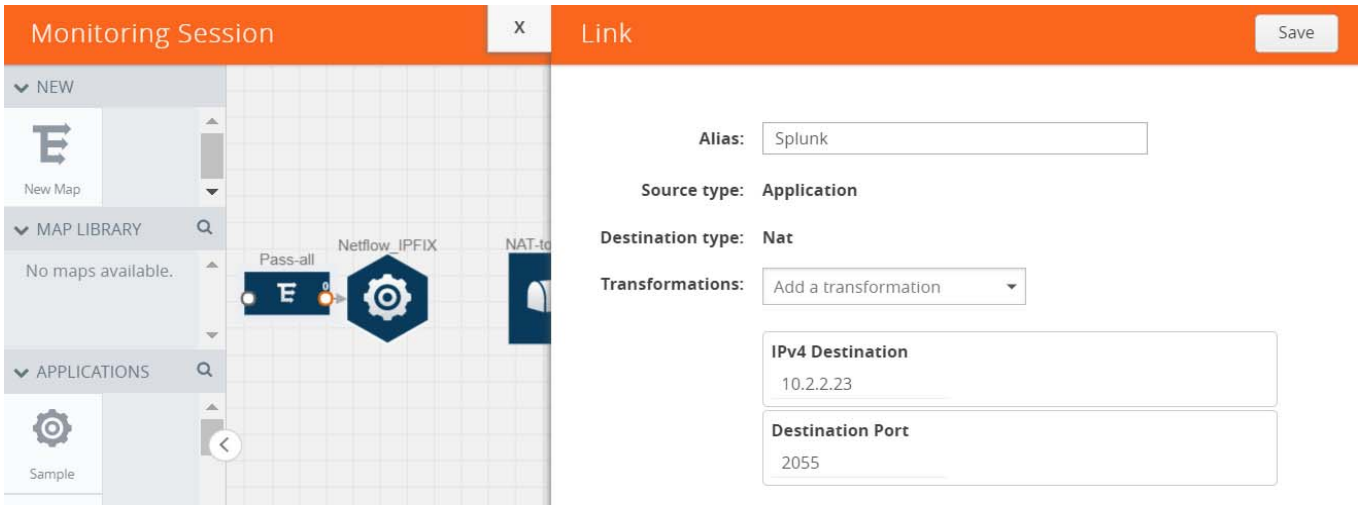

Figure 3-34: Creating a Link from NetFlow to NAT

2. In the **Alias** field, enter a name for the link.
3. From the **Transformations** drop-down list, select any one of the header transformations:
   - IPv4 Destination
   - ToS
   - Destination Port

NOTE: Only the above three header transformations are allowed on the link from the NetFlow application to a NAT device.

4. In **IPv4 Destination**, enter the IP address of the NetFlow collector.
5. (Optional) By default, the Destination Port is 2055. To change the destination port, enter a port number.
6. Click **Save**. The transformed link is displayed in Orange.

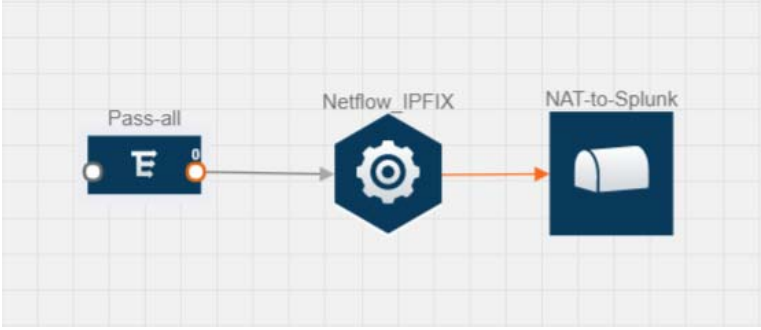

Figure 3-35: Linking NetFlow to NAT

7. Repeat steps 7 to 10 to send additional NetFlow records to NAT.

NetFlow Examples

This section provides some examples demonstrating the NetFlow application configuration in the GigaVUE V Series nodes.

- *Example 1* on page 80
- *Example 2* on page 84

Example 1

In this example, a pass all map is created and the entire traffic from a VPC is sent to a tool for full packet inspection. At the same time, a NetFlow application is added to generate flow records for flow inspection.

1. Create a monitoring session. For steps, refer to *Creating a Monitoring Session* on page 55.

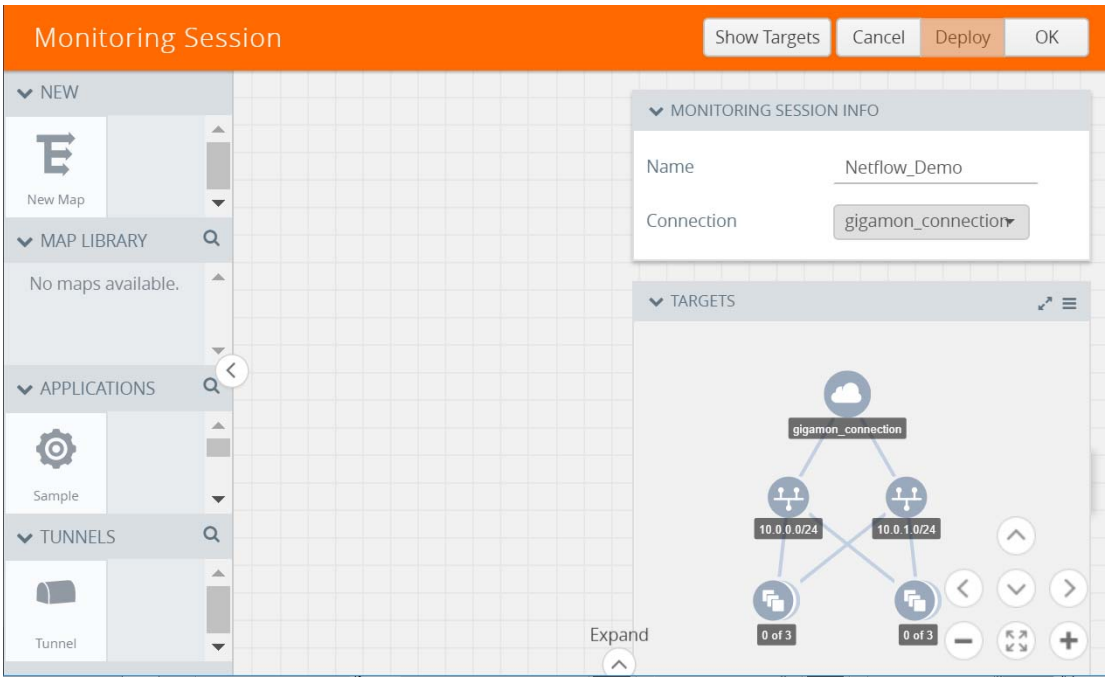

Figure 3-36: Creating a Monitoring Session

2. In the monitoring session, create a Pass all map. A pass all map sends all the traffic received from the G-vTAP agents to the tunnel endpoint or NAT. For steps, refer to *Creating a Map* on page 57.

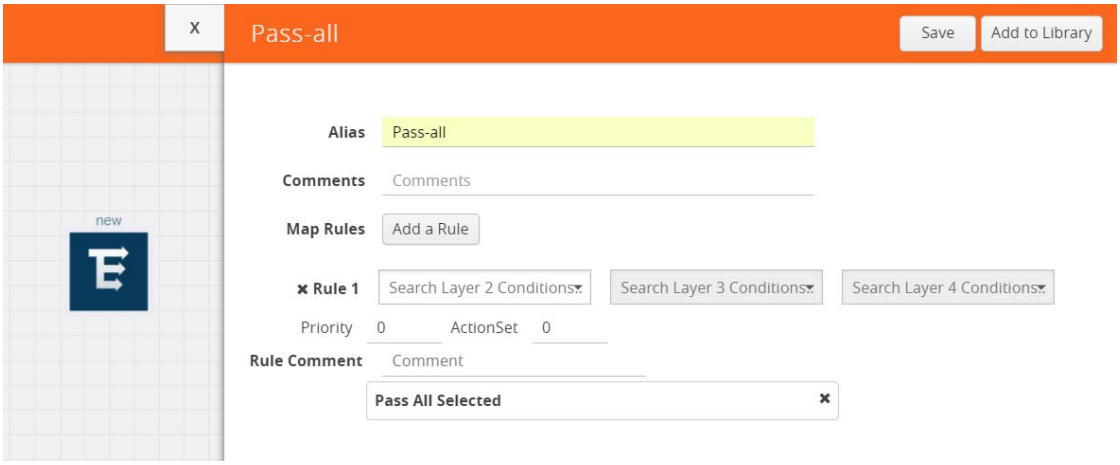

Figure 3-37: Creating a Pass All Map

3. Drag and drop a tunnel from **Tunnels**. A tunnel encapsulates the flow records and then sends them to the tools for full packet inspection.

Figure 3-38: Adding a Tunnel

4. Create a link from the Pass-all map to the tunnel endpoint. The traffic from the Pass-all map is forwarded to the tunnel endpoint that is connected to a tool.

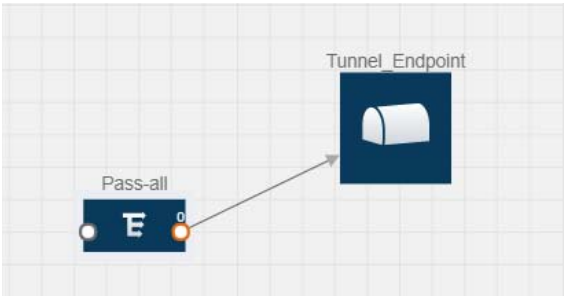

Figure 3-39: Creating a Link from Pass-all Map to Tunnel_Endpoint

5. Drag and drop a v5 NetFlow application.

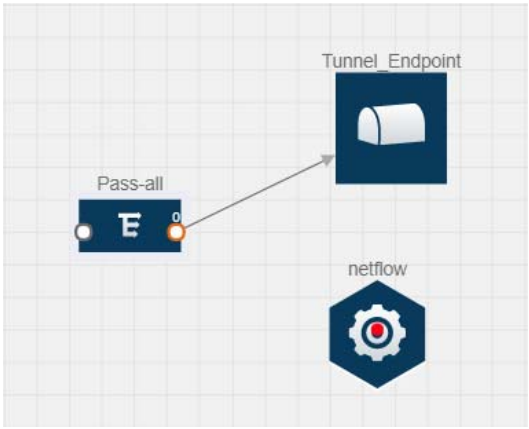

Figure 3-40: Adding a link from Pass-all Map to Tunnel_Endpoint

6. Click the NetFlow application and select **Details**. The Application quick view is displayed. For steps to configure the v5 NetFlow application, refer to *Adding a Version 5 NetFlow Application* on page 74.

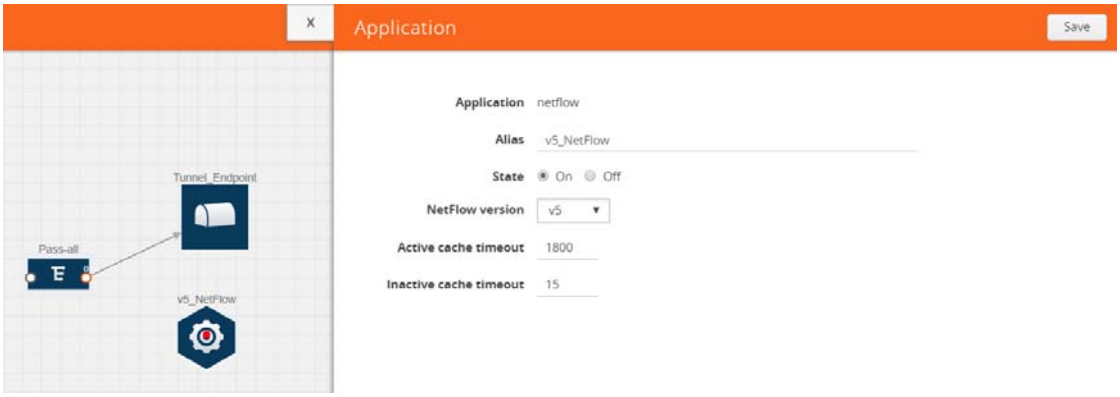

Figure 3-41: Configuring the NetFlow Application

7. Create a link from the Pass all map to the v5 NetFlow application.

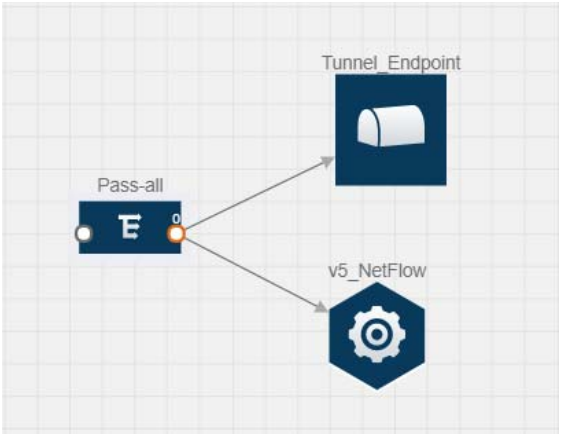

Figure 3-42: Adding a link from Pass-all Map to v5_NetFlow

8. Drag and drop **NAT** to the graphical workspace. A quick view to configure the NAT device is displayed. For steps to configure the NAT device, refer to *Adding NAT* on page 77.

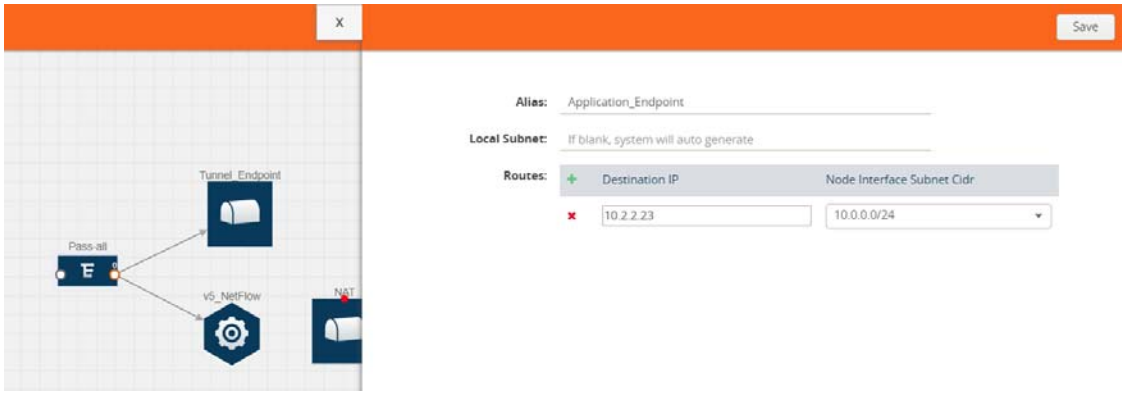

Figure 3-43: Adding a NAT Device

9. Create a link from the v5 NetFlow application to NAT. The link must be configured with the destination IP address of the NetFlow collector and the GigaVUE V Series node interface. For steps to configure the link, refer to *Linking a NetFlow Application to NAT* on page 79.

