FIREEYE TECHNICAL DOCUMENTATION

DISTRIBUTED NETWORK SECURITY

MVX SMART GRID GUIDE

RELEASE 2021.2

Distributed Network Security MVX Smart Grid Guide

Software Release 2021.2

Revision 1

FireEye Contact Information:

Website: www.fireeye.com

Technical Support: https://csportal.fireeye.com

Phone (US):

1.408.321.6300

1.877.FIREEYE

Contents

PART I: Overview

CHAPTER 1: Introduction

A standard (or *integrated*) appliance performs both monitoring and analysis. FireEye Distributed Network Security separates these two functions. Appliances that function as *sensors* extract objects and URLs from the traffic they monitor, and send submissions to an MVX *cluster* for inspection and analysis. A sensor and an integrated appliance have identical features and detection efficacy.

An appliance running in MVX *hybrid* mode can send submissions to an MVX cluster, but only when a predefined capacity threshold is reached. This offloads the analysis function from the appliance to the MVX cluster, which prevents delays and reduced efficacy when volume and other processing demands are high. When the capacity falls below this threshold, the appliance resumes sending submissions to its on-board analysis engine.

Sensors can be managed by the Central Management appliance that manages the MVX cluster or by another Central Management appliance. The sensors can also be standalone appliances that are not managed by a Central Management appliance.

Hybrid appliances must be managed by the Central Management appliance that manages the MVX cluster. They cannot be standalone appliances.

The MVX cluster contains *compute nodes*, which are VX Series appliances with MVX analysis engines. One or two compute nodes are designated as *brokers*. The brokers receive the submissions from the sensors and manage them in a queue that is distributed across the brokers in the cluster. The compute nodes pull submissions from the queue, perform the analysis, and send the verdict to the sensors through the brokers.

The sensors generate alerts based on the verdict. A managed sensor sends the alerts to its managing Central Management appliance, which aggregates the alerts and displays them on a single interface. A standalone sensor displays its own alerts.

Sensor Enrollment and Submission

A sensor or hybrid appliance must be enrolled in a cluster. The Central Management appliance that manages the MVX cluster provides the enrollment service for the sensors and hybrid appliances and the cluster. A cluster uses the enrollment service to publish its identity, capacity, and capabilities. A sensor or hybrid appliance uses the enrollment

service to enroll with a cluster based on matching criteria and available capacity. The sensors and hybrid appliances and the brokers use the enrollment service to authenticate with each other for secure communication.

Automatic enrollment is enabled by default. After a sensor or hybrid appliance is connected to a Central Management appliance with at least one cluster, the sensor is automatically enrolled with the cluster that has the most capacity, and is connected to a dedicated broker in the cluster. The enrollment is permanent—submissions are distributed between brokers in the same cluster, but not between clusters. If sensors are not managed by the Central Management appliance that manages the MVX cluster, you must configure the enrollment service as described in Enrolling a Managed Sensor Directly on page 64.

Queue Management and Analysis

A Virtual Execution appliance is ready to be a compute node as soon as it is added to an MVX cluster. The cluster interface that the compute node uses to communicate with the brokers and other compute nodes is defined in the configuration wizard during the initial configuration of the Virtual Execution appliance. The submission interface that the broker uses to communicate with sensors and hybrid appliances is also defined in the configuration wizard. You must manually designate a compute node as a broker. A broker can perform analysis as well as managing the queue.

The detection-related configuration settings on all compute nodes must match, because any compute node in a cluster can process submissions from any sensor or hybrid appliance. One broker must be designated as the master configuration. The Central Management appliance that manages the MVX cluster synchronizes relevant configuration settings on other compute nodes (including brokers) with the master configuration.

An Network Security sensor or an Network Security hybrid appliance applies its policies and filters to the traffic it receives on its monitoring ports, and extracts suspicious or malicious files and URLs that need to be inspected by the MVX cluster. An Email Security — Server Edition sensor or hybrid appliance applies its policies and filters to the emails it receives on its network interfaces, and extracts suspicious or malicious files or URLs that need to be inspected by the MVX cluster. An Network Security sensor or hybrid appliance applies its policies and filters to the scans it runs on network shares, and extracts suspicious or malicious files that need to be inspected by the MVX cluster. The sensor or hybrid appliance then sends a submission to its dedicated broker.

The broker receives the submissions from the sensor or hybrid appliance and places them in the queue. Each compute node is connected to all brokers in its cluster. When a compute node has processing capacity, it pulls submissions from a broker and performs the analysis. The compute node sends the verdict and malware artifacts to the sensor or hybrid appliance through the broker. The sensor or hybrid appliance generates alerts and takes action based on its deployment and operational mode and policies.

Alerts and Notifications

Alerts from all managed sensors and hybrid appliances and any other managed appliance types are aggregated on the Central Management appliance. Alert types and notification types for sensors are the same as for an integrated appliance.

Detection Features

A sensor has the same detection features and requires the same configuration as an integrated appliance, with the exception of guest images, YARA rules, and static analysis tools.

IMPORTANT: Do not configure YARA rules and static analysis tools using the File Protect appliance settings. They are configured on and managed by the Virtual Execution appliances (nodes).

Hybrid appliances require the same configuration as integrated appliances, because they operate as integrated appliances until their capacity reaches the predefined threshold for sending submissions to an MVX cluster.

A compute node or broker requires little configuration:

- The values you define and the default values you accept in the configuration wizard during the initial configuration are sufficient to deploy the Virtual Execution appliance in your network and to enable it to function as a broker or compute node.
- Guest images are pre-installed on the Virtual Execution appliance; the Central Management Web UI warns you when they are out-of-date and provides an easy way to update them.
- Static analysis tools are enabled by default.
- FireEye provides a set of YARA rules. You can define and upload custom YARA rules on a Virtual Execution appliance. The "Write changes to group" feature in the Central Management Web UI allows you to apply the rules to all brokers and compute nodes at once.

Health Monitoring

The Central Management appliance continuously monitors the health and capacity of the MVX cluster. If the cluster loses integrity, the Central Management appliance sends email notifications to configured recipients and displays warnings in its Dashboard and other Web UI pages, and in the output of CLI commands. The warnings persist until you explicitly acknowledge them.

You can check the health of the cluster after you deploy it, as described in Viewing Cluster Utilization on page 133 and Viewing Submission Statistics on page 137.

Advanced Configurations

The following advanced configurations can be used in an MVX cluster deployment in which a Central Management appliance manages the sensors.

- **Network Address Translation (NAT)**—The Central Management appliance in an MVX cluster deployment can be deployed in an internal network behind a NAT gateway. In certain scenarios, you must configure an accessible address for the Central Management management interface so the managed sensors and nodes can reach it. For details, see *Administration Guide* or *System Administration Guide* for the appliance.
- **Client-Initiated Connections**—This guide shows how to add the appliances using a server-initiated connection, in which you use the Central Management appliance to directly connect appliances. Alternatively, you can use a client-initiated connection, in which the appliance administrator initiates a request to be managed, and a Central Management administrator explicitly accepts the request. For information about using a client-initiated connection, see the *Administration Guide* or *System Administration Guide* for the appliance.
- **Multiple DTI Sources**—By default, a standalone sensor and the Central Management appliance use **cloud.fireeye.com** (redirected to **download.fireeye.com**) to download software updates from the DTI network. You can change the DTI source server to **staticcloud.fireeye.com**. By default, managed appliances use the Central Management appliance as their DTI source server. In certain configurations, you can change the DTI source server for sensors and nodes to **cloud.fireeye.com** or **staticcloud.fireeye.com**. For details, see the *Administration Guide* or *System Administration Guide* for the appliance.
- **Secure Shell (SSH) Authentication**—The Secure Shell (SSH) protocol is used for secure communication between the Central Management appliance and the sensors and nodes. You can configure advanced options for SSH user authentication, which verifies the identity of the remote user attempting the connection. You can also configure advanced options for SSH host authentication, which verifies the identity of the Central Management appliance to the sensor or node, and verifies the identity of the sensor or node to the Central Management appliance. For details, see the *FireEye System Security Guide*.

CHAPTER 2: Product Terminology

Some FireEye MVX Smart Grid components are referred to differently in the user interface and documentation. The following table maps the component name to the user interface and documentation term.

Product or Component Name	User Interface Term
Network Smart Node	Sensor or Hybrid Appliance
MVX Smart Grid Broker	Broker Node
MVX Smart Grid Element	Compute Node

Other Terms You Should Know

Physical Sensor

A hardware Network Security appliance that has been converted to a sensor (which disables its on-board analysis engine) or a hardware appliance that has no on-board analysis engine. These appliances send submissions to an MVX cluster for analysis. A physical sensor operates in MVX sensor mode.

Virtual Sensor

A virtual appliance that has no on-board analysis engine and instead sends submissions to an MVX cluster for analysis. A virtual sensor operates in MVX sensor mode.

Integrated Appliance

A hardware appliance that uses its on-board analysis engine instead of sending submissions to an MVX cluster for analysis. An integrated appliance operates in MVX local mode.

Hybrid Appliance

An integrated hardware appliance that can send submissions to an MVX cluster, but only when a predefined capacity threshold is reached. When the capacity is below this threshold, the appliance uses its on-board analysis engine for inspection and analysis. A hybrid appliance operates in MVX hybrid mode.

Client-Initiated Connection

A method in which an appliance administrator (for example, an Network Security administrator) sends a request for management to the Central Management appliance, and a Central Management administrator accepts the request.

Server-Initiated Connection

A method in which a Central Management administrator directly adds an appliance to the Central Management appliance for management.

Central Management-Managed Appliance

An appliance that is under Central Management management.

Standalone Appliance

An appliance that is not under Central Management management.

MSSP Deployment

A Distributed Network Security environment that is deployed and maintained by a managed security service provider (MSSP).

CHAPTER 3: Supported Releases

This release of the *MVX Smart Grid Guide* supports the following releases:

- **Virtual Execution:** Version 8.3.0 or later for cluster nodes
- **Central Management:** Version 8.5.0 or later for cluster upgrade and configuration

For information on supported sensors, see System Requirements on page 35.

CHAPTER 4: Architecture

FireEye provides three Distributed Network Security deployment options:

- The same Central Management appliance manages both the sensors and hybrid appliances and the MVX cluster.

  With this option, a single physical or virtual Central Management appliance or a physical Central Management High Availability (HA) pair manages the MVX cluster components. All sensors, brokers, and compute nodes must be connected to the same Central Management appliance.

  NOTE: This is the only option for an MVX hybrid appliance.

- The sensors are managed by a local Central Management appliance and submit to a remote MVX cluster that is managed by a different Central Management appliance. With this option, the brokers and compute nodes must be connected to the same Central Management appliance.
- The sensors are standalone appliances, and submit directly to an MVX cluster.

A Central Management appliance can manage both the MVX components and other appliances, such as integrated Network Security appliances and Email Security — Server Edition appliances.

The MVX cluster components (brokers and compute nodes) must be deployed on the same LAN. The MVX cluster, the sensors, and the Central Management appliance can be in different physical locations.

IMPORTANT: FireEye does not recommend transcontinental deployments due to throughput, reliability, and latency issues.

A sensor can be a physical or virtual Network Security appliance, a virtual Email Security — Server Edition appliance, or a virtual File Protect appliance. Some Network Security appliance models can function only as sensors, because they do not include an MVX analysis engine. Some physical Network Security appliances can be enabled as sensors, in which case the analysis engine is disabled. A virtual appliance can function only as a sensor.

A hybrid appliance is a physical Network Security, Email Security — Server Edition, or File Protect appliance. A hybrid appliance has an on-board analysis engine, but stops using it when a predefined threshold is exceeded. The appliance then essentially acts as a sensor, because it sends all submissions to an MVX cluster until the capacity falls below this threshold.

A broker or node is a physical Virtual Execution appliance. A Virtual Execution appliance serves no purpose until it is added to a cluster.

The following diagrams show an architecture in which the Central Management appliance manages both physical and virtual sensors, one MVX cluster, and integrated appliances that are not sensors or MVX cluster components.