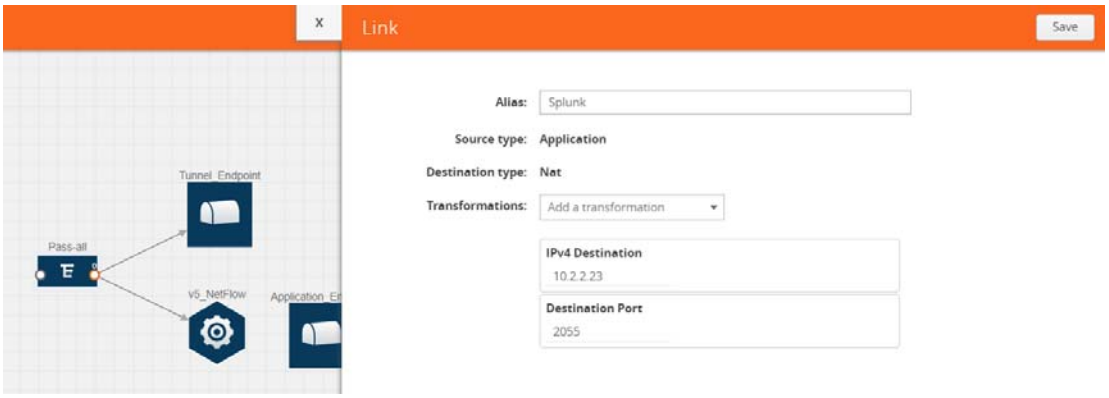

Figure 3-44: Adding a Link from v5 NetFlow Application to NAT

10. Click on the link created from the v5 NetFlow application to NAT. The information about the NetFlow collector destination IP and port is displayed.

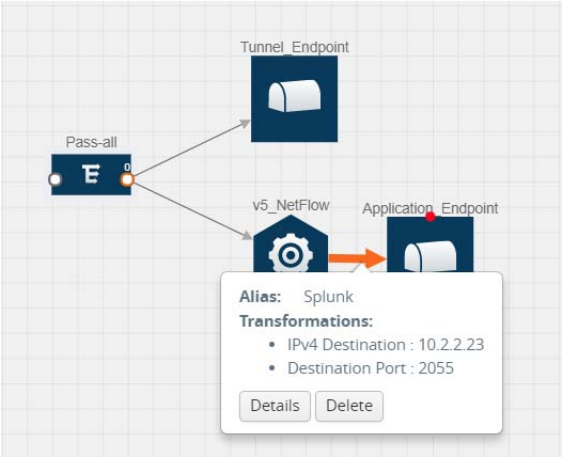

Figure 3-45: Viewing the Transformation Dialog Box

Example 2

In this example, two different versions of NetFlow applications are created. One map is configured to send the TCP packets to the v9 NetFlow application. Another map is configured to send the UDP packets to the IPFIX NetFlow application. The flow records generated from v9 and IPFIX NetFlow application are sent to NAT.

1. Create a monitoring session. For steps, refer to *Creating a Monitoring Session* on page 55.

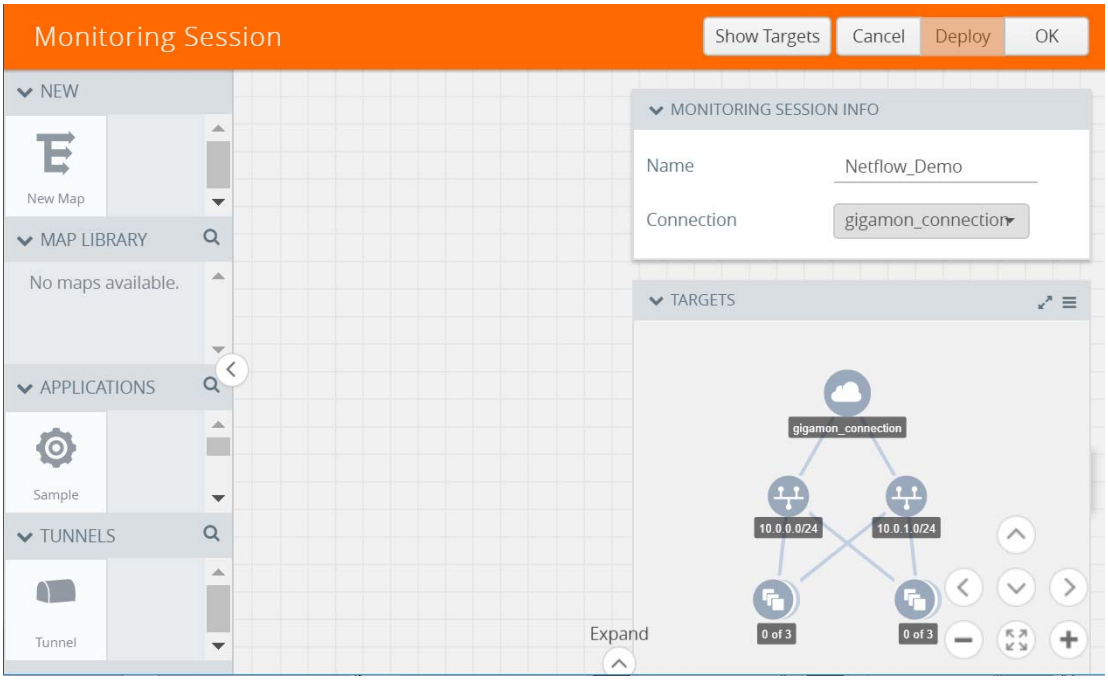

Figure 3-46: Creating a Monitoring Session

2. Create a map rule to filter the TCP packets. For steps on creating a map, refer to *Creating a Map* on page 57.

Figure 3-47: Creating a TCP Map

3. Create another map rule to filter the UDP packets.

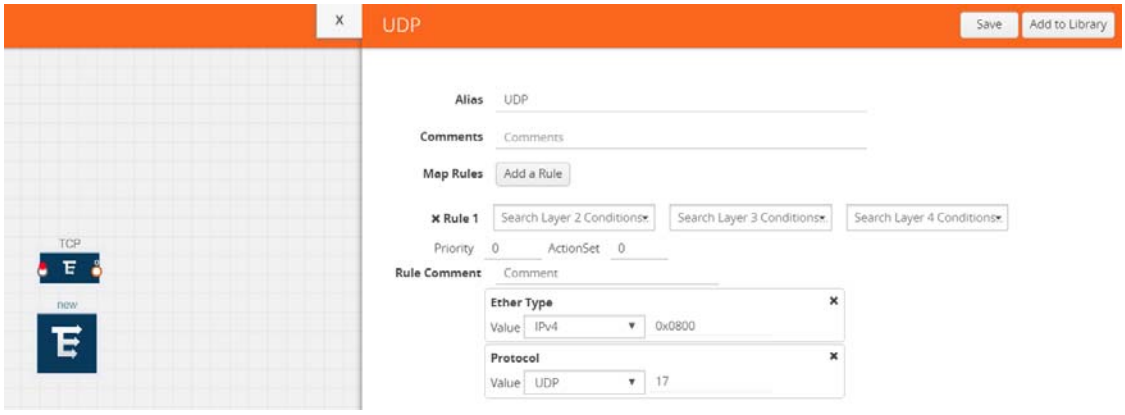

Figure 3-48: Creating a UDP Map

4. Create another map rule to filter the UDP packets.

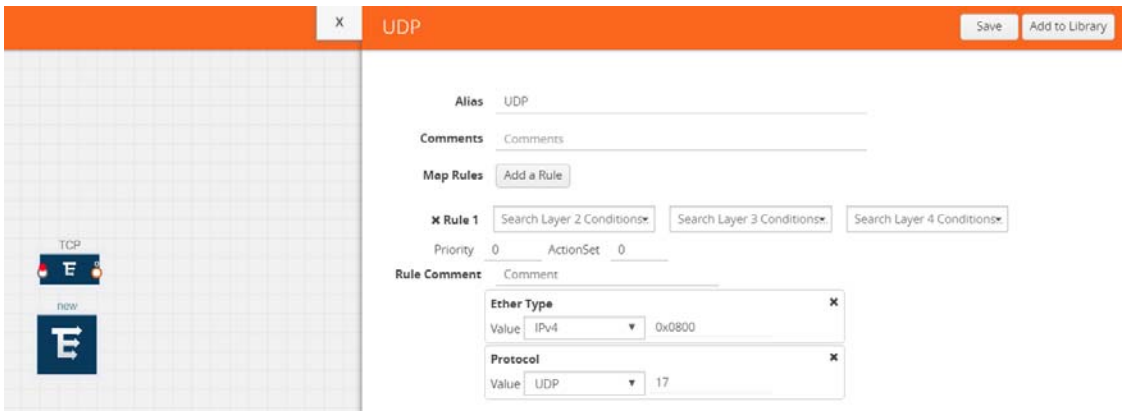

Figure 3-49: Creating a Map to Filter UDP Packets

5. Drag and drop a NetFlow application. Choose v9 as the NetFlow version. Select the match and the collect fields. For steps to configure the v9 NetFlow application, refer to *Adding a Version 9 and IPFIX NetFlow Application* on page 75.

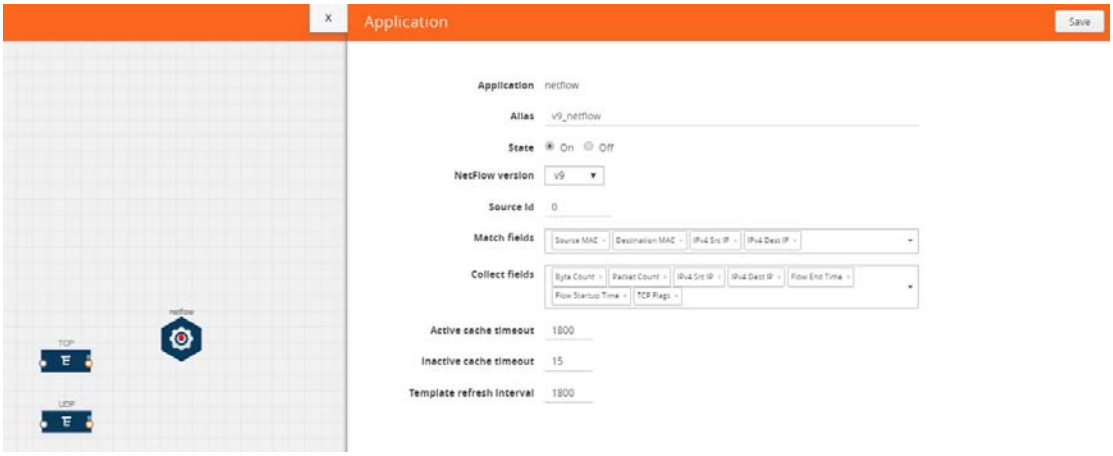

Figure 3-50: Adding a v9 NetFlow Application

6. Create a link from the TCP map to the v9 NetFlow application.

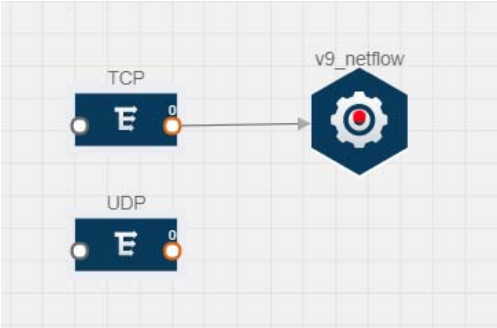

Figure 3-51: Adding a Link from the TCP Map to v9_netflow Application

7. Drag and drop a NetFlow application. Choose IPFIX as the NetFlow version. Select the match and the collect fields. For steps to configure the IPFIX NetFlow application, refer to *Adding a Version 9 and IPFIX NetFlow Application* on page 75.

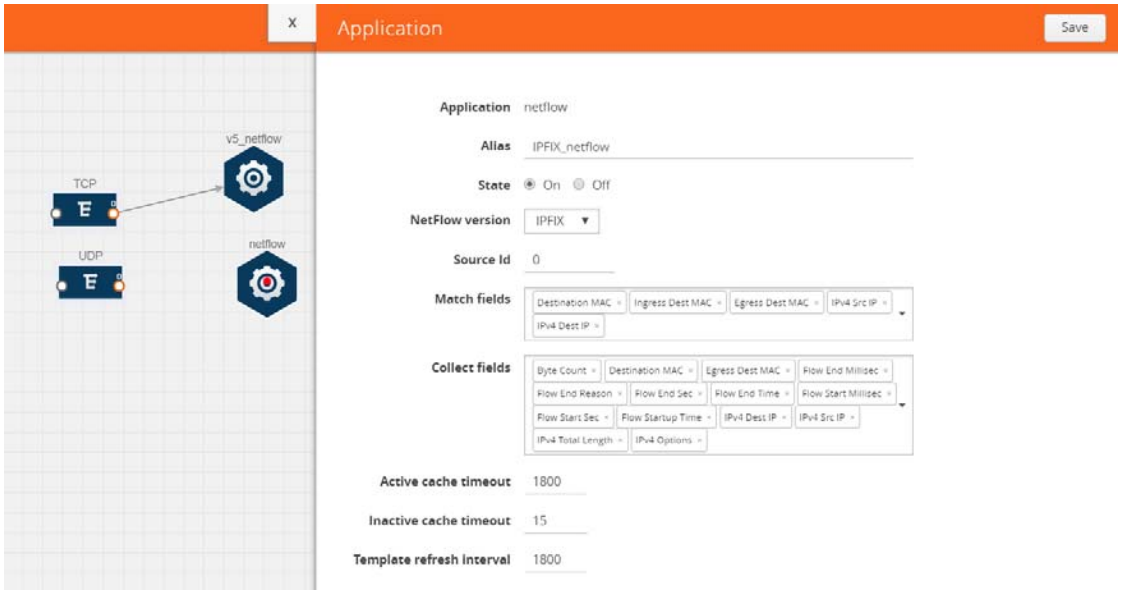

Figure 3-52: Adding a IPFIX NetFlow Application

8. Create a link from the UDP map to the IPFIX NetFlow application.

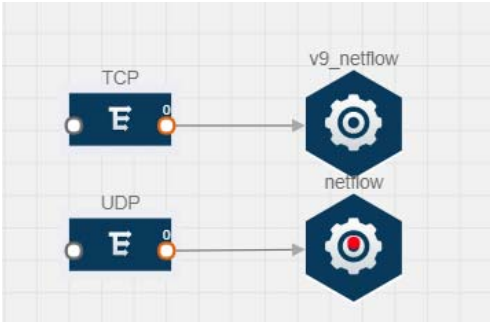

Figure 3-53: Adding a Link from the UDP Map to the IPFIX NetFlow Application

9. Drag and drop a NAT. A quick view to configure the NAT is displayed. For steps to configure a NAT, refer to *Adding NAT* on page 77.

Figure 3-54: Adding a NAT

10. Create a link from the v9 NetFlow application to the NAT.

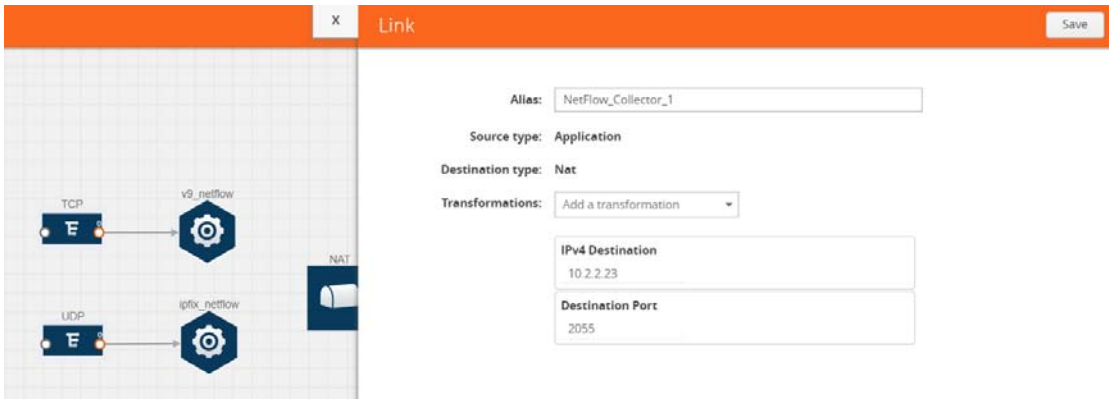

Figure 3-55: Adding a Link from NetFlow Application to NAT

11. Create another link from the IPFIX NetFlow application to the NAT.

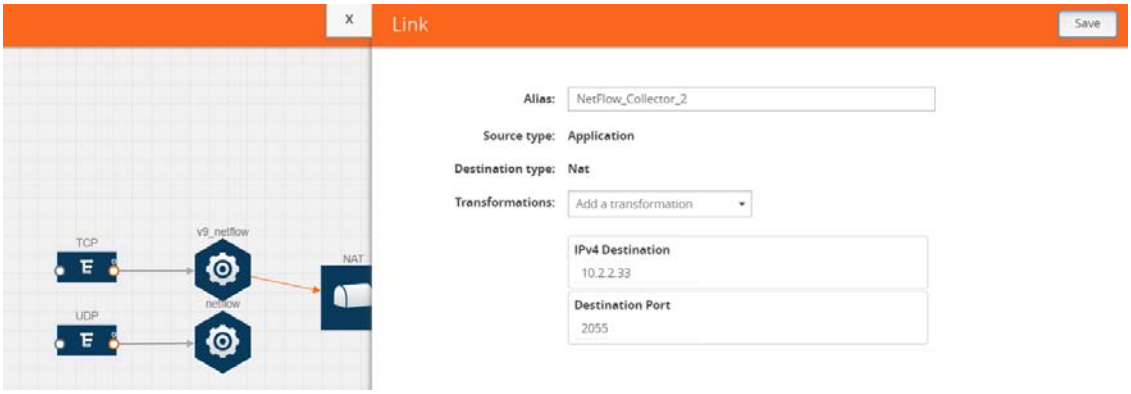

Figure 3-56: Adding a Link from NetFlow Application to NAT

12. Click on the link connecting the NetFlow application to the NAT.

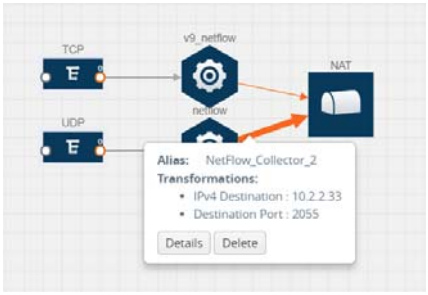

Figure 3-57: Viewing the Header Transformation

Deploying the Monitoring Session

To deploy the monitoring session:

1. Drag and drop one or more maps from the **MAP Library** to the workspace.
2. (Optional) To add Inclusion and Exclusion maps, drag and drop the maps from the Map Library to their respective section at the bottom of the workspace.

3. (Optional) Drag and drop one or more applications from the APPLICATIONS section to the workspace.

   NOTE: For information about adding applications to the workspace, refer to *Adding Applications to the Monitoring Session* on page 62.

4. Drag and drop one or more tunnels from the TUNNELS section.

   Figure 3-58 on page 89 illustrates three maps, one exclusion map, one application, and two tunnel endpoints dragged and dropped to the workspace.

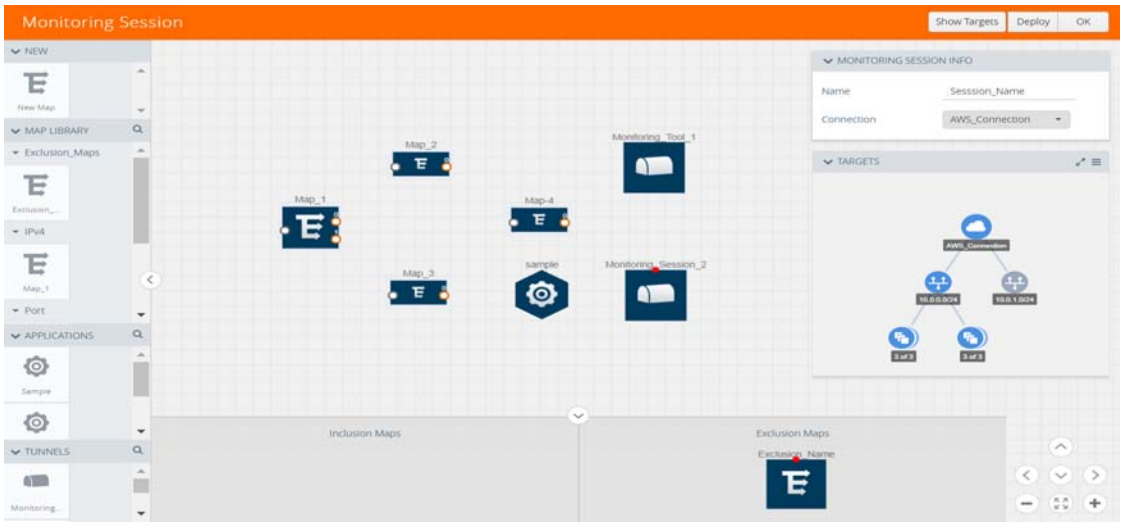

Figure 3-58: Dragging and Dropping the Maps, Applications, and Monitoring Tools

   NOTE: You can add up to 8 links from a single map to different maps, applications, or monitoring tools.

5. Hover your mouse on the map, click the red dot, and drag the link over to another map, application, or tunnel. You can drag more than one link from a map to the destination. On these links, you can apply link transformation to alter the packets. Refer to Figure 3-59 on page 90. For information about adding link transformation, refer to *Adding Header Transformations* on page 91.

6. Hover your mouse on the application, click the red dot, and drag the link (arrow) over to the tunnel endpoints.

In Figure 3-59 on page 90, the traffic matching the rules in each action set is routed to maps, applications, or monitoring tools.

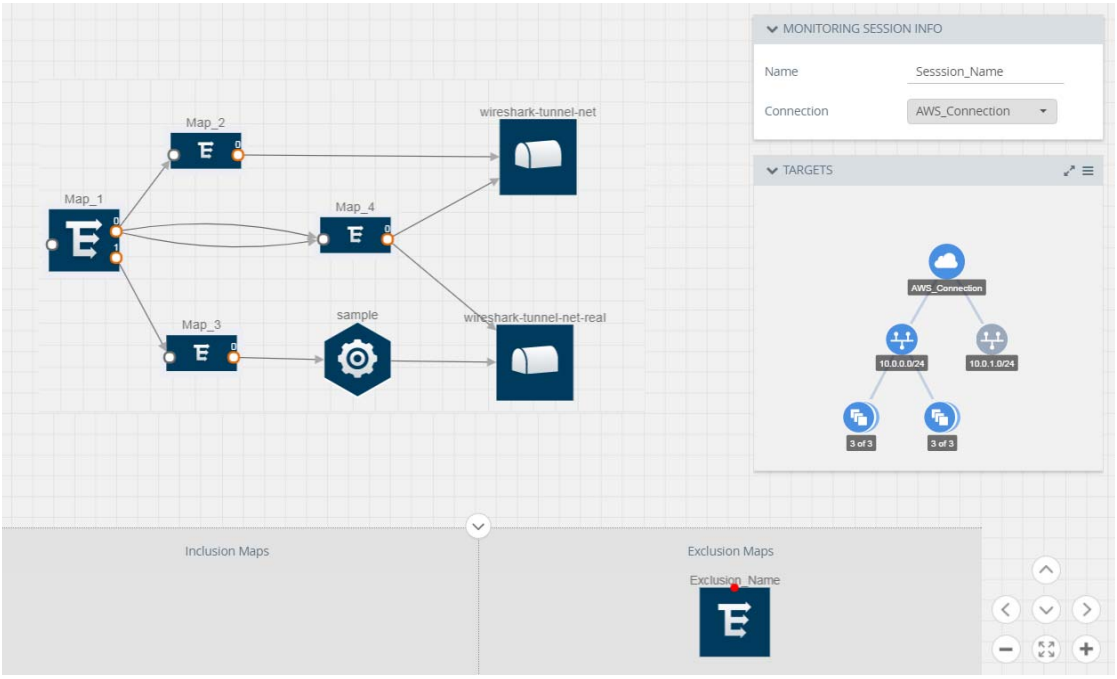

Figure 3-59: Connecting the Maps, Applications, and Monitoring Tools

7. Click **Show Targets** to view details about the subnets and monitoring instances.

   The instances and the subnets that are being monitored are highlighted in blue.

8. Click **Deploy** to deploy the monitoring session.

   The status is displayed as **Success** in the Monitoring Sessions page. The session is successfully deployed on all GigaVUE V Series nodes and G-vTAP agents.

   If the monitoring session is not deployed properly, then one of the following errors is displayed:

   - Partial Success—The session is not deployed on one or more instances due to G-vTAP or GigaVUE V Series node failure.
   - Failure—The session is not deployed on any of the GigaVUE V Series nodes and G-vTAP agents.

Click on the status link to view the reason for the partial success or failure. Refer to Figure 3-60 on page 91.

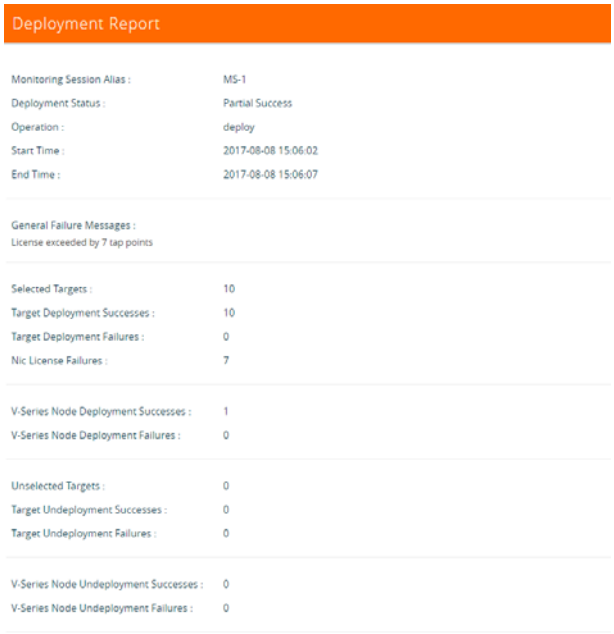

Figure 3-60: Deployment Status

9. Click **View** under Statistics to view and analyze the incoming and outgoing traffic.

You can also do the following in the Monitoring Session page:

- Use the **Redeploy** button to redeploy a monitoring session that is not deployed or partially successful.
- Use the **Undeploy** button to undeploy the selected monitoring session.
- Use the **Clone** button to duplicate the selected monitoring session.
- Use the **Edit** button to edit the selected monitoring session.
- Use the **Delete** button to delete the selected monitoring session.

Adding Header Transformations

Header transformation is performed on a link in a monitoring session. You can select a link and modify the packet header before they are sent to the destination. The header transformation feature is supported only with GigaVUE V Series node version 1.3-1 and above.

Header transformations are used to perform many simple operations on the network packets. The source and destination MAC addresses, port numbers, and IP addresses can be masked to prevent the information from being exposed to the monitoring tools.

You can also add VLAN ID, VLAN priority, and DSCP bits to distinguish the traffic coming from multiple VPCs with the same subnet range. The monitoring tools cannot always distinguish the traffic coming from multiple VPCs with the same subnet range.

In addition to header transformation, GigaVUE V Series node allows you to add multiple links to the same destination. Using multiple links, you can send duplicate packets or various transformed packets to the same destination. For example, you can add different L2GRE or VXLAN tunnel IDs to the packets and send them to different applications within the same tool.

In Figure 3-61 on page 92, the filtered packets from the ICMP map are sent to the same tunnel endpoint in four different links. In each link, you can apply one or more header transformations. A link with the header transformation applied is displayed in orange. When you mouse over the orange link, a detailed information about the alias and the type of transformation is displayed.

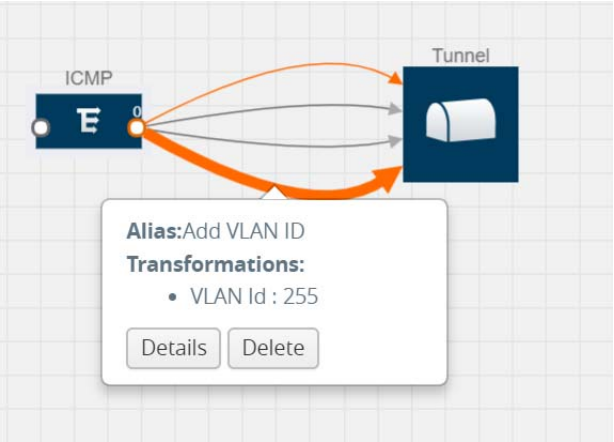

Figure 3-61: Action Set with Multiple Links

GigaVUE V Series node supports the following header transformations:

Table 3-8: Header Transformations

Option	Description
MAC Source	Modify the Ethernet source address.
MAC Destination	Modify the Ethernet destination address.
VLAN Id	Specify the VLAN ID.
VLAN PCP	Specify the VLAN priority.
Strip VLAN	Strip the VLAN tag.
IPv4 Source	Specify the IPv4 source address.
IPv4 Destination	Specify the IPv4 destination address.
ToS	Specify the DSCP bits in IPv4 traffic class.
Source Port	Specify the UDP, TCP, or SCTP source port.