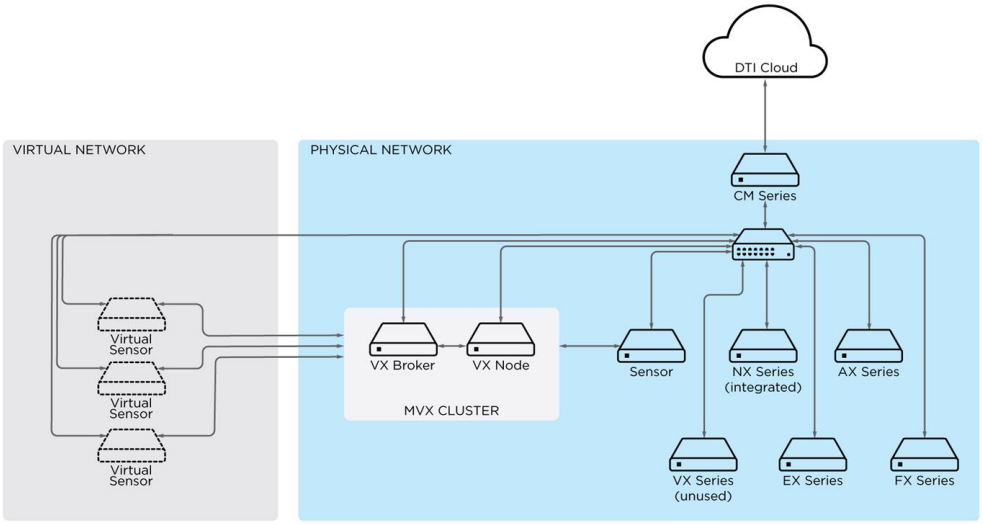

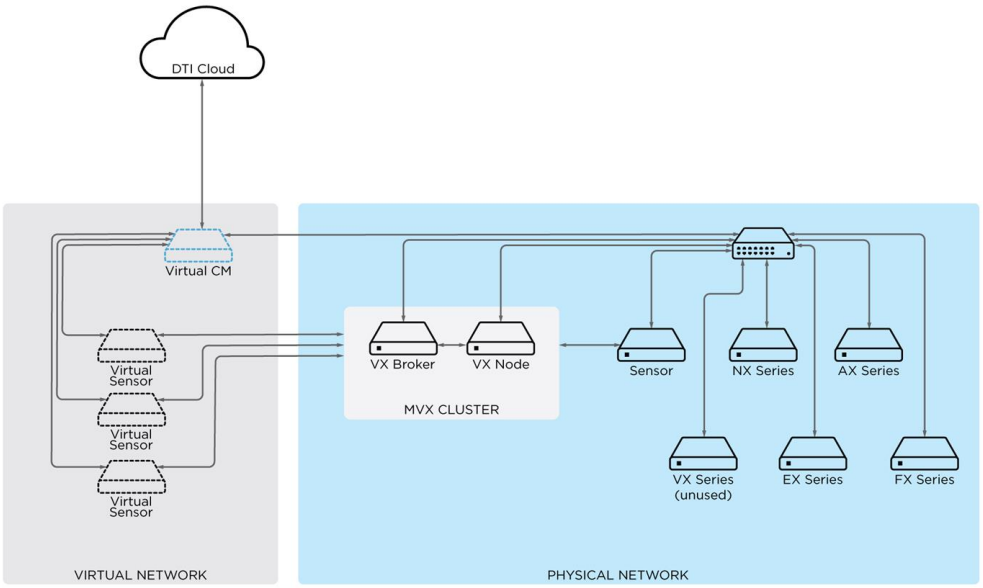

Communication Paths

The following diagram shows the communication paths between the MVX cluster components.

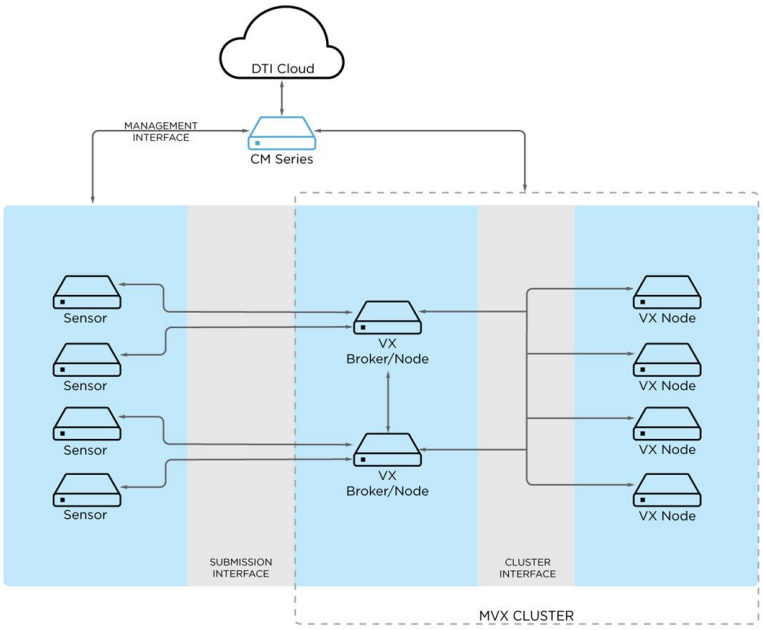

During the initial configuration, you can accept the default interfaces provided by the configuration wizard, or you can change them.

- **Management Interface**—In the preceding diagrams, the enrolled sensors and the brokers and nodes are connected to the same Central Management appliance. The connection is established through the management interface the same way as it is for other managed appliances. By default, both SSH and HTTPS traffic use the SSH port (port 22).

  

  NOTE: Alternatively, the sensors can be connected to another Central Management appliance or can be standalone appliances.

- **Submission Interface**—Sensors and brokers communicate with each other through the submission interface using SSH. The default submission interface is ether1. A second interface (ether2) can also be configured.
- **Cluster Interface**—Brokers and nodes communicate with each other through the cluster interface using SSH and other protocols. The default cluster interface is ether1, but another interface (for example, ether2) can be configured instead.

Brokers expose ports 25672 and 4369 for inter-broker communication. They also expose port 5671 for communication with compute nodes. Ports 25672 and 5671 are SSL-encrypted. Port 4369 is protected by key hash. Sensors and compute nodes connect to brokers using SSH (port 22). The cluster database uses TCP port 7001. Cluster management communication uses TCP and UDP ports 18300 through 18303.

CHAPTER 5: Standard Deployment Scenarios

This section illustrates example deployment scenarios.

Sensors and Clusters Managed by Same Central Management Appliance

In the following two deployment scenarios, the sensors and MVX clusters are managed by the same Central Management appliance.

Distributed Sensors

In the following example, the Central Management appliance and two MVX clusters are deployed in one data center. Each MVX cluster analyzes objects submitted by sensors that are in two geographic regions.

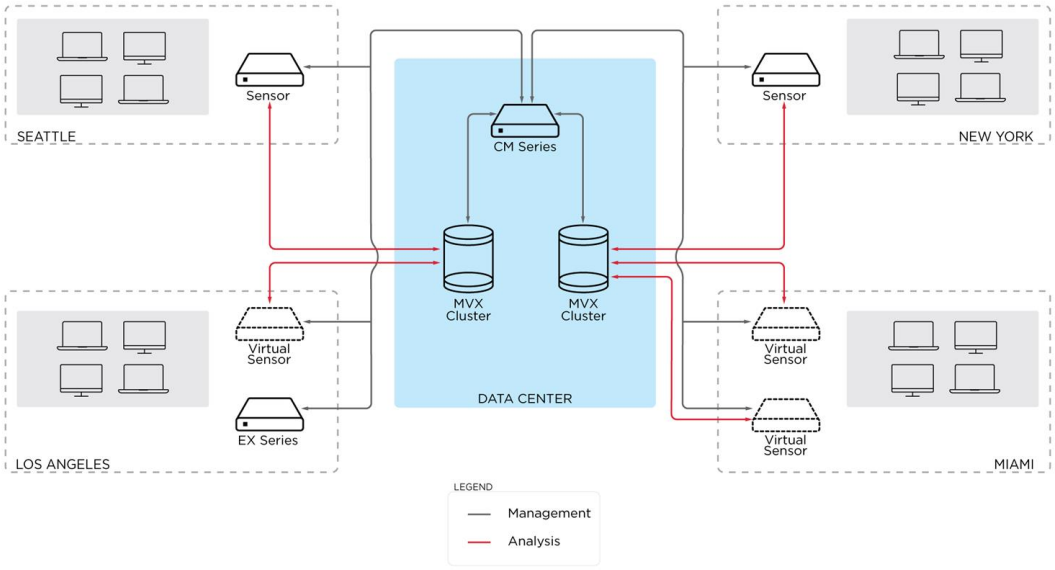

Distributed Clusters and Sensors

In the following example, the Central Management appliance and one MVX cluster are deployed in the same data center, and the other MVX cluster is deployed in another data center.

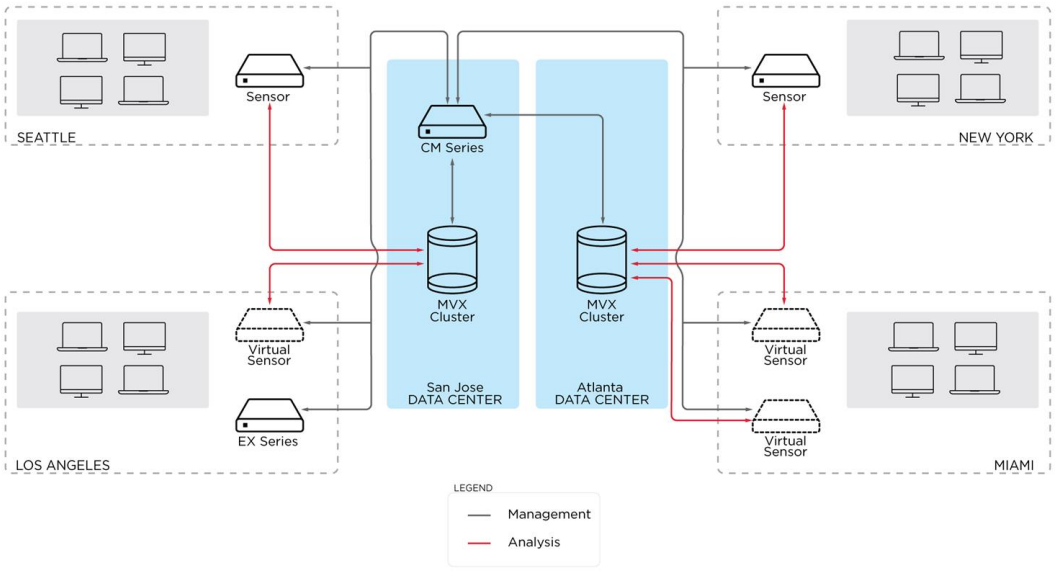

Sensors and Cluster Managed by Different Central Management Appliances

In the following two deployment scenarios, the sensors are managed by one Central Management appliance, and the cluster is managed by another Central Management appliance.

Sensor Enrolls Directly with Cluster

The enrollment service address type is DTI and its address is the IP address of the Central Management appliance that manages the MVX cluster. You must prevent the Central Management appliance that manages the sensor from overriding the configuration. See Enrolling a Managed Sensor Directly on page 64 for configuration information.

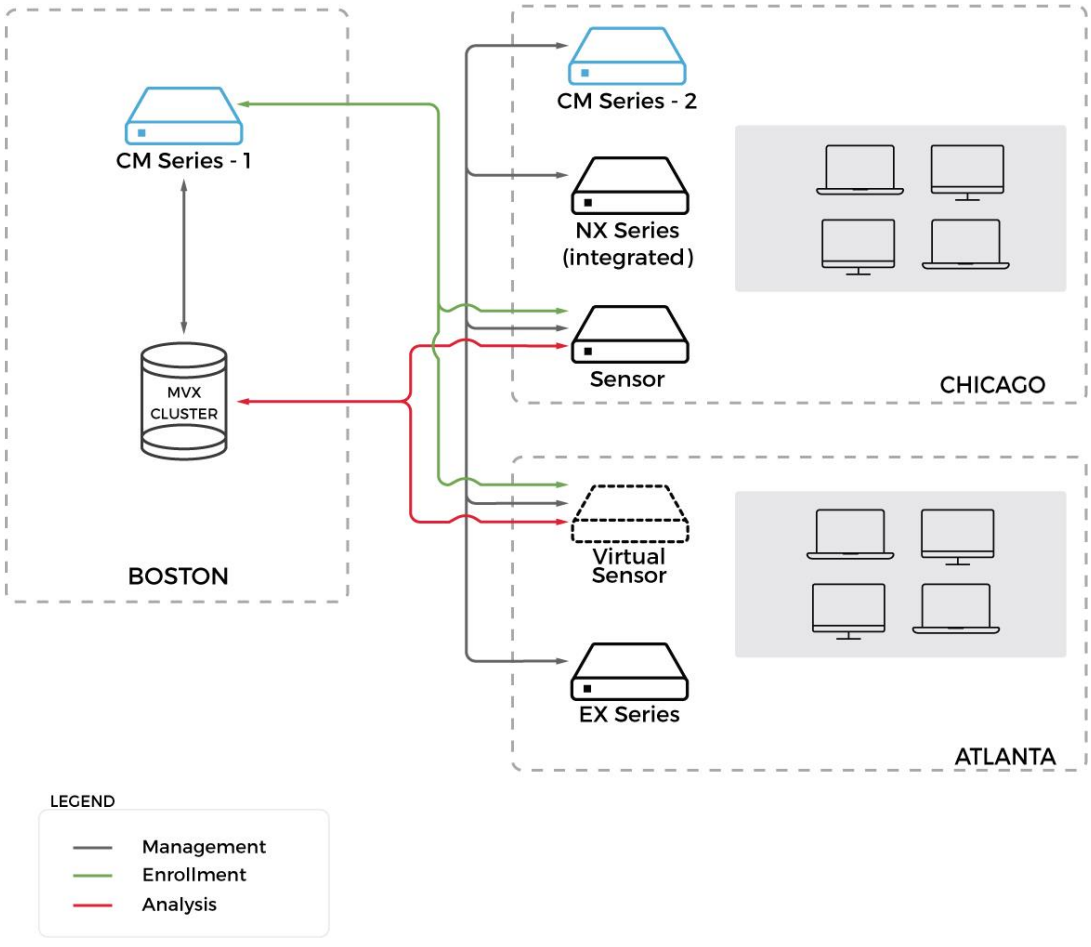

Sensor Uses Proxy Enrollment

On the sensor, the enrollment service type is CMS and its address is the IP address of the Central Management appliance that manages the sensor. On the Central Management appliance that manages the sensor, the enrollment service type is DTI and its address is the IP address of the Central Management appliance that manages the MVX cluster. See Enrolling a Managed Sensor Through a Proxy on page 66 for configuration information.