Table 3-8: Header Transformations

Option	Description
Destination Port	Specify the UDP, TCP, or SCTP destination port.
Tunnel ID	Specify the tunnel ID. The tunnel ID header transformation can only be applied on the links with the tunnel endpoint destination. Using Tunnel ID header transformation, the filtered packets can be sent to different applications or programs within the same monitoring tool.

To add a header transformation:

1. On the Monitoring Session, click the link and select **Details.** The Link quick view is displayed.

Figure 3-62: Opening the Link Quick View

2. From the **Transformations** drop-down list, select one or more header transformations.

NOTE: Do not apply VLAN Id and VLAN PCP transformation types with the Strip VLAN ID transformation type on the same link.

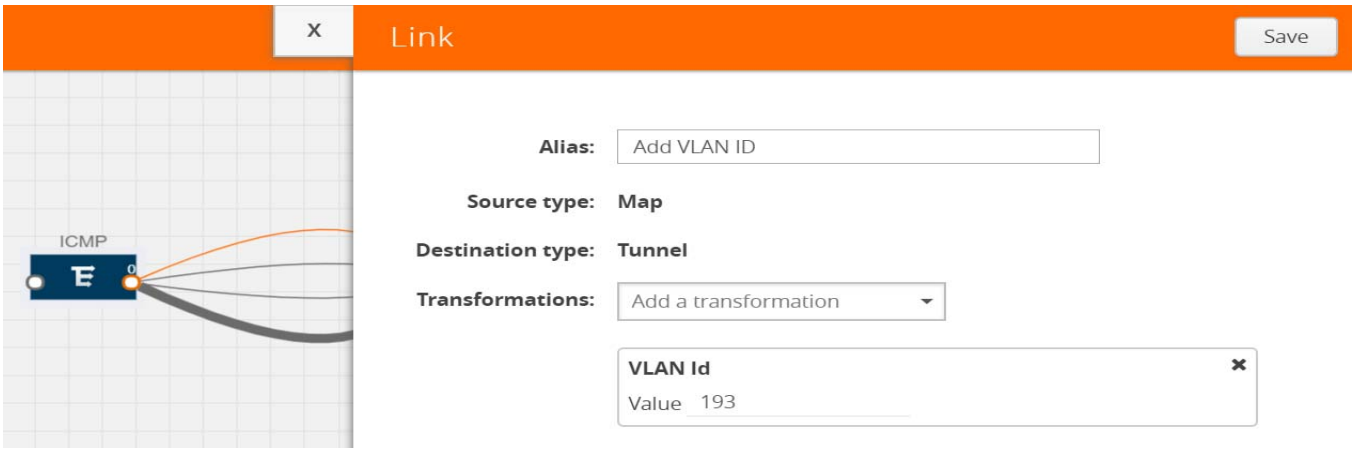

Figure 3-63: Adding Transformation

3. Click **Save**. The selected transformation is applied to the packets passing through the link.
4. Click **Deploy** to deploy the monitoring session.

Viewing the Statistics

The Monitoring Session Statistics page lets you analyze the incoming and outgoing traffic on an hourly, daily, weekly, and monthly basis. The traffic can be viewed based on kilobits/second, megabits/second or gigabits/second.

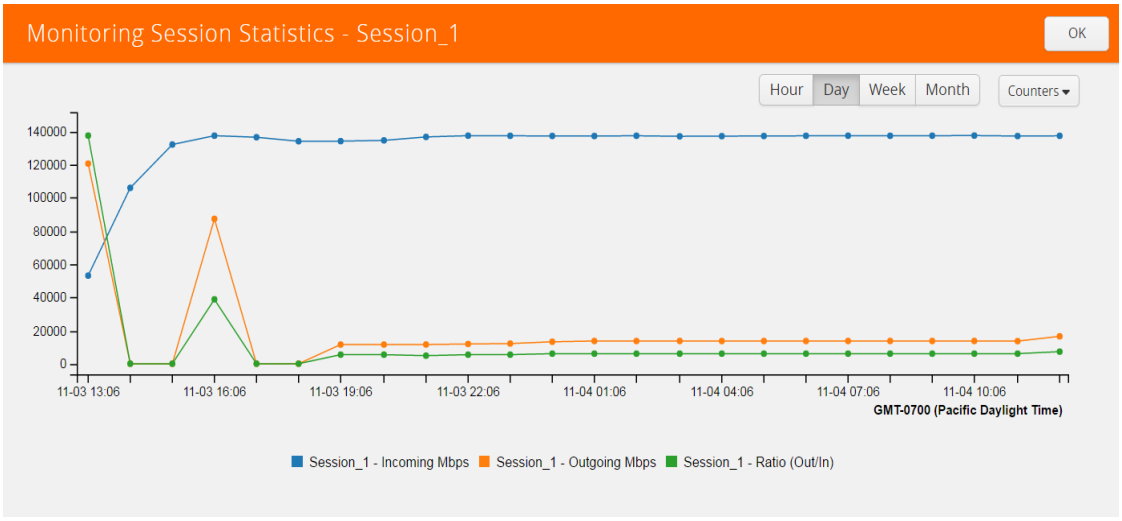

Figure 3-64: Viewing the Monitoring Session Statistics

You can click on Incoming Maps, Outgoing Maps, and Ratio at the bottom of the graph to view the statistics individually.

You can expand the **View Monitoring Session Diagram** and click on each individual map, application, and tunnel to view more details about the incoming and outgoing traffic on the selected statistics page. The Map Statistics page lets you choose the map rules to view the traffic matching the selected rule.

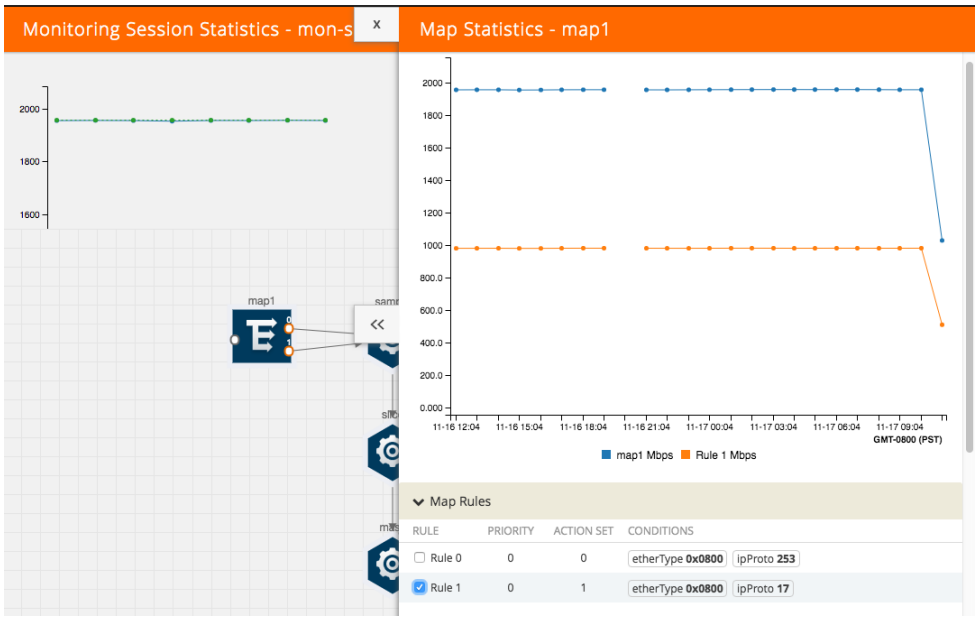

Figure 3-65: Viewing the Map Statistics

Viewing the Topology

You can have multiple VPC connections in GigaVUE-FM. Each VPC can have multiple monitoring sessions configured within them. You can select the connection and the monitoring session to view the selected subnets and instances in the topology view.

To view the topology diagram:

1. Select **AWS > Topology**.
2. Select a connection from the **Select connection...** list. The topology view of the subnets and instances is displayed.
3. (Optional) Select a monitoring session from the **Select Monitoring Session...**list. The topology view of the monitored subnets and instances in the selected session are displayed.
4. Select one of the following check boxes:
   - **Source**— Displays the topology view of the source target interfaces that are being monitored.
   - **Destination**—Displays the topology view of the destination target interfaces where the traffic is being mirrored.

- **Other**—Displays the topology view of the non-G-vTAP agents such as GigaVUE V Series Controllers, G-vTAP Controllers, monitoring tools, and instances that are being used in the connection.

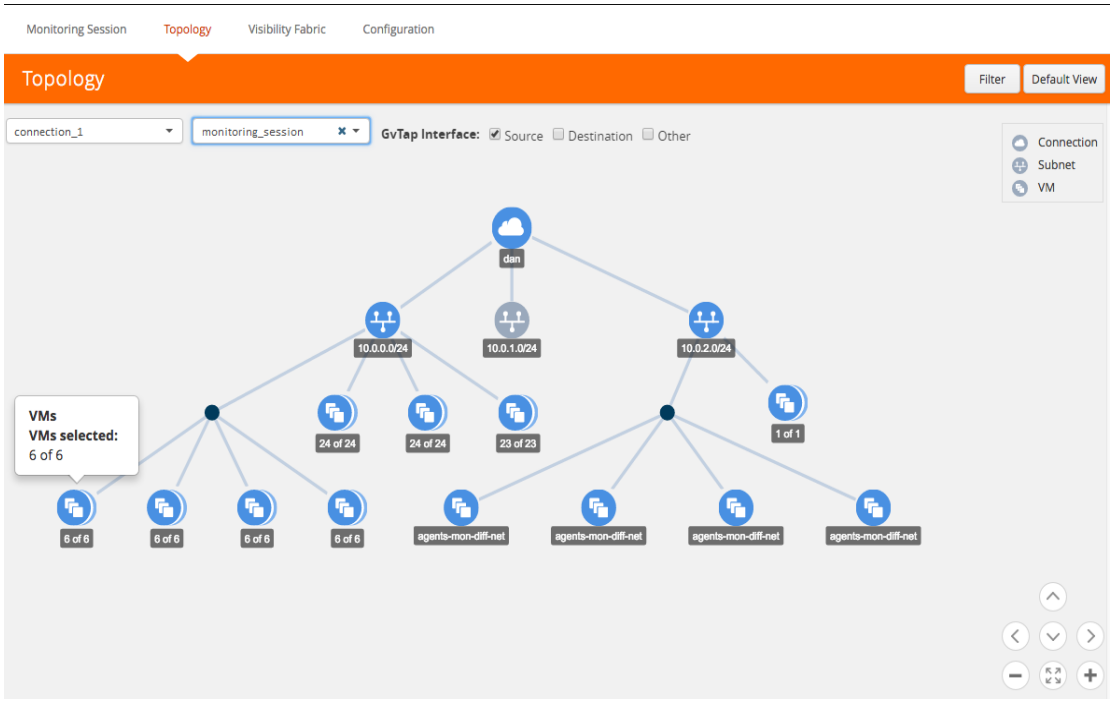

Figure 3-66: Viewing the Topology

5. (Optional) Hover over or click the subnet or VM Group icons to view the subnets or instances present within the group.

In the topology page, you can also do the following:

- Use the **Filter** button to filter the instances based on the VM name, VM IP, Subnet ID, or Subnet IP, and view the topology based on the search results. Refer to Figure 3-67 on page 97.

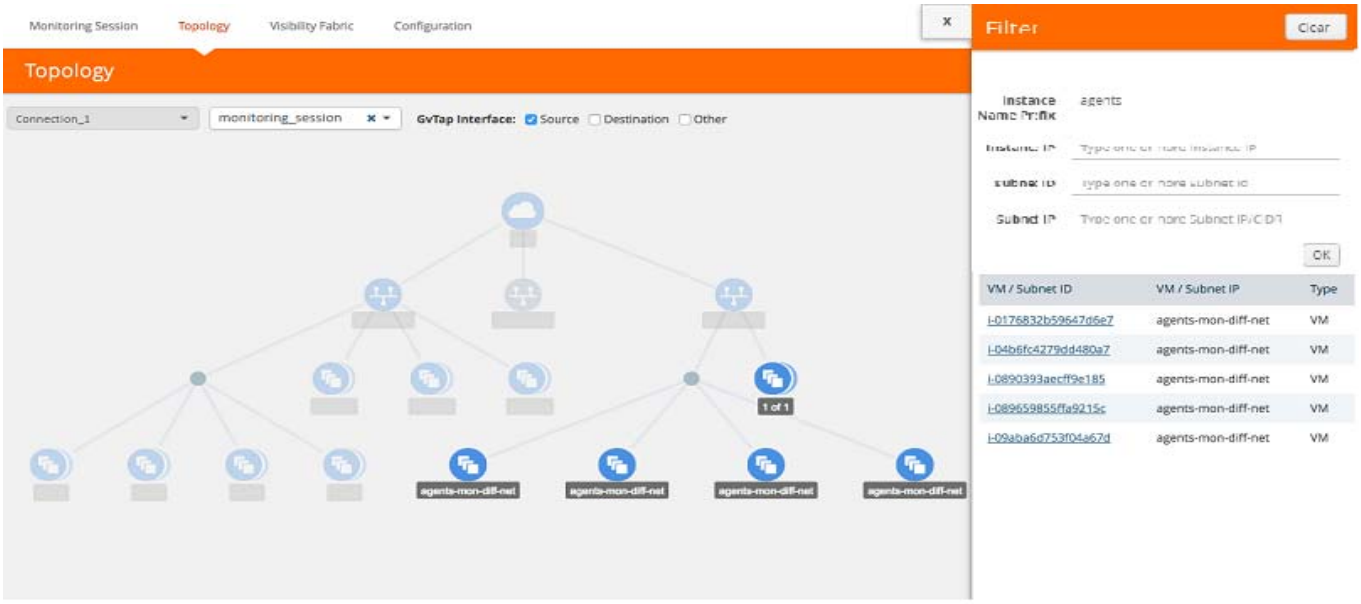

Figure 3-67: Filtering in Topology View

- Use the **Default View** button to view the topology diagram based on the source interfaces of the monitoring instances.
- Use the arrows at the right-bottom corner to move the topology page up, down, left, or right. Click the **Fit-to-Width** icon to fit the topology diagram according to the width of the page.
- Use **+** or **-** icons to zoom in and zoom out the topology view.

Configuring the AWS Settings

This section provides information on how to configure the maximum number of connections, refresh intervals for instance and non-instance inventory, and maximum batch size for monitoring session updates. It also provides information on how to enable CloudWatch events.

Use the **AWS** > **Configurations** > **AWS Settings** to edit these AWS settings. Refer to Table 3-9 on page 98 for more information about the settings:

Table 3-9: AWS Settings

Settings	Description
Maximum number of connections allowed	Specifies the maximum number of VPC connections you can establish in GigaVUE-FM.
Refresh interval for instance inventory (secs)	Specifies the frequency for updating the state of EC2 instances in AWS.
Refresh interval for non-instance inventory (secs)	Specifies the frequency for updating the state of non-instance information such as subnets, security groups, images, key pairs, VPCs, and elastic IP addresses.
Number of instances per GigaVUE V Series Node	Specifies the maximum number of instances that can be assigned to the GigaVUE V Series node.
Refresh interval for G-vTAP agent inventory (secs)	Specifies the frequency for discovering the G-vTAP agents available in the VPC.
AWS CloudWatch event-based inventory refresh	Enables or disables the AWS CloudWatch event-based inventory refresh. If enabled, CloudWatch event rules updates GigaVUE-FM with EC2 instance state changes.

Configuring the Proxy Server

Sometimes, the VPC in which the GigaVUE-FM is launched may not have access to the Internet. Without Internet access, GigaVUE-FM cannot connect to the AWS API endpoints. For GigaVUE-FM to connect to AWS, a proxy server must be configured.

To create a proxy server:

1. Select **AWS > Configuration > Proxy Server**.
2. Click **Add**. The Add Proxy Server page is displayed as shown in Figure 3-68 on page 99.

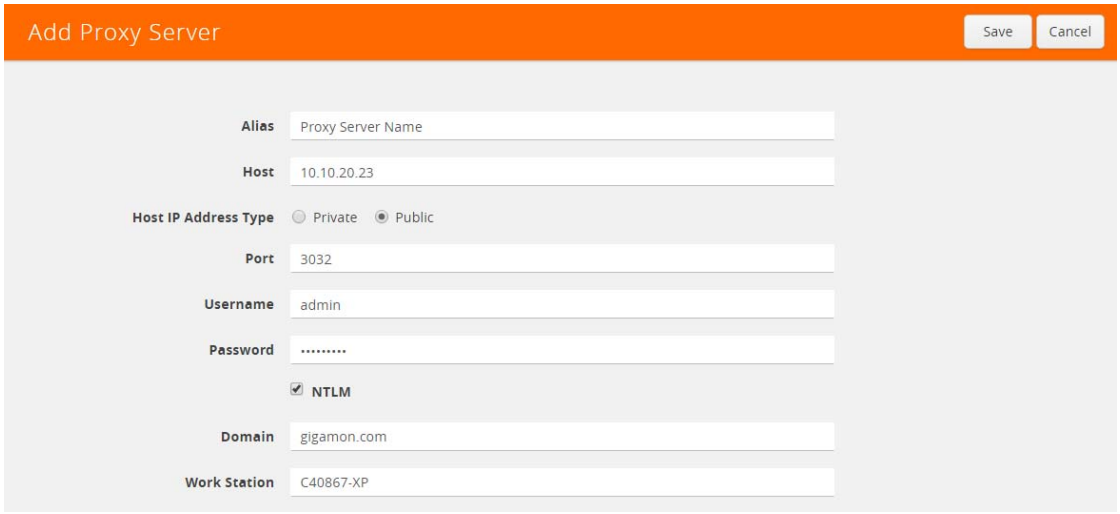

Figure 3-68: Adding a Proxy Server

3. Select or enter the appropriate information as shown in Table 3-10 on page 99.

Table 3-10: Fields for Proxy Sever Configuration

Field	Description
Alias	The name of the proxy server.
Host	The host name or the IP address of the proxy server.
Host IP Address Type	The type of the Host IP address that indicate whether the proxy server IP address is private or public to the VPC.
Port	The port number used by the proxy server for connecting to the Internet.
Username	(Optional) The username of the proxy server.
Password	The password of the proxy server.
NTLM	(Optional) The type of the proxy server used to connect to the VPC.
Domain	The domain name of the client accessing the proxy server.
Workstation	(Optional) The name of the workstation or the computer accessing the proxy server.

4. Click **Save**.

   The new proxy server configuration is added to the Proxy Server Configuration page. Refer to Figure 3-69. The proxy server is also listed in the AWS Connection page. Refer to *Connecting to AWS* on page 39.

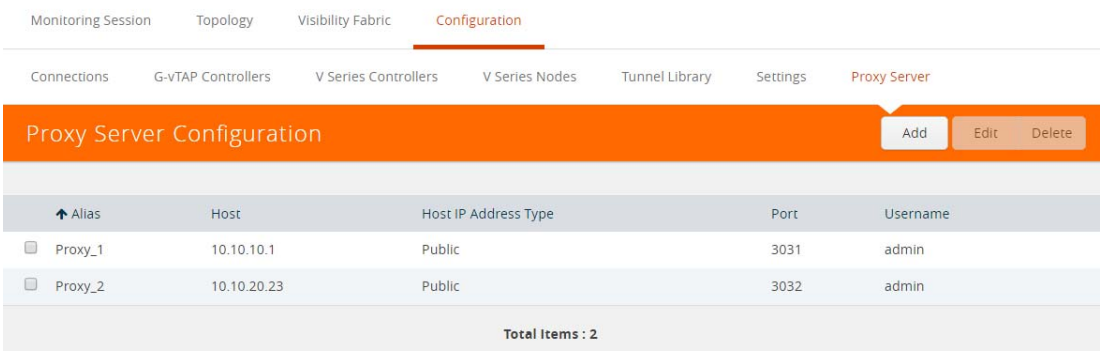

Figure 3-69: Proxy Server Configuration Page

Setting Up Email Notifications

Notifications are triggered by a range of events such as AWS license expiry, VM instance terminated, and so on. You can setup the email notification for a particular event or a number of events and the recipient or recipients to whom the email should be sent.

Gigamon strongly recommends you to enable email notifications so there is immediate visibility of the events affecting node health.

The following are the events for which you can setup the email notifications:

- AWS License Expire
- Fabric Node Down
- Fabric Node Reboot Failed
- Fabric Node Rebooted
- Fabric Node Replacement Launch Failed
- Fabric Node Replacement Launched
- Fabric Node Restart Failed
- Fabric Node Restarted
- Fabric Node Unreachable
- Fabric Node Up

Configuring the Email Notifications

To configure the automatic email notifications:

1. Click **Administration** on the top navigation link.

2. On the left navigation pane, select **System** > **Notifications**.

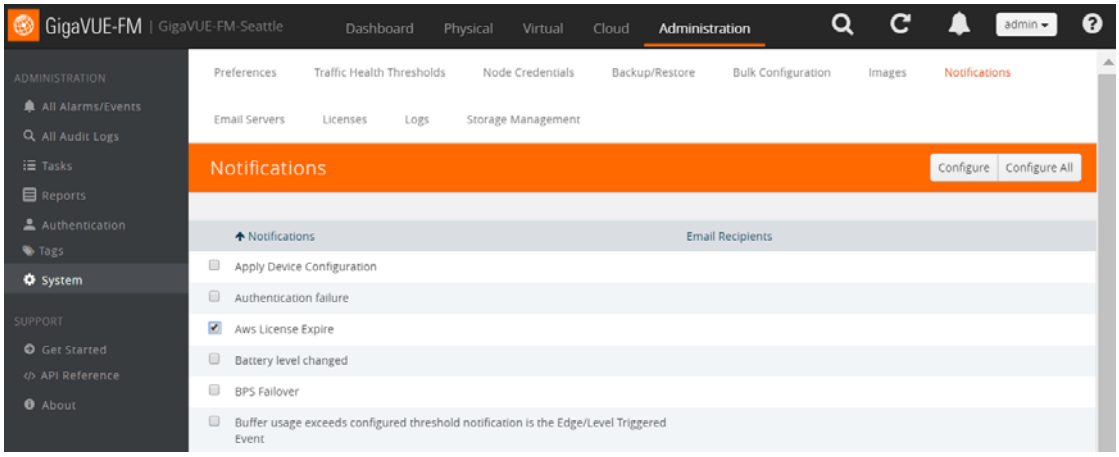

Figure 3-70: Setting up Email Notifications

3. In the Notifications page, select the event and click **Configure**.

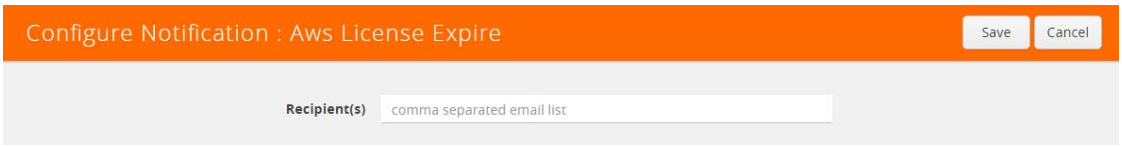

Figure 3-71: Configure Email Notification

4. In the Recipient(s) box, enter one or multiple email IDs separated by a comma.
5. Click **Save**.

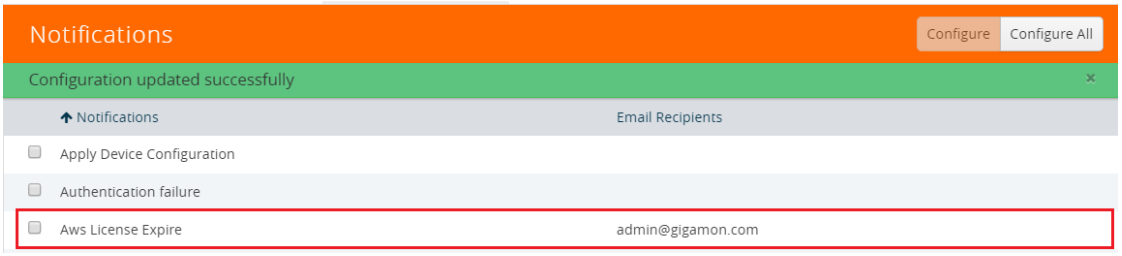

Figure 3-72: Email Notification Configured

Alarms and Events

The Alarms and Events page displays all the events occurring in the virtual fabric node, VM Domain, and VM manager. An event is an incident that occur at a specific point in time. Examples of events include:

- AWS License Expire
- G-vTAP Agent Inventory Update Completed
- AWS 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 could be AWS license expiry.

The alarms and events broadly fall into the following categories: Critical, Major, Minor, or info.

Click **Cloud** on the top navigation link. On the left navigation pane, click **Alarms/ Events**.

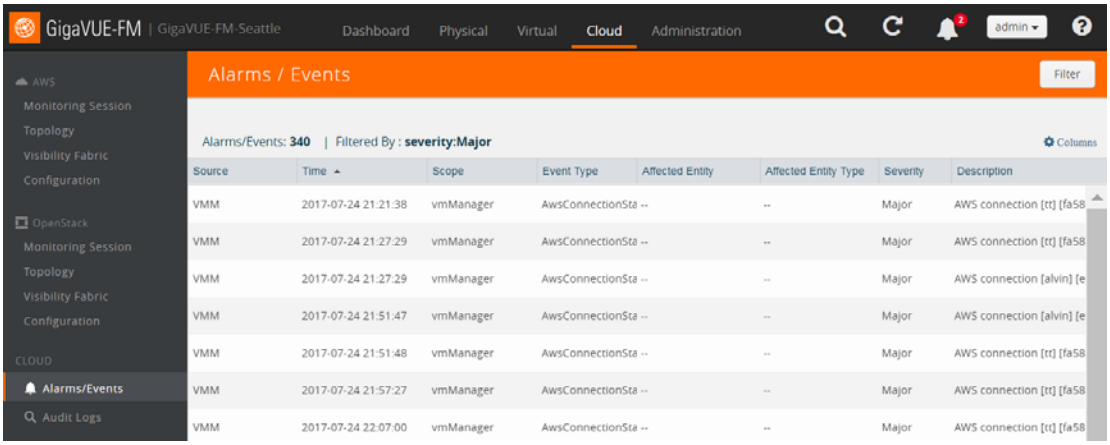

Figure 3-73: Alarms and Events

Table 3-11 describes the parameters recording for each alarm or event. You can also use filters to narrow down the results. Refer to *Filtering Alarms/Events* on page 103.

Table 3-11: All Alarm/Event Parameters

Controls/ Parameters	Description
Source	The source from where the alarms and events are generated.
Time	The timestamp when the event occurred. **IMPORTANT:** Timestamps are shown in the time zone of the client browser's computer and not the timezone 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 timezone.
Scope	The category to which the alarms or events belong. Alarms and events can belong to the following category: Virtual Fabric Node, VM Domain, VM Manager.
Event Type	The type of event that generated the alarms and events.
Severity	The severity is one of Critical, Major, Minor, or Info. Info is informational messages. For example, when GigaVUE V Series nodes are installed, such a message is displayed as Info.
Affected Entity Type	The resource type associated with the alarm or event.
Affected Entity	The resource ID of the affected entity type.
Description	The description of the event, which includes any of the possible notifications with additional identifying information where appropriate.

Table 3-11: All Alarm/Event Parameters

Controls/ Parameters	Description
Device IP	The IP address of the device.
Host Name	The host name of the device.

Filtering Alarms/Events

To filter the alarms and event:

1. Click **Filter**.

   The Filter quick view is displayed.

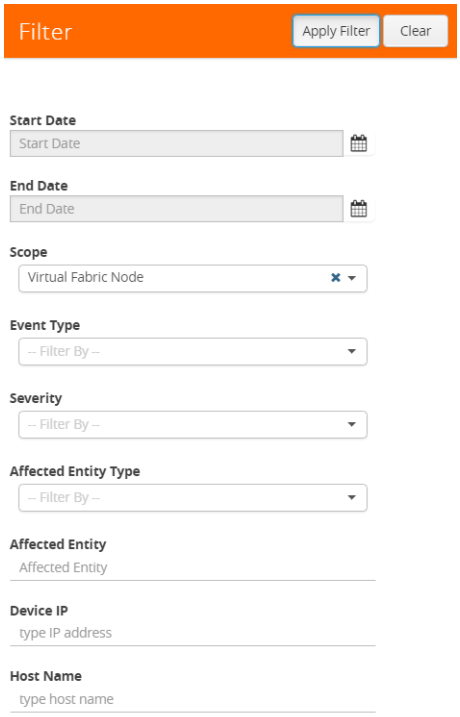

Figure 3-74: Filtering Alarms/Events

2. Select the filtering criteria, then click **Apply Filter**. The results are displayed in the Alarms/Events page.

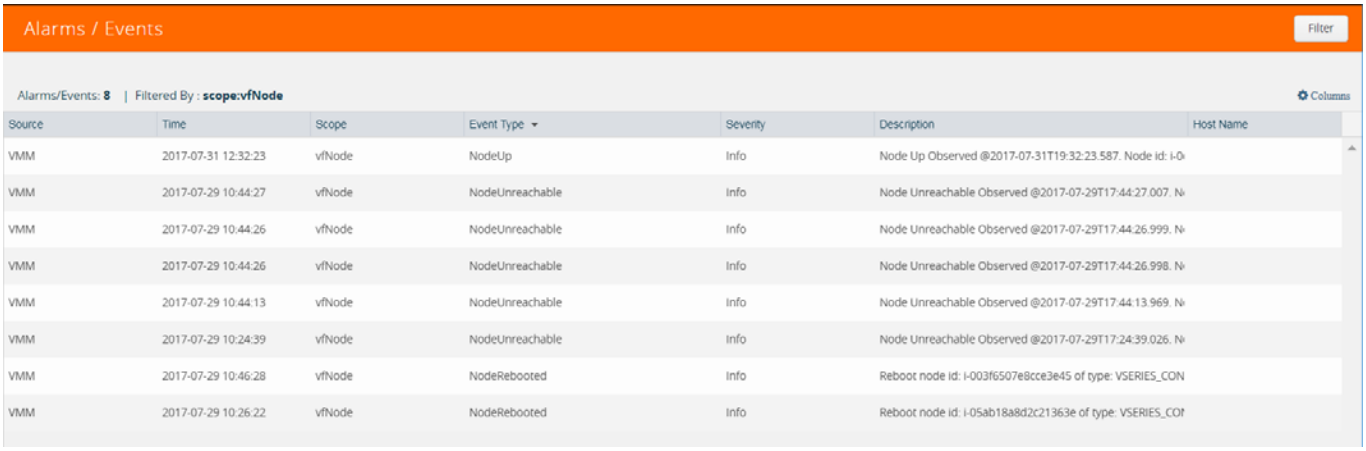

Figure 3-75: Alarms/Events Filter Results

Audit Logs

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

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 are logged by the system such as: • Log in and Log out based on users. • Create/Delete/Edit tasks, GS operations, maps, virtual ports, and so on.
Source	Provides details on whether the user was in FM or on the node when the event occurred.
Status	Success or Failure of the event.
Description	In the case of a failure, provides a brief update on the reason for the failure.