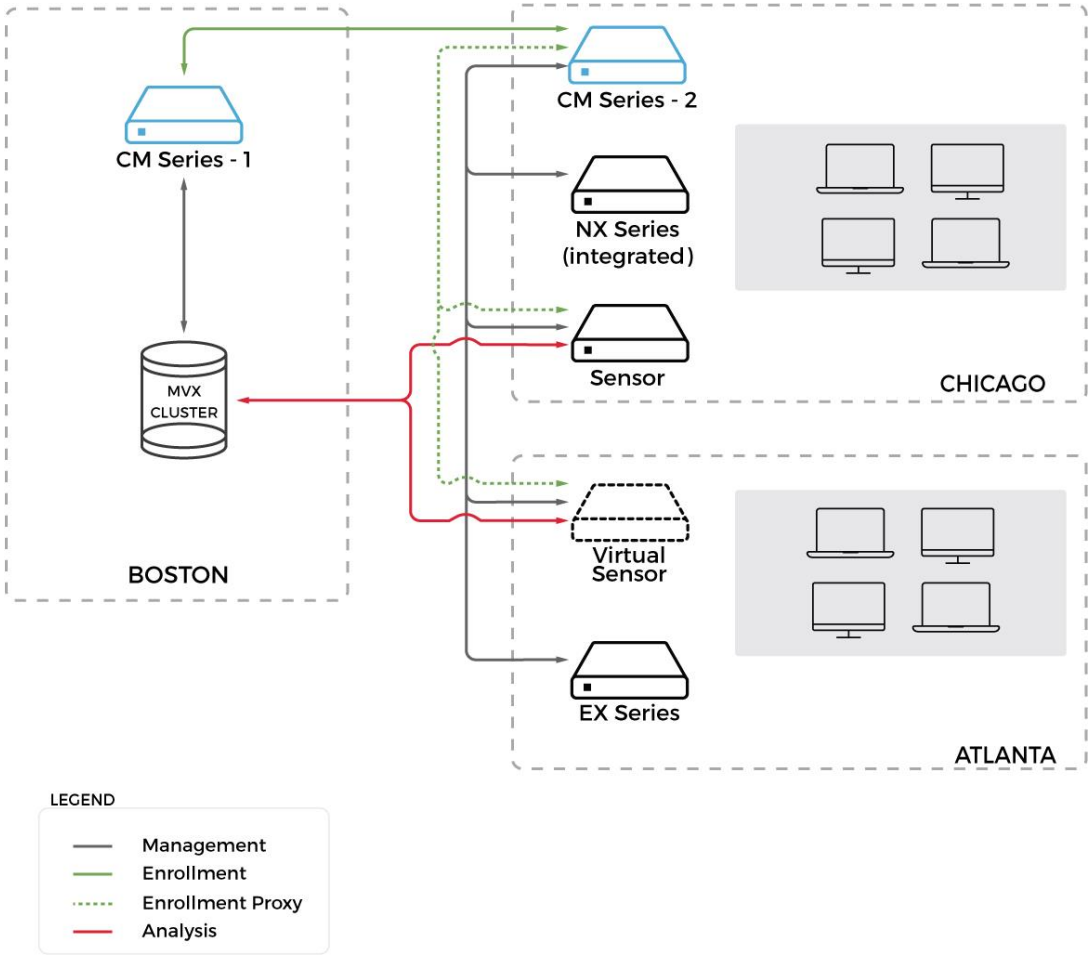

Standalone Sensors

For sensors that are not managed by a Central Management appliance, the enrollment service address type is DTI and its address is the IP address of the Central Management appliance that manages the MVX cluster. See Enrolling a Standalone Sensor on page 62 for configuration information.

Evidence Collector

In the following example, the Central Management appliance and two MVX clusters are deployed in the same data center. One cluster analyzes objects submitted by a physical Network Security sensor that is deployed to monitor network traffic in the Seattle region. The other cluster analyzes objects submitted by two virtual sensors in the Portland region.

The physical Network Security sensor has Evidence Collector enabled. Evidence Collector collects network event logs generated by the sensor, and sends them to Helix for analysis. You can use the data to correlate Network Security alerts with a detailed event log analysis. (For information about enabling Evidence Collector, see the *Network Security User Guide*.)

In a Helix deployment, the sensor submits objects to the MVX cluster on one path, and sends the events to Helix on a different path.

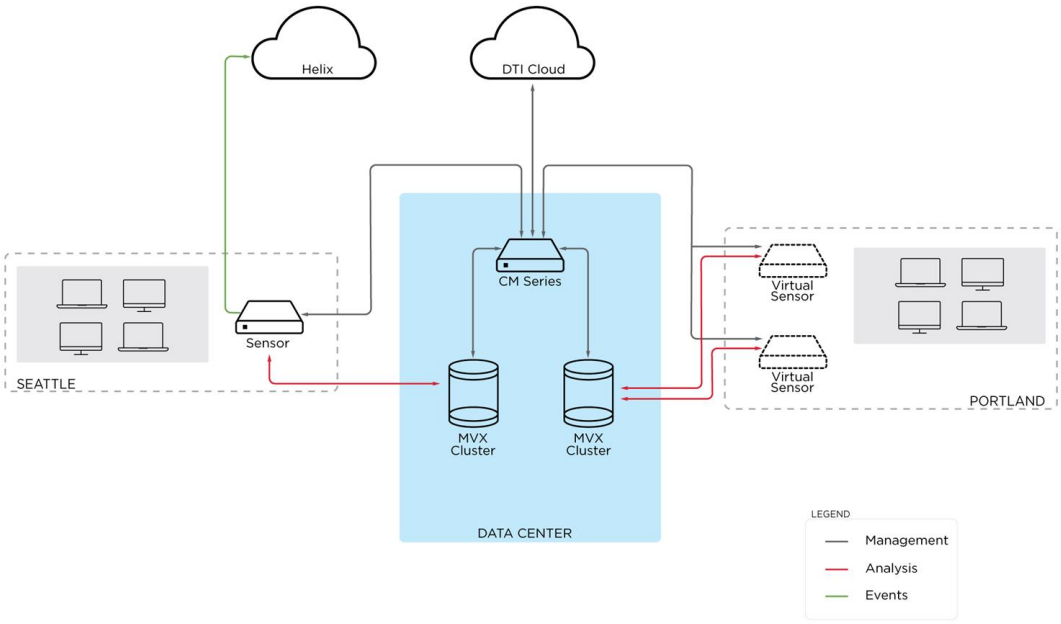

CHAPTER 6: MSSP Deployment Scenario

In a managed security service provider (MSSP) deployment, the MSSP deploys and maintains the MVX cluster in its premises (also referred to as a *private cloud*). A single physical or virtual Central Management appliance or Central Management High Availability (HA) pair manages the cluster. The MSSP also deploys and maintains the Central Management appliance. The MVX cluster and the Central Management appliance that manages the MVX cluster can be in different data centers.

MSSP customers configure their on-premises sensors to enroll with the cluster. The sensors can be standalone appliances or be managed by an on-premises Central Management appliance. The sensors can be in different locations.

Although the MVX cluster components (brokers and compute nodes) only need reliable IP connectivity, FireEye recommends that they be deployed on the same LAN.

IMPORTANT: FireEye does not recommend transcontinental deployments due to throughput, reliability, and latency issues.

NOTE: The Central Management appliance that manages sensors that submit to a cluster in a private cloud must be running a release that is compatible with the managed sensors.

A sensor can be a physical Network Security appliance or a virtual Network Security, Email Security — Server Edition, or File Protect appliance. The physical NX 1500 model can function only as a sensor, because it does not include an on-board MVX analysis engine. Some physical Network Security appliances can be enabled as sensors, in which case their on-board MVX analysis engines are disabled. A virtual appliance can function only as a sensor.

A broker or node is a physical Virtual Execution appliance. A Virtual Execution appliance serves no purpose until it is added to a cluster.

The following diagram illustrates the basic components of an MSSP deployment.

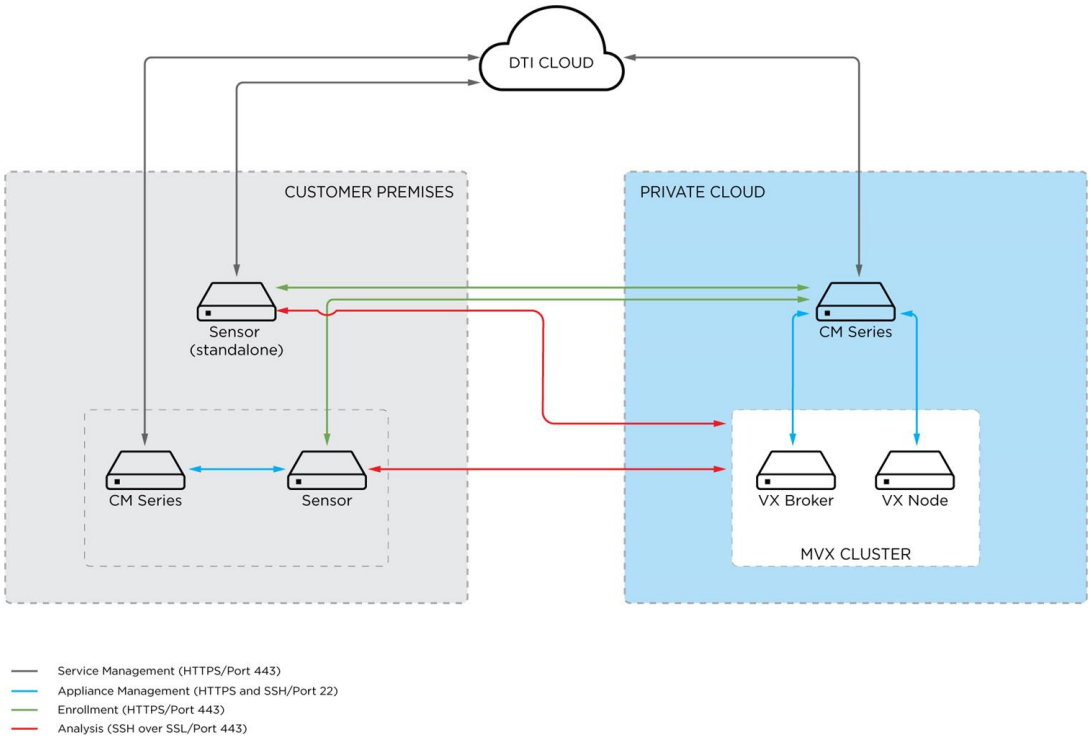

Communication Paths

The following diagram shows the communication paths between the MVX cluster components in an MSSP deployment.

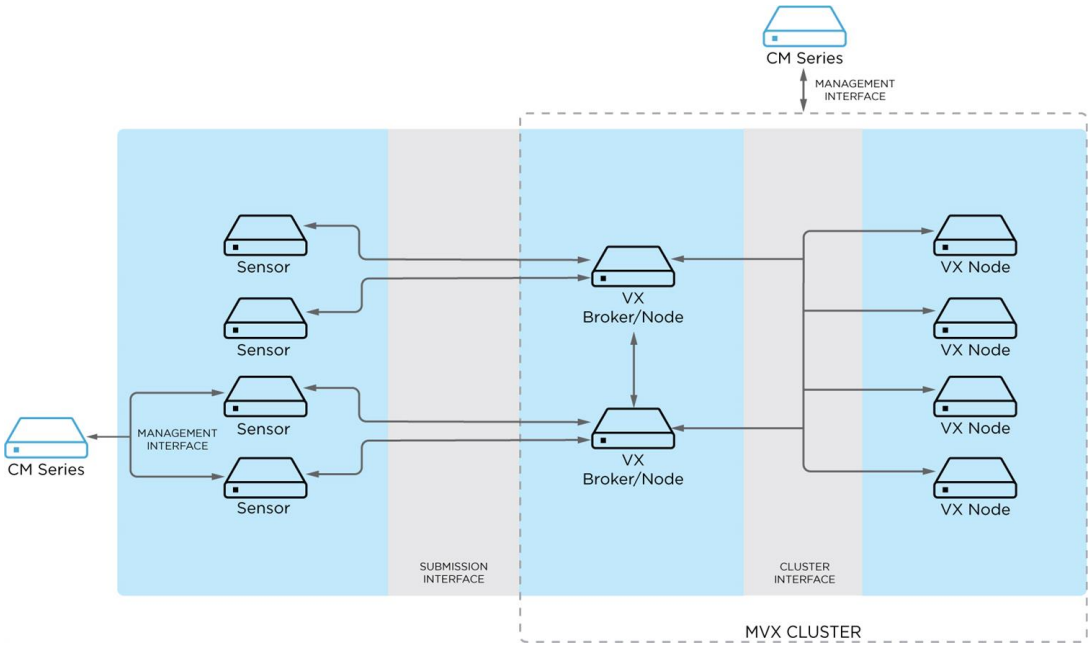

During the initial configuration, you can accept the default interfaces provided by the configuration wizard, or you can change them.

- **Management Interface**—In the preceding diagram, the connection between the Central Management appliance in the private cloud and the brokers and nodes is established through the management interface the same way as it is for other managed appliances. By default, both SSH and HTTPS traffic use the SSH port (port 22). The connection between the on-premises Central Management appliance and the sensor is established the same way.
- **Submission Interface**—Sensors and brokers communicate with each other through the submission interface using SSH. The default submission interface is ether1. A second interface (ether2) can also be configured.
- **Cluster Interface**—Brokers and nodes communicate with each other through the cluster interface using SSH and other protocols, secured through TLS (SSL). The default cluster interface is ether1, but another interface (for example, ether2) can be configured instead.

Brokers expose ports 25672 and 4369 for inter-broker communication. They also expose port 5671 for communication with compute nodes. Ports 25672 and 5671 are SSL-encrypted. Port 4369 is protected by key hash. Sensors and compute nodes connect to brokers using SSH (port 22). The cluster database uses TCP port 7001. Cluster management communication uses TCP and UDP ports 18300 through 18303.

Submission traffic from sensors to the broker in the MVX cluster in the private cloud must be exempted from any Web filtering product or man-in-the-middle SSL proxy you have deployed in your network. (The submission interface uses SSH over SSL (port 443); it does not carry HTTPS traffic.)

NOTE: For deployment steps, see MSSP Deployment Tasks on page 47.