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

Filtering Audit Logs

Filtering the audit logs allows you to display specific type of logs. You can filter based on any of the following:

- When—display logs that occurred within a specified time range.
- Who—display logs related a specific user or users.

- What—display logs for one or more operations, such as Create, Read, Update, and so on.
- Where—display logs for GigaVUE-FM or devices.
- Result—display logs for success or failure.

To filter the audit logs, do the following:

1. Click **Filter**.

   The quick view for Audit Log Filters displays.

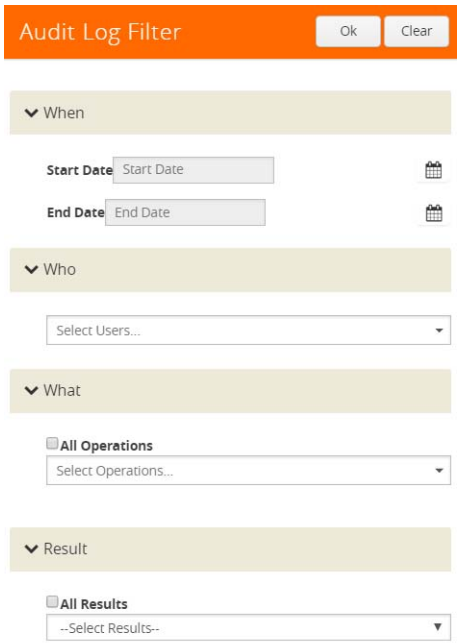

Figure 3-76: Audit Logs Filter

2. Specify any 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, either FM or device. Select **All Systems** apply both 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. Click **OK** to apply the selected filters to the Audit Logs page.

Chapter 4

Upgrading the GigaVUE-FM Instance

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

Refer to the following sections for details:

At a Glance

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

Step 1: Stop the existing version of the GigaVUE-FM instance.

Step 2: Create a snapshot of the second disk (dev/sdb) of the FM instance.

Step 3: Make a note of the snapshot ID.

Step 4: Launch the latest version of the GigaVUE-FM instance. While launching the latest version, enter the snapshot ID of the old version of the GigaVUE-FM instance in **Add Storage** > **Add New Volume**.

Step 5: Complete the launch.

Step 6: Verify if the data from the previous GigaVUE-FM instance is restored in the new instance.

Step 7: Terminate the old FM instance.

Stopping the GigaVUE FM Instance

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

NOTE: Do not terminate the GigaVUE-FM instance.

To stop the GigaVUE-FM instance:

1. Login to the AWS account and select **Services > EC2**.
2. In the left navigation pane, select **Instances**. Refer to Figure 4-1 on page 108.

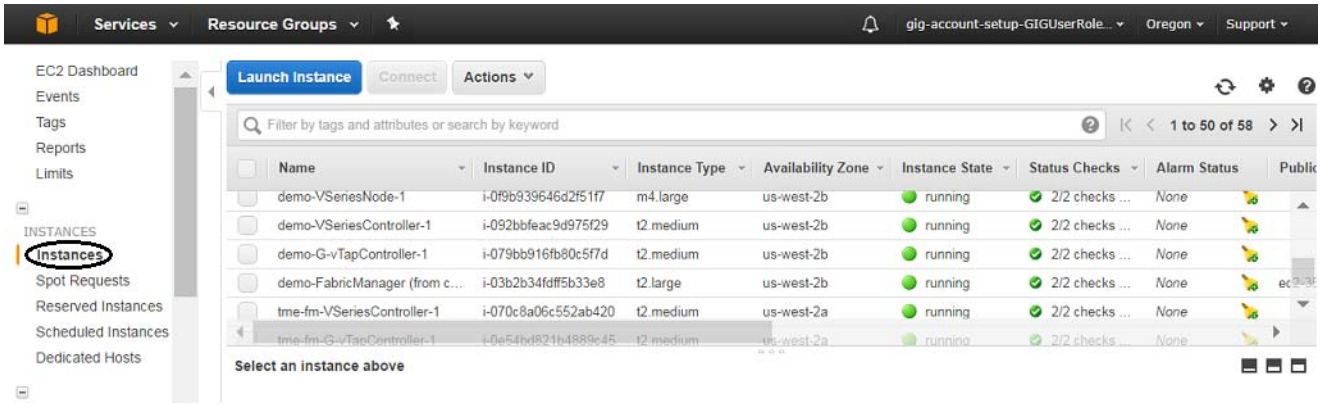

Figure 4-1: Selecting Instances

3. In the search field, enter the name of the existing GigaVUE-FM instance and select the Instance ID.

NOTE: If the instance ID is the password for logging in to the existing GigaVUE-FM, make note of this instance ID. This instance ID will be used as the password for logging in to the upgraded GigaVUE-FM as well. If the password is changed, use the changed password to login to the upgraded GigaVUE-FM.

4. Go to **Actions** > **Instance State** > **Stop**.

Creating a Snapshot of the GigaVUE-FM Instance

You must create a snapshot of the volume of the existing version (dev/sdb) of the GigaVUE-FM instance. Snapshots capture data that are written to your Amazon EBS volume at the time the snapshot is taken. This excludes any data that are cached by any applications or the operating system.

To create a snapshot:

1. Select the GigaVUE-FM instance and click the **Description** tab. Refer to Figure 4-2 on page 109.

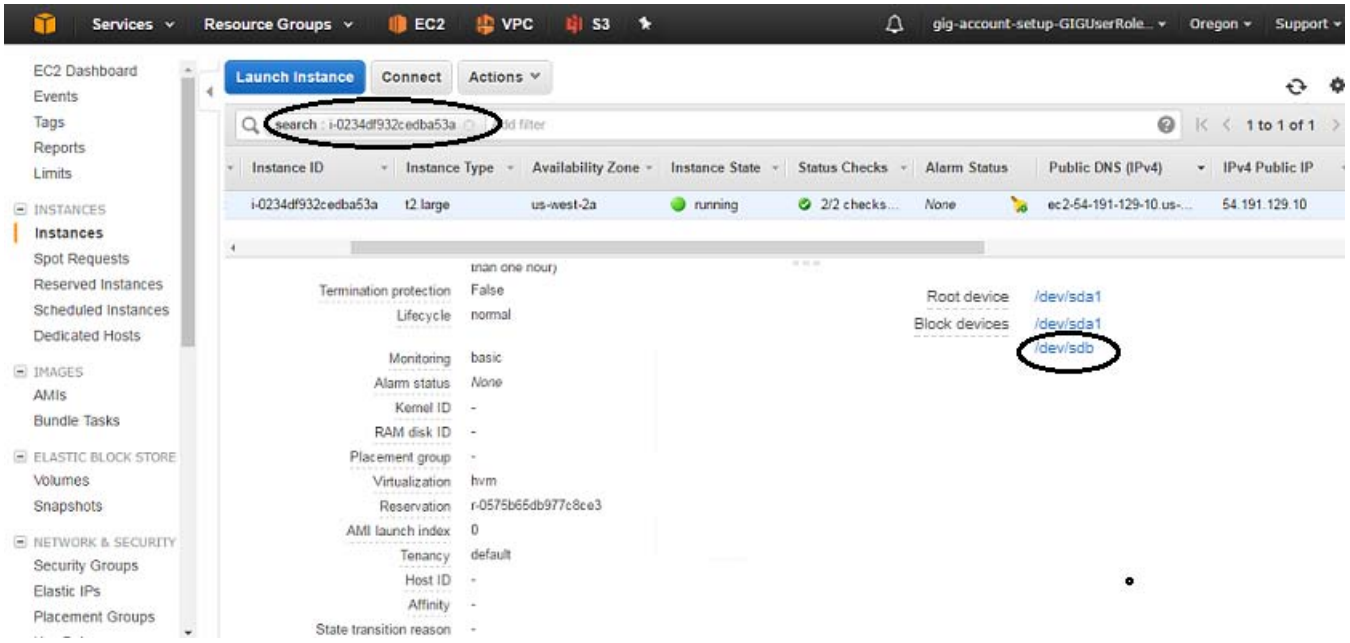

Figure 4-2: Searching for the GigaVUE-FM Instance

2. Scroll down and locate Block Devices. Refer to Figure 4-2 on page 109.
3. Click the **/dev/sdb** link. The Block Device dialog box is displayed with the volume ID link. Refer to Figure 4-3 on page 109.

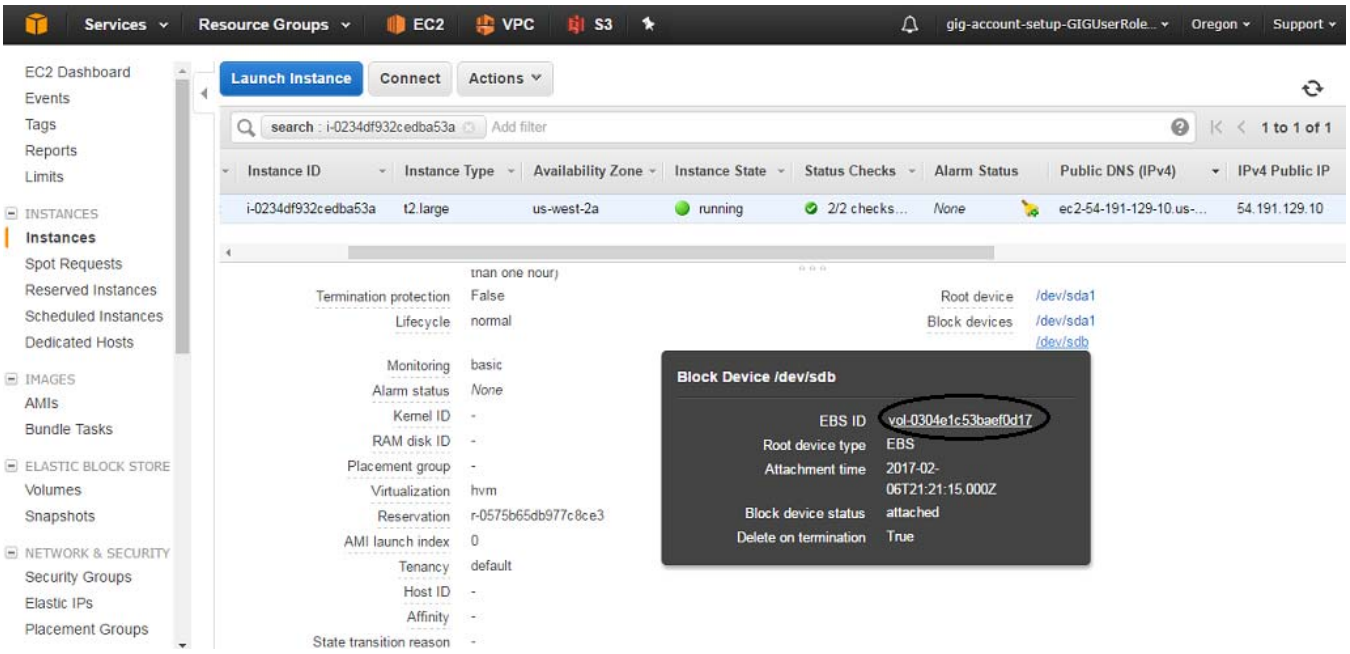

Figure 4-3: Opening Block Device Dialog Box

4. In the Block Device dialog box, click the volume ID link. The Volumes page is displayed. Refer to Figure 4-4 on page 110.

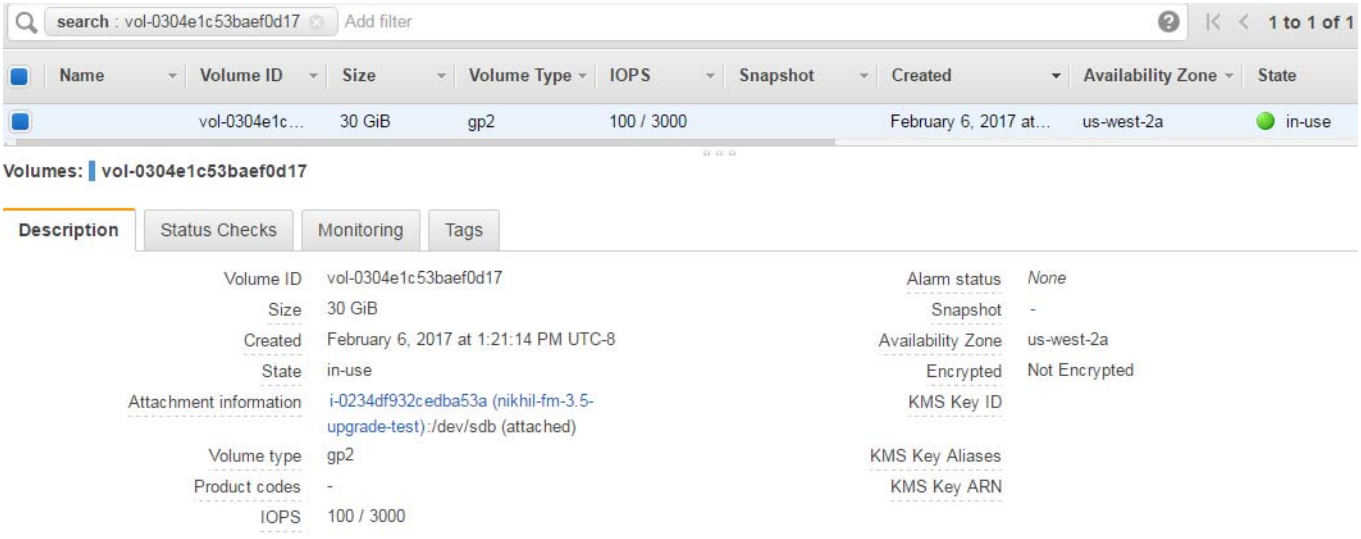

Figure 4-4: Viewing the Volumes Page

5. Click **Actions** and select **Create Snapshot.** Refer to Figure 4-5 on page 110.

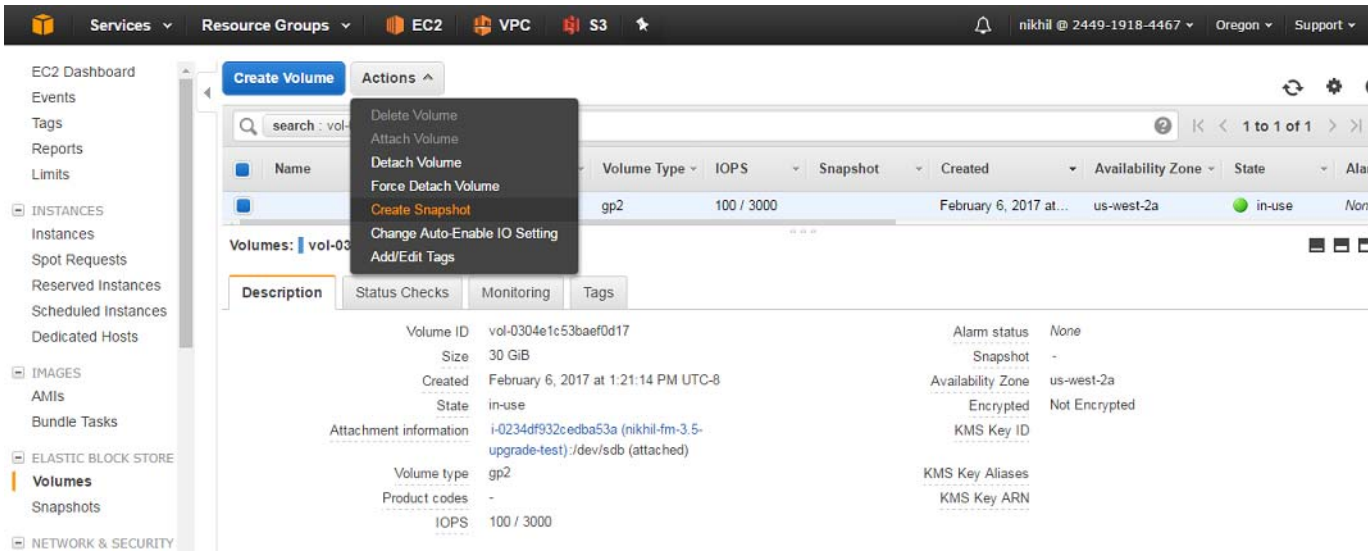

Figure 4-5: Selecting Create Snapshot

The Create Snapshot dialog box is displayed. Refer to Figure 4-6 on page 111.

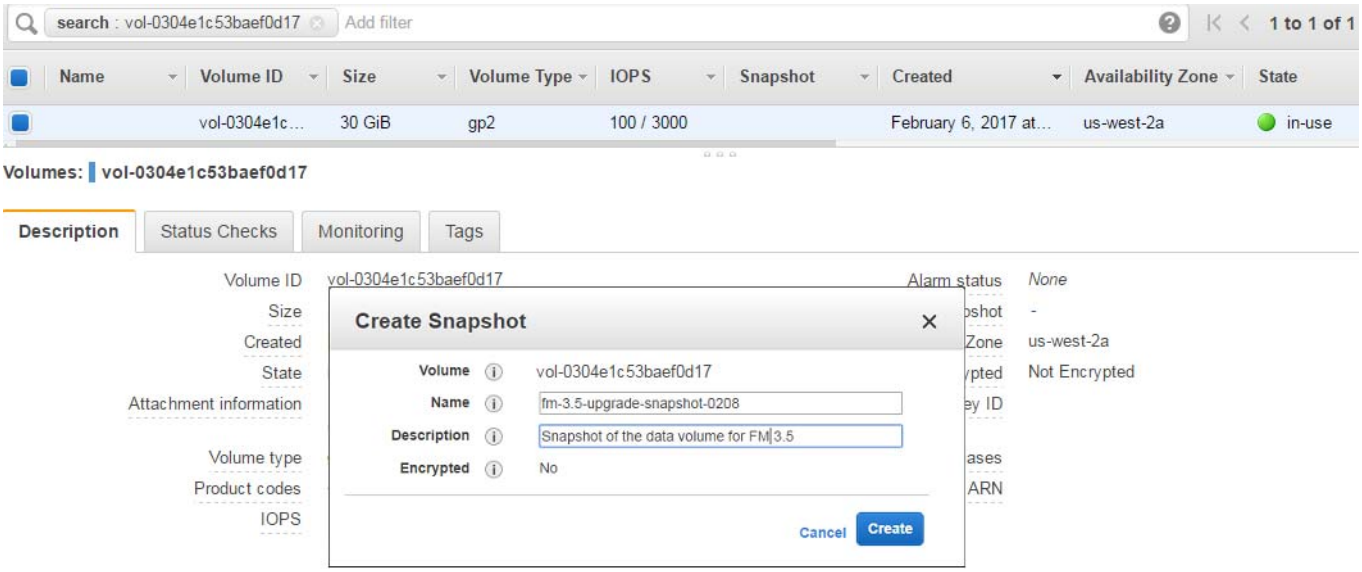

Figure 4-6: Creating a Snapshot

6. In the Create Snapshot dialog box, enter the following information:

Table 4-1: Fields for Creating a Snapshot

Field	Description
Name	The name of the snapshot.
Description	The description of the snapshot.

7. Click **Create**. It will take several minutes for the snapshot to be created. Refer to Figure 4-7 on page 111.

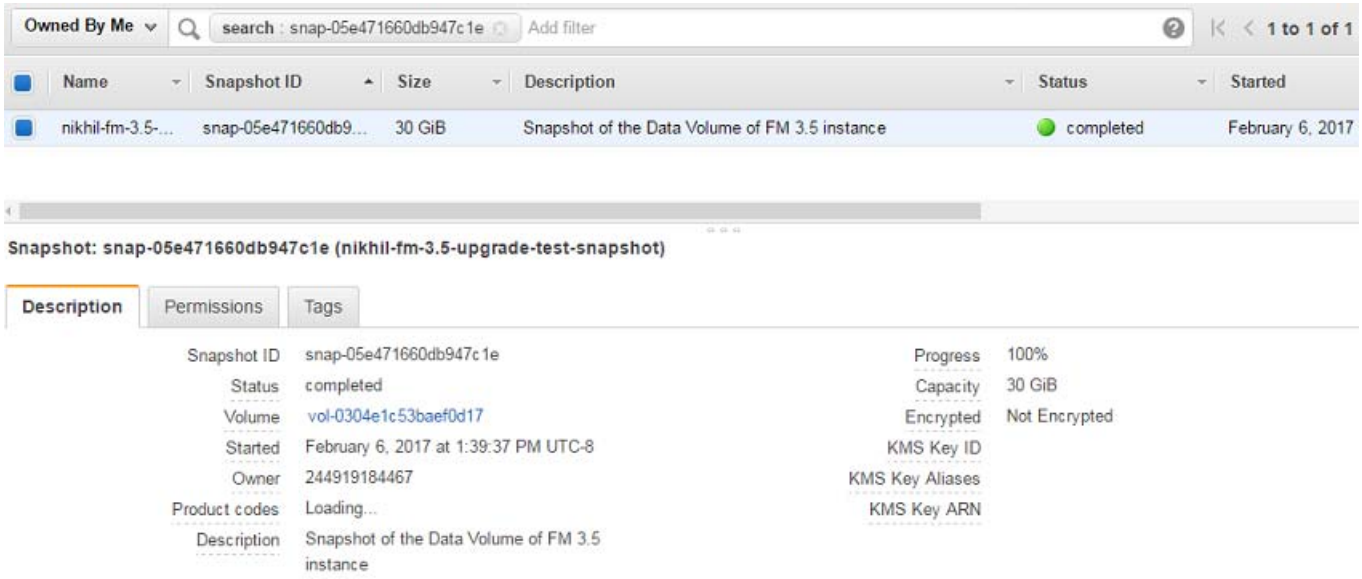

Figure 4-7: Viewing the Snapshot Page

NOTE: Make a note of the snapshot ID. This snapshot ID will be used to find the snapshot and add the volume while upgrading the GigaVUE-FM instance.

Upgrading the GigaVUE-FM Instance

While upgrading the GigaVUE-FM instance, the Amazon EBS volume must be restored with the data from the snapshot that is created in *Creating a Snapshot of the GigaVUE-FM Instance* on page 108.

To upgrade the GigaVUE-FM instance:

1. Select **Services > EC2**.
2. Click **Launch Instance** and go to **AWS Marketplace** or **Community AMIs**.
3. Search for **Gigamon**, locate the latest version of the GigaVUE-FM AMI, and click **Select**.
4. Choose the Instance Type. The recommended instance type is **m4.xlarge**.

NOTE: Do not select the t2 instance types as they are not supported.

5. Click **Next: Configure Instance Details.** Refer to Figure 4-8 on page 112.

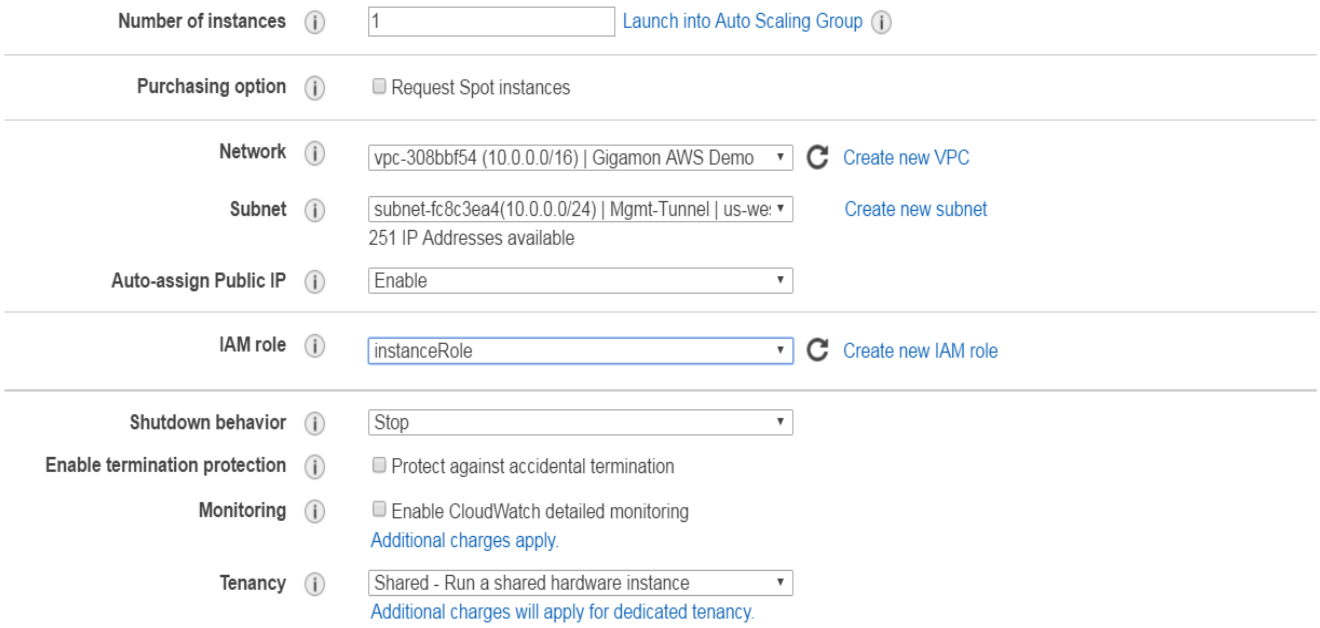

Figure 4-8: Configuring an Instance

6. Enter the following information.
   - **Network**— Select the VPC where you want to launch the AMI.
   - **Subnet**— Select the management subnet that the instance will use after launch.
   - **Auto-assign Public IP**— Select **Enable**.
   - **IAM role**—Select an existing IAM role to associate with the instance.
7. Click **Next: Add Storage** and click **Add New Volume**. Refer to Figure 4-9 on page 113.

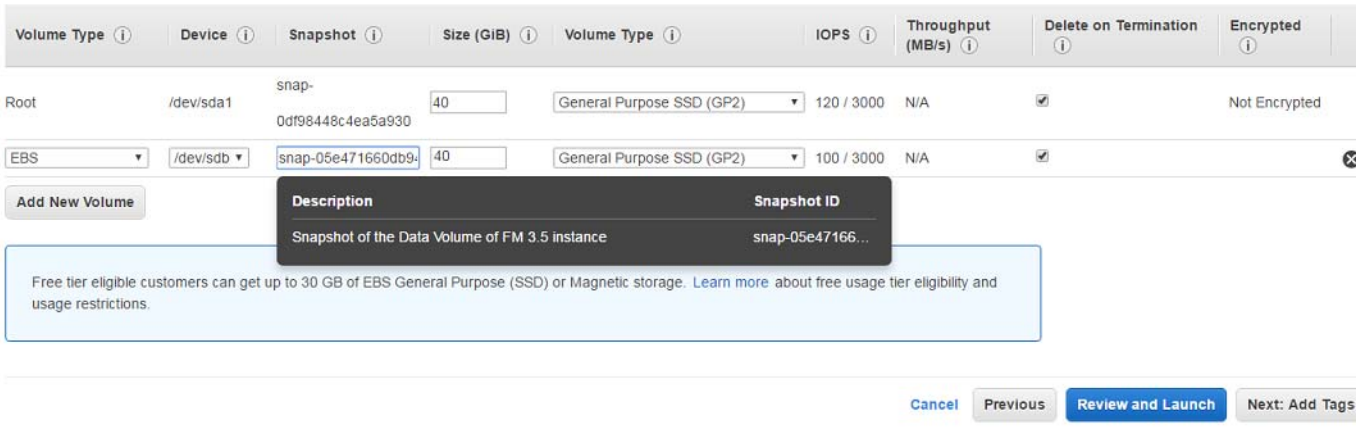

Figure 4-9: Adding New Volume

8. Enter the following storage device settings as shown in Figure 4-9 on page 113:
   - **Snapshot**—Enter the name of the snapshot that is created in step 9 in the section *Creating a Snapshot of the GigaVUE-FM Instance* on page 108.
   - **Size (GiB)**— Enter a minimum of 40Gb of storage. The size of the volume must be same as the volume selected while launching the previous version of the GigaVUE-FM instance.
   - **Volume Type**— Select a volume type. The recommended volume is General Purpose SSD (GP2).
   - **Delete on Termination**— Select this check box to make sure the volumes are cleaned up when the GigaVUE-FM instance is removed.
9. Click **Next: Tag Instance,** and then add a key-value pair to identify the instance. Refer to Figure 4-10 on page 113.