PART II: Planning

CHAPTER 7: System Requirements

Before you deploy your FireEye MVX cluster and related components, make sure the following requirements are met.

NOTE: This guide does not provide information about expected throughput, performance, and capacity. See your FireEye representative for guidance with this.

Supported Sensor Models

Virtual Sensors

All virtual Network Security, Email Security — Server Edition, and File Protect models listed in the *FireEye Device Deployment Guide* are sensors.

Physical Network Security Sensors

NX 1500 is a sensor-only model with no on-board MVX analysis engine. The other physical NX models are integrated appliances with an on-board MVX analysis engine, but the engine is disabled after MVX sensor mode is enabled on the appliance.

- NX 900
- NX 1300/NX 1310
- NX 1400
- NX 1500
- NX 2300/NX 2310
- NX 2400
- NX 2500
- NX 2550
- NX 3500
- NX 4310/NX 4320
- NX 4400
- NX 4500
- NX 5500
- NX 6500
- NX 7300/NX 7320
- NX 7400
- NX 7420
- NX 7500
- NX 9450
- NX 10000
- NX 10450
- NX 10550

Supported Hybrid MVX Models

You can use the following physical appliance models as hybrid appliances.

- Network Security: All physical NX models
- Email Security — Server Edition: All physical EX models.
- File Protect: All physical FX models.

Supported Central Management Models

You can use the following CM models to manage the Virtual Execution appliances that form an MVX cluster.

Virtual Central Management Models

All virtual Central Management models listed in the *FireEye Device Deployment Guide* can manage an MVX cluster.

Physical Central Management Models

- CM 4400
- CM 4500
- CM 7400
- CM 7500
- CM 9400
- CM 9500

Supported Virtual Execution Models

You can use the following physical VX models in an MVX cluster.

- VX 5500 (physical)
- VX 12500 (physical)
- VX 12550 (physical)
- VX 12550V (virtual AWS model called the *FireEyeVX12550CloudEc2c5metal*)

NOTE: A cluster can contain either all physical Virtual Execution appliances or all virtual Virtual Execution appliances. A cluster cannot contain a combination of physical and virtual Virtual Execution appliances.

Virtual Machine Requirements

See the *FireEye Device Deployment Guide* for requirements and specifications.

Network Requirements

- One of the following:
  - Sensors, hybrid appliances (if any), brokers, and compute nodes connected to the same Central Management appliance.
  - Brokers and compute nodes connected to the same Central Management appliance; and sensors connected to another Central Management appliance, or standalone sensors not connected to a Central Management appliance. For this option, the enrollment service must be configured as described in Enrolling with an MVX Cluster on page 61.
- Connectivity with the DTI network (one-way or two-way sharing).
- Network access to the ports listed in the "Multi-Vector Execution (MVX) Platforms" section of the *Ports and Protocols Guide*.
- Reliable IP connectivity between brokers and compute nodes.

  NOTE: All brokers and compute nodes must be deployed on the same LAN.
- Communication between brokers not blocked by a firewall (to allow the dynamic port allocation described in Communication Paths on page 21.

IMPORTANT: Do not configure the submission and cluster interfaces on the same VLAN without guidance from FireEye Technical Support.

Management Path Requirements

See the *FireEye Device Deployment Guide* for management path requirements, including information about environments that restrict outbound access to certain IP addresses, and domain-based Proxy ACL rules.

Software Requirements

- Release 7.9.1 or later of the Network Security software image running on the Network Security sensors.
- Release 8.0.0 or later of the File Security software running on the File Security sensors.
- Release 8.0.0 or later running on the Email Security — Server Edition, File Security, and Network Security hybrid appliances.
- Release 8.1.4 or later running on virtual Email Security — Server Edition sensors.
- *Central Management appliance that manages a cluster*: Release 8.1.0 or later of the Central Management system image. If virtual Email Security — Server Edition sensors submit to the cluster, CM Series Release 8.3.0 is required.
- *Central Management appliance that manages sensors that submit to a cluster managed by another Central Management appliance*: A Central Management release that is compatible with the managed sensors.
- The same major and minor version (Release 8.0.0 or later) of the VX Series system image running on all nodes. If virtual Email Security — Server Edition sensors submit to the cluster, VX Series Release 8.2.0 is required.
- MVX sensor mode or MVX hybrid mode enabled on integrated Network Security appliances (see Enabling and Disabling MVX Sensor or Hybrid Mode on page 55).
- MVX hybrid mode enabled on integrated Email Security — Server Edition and File Security appliances (see Enabling and Disabling MVX Sensor or Hybrid Mode on page 55).
- The same guest images (profile and version) running on all nodes.
- The same security content version running on all nodes.
- The same configuration settings for relevant features on all nodes.

Licensing Requirements

The following table shows the licenses that must be installed on each sensor, compute node, and Central Management appliance. (Hybrid appliances have the same licensing requirements as integrated appliances.)

License	NX Sensor	FX Sensor	EX Sensor	Compute Node	CM Series
FIREEYE_APPLIANCE	✓	✓	✓	✓	✓
CONTENT_UPDATES	✓	✓	✓	✓	✓
FIREEYE_SUPPORT	✓	✓	✓	✓	✓
CLOUD_MVX	✓	✓	✓		
EMPS_ATTACHMENT_SCAN			✓		
EMPS_URL_SCAN			✓		
AV_ENGINE_SOPHOS	✓			✓	
IPS	✓				
ATI	✓		✓		
MD_ACCESS	✓	✓	✓		

NOTE: The AV_ENGINE_SOPHOS, IPS, ATI, and MD_ACCESS licenses are required if the associated optional features are enabled on the appliance.

Limitations

Note the following deployment limitations in this release.

- To maintain high availability and data consistency, FireEye recommends that you create clusters with three broker nodes.
- Clusters with two broker nodes are not recommended.
- DHCP is not supported on the submission or cluster interface.
- IPv6 is not supported on the management interface, cluster interface, or submission interface.

- On an Email Security — Server Edition sensor, ether2 cannot be used as both the submission interface and the URL Dynamic Analysis interface.
- A node must be removed from the cluster before you change its cluster interface or hostname.
- In the **Create New Cluster** or **Edit Cluster** wizard in the Central Management Web UI, it is recommended that you perform the "move node to cluster" and "remove node from cluster" actions on no more than three nodes at a time.
- Do not upgrade an individual node if it is currently part of a cluster. Use the upgrade orchestration procedure described in Upgrades on page 147 to ensure that at least one broker and one compute node are running during the upgrade.
- Transcontinental deployments are not recommended due to throughput, reliability, and latency issues.
- Physical or virtual Network Security sensor limitations:
  - Jumbo frames are not supported.
  - Network Security High Availability is not supported.
  - The Essentials edition is not supported.
- Virtual sensor limitations: See the *FireEye Device Deployment Guide*.

Best Practices

The following best practices can help you achieve a successful FireEye MVX Smart Grid deployment.

- Add sensors one at a time, and monitor the cluster utilization. This allows you to determine whether additional nodes are needed.
- Initially deploy sensors in an out-of-band mode to get a benchmark on performance. If the performance is good, then change to an inline deployment mode.

CHAPTER 8: Standard Deployment Tasks

You must perform the following tasks to deploy FireEye Distributed Network Security.

1. Plan below
2. Gather Items from Your Network Administrator on the next page
3. Gather Information from FireEye on the next page
4. Deploy Virtual Appliances on page 43
5. Deploy Physical Appliances on page 43
6. Configure the Enrollment Service on page 43
7. Add Appliances to the Central Management Appliance on page 44
8. Complete the Configuration on page 44
9. Create the Cluster on page 45
10. Check the Cluster Health and Performance on page 45

Plan

Perform the following steps before you begin the deployment.

1. Decide where you want to deploy sensors, hybrid appliances, and clusters in your network.
2. Decide the deployment and operational mode for each Network Security sensor and each Network Security integrated appliance that will operate in sensor or hybrid mode.

   Sensors and hybrid appliances can be deployed in the same modes as integrated appliances. See the *Network Security System Administration Guide* and *Network Security User Guide* for details about the modes.

3. Decide the deployment and operational mode for each Email Security — Server Edition sensor and integrated Email Security — Server Edition appliance that will operate in hybrid mode.

   Sensors and hybrid appliances can be deployed in the same modes as integrated appliances. See the *Email Security — Server Edition User Guide* for information about the modes.

4. Decide how File Protect sensors and File Protect integrated appliances that will operate in hybrid mode will access network shares.

   Sensors, hybrid appliances, and integrated appliances can access network shares in the same way. See the *File Protect User Guide* for information about network share access.

Gather Items from Your Network Administrator

Have the following items ready before you begin the deployment.

- Static or reserved IP address, subnet mask, and default gateway address for the management interface.
- IP address for each Domain Name System (DNS) server.
- IP address for each Network Time Protocol (NTP) server.
- Telnet or SSH client on the remote system (if the component will be managed remotely).
- *Physical appliances*: If you plan to configure initial settings using the serial console port and a Windows or Mac laptop, obtain a USB-to-serial cable.

Gather Information from FireEye

Have the following items ready before you begin the deployment.

- License keys (if the license update service is not enabled).
- *Virtual appliances*:
    - Activation code, which gives the virtual appliance a unique identity (its appliance ID), activates the product (FIREEYE_APPLIANCE) license, allows access to the license token server, provides access to the DTI network, protects against fraudulent use of the appliance, and allows the appliance to initialize.
    - Link to an OVA file containing your customer-specific system image.

Deploy Virtual Appliances

See the *FireEye Device Deployment Guide* for information about deploying virtual appliances.

Deploy Physical Appliances

Perform the following steps to deploy physical appliances in your network.

1. Install the appliances in your network.

   See the *Hardware Administration Guide* for the appliance and the *FireEye Device Deployment Guide*.
2. Enable sensor mode on integrated Network Security appliances, as described in Enabling and Disabling MVX Sensor or Hybrid Mode on page 55.
3. Enable hybrid MVX mode on integrated Network Security, Email Security — Server Edition, and File Protect appliances, as described in Enabling and Disabling MVX Sensor or Hybrid Mode on page 55.

Configure the Enrollment Service

Enrollment is automatic for sensors and hybrid appliances that are managed by the same Central Management appliance that manages the MVX cluster. Perform the procedure below based on your enrollment scenario.

- If you want managed sensors to enroll directly with an MVX cluster on another Central Management appliance, configure the enrollment service on the sensors to point to the other Central Management appliance. See Enrolling a Managed Sensor Directly on page 64.

- If you want sensors to enroll through the Central Management appliance that manages the sensors, configure the enrollment service on that Central Management appliance to point to the other Central Management appliance. See Enrolling a Managed Sensor Through a Proxy on page 66.
- If your sensors are standalone appliances, configure the enrollment service to point to the Central Management appliance that manages the MVX cluster. See Enrolling a Standalone Sensor on page 62.

Add Appliances to the Central Management Appliance

1. Add the Virtual Execution appliances, as described in Adding Nodes to a Central Management Appliance on page 73.
2. *If the sensors will be managed*: Add the appliances, as described in Adding Sensors and Hybrid Appliances to a Central Management Appliance on page 76.
3. *If you are deploying hybrid appliances*: Add the appliances, as described in Adding Sensors and Hybrid Appliances to a Central Management Appliance on page 76.

Complete the Configuration

Perform the following steps to complete the configuration.

1. *If the license update feature is disabled:* Install FIREEYE_SUPPORT and feature licenses.

   The license update feature enables your appliance to automatically download and apply licenses to which you are contractually entitled. This feature is enabled with the configuration wizard during the initial configuration, and is fully functional after the configuration wizard is completed.

   See the *Administration Guide* or *System Administration Guide* for the appliance.
2. Configure other system administration features such as AAA, SSL certificates, SNMP, email notification recipients, and so on.

   See the *Administration Guide* or *System Administration Guide* for the appliance.
3. Configure detection settings, such as operational mode, policies, notifications, reports, and so on. See the *Network Security User Guide* and *Email Security — Server Edition User Guide*.
4. Add storage and configure detection settings, such as scans, notifications, reports, and so on. See the *File Protect User Guide*.

5. (*Optional*) Define and upload custom YARA rules. See the *Virtual Execution Administration Guide*.

Create the Cluster

Create the cluster, as described in Creating a Cluster on page 51.

Check the Cluster Health and Performance

Check that the cluster and its components are healthy and that the cluster utilization and performance are acceptable, as described in Viewing Cluster and Node Status on page 103 and Viewing Cluster Utilization on page 133.

CHAPTER 9: MSSP Deployment Tasks

This section describes the steps a Managed Service Security Provider (MSSP) must perform to deploy an MVX cluster in the MSSP premises and to allow sensors in MSSP customer premises to enroll with the cluster.

NOTE: This topic assumes that the appliances in the MSSP premises and in the customer premises are already configured and deployed.

Basic MSSP Tasks

The basic MSSP premises tasks are described in the following table.

Task	Instructions
Add the Virtual Execution appliances to the Central Management appliance in the private cloud.	1. Select **Appliances** > **Nodes** in the Central Management Web UI. 2. Click **Create New Node** and complete the fields in the dialog box. 3. Click **Add**. For details, see Adding Nodes to a Central Management Appliance on page 73.
Create the MVX cluster.	1. Select **Appliances** > **Clusters** in the Central Management Web UI. 2. Click **Create New Cluster** and follow the instructions in the wizard. For details, see Creating a Cluster on page 51.

Task	Instructions
Configure the Central Management appliance in the private cloud as the enrollment server.	Run the `fenet dti enrollment service default LOCAL` command in the private cloud Central Management CLI.

Basic Sensor Tasks

The basic sensor tasks are described in the following table.

Task	Instructions
Install a CLOUD_MVX license.	1. Select **Settings** > **Appliance Licenses** in the sensor Web UI. 2. Click **Add License**, enter the license key, and then click **Add**. For details, see the *System Administration Guide* for the sensor.
Configure the enrollment service.	**Standalone Sensor** Run the following command in the sensor CLI: `fenet dti enrollment service type DTI address <cluster CM address>` where `cluster CM address` is the IP address of the Central Management appliance that manages the MVX cluster.

	Managed Sensor
	Run the following three commands in the sensor CLI:
	`no fenet dti enrollment service override enable`
	`fenet dti enrollment service type DTI address <cluster CM address>`
	`fenet dti enfollment service default DTI`
	where `cluster CM address` is the IP address of the Central Management appliance that manages the MVX cluster.

PART III: Deployment

CHAPTER 10: Creating a Cluster

The following sections describe how to create an MVX cluster.

Prerequisites

- Operator or Admin access.
- Requirements listed in System Requirements on page 35 are met.
- Managing Central Management appliance and Virtual Execution appliances are fully configured.
- The Virtual Execution appliances in a cluster must be all physical appliances or all virtual appliances. A combination of physical and virtual Virtual Execution appliances in the same cluster is not supported.
- If the MVX cluster is in an internal network behind a NAT gateway and your sensors are in an external network: an accessible IP address for each broker (see Configuring an Accessible Broker Address on page 91).

Creating a Cluster Using the Web UI

This section describes how to create a cluster using the Central Management Web UI.

To create a cluster:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click **Add Cluster**. The **Create New Cluster** box opens.
4. Enter a unique name for the cluster in the **Cluster Name** box and select a node to add to the cluster. You can add only one node at a time.
5. Click **Create**. A message is displayed while the cluster is being created.

   After the cluster is created, its information is displayed on the **Clusters** tab.

Creating a Cluster Using the CLI

Use the commands in this section to create an MVX cluster.

Create the cluster framework:

1. Log in to the Central Management CLI.
2. Go to CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Create the cluster:

   ```
   cm-hostname (config) # cmc mvx cluster <cluster name>
   ```

   where `<cluster name>` is a unique name identifying the cluster.

Add nodes to the cluster:

1. Add a node to the cluster:

   ```
   cm-hostname (config) # cmc mvx cluster <cluster name> node <nodeName>
   ```

2. Repeat the previous step for each node you want to add.

Enable brokers:

1. Enable broker mode on a node:

   ```
   cm-hostname (config) # cmc mvx cluster <cluster name> broker <node
   name> enable
   ```

2. Repeat the previous step for each additional broker (if any).

Verify and save your changes:

1. Verify your changes:

   ```
   cm-hostname (config) # show cmc mvx cluster detail
   ```

2. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

Example

The following example creates a cluster named Cluster-01 on the cm-1 Central Management appliance, adds vx-1 and vx-2 as cluster nodes, enables broker mode on vx-1, and designates vx-1 as the master configuration.

```
cm-1 (config) # cmc mvx cluster Cluster-01
cm-1 (config) # cmc mvx cluster Cluster-01 node vx-1
cm-1 (config) # cmc mvx cluster Cluster-01 node vx-2
cm-1 (config) # cmc mvx cluster Cluster-01 broker vx-1 enable
cm-1 (config) # cmc mvx cluster Cluster-01 master vx-1
cm-1 (config) # show cmc mvx cluster detail
```

```
MVX Cluster: Cluster-01

Version             : 8.3.0
Utilization         : 0 %
Status              : ready
Total Nodes         : 2

Member Status:
       Brokers:
               vx-1                    : 10.11.121.12 - ready

       Compute Nodes:
               vx-2                    : 10.11.121.18 - ready
```

CHAPTER 11: Enabling and Disabling MVX Sensor or Hybrid Mode

You can enable MVX sensor or hybrid mode on an integrated appliance using either the managing Central Management Web UI or the appliance CLI.

NOTE: For information about restoring local mode on a hybrid-enabled integrated appliance, see Restoring Local Mode on page 58.

Prerequisites

- Admin access

Enabling MVX Sensor Mode

The following sections describe how to enable MVX sensor mode.

Enabling MVX Sensor Mode Using the Web UI

Use the **Sensors** page to enable MVX sensor mode (also known as "cluster" mode) on a managed integrated appliance. The appliance must reload for the mode change to take effect.

When sensor mode is enabled, the **Cluster Enrollment** column shows the cluster and broker node to which the appliance submits objects.

To enable MVX sensor mode:

1. Log into the Web UI of the Central Management appliance that manages the appliance.
2. Select **Appliances** > **Sensors**.

3. Locate the integrated appliance.
4. In the **Action** column, click **Select** > **Enable Cluster Mode**.
5. When prompted, click **OK** to allow the system to reload.

Enabling MVX Sensor Mode Using the CLI

Use the commands in this section to enable MVX sensor mode on an integrated appliance.

To enable MVX sensor mode:

1. Log in to the integrated appliance CLI.
2. Go to CLI configuration mode:

```
hostname > enable
hostname # configure terminal
```

3. Enable sensor mode:

```
hostname (config) # mvx mode sensor
```

4. Reload the appliance:

```
hostname (config) # reload
```

5. Press Enter when prompted to save changes.

Example

The following example enables nx-3 to function as a sensor.

```
nx-3 (config) # mvx mode sensor
MVX sensor mode configuration successful. Please reload the appliance to come
up in the configured mode
nx-3 (config) # reload
Configuration has been modified; save first? [yes]

nx-3 # show mvx status
MVX Mode Status:
    Sensor Config Enabled:   yes
    Mode Reboot Required:    no
    Current Operating Mode:  sensor
```

Enabling Hybrid Mode

You can enable MVX hybrid mode on an integrated appliance using the Central Management Web UI or the integrated appliance CLI.

NOTE: For information about restoring local mode on a hybrid-enabled integrated appliance, see Restoring Local Mode on page 58.

Prerequisites

- Admin access

Enabling Hybrid Mode Using the Web UI

Use the **Appliances** > **Sensors** page of the Central Management Web UI to enable MVX hybrid mode on an integrated appliance.

NOTE: You cannot use the Central Management Web UI to enable hybrid mode on an Network Security or File Protect appliance. Instead, use the procedure in Enabling Hybrid Mode Using the CLI below.

To enable hybrid mode:

1. Log into the Central Management Web UI.
2. Select **Appliances** > **Sensors** tab.
3. Locate the appliance row.
4. Click **Select** in the **Action** column.
5. Click **Enable Hybrid Mode**.
6. Click **OK** to confirm the action.

Enabling Hybrid Mode Using the CLI

Use the commands in this section to enable an integrated appliance to function in MVX hybrid mode.

To enable hybrid mode:

1. Go to CLI configuration mode:

   ```
   hostname > enable
   hostname # configure terminal
   ```

2. Enable hybrid mode:

   ```
   hostname (config) # mvx mode hybrid
   ```

3. Verify your change:

   ```
   hostname (config) # show mvx status
   ```

4. Save your change:

   ```
   hostname (config) # write memory
   ```

Example

The following example enables hybrid mode on FX-04.

```
FX-04 (config) # mvx mode hybrid
MVX hybrid mode configuration successful.
Submission interface is ether1. To change it, use:
'mvx sensor config submission-if' or 'configuration jump-start'.
FX-04 (config) # show mvx status
MVX Mode Status:
    Current Operating Mode:  hybrid
    ...
```

Restoring Local Mode

You can restore MVX local mode on an appliance using either the managing Central Management Web UI or the integrated appliance CLI. After integrated mode is restored, the appliance uses its on-board MVX analysis engine for all submissions.

Prerequisites

- Admin access

Restoring MVX Local Mode Using the Web UI

Use the **Appliances** > **Sensors** page to restore MVX local mode on an appliance.

To restore local mode on a sensor-enabled appliance:

1. Log into the Central Management Web UI.
2. Select **Appliances** > **Sensors**.
3. In the **Action** column, click **Select** and then click **Disable Cluster Mode**.
4. When prompted, click **OK** to confirm the mode change.

To restore local mode on a hybrid-enabled appliance:

1. Log into the Central Management Web UI.
2. Select **Appliances** > **Sensors**.
3. In the **Action** column, click **Select**, and then click **Disable Hybrid Mode**.
4. When prompted, click **OK** to confirm the mode change.

After you restore local mode, the appliance submits all objects to its on-board MVX analysis engine. The value of the **Cluster Enrollment** column becomes **Local**.

NOTE: You cannot currently use the Central Management Web UI to disable hybrid mode on a managed Network Security or File Protect appliance. Use the procedure in Restoring Local Mode Using the CLI on the facing page instead.

Restoring Local Mode Using the CLI

Use the commands in this section to restore MVX local mode on an appliance.

To restore local mode:

1. Log in to the integrated appliance CLI.
2. Go to CLI configuration mode:

   ```
   hostname > enable
   hostname # configure terminal
   ```

3. Enable local mode:

   ```
   hostname (config) # mvx mode local
   ```

4. Verify your change:

   ```
   hostname (config) # show mvx status
   ```

Example

The following example restores local mode on the nx-3 appliance.

```
nx-3 (config) # mvx mode local

nx-3 (config) # show mvx status
MVX Mode Status:
    Current Operating Mode:  integrated
    ...
```

CHAPTER 12: Enrolling with an MVX Cluster

No manual configuration is required if the same Central Management appliance manages both the sensors and the MVX cluster. This section describes how to enroll sensors in scenarios that require manual configuration. In these scenarios, either the MVX cluster is not managed by the Central Management appliance that manages the sensors, or the sensors are standalone appliances. For example:

- The MVX cluster is managed by another Central Management appliance. The sensor can either enroll directly, or enroll through its managing Central Management appliance. In the second case, the Central Management appliance acts as a proxy for the enrollment. (The proxy scenario is convenient if the Central Management appliance manages multiple sensors, because the configuration can be done once on the Central Management appliance instead of being done on each sensor.)
- The sensor is a standalone Network Security appliance (it is not managed by a Central Management appliance).

No manual configuration is required to enroll hybrid appliances, because they must be managed by the same Central Management appliance that manages the MVX cluster.

IMPORTANT! Information about standalone sensors is not displayed in the Web UI of the Central Management appliance that manages the cluster.

NOTE: For examples of enrollment scenarios, see Standard Deployment Scenarios on page 23.

Prerequisites

- Operator or Admin access
- Requirements in System Requirements on page 35

Enrolling a Standalone Sensor

Use the commands in this section to enroll a standalone sensor with an MVX cluster.

To enroll a standalone sensor:

1. Log in to the sensor CLI.
2. Go to CLI configuration mode:

   ```
   sensor-hostname > enable
   sensor-hostname # configure terminal
   ```

3. Configure the enrollment service address (the IP address of the Central Management appliance that manages the cluster):

   ```
   sensor-hostname (config) # fenet dti enrollment service type DTI
   address <cluster CM address>
   ```

   where `<cluster CM address>` is the IP address of the Central Management appliance that manages the MVX cluster.

4. Verify your changes:

   ```
   sensor-hostname (config) # show fenet dti configuration
   ```

5. Verify the enrollment service settings:

   ```
   sensor-hostname (config) # show mvx cluster enrollment service
   ```

6. Save your changes:

   ```
   sensor-hostname (config) # write memory
   ```

Example

In this example, the standalone nx-5 sensor is enrolled with a cluster that is managed by a Central Management appliance with an IP address of 10.11.10.11.

```
nx-5 (config) # fenet dti enrollment service type DTI address 10.11.10.11

nx-5 (config) # show fenet dti configuration
DTI CLIENT CONFIGURATIONS:

  ACTIVE SETTINGS:

   Mode               : online
   Download source    : DTI (User8@cloud.fireeye.com)
   Upload destination : DTI (User8@up-cloud.fireeye.com)
   Mil service        : DTI (User8@up-cloud.fireeye.com)
   Enrollment service : DTI (User8@10.11.10.11)
 ...

nx-5 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status
```

```
 Enrollment Client :
  Status ok             : yes
  Status description    : enrolled
  Last checked at       : 2019/08/14 14:56:01

Enrollment Service :
 Auto enabled           : yes
 Service address        : DTI (10.11.10.11)
 Preferred cluster      : any (less loaded)
 Cloud enabled          : no
 Cloud License enabled  : no
 Connect on demand      : no

Broker Info :
 Cluster Name           : Cluster-02
 Broker Name            : vx-4
 Broker Address         : 10.11.12.13
 Broker ID              : 002XXXXXXXXX
 Broker State           : Connected
 Failure Reason         : None
 Last Connection Attempt: 2019/08/12 15:01:11
 Connection Last Formed : 2019/08/12 15:01:12
 Connection Last Broken :
```

Enrolling a Managed Sensor Directly

Use the commands in this section to enroll a managed sensor with an MVX cluster managed by another Central Management appliance.

To enroll a managed sensor directly:

1. Log in to the sensor CLI.
2. Go to CLI configuration mode:

   ```
   sensor-hostname > enable
   sensor-hostname # configure terminal
   ```

3. Prevent the local Central Management appliance (the one that manages the sensor) from overriding the settings you are configuring in this procedure:

   ```
   sensor-hostname (config) # no fenet dti enrollment service override
   enable
   ```

4. Change the enrollment service type to `DTI`, and configure the enrollment service address (the IP address of the Central Management appliance that manages the cluster):

   ```
   sensor-hostname (config) # fenet dti enrollment service type DTI
   address <cluster CM address>
   ```

   where `<cluster CM address>` is the IP address of the Central Management appliance that manages the MVX cluster.

5. Set `DTI` as the default enrollment service type:

   ```
   sensor-hostname (config) # fenet dti enrollment service default DTI
   ```

6. Verify your changes:

   ```
   sensor-hostname (config) # show fenet dti configuration
   ```

7. Verify the enrollment service status:

   ```
   sensor-hostname (config) # show mvx cluster enrollment status
   ```

8. Save your changes:

   ```
   sensor-hostname (config) # write memory
   ```

Example

In this example, the nx-1 sensor that is managed by a Central Management appliance with an IP address of 172.1.2.3 will be enrolled with a cluster that is managed by another Central Management appliance with an IP address of 10.11.10.11.

```
nx-1 (config) # no fenet dti enrollment service override enable
nx-1 (config) # fenet dti enrollment service type DTI address 10.11.10.11
```

```
nx-1 (config) # fenet dti enrollment service default DTI
nx-1 (config) # show fenet dti configuration
DTI CLIENT CONFIGURATIONS:

  ACTIVE SETTINGS:

   Mode               : online
   Download source    : CMS (User8@172.1.2.3 : singleport) - Managed by CMS
   Upload destination : CMS (User8@172.1.2.3 : singleport) - Managed by CMS
   Mil service        : CMS (User8@172.1.2.3 : singleport) - Managed by CMS
   Enrollment service : DTI (User8@10.11.10.11) - Managed by Appliance
 ...

nx-1 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client :
   Status ok             : yes
   Status description    : enrolled
   Last checked at       : 2019/08/14 14:56:01

 Enrollment Service :
  Auto enabled           : yes
  Service address        : DTI (10.11.10.11)
  Preferred cluster      : any (less loaded)
  Cloud enabled          : no
  Cloud License enabled  : no
  Connect on demand      : no

 Broker Info :
  Cluster Name           : Cluster-02
  Broker Name            : vx-4
  Broker Address         : 10.11.12.13
  Broker ID              : 002XXXXXXXXX
  Broker State           : Connected
  Failure Reason         : None
  Last Connection Attempt: 2019/08/12 15:01:11
  Connection Last Formed : 2019/08/12 15:01:12
  Connection Last Broken :
```

Enrolling a Managed Sensor Through a Proxy

Use the commands in this section to enroll a managed sensor with an MVX cluster managed by another Central Management appliance. In this scenario, the local Central Management appliance that manages the sensor will act as a proxy for the enrollment.

This procedure changes the enrollment service address for the local Central Management appliance to the IP address of the Central Management appliance that manages the MVX cluster. The enrollment service address for the sensor will continue to be the IP address of the local Central Management appliance.

To enroll a managed sensor through a proxy:

1. Log in to the local Central Management CLI.
2. Go to CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Change the enrollment service type to `DTI`, and configure the enrollment service address (the IP address of the Central Management appliance that manages the cluster):

   ```
   cm-hostname (config) # fenet dti enrollment service type DTI address
   <cluster CM address>
   ```

   where `<cluster CM address>` is the IP address of the Central Management appliance that manages the MVX cluster.

4. Set `DTI` as the default enrollment service type:

   ```
   cm-hostname (config) # fenet dti enrollment service default DTI
   ```

5. Verify your changes:

   ```
   cm-hostname (config) # show fenet dti configuration
   ```

6. Save your changes.

   ```
   cm-hostname (config) # write memory
   ```

7. Log in to the sensor CLI.
8. Verify the enrollment service settings:

   ```
   sensor-hostname (config) # show mvx cluster enrollment service
   ```

Example

In this example, the nx-1 sensor is managed by the local Central Management appliance (cm-1) with an IP address of 172.1.2.3. The local Central Management appliance will act as a proxy to enroll the sensor with Cluster-02, which is managed by another Central Management appliance (cm-2) with an IP address of 10.11.10.11.

```
cm-1 (config) # fenet dti enrollment service type DTI address 10.11.10.11
cm-1 (config) # fenet dti enrollment service default DTI
cm-1 (config) # show fenet dti configuration
DTI CLIENT CONFIGURATIONS:

  ACTIVE SETTINGS:

   Mode               : online
   Download source    : DTI (User8@cloud.fireeye.com)
   Upload destination : DTI (User8@up-cloud.fireeye.com)
   Mil service        : DTI (User8@mil-cloud.fireeye.com)
   Enrollment service : DTI (User8@10.11.10.11)
 ...

nx-1 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client :
   Status ok             : yes
   Status description    : enrolled
   Last checked at       : 2019/08/14 14:56:01

 Enrollment Service :
  Auto enabled           : yes
  Service address        : CMS (172.1.2.3)  Preferred cluster      : any
(less loaded)
  Cloud enabled          : no
  Cloud License enabled  : no
  Connect on demand      : no

 Broker Info :
  Cluster Name           : Cluster-02
  Broker Name            : vx-4
  Broker ID              : 002XXXXXXXXX
  Broker Address         : 10.11.12.13
  Broker State           : Connected
  Failure Reason         : None
  Last Connection Attempt: 2019/08/12 15:01:11
  Connection Last Formed : 2019/08/12 15:01:12
  Connection Last Broken :
```

Enrolling with a Preferred Cluster

You can specify a preferred cluster in a deployment with multiple clusters. When you do so, the sensor or hybrid appliance submits to this cluster instead of the one it was automatically enrolled with when it was added to the Central Management appliance that also manages the cluster.

Reasons to specify a preferred cluster include:

- **Geographic proximity**—Sensors submit to clusters that are physically close to them to reduce network latency.
- **Distributed guest-images profiles**—Sensors monitoring Windows devices submit to clusters running Windows guest images, and sensors monitoring Mac devices submit to clusters running OS X guest images.

To configure a preferred cluster:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI configuration mode:

   ```
   appl-hostname > enable
   appl-hostname # configure terminal
   ```

3. Unsubscribe the sensor or hybrid appliance from the cluster with which it is currently subscribed (if any):

   ```
   appl-hostname (config) # mvx cluster unenroll now
   ```

4. Configure the preferred cluster.

   ```
   appl-hostname (config) # mvx cluster enrollment-service preferred name
   <clusterName>
   ```

   where `<clusterName>` is the name of the cluster to enroll the sensor with.

5. Verify your change:

   ```
   appl-hostname (config) # show mvx cluster enrollment status
   ```

6. Save your change:

   ```
   appl-hostname (config) # write memory
   ```

To remove the preferred cluster enrollment:

1. Remove the preferred cluster enrollment:

   ```
   appl-hostname (config) # no mvx cluster enrollment-service preferred
   name
   ```

2. Verify your change:

   ```
   appl-hostname (config) # show mvx cluster enrollment status
   ```

3. Save your change:

```
appl-hostname (config) # write memory
```

NOTE: The sensor or hybrid appliance will be automatically enrolled with the cluster with the most capacity after you remove the preferred cluster enrollment.

Example

The following example unsubscribes the nx-1 sensor from the cluster it is currently enrolled with, and then configures Cluster02 as its preferred cluster.

```
nx-1 (config) # mvx cluster unenroll now
nx-1 (config) # mvx cluster enrollment-service preferred name Cluster02
nx-1 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

Enrollment Client:
Status OK                : yes
STatus description       : enrolled
Last checked at          : 2019/08/19 0023:02

Enrollment Service :
  Auto enabled           : yes
  Service Address        : CMS (DTIuser@10.11.121.13 : singleport)
  Preferred cluster      : Cluster02
...
```

Enrolling a Sensor Using the Central Management CLI

Use the commands in this section to enroll a sensor using the Central Management CLI.

NOTE: You can also use the `cmc execute` command from the Central Management CLI to perform a sensor operation on a named group of sensors. In this scenario, the Central Management appliance executes the command individually on each member of the group. See Enrolling a Named Group of Sensors on the facing page for an example. For details about using the `cmc execute` command, see the *Central Management Administration Guide*.

To enroll a sensor:

1. Log in to the CM Series CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Enroll the sensor:

   ```
   cm-hostname (config) # cmc mvx sensor enrollment enroll <sensor name>
   ```

   where `<sensor name>` is the name of the sensor (the name shown in the `show cmc appliances` command output).

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx cluster enrollment status
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

To unsubscribe a sensor:

1. Log in to the CM Series CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Unsubscribe the sensor:

   ```
   cm-hostname (config) # cmc mvx sensor enrollment unenroll <sensorName>
   ```

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx cluster enrollment status
   ```


5. Save your change:

```
cm-hostname (config) # write memory
```

Examples

Enrolling a Single Sensor

The following example enrolls the nx-6 sensor with Cluster-Acme. In this example, the sensors are connected to the vx-1 broker.

```
cm-5 (config) # cmc mvx sensor enrollment enroll nx-6
Sensor has been successfully enrolled
cm-5 (config) # show cmc mvx cluster enrollment status
SENSOR NAME        CLUSTER NAME     BROKER NAME     BROKER ADDRESS
===========        ============     ===========     ==============
nx-1               Cluster-Acme     vx-1            10.11.121.12
nx-2               Cluster-Acme     vx-1            10.11.121.12
nx-6               Cluster-Acme     vx-1            10.11.121.12
```

Enrolling a Named Group of Sensors

The following example enrolls each sensor in the Tokyo group with the cluster that has the least load at the time the individual command is executed.

```
cm-4 (config) # cmc execute group Tokyo command "mvx cluster enroll now"
============ Appliance nx-6 =============
Execution was successful.
Execution output:
Operation initiated in the background.
  Run 'show mvx cluster enrollment status'

============ Appliance nx-7 =============
Execution was successful.
Execution output:
Operation initiated in the background.
  Run 'show mvx cluster enrollment status'
cm-4 (config) # show cmc mvx cluster enrollment status
```

NOTE: Use the Central Management command (`show cmc mvx cluster enrollment status`) to view the cluster enrollment status, not the Network Security command (`show mvx cluster enrollment status`) that is shown in the command output.

CHAPTER 13: Adding Nodes and Sensors to a Central Management Appliance

You can use the Central Management Web UI or the CLI to add nodes and sensors to a Central Management appliance.

Prerequisites

- Operator or Admin access

Adding Nodes to a Central Management Appliance

You can use the Central Management Web UI or the CLI to add nodes to a Central Management appliance.

Prerequisites

- Operator or Admin access

Adding a Node to the Central Management Appliance Using the Web UI

Use the **Add New Node** dialog box to add Virtual Execution appliances that will function as nodes.

Field	Description
Node Name	The name of the Virtual Execution appliance that will be a node in the MVX cluster. (The name will be displayed in the `show cmc appliances` command output.)
IP Address	The IPv4 or IPv6 address of the management interface on the appliance.
Username	A user account on the appliance with the admin role.
Password	The password for the user account specified in the **Username** field.
Host Key	(*Required if host-key authentication is enabled*) The RSA v2 host key for the VX Series appliance. For details, see the *FireEye System Security Guide*.
Comments	(*Optional*) A comment that describes the node.

To add a node:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Nodes**.
3. Click **Add Node**. The **Add New Node** dialog box opens.
4. Enter values for the fields described in the preceding table.
5. Click **Add**.

Adding a Node to the Central Management Appliance Using the CLI

Use the commands in this section to add Virtual Execution appliances that will function as nodes.

To add a node:

1. Log in to the Central Management CLI.
2. Go to CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. Add the first Virtual Execution appliance:

   ```
   cm-hostname (config) # cmc appliance <appl ID> address <IP address>
   cm-hostname (config) # cmc appliance <appl lD> auth password username
   <username>
   cm-hostname (config) # cmc appliance <appl ID> auth password password
   <password>
   ```

   where:

   `appl ID` is the name of the Virtual Execution appliance that will be a node in the MVX cluster. (The name is displayed in the `show cmc appliances` command output.)

   `IP address` is the IPv4 or IPv6 address of the management interface on the appliance.

   `username` is the username of a user account on the appliance with the admin role.

   `password` is the password of the user specified by `username`.

4. Repeat the previous step for each additional Virtual Execution appliance.
5. Check the connection status:

   ```
   cm-hostname (config) # show cmc appliances
   ```

Example

The following example adds the vx-1 and vx-2 appliances to the Central Management appliance.

```
cm-hostname (config) # cmc appliance vx-1 address 172.16.1.1
cm-hostname (config) # cmc appliance vx-1 auth password username admin
cm-hostname (config) # cmc appliance vx-1 auth password password admin123
cm-hostname (config) # cmc appliance vx-1 address 172.16.2.2
cm-hostname (config) # cmc appliance vx-2 auth password username admin
cm-hostname (config) # cmc appliance vx-2 auth password password admin123
cm-hostname (config) # show cmc appliances
Appliance vx-1:
  Address:                    172.16.1.1
  Enabled:                    yes
  Connected:                  yes (server-initiated)
  Status check OK:            no
  Version compatible:         yes

Appliance vx-2:
  Address:                    172.16.2.2
  Enabled:                    yes
  Connected:                  yes (server-initiated)
  Status check OK:            no
  Version compatible:         yes
```

Adding Sensors and Hybrid Appliances to a Central Management Appliance

You can use the Central Management Web UI or CLI to add sensors and hybrid appliances to an on-premises Central Management appliance for management.

Prerequisites

- Operator or Admin access

Adding Sensors and Hybrid Appliances to the Central Management Appliance Using the Web UI

Use the **Appliances** > **Sensors** page to add appliances that will function as sensors or hybrid appliances to a Central Management appliance.

Field	Description
Display Name	The name of the appliance. (The name will be displayed in the `show cmc appliances` command output.)
IP Address	The IPv4 or IPv6 address of the management interface on the appliance.
Username	A user account on the appliance with the admin role.
Password	The password for the user account specified in the **Username** field.
Appliance Host Key	(*Required if host-key authentication is enabled*) The RSA v2 host key for the appliance. For details, see the *FireEye System Security Guide*.
Comment	(*Optional*) A comment that describes the appliance.

To add an appliance:

1. Log into the Central Management Web UI.
2. Select **Appliances** > **Sensors**.
3. Click **Actions** > **Add Sensor**. Enter values for the fields described in the preceding table.
4. Click **Add**. Information about each appliance is displayed. To see details about a specific appliance, click the appliance name.

Adding Sensors and Hybrid Appliances to the Central Management Appliance Using the CLI

Use the commands in this section to add appliances that will function as sensors or hybrid appliances to a Central Management appliance.

To add an appliance:

1. Log in to the Central Management CLI.
2. Go to CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Add the appliance:

   ```
   cm-hostname (config) # cmc appliance <appl ID> address <IP address>
   cm-hostname (config) # cmc appliance <appl lD> auth password username
   <username>
   cm-hostname (config) # cmc appliance <appl ID> auth password password
   <password>
   ```

   where:

   - `appl ID` is the name of the appliance. (The name is displayed in the `show cmc appliances` command output.)
   - `IP address` is the IPv4 or IPv6 address of the management interface on the appliance.
   - `username` is the username of a user account on the appliance with the admin role.
   - `password` is the password of the user specified by `username`.

4. Repeat the previous step for each additional appliance.
5. Check the connection status:

   ```
   cm-hostname (config) # show cmc appliances
   ```

6. Save your changes:

   ```
   cm-hostname (config) # write memory
   ```

Example

The following example adds the nx-1 and nx-2 appliances to the Central Management appliance.

```
cm-1 (config) # cmc appliance nx-1 address 172.17.74.50
cm-1 (config) # cmc appliance nx-1 auth password username admin
cm-1 (config) # cmc appliance nx-1 auth password password admin123
cm-1 (config) # cmc appliance nx-2 address 10.13.65.14
cm-1 (config) # cmc appliance nx-2 auth password username admin
cm-1 (config) # cmc appliance nx-2 auth password password admin123
cm-1 (config) # show cmc appliances
```

```
Appliance nx-1:
   Address:                    172.17.74.50
   Enabled:                    yes
   Connected:                  yes (server-initiated)
   Status check OK:            yes
   Version compatible:         no
Appliance nx-2:
   Address:                    10.13.65.14
   Enabled:                    yes
   Connected:                  yes (server-initiated)
   Status check OK:            yes
   Version compatible:         no
```

CHAPTER 14: Defining the Interfaces

The submission interface is used for communication between sensors or hybrid appliances and brokers. The cluster interface is used for communication between broker nodes and compute nodes, and between compute nodes. By default, both the submission interface and cluster interface are defined to be *ether1*.

FireEye recommends that you configure another management interface (such as ether2), not a monitoring interface, for the submission and cluster interfaces. This prevents the ether1 management interface from becoming too busy, and keeps management and data traffic separate. However, ether2 cannot be used as both the submission interface and the URL Dynamic Analysis interface on an Email Security — Server Edition sensor.

If the sensor or hybrid appliance and broker are in different subnets, you must define the default gateway for the interface.

NOTE: DHCP is not currently supported on the submission or cluster interface.

NOTE: Ether1 is the only supported submission interface on File Protect sensors and hybrid appliances.

Prerequisites

- Operator or Admin access

Defining the Interfaces Using the CLI

Use the commands in this section to define the submission interface on a sensor or hybrid appliance and broker node, and to define the cluster interface on a broker node and compute node.

NOTE: Alternatively, you can change the interfaces for a broker and node in the configuration wizard. Use the `configuration jump-start` command to start the wizard.

Defining the Submission Interface

To define the submission interface on a sensor or hybrid appliance:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI configuration mode:

   ```
   hostname > enable
   hostname # configure terminal
   ```

3. Specify the interface:

   ```
   hostname (config) # mvx sensor config submission-if <interface name>
   ```

4. *If the sensor or hybrid appliance and broker are in different subnets*: Specify the IPv4 address of the default gateway for the interface:

   ```
   hostname (config) # mvx sensor config submission-if <interface name>
   default-gateway ipv4 <ipv4 address>
   ```

5. Verify your change:

   ```
   hostname (config) # show mvx status
   ```

6. Save your change:

   ```
   hostname (config) # write memory
   ```

To define the submission interface on a broker node:

1. Log in to the broker (Virtual Execution) CLI.
2. Enable the CLI configuration mode

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Specify the interface name:

   ```
   vx-hostname (config) # mvx node config submission-if <interface name>
   ```

4. *If the broker and any sensors are in different subnets*: Specify the IPv4 address of the default gateway for the interface:

   ```
   vx-hostname (config) # mvx node config submission-if <interface name>
   default-gateway ipv4 <ipv4 address>
   ```

5. Verify your changes:

   ```
   vx-hostname (config) # show mvx node status
   ```

6. Save your changes:

   ```
   vx-hostname (config) # write memory
   ```

Defining the Cluster Interface

To define the cluster interface on a broker node or compute node:

1. Log in to the Virtual Execution CLI.
2. Enable the CLI configuration mode:

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Specify the interface:

   ```
   vx-hostname (config) # mvx node config cluster-if <interface name>
   ```

4. Verify your change:

   ```
   vx-hostname (config) # show mvx node status
   ```

5. Save your change:

   ```
   vx-hostname (config) # write memory
   ```

Examples

The following example defines ether2 as the cluster interface on the vx-3 node.

```
vx-3 (hostname) # mvx node config cluster-if ether2
vx-3 (hostname) # show mvx node status

MVX Cluster: Node Status

Broker role:

     Enabled              : no
     Ready:               : yes
     SSH port             : 22
     Submission Interface : ether1
     Cluster Interface    : ether2
     Key Hash             : f8:xx:xx:...
...
```

The following example defines ether2 as the submission interface on the nx-4 sensor.

```
nx-4 (hostname) # mvx node config submission-if ether2
nx-4 (hostname) # show mvx status
MVX Mode Status:
    Sensor Config Enabled:   yes
    Current Operating Mode:  sensor
    Mode Reboot Required:    no
    Submission Interface:    ether2
    Modes Supported:         mvx configurable
    ...
```

PART IV: Configuration

CHAPTER 15: Working with Brokers and Compute Nodes

The following topics describe how to work with brokers and compute nodes.

Adding a Node to a Cluster

A Virtual Execution appliance serves no purpose until it is added to a cluster. After you add a Virtual Execution appliance to the Central Management appliance, it is an available node that can be added to an MVX cluster. You can determine whether a Virtual Execution appliance is an available node from the Central Management Web UI or CLI and the Virtual Execution CLI.

The **Cluster** column in the Central Management Web UI indicates whether a node is available.

The `Node information` section in the following CLI command output is empty because the vx-2 node is not currently part of a cluster, so it is available.

```
vx-2 # show mvx node status full
MVX Cluster: Node Status
...
Node information:

vx-2 #
```

Prerequisites

- Operator or Admin access to the Central Management appliance

Adding a Node to a Cluster Using the Web UI

Use the **Appliances** > **Clusters** page to add a node to an MVX cluster.

To add a node to a cluster:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click the **Actions** menu, and then select **Edit**.
4. Click the arrow to move the node from the **Available** list to the **Selected** list.
5. Click **Next**.
6. Enable at least one node as a broker and define one node as the master configuration.
7. Click **Update**.

Adding a Node to a Cluster Using the CLI

Use the commands in this section to add a node to an MVX cluster.

To add a node to a cluster:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. Add the node:

```
cm-hostname (config) # cmc mvx cluster <Cluster name> node <node name>
```

4. Verify that the node was added:

```
cm-hostname (config) # show cmc mvx cluster Cluster-Acme nodes
NODES: Cluster-Acme      CONNECTED    HEALTHY      ADDRESS
-------------            ---------    -------      -------
    vx-1                 yes          yes          10.11.121.12
    vx-2                 yes          yes          10.11.121.18
...
```

Enabling and Disabling Broker Mode

You can enable a compute node to function as a broker. You can also remove this functionality from a broker, which restores it as a compute node.

Prerequisites

- Operator or Admin access to the Central Management appliance that manages the cluster.

- Submission and cluster interfaces are configured on the node to be enabled as a broker node.

Enabling and Disabling Broker Mode Using the Web UI

Use the **Appliances** > **Nodes** page to enable or disable broker mode.

To enable or disable broker mode:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Nodes**.
3. To enable broker mode:
    a. Click the **Actions** menu, and then select **Enable Broker**.
    b. If you are prompted to make this the master configuration and want to do so, select the **Enable Master Config** check box.
    c. Click **ENABLE** to confirm your action.
4. To disable broker mode:
    a. To disable broker mode, click the **Actions** menu, and then select **Disable Broker**.
    b. When prompted, click **DISABLE** to confirm your action.

Enabling and Disabling Broker Mode Using the CLI

Use the commands in this section to enable or disable broker mode.

To enable or disable broker mode:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. To enable broker mode:

```
cm-hostname (config) # cmc mvx cluster <cluster name> broker <node
name> enable
```

4. To disable broker mode:

```
cm-hostname (config) # no cmc mvx cluster <cluster name> broker <broker
name> enable
```

5. Verify your changes:

```
cm-hostname (config) # show cmc mvx cluster <cluster name> nodes
```

6. Save your changes:

```
cm-hostname (config) # write memory
```

Example

The following example enables vx-3 to function as a broker node.

```
cm-1 (config) # cmc mvx cluster Cluster-Acme broker vx-1 enable
Broker role enabled successfully on vx-1
```

Removing a Node from a Cluster

After a broker node or compute node is removed from an MVX cluster, it remains connected to the Central Management appliance as an available node. The node serves no purpose until it is added to another MVX cluster.

Prerequisites

- Operator or Admin access to the Central Management appliance
- If the node you are removing is the only broker node, enable another broker node.

Removing a Node from a Cluster Using the Web UI

Use the **Appliances** > **Clusters** page or the **Appliances** > **Nodes** page to remove a broker node or compute node from an MVX cluster.

Using the Clusters Page to Remove a Node

To remove a broker or compute node:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click the **Actions** menu, and then click **Edit**.
4. Click the icon to move the node from the **Selected** list to the **Available** list.
5. Click **Next**.
6. Modify broker assignments if necessary, and then click **Update**.

Using the Nodes Page to Remove a Node

To remove a broker or compute node:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Nodes**.
3. Click the **Actions** menu, and then click **Remove**.
4. When prompted, click **REMOVE** to confirm that you want to remove the node.

Removing a Node from a Cluster Using the CLI

Use the commands in this section to remove a broker node or compute node from an MVX cluster.

To remove a broker or compute node:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Remove the node:

   ```
   cm-hostname (config) # no cmc mvx <cluster name> node <node name>
   ```

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx cluster <cluster name> nodes
   ```

Example

The following example removes the vx-2 from Cluster-Acme.

```
cm-1 (config) # no cmc mvx cluster Cluster-Acme node vx-2
```

Removing a Node from a Cluster on an Offline Central Management Appliance Using the CLI

Use the commands in this section to remove a broker node or compute node from an MVX cluster that is managed by an offline Central Management appliance.

To remove a broker or compute node:

1. Log in to the Virtual Execution CLI.
2. Enable the CLI configuration mode:

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Remove this node:

   ```
   vx-hostname (config) # mvx node detach
   ```

4. Remove this node or another node in the cluster:

   ```
   vx-hostname (config) # mvx node detach <node>
   ```

Example

The following example removes the vx-2 node from the cluster.

```
vx-1 (config) # mvx node detach vx-2
```

Deleting a Node from the Central Management Appliance

You must remove a broker node or compute node (Virtual Execution appliance) from an MVX cluster before you can delete it from the Central Management appliance.

Prerequisites

- Operator or Admin access to the Central Management appliance
- Node is an available node (that is, it is not currently part of an MVX cluster).

Deleting a Node from the Central Management Appliance Using the Web UI

Use the **Appliances** > **Nodes** page to delete a node from the Central Management appliance.

IMPORTANT: The node must be removed from the cluster before it can be removed from the Central Management appliance.

To delete a node:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Nodes**.
3. Click the **Actions** menu, and then select **Delete**.
4. When prompted, click **DELETE** to confirm that you want to delete the node.

Deleting a Node from the Central Management Appliance Using the CLI

Use the commands in this section to delete a node (Virtual Execution appliance) from the Central Management appliance.

To delete a node from the Central Management appliance:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Delete the node:

   ```
   cm-hostname (config) # no cmc appliance <appliance ID>
   ```

4. Verify your change:

   ```
   cm-hostname (config) # show cmc appliances brief
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

Example

The following example deletes the vx-3 node from the Central Management appliance.

```
cm-1 (config) # no cmc appliance vx-3
```

Configuring an Accessible Broker Address

Sensors cannot communicate with brokers using the default submission interface if your MVX cluster is in an internal network behind a NAT gateway and your sensors are in an external network. In this scenario, you must configure an accessible (*public*) IP address for each Virtual Execution appliance that functions as a broker. This is the virtual NAT IP address and port that a network administrator must map to the Virtual Execution internal IP address and port 22.

The enrollment service on the Central Management appliance that manages the MVX cluster uses the accessible IP address (instead of the submission interface IP address) for sensor enrollment. The `show mvx node status` command output appends `(public)` to the broker address when an accessible IP address is configured.

Prerequisites

- Admin access

Configuring an Accessible Broker Address

Use the commands in this section to configure an accessible IP address for a broker.

To configure an accessible IP address for a broker:

1. Log in to the Virtual Execution (broker) CLI.
2. Enable the CLI configuration mode:

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Configure the accessible IP address:

   ```
   vx-hostname (config) # mvx node broker public-address <public IP addr>
   ```

4. Verify your change:

   ```
   vx-hostname (config) # show mvx node status
   ```

5. Save your change:

   ```
   vx-hostname (config) # write memory
   ```

Example

The following example configures 172.2.3.4 as the accessible IP address for the vx-1 broker whose internal IP address is 10.1.2.3. The accessible IP address is displayed in the `Broker Role` section, and the internal IP address is displayed in the `Node information` section.

```
vx-1 (config) # mvx node broker public-address 172.2.3.4
vx-1 (config) # show mvx node status

MVX Node: vx-2

       Cluster IP Address    : 172.2.3.4 (Public)
       IP Address            : 10.13.65.64
       Role                  : broker
       Status                : ready
       Version               : 8.3.0

       queue:
              Status          : ready
              utilization     : 0

storage:
       Status                : ready

       compute:
              Status          : ready
              total           : 30
              running         : 0
              utilization     : 0
```

```
MVX Node: vx-1

       Cluster IP Address    : 10.13.65.63
       IP Address            : 10.13.65.63
       Role                  : compute
       Status                : ready
       Version               : 8.3.0

       compute:
              Status          : ready
              total           : 30
              running         : 0
              utilization     : 0
```

Removing an Accessible Broker Address

Use the commands in this section to remove the accessible IP address from a broker. This restores the default submission interface as the IP address that sensors use to communicate with the broker.

To remove an accessible IP address from a broker:

1. Log in to the Virtual Execution (broker) CLI.
2. Enable the CLI configuration mode:

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Remove the accessible IP address:

   ```
   vx-hostname (config) # no mvx node broker public-address
   ```

4. Verify your change:

   ```
   vx-hostname (config) # show mvx node status
   ```

5. Save your change:

   ```
   vx-hostname (config) # write memory
   ```

Example

The following example removes the accessible IP address from the vx-1 broker.

```
vx-1 (config) # no mvx node broker public-address
vx-1 (config) # show mvx node status

MVX Node: vx-2

       Cluster IP Address    : 10.13.65.64
       IP Address            : 10.13.65.64
       Role                  : broker
       Status                : ready
       Version               : 8.3.0

       queue:
              Status          : ready
              utilization     : 0
```

```
        storage:
                Status                  : ready

        compute:
                Status          : ready
                total           : 30
                running         : 0
                utilization     : 0

MVX Node: vx-1

Cluster IP Address    : 10.13.65.63
       IP Address            : 10.13.65.63
       Role                  : compute
       Status                : ready
       Version               : 8.3.0

       compute:
               Status          : ready
               total           : 30
               running         : 0
               utilization     : 0
```

CHAPTER 16: Changing Utilization Data Reporting

The following topics describe ways you can change the way utilization data is reported.

Prerequisites

- Operator or Admin access to the Central Management appliance that manages the cluster

Changing Alert Levels

As described in Viewing Cluster Utilization on page 133, by default warning alerts are generated when the cluster utilization reaches 60% of capacity, and critical alerts are generated when the cluster utilization exceeds 85% of capacity.

You can change the levels at which warning and critical alerts will be generated. The warning level cannot be set to a higher value than the critical level.

To change the warning level:

1. Log in to the Central Management CLI.
2. Go to CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Specify the percentage:

   ```
   cm-hostname (config) # cmc mvx status cluster-sizing threshold warning
   <percentage>
   ```

   where `<percentage>` can be a value from 10–90 and must be less than the "critical" percentage.

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx status cluster-sizing config
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

To change the critical level:

1. Log in to the Central Management CLI.
2. Go to CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Specify the percentage:

   ```
   cm-hostname (config) # cmc mvx status cluster-sizing threshold critical
   <percentage>
   ```

   where `<percentage>` can be a value from 20–100 and must be a greater than the "warning" value.

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx status cluster-sizing config
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

Example

The following example changes the warning level to 75% of total capacity.

```
cm-1 (config) # cmc mvx status cluster-sizing threshold warning 75
cm-1 # show cmc mvx status cluster-sizing config

MVX Cluster Sizing Configurations:
Enabled:                        yes
Utilization Warning Threshold:  75%
Utilization Critical Threshold: 95%
```

Disabling Cluster Utilization Statistics

You can stop the collection and display of MVX cluster utilization statistics.

NOTE: The **Cluster Utilization** panel on the Central Management Dashboard is removed when you disable cluster utilization statistics.

To disable cluster utilization statistics:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Disable the feature:

   ```
   cm-hostname (config) # no cmc mvx status cluster-sizing enable
   ```

4. Verify your change:

   ```
   cm-hostname (config) # show cmc mvx status cluster-sizing config
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

NOTE: To re-enable cluster sizing statistics, use the `cmc mvx status cluster-sizing enable` command.

Example

The following example stops the collection and display of MVX cluster utilization statistics.

```
cm-1 (config) # no cmc mvx status cluster-sizing enable

MVX Cluster Sizing Configurations:
Enabled:                        no
Utilization Warning Threshold:  80%
Utilization Critical Threshold: 95%
```

CHAPTER 17: Deleting a Cluster

You can delete an MVX cluster. The broker nodes and compute nodes that were in the cluster remain connected to the Central Management appliance as available nodes that are ready to be added to another cluster.

Prerequisites

- Operator or Admin access to the Central Management appliance

Deleting a Cluster Using the Web UI

Use the **Appliances** > **Clusters** page to delete an MVX cluster.

To delete an MVX cluster:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click the **Actions** menu, and then select **Delete**.
4. When prompted, click **DELETE** to confirm that you want to delete the cluster.

Deleting a Cluster Using the CLI

Use the CLI commands in this topic to delete a cluster from the Central Management appliance.

To delete a cluster:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. Delete the cluster:

   ```
   cm-hostname (config) # no cmc mvx cluster <cluster name>
   ```

4. Verify that the cluster is no longer configured:

   ```
   cm-hostname (config) # show cmc mvx cluster
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

Example

The following command deletes Cluster-Acme from the Central Management appliance.

```
cm-1 (config) # no cmc mvx cluster Cluster-Acme
cm-1 # show cmc mvx cluster
No MVX cluster configured.
```

PART V: Monitoring

CHAPTER 18: Viewing Cluster and Node Status

The following topics describe how to view status for a cluster and individual nodes.

Viewing Cluster Status

The Central Management Dashboard provides instant visibility into MVX cluster status. You can drill down to other pages in the Central Management Web UI to view details. You can also use Central Management and Virtual Execution CLI commands to view status information.

Prerequisites

- Monitor, Operator, or Admin access

Viewing the Cluster Status Using the Central Management Dashboard

Use the Central Management Dashboard to check the health of the cluster and its components.

A healthy cluster has the following characteristics:

- Virtual Execution appliances (nodes) are connected to the Central Management appliance.
- The cluster status in Ready status.
- The cluster utilization percentage is in the normal range.
- The nodes are in Ready status.
- There is at least one broker.

- Sensors that are managed by the Central Management appliance that manages the cluster are connected and healthy.
- Hybrid appliances that are connected to the Central Management appliance that manages the cluster are healthy.

To view cluster status:

1. Log in to the Central Management Web UI.
2. Click the **Dashboard** tab, if it is not already displayed.
3. Select the cluster from the group menu.
4. Review the information in the **Summary** pane.
5. Click the cluster name, and review the high-level information about the cluster.
6. Click the cluster name again to open the **Details** dialog box. View more detailed information on the **Summary** tab.
7. Click the **Nodes** tab to view information about the nodes.
8. Click the node name to view more detailed information about the node.
9. Click the **Enrolled Sensors** tab to view information about the sensors and hybrid appliances.

Viewing Cluster Status Using the Central Management Web UI

Use the **Appliances** > **Clusters** page in the Central Management Web UI to view a list of the clusters this Central Management appliance manages and their overall status. You can then click a cluster name for more detail.

The following table describes the cluster status fields.

Field	Description
Name	The name of the cluster.
Connection	The state of the connection between the nodes and the cluster. • **OK** or **Connected**—At least one node is connected to the cluster. • **Disconnected**—No node is connected to the cluster.
Health	• **OK** or **Healthy**—The cluster and the individual nodes are healthy. • **Warning**—Configuration settings are not the same on all nodes. • **Critical**—The cluster does not have a broker node.

Field	Description
Utilization (hourly)	The average cluster utilization level over the last hour: **Healthy**, **Warning**, **Critical**. See Viewing Cluster Utilization on page 133 for information about the three levels.
OS	The Virtual Execution system image version that is running on the nodes.
Guest Image	The guest images version installed on the nodes.

To view cluster status:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click the cluster name in the **Name** column to view additional details about the cluster, the nodes, and the enrolled sensors.

Viewing Cluster Status Using the Central Management CLI

Use the commands in this section to monitor the status of the cluster and its components.

To view cluster status:

1. Log in to the Central Management CLI.
2. Enable the CLI enable mode:

   ```
   cm-hostname > enable
   ```
3. To view summary information:

   ```
   cm-hostname # show cmc mvx cluster
   ```
4. To view brief information:

   ```
   cm-hostname # show cmc mvx cluster brief
   ```
5. To view detailed information about a specific cluster:

   ```
   cm-hostname # show cmc mvx cluster <cluster name>
   ```
6. To view information about the nodes in a specific cluster:

   ```
   cm-hostname # show cmc mvx cluster <cluster name> nodes
   ```

Examples

The following example shows summary information about the Cluster-Acme cluster (the only cluster being managed by this Central Management appliance).

```
cm-1 # show cmc mvx cluster

MVX Cluster: Cluster-Acme

  Health OK:           yes
  Health severity:     OK
  Master broker:       vx-1
  Member node count:   2
  All connected:       yes
  Description:
```

The following example shows brief information about the Cluster-Acme cluster (the only cluster being managed by this Central Management appliance).

```
cm-1 # show cmc mvx cluster brief
CLUSTER NAME          HEALTHY   CONNECTED NODES MASTER
------------          -------   --------- ----- ------
Cluster-Acme          yes       yes       2     vx-1
```

The following example shows detailed information about the Cluster-Acme cluster.

```
cm-1 (config) # show cmc mvx cluster detail

MVX cluster: Cluster-Acme

   Health OK:          yes
   Health severity:    OK
   Master broker:      vx-1
   Member node count:  2
   All connected:      yes
   Description:

   Health Status:
      Nodes connected all:              yes
      System configuration in sync:     yes
      System software version match:    yes
      Security content version match:   yes
      Guest-images version match:       yes
      Master Node Selected:             yes
      Broker selected:                  yes

   Update Status:
      Latest OS version installed:      yes
      GI update available:              no

   Member Status (Total 2 Nodes):
      Brokers:
         vx-1 (master)                  10.11.121.12   ok

      Compute Nodes:
         vx-2                           10.11.121.18   ok
```

The following information shows information about the nodes in the Cluster-Acme cluster.

```
cm-1 # show cmc mvx cluster Cluster-Acme nodes
NODES: Cluster-Acme     CONNECTED   HEALTHY   ADDRESS
--------                ---------   -------   -------
    vx-1                yes         yes       10.11.121.12
    vx-2                yes         yes       10.11.121.18
Brokers (active)
    vx-1                yes         yes       10.11.121.12
Brokers (ready)
    vx-2                yes         yes       10.11.121.18
```


Output Fields

Field	Description
Health OK	Whether the cluster is healthy.
Health severity	Severity of the cluster's health condition: • `OK`—The cluster is healthy. • `NOT READY`—The cluster has no master configuration. • `WARNING`—The configuration settings on the nodes do not match and need to be synchronized with the master configuration, the security content versions on the nodes do not match, or the node is not healthy. • `DEGRADED`—One or more nodes in the cluster are not connected to the Central Management appliance, or the nodes are not running the same system image. • `CRITICAL`—The cluster has no broker.
Master broker	Hostname of the master broker.
Member node count	Number of nodes in the cluster (including both brokers and compute nodes).
All connected	Whether all members of the cluster are connected.
Description	Description of the cluster (if one was provided).
Health Status	
Nodes connected all	Whether all nodes are connected.
System configuration in sync	Whether relevant detection-related configuration settings match on all nodes.
System software version match	Whether the same Virtual Execution software version is running on all nodes.

Field	Description
Security content version match	Whether the same security content version is running on all nodes.
Guest-images version match	Whether the same versions of guest images are installed on all nodes.
Master Node Selected	Whether the master configuration is defined.
Broker selected	Whether a broker is defined.
Update Status	
Latest OS version installed	Whether the latest Virtual Execution software version is running on all nodes.
GI update available	Whether the latest version of guest images is running on all nodes.
Member Status	
Brokers	The hostname, IP address, and health status of each broker node. The health status can be `ok`, `WARN`, or `CRIT`.
Compute Nodes	The hostname, IP address, and health status of each compute node. The health status can be `ok`, `WARN`, or `CRIT`.

Viewing Cluster Status Using the Virtual Execution CLI

Use the commands in this section to monitor the status of a cluster and its member nodes.

To view cluster status:

1. Log in to the Virtual Execution CLI.
2. Enable the CLI enable mode:

   ```
   vx-hostname > enable
   ```

3. To view cluster status:

   ```
   vx-hostname # show mvx cluster status
   ```

4. To view detailed information about each node in the cluster:

   ```
   vx-hostname # show mvx cluster status detail
   ```

Examples

The following example shows the status of the Cluster-Acme cluster from the vx-1 member node.

```
vx-1 # show mvx cluster status

MVX Cluster Status

    Cluster Name         : Cluster-Acme
    Cluster Health       : ok
    Cluster Size         : 3
    Cluster Utilization  : 48 %
```

The following example shows the status of the Cluster-Acme cluster and its nodes from the vx-2 member node.

```
vx-2 # show mvx cluster status detail

MVX Cluster Status

    Cluster Name         : Cluster-Acme
    Cluster Health       : ok
    Cluster Size         : 3
    Cluster Utilization  : 57 %

    MVX Cluster Nodes

        Node Name        : vx-1
        Node Health      : ok
        Node Address     : 10.1.1.1
        Node Role        : broker
        Node Utilization : 65 %

        Node