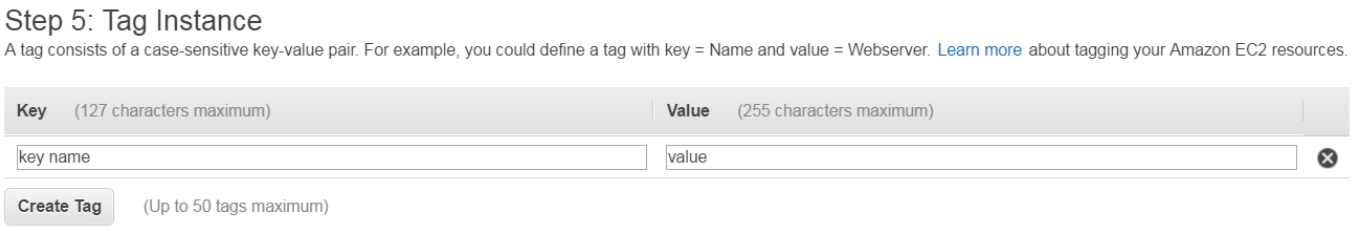

Figure 4-10: Adding a Tag to an Instance

10. Click **Next: Add Security Group**. Click the **Select an existing security group** check box if the security group is already created. Otherwise, select the **Create a new security group** check box and click **Add Rule**. For more information on creating a security group, refer to *Security Group* on page 20.
11. Click **Review and Launch**. Review the instance launch details and click **Launch**.
12. Select the SSH key pair, check the acknowledgment check box, and click **Launch Instances** as shown in Figure 4-11 on page 114.

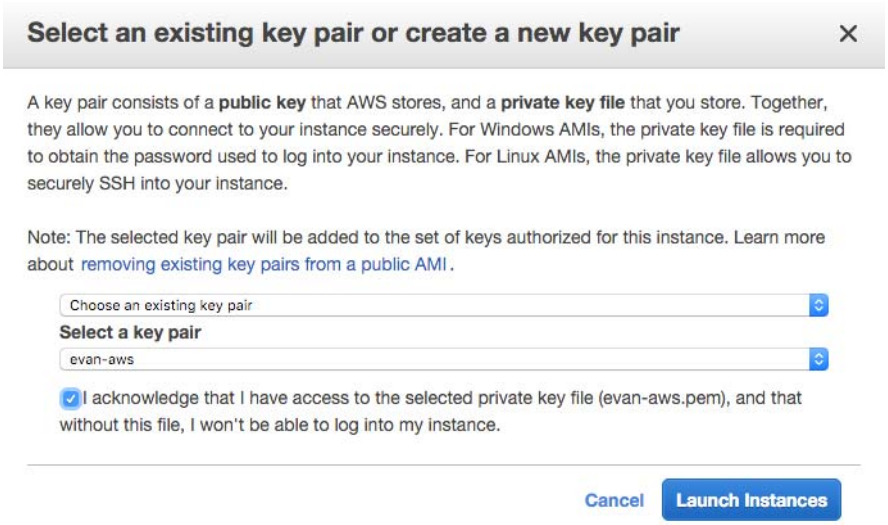

Figure 4-11: Selecting an SSH Key Pair

13. It will take several minutes for the instance to initialize. After the initialization is completed, verify the instance through the Web interface as follows:
    a. Find the instance and expand the page in the **Descriptions** tab to view the instance information, if necessary.
    b. Copy the Public DNS value and paste the value into a new browser window or tab.
    c. Copy the Instance ID of the previous version of the GigaVUE-FM. If the password is changed, use the changed password to login to the upgraded GigaVUE-FM.

NOTE: Do not have multiple versions of GigaVUE-FM instances monitoring the same AWS connection.

Launch the new version of the GigaVUE-FM instance. Verify if the data from the previous GigaVUE-FM instance is restored in the new instance. Once the data is verified, terminate the old version of the GigaVUE-FM instance.

Chapter 5

Upgrading the Virtual Fabric

This chapter describes how to upgrade GigaVUE V Series Controllers and GigaVUE V Series nodes.

NOTE: G-vTAP Controllers cannot be upgraded. Only a new version that is compatible with the G-vTAP agents' version can be added during the G-vTAP configuration.

Prerequisite

Before you upgrade the GigaVUE V Series Controllers and GigaVUE V Series nodes, you must upgrade GigaVUE-FM to software version 5.1 or above. For information about upgrading the GigaVUE-FM instance, refer to *Upgrading the GigaVUE-FM Instance* on page 107.

NOTE: The older version of virtual fabric is compatible with GigaVUE-FM 5.1. For better performance, Gigamon recommends you to upgrade to the latest version.

Upgrading the GigaVUE V Series Controllers and Nodes

GigaVUE-FM lets you upgrade GigaVUE V Series Controllers and GigaVUE V Series nodes at a time.

There are multiple ways to upgrade the GigaVUE V Series Controllers 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 Controller and 10 GigaVUE V Series nodes in your VPC, you can upgrade all of them at once. First, the new version of GigaVUE V Series controller is launched. Next, the new version of GigaVUE V Series nodes are launched. Then, the old version of V Series controller and nodes are deleted from the VPC.

  NOTES:

  - When the new version of nodes and controllers are launched, the old version still exists in the VPC until they are deleted. Make sure the instance type determined during the configuration can accommodate the total number of new and old instances present in the VPC. If the

instance type cannot support so many instances, 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 VPC, the new version of the fabric is immediately deleted and the old version of the fabric is retained as before.

- 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 Controllers and GigaVUE V Series Nodes:

1. Click **Cloud** in the top navigation link.
2. In the left navigation pane, select **Visibility Fabric** > **V Series Controllers**. Refer to Figure 5-1 on page 116.

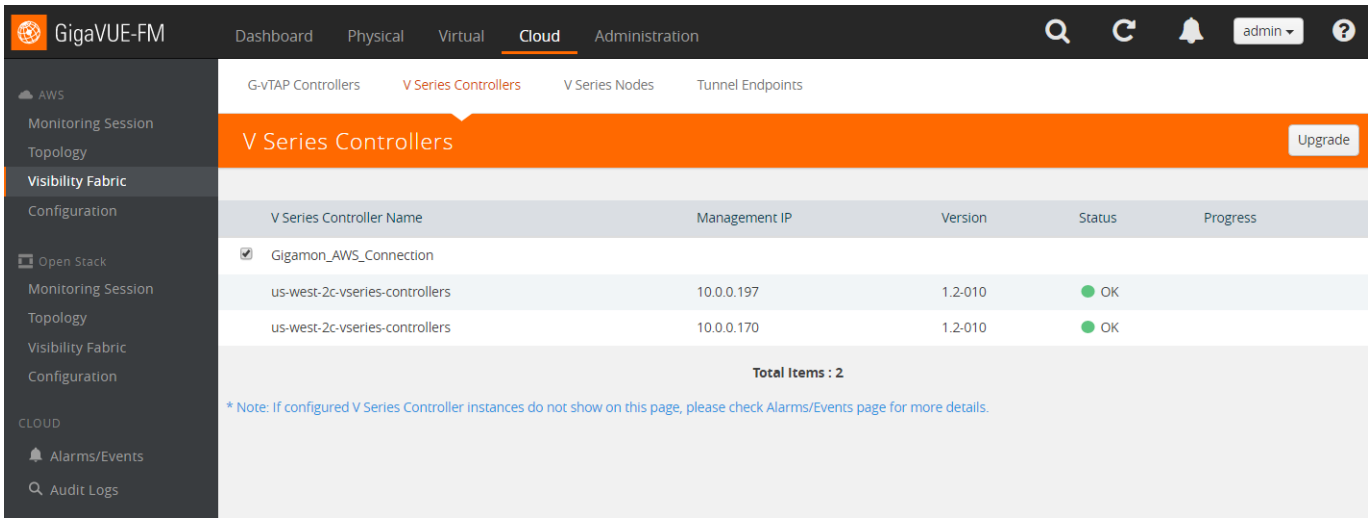

Figure 5-1: Gigamon Virtual Fabric Upgrade

3. Select the connection name check box and click **Upgrade**. The V Series Controller and Node Upgrade page is displayed. Refer to Figure 5-2 on page 116.

Figure 5-2: GigaVUE V Series Controller and Node Upgrade

4. From the **Version** drop-down list, select the latest version of the GigaVUE V Series Controller.

5. To upgrade the GigaVUE V Series Controllers, specify the batch size in the **Batch Size for V Series Controller** box.

   For example, if there are 4 GigaVUE V Series Controllers in your VPC, you can specify 4 as the batch size and upgrade all of them at once or specify 2 as the batch size and upgrade 2 GigaVUE V Series Controllers in each batch.

6. To upgrade the GigaVUE V Series nodes, specify the batch size in the **Batch Size for V Series Nodes** 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.

7. Click **Upgrade**.

   The upgrade process takes a while depending on the number of GigaVUE V Series controllers and nodes upgrading in your AWS environment. First, the new version of the GigaVUE V Series Controllers is launched. Next, the new version of GigaVUE V Series nodes is launched. Then, the older version of both is deleted from the VPC. In the V Series Controllers page, click the link under Progress to view the upgrade status. Refer to Figure 5-3 on page 117.

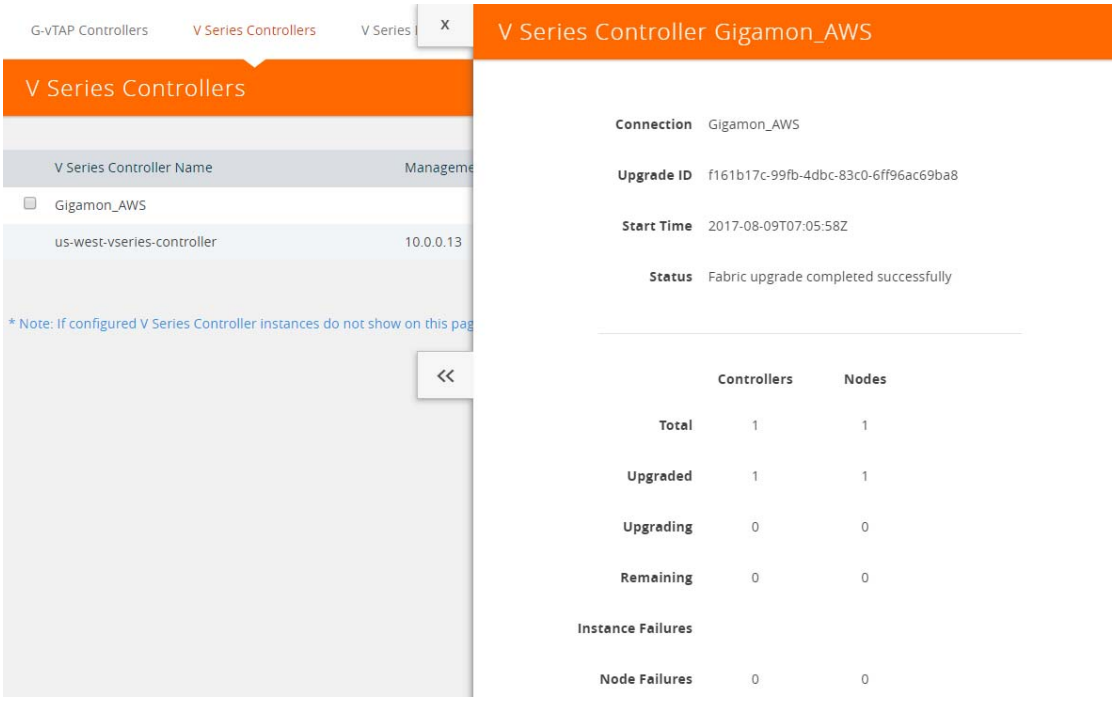

Figure 5-3: GigaVUE Fabric Upgrade Status

Appendix A

Glossary

This appendix lists the AWS terminologies used in this document. To find a brief definition of these terms, refer to AWS Glossary.

- Access Key
- Access key ID
- Amazon API Gateway
- Amazon CloudWatch
- Amazon CloudWatch Events
- Amazon Elastic Compute Cloud (Amazon EC2)
- Amazon VPC
- AMI
- AWS
- AWS Identity and Access Management (IAM)
- CIDR block
- EC2 Instances
- Elastic IP address
- Endpoint
- Instance
- Instance type
- Internet gateway
- Key pair
- Secret access key
- Subnet
- Tag
- Target Instance
- Tunnel

Appendix B

Compatibility Matrix

This appendix provides the compatibility matrix information for the Gigamon Visibility Platform components. The following table lists the GigaVUE-FM instance version and the supported versions of its components.

GigaVUE-FM	G-vTAP Agent Version	G-vTAP Controller Version	GigaVUE V Series Controller	GigaVUE-V Series Nodes
5.0	v1.0	v1.0	v1.0	v1.0
5.0.1	v1.0	v1.0	v1.0, v1.1	v1.0, v1.1
5.1	v1.0	v1.0	v1.0, v1.1	v1.0, v1.1
5.2	v1.2, v1.3	v1.2, v1.3	v1.0, v1.1, v1.2, v1.3	v1.0, v1.1, v1.2, v1.3

The following table lists the features and the supported version of GigaVUE V Series node:

Features	GigaVUE V Series v1.0	GigaVUE V Series v1.2	GigaVUE V Series v1.3
Header Transformation	No	No	Yes
Multi-link Support	No	Yes	Yes
NetFlow Application	No	Yes	Yes
NAT Support	No	Yes	Yes

The following table lists the features and the supported version of G-Tap Agents:

Features	G-vTAP Agent v1.2	G-vTAP Agent v1.3
Single ENI Support	No	Yes
VXLAN Support	No	Yes

Appendix C

Additional Sources of Information

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

Documentation

Gigamon provides additional documentation for this solution on the Gigamon Customer Portal.

Document	Summary
GigaVUE-FM and GigaVUE VM User's Guide	Describes how to install, deploy, and operate the GigaVUE® Fabric Manager (GigaVUE-FM) and GigaVUE® Virtual Manager (GigaVUE-VM) from Gigamon® Inc.

Documentation Feedback

To send feedback and report issues in our documentation, complete the short survey at the following link:

https://www.surveymonkey.com/r/gigamondocumentationfeedback

Contacting Technical Support

Refer to http://www.gigamon.com/support-and-services/contact-support for Technical Support hours and contact information. You can also email Technical Support at support@gigamon.com.

Contacting Sales

Table i shows how to reach the Sales Department at Gigamon.

Table i: Sales Contact Information

Telephone	+1 408.831.4025
Sales	inside.sales@gigamon.com

Premium Support

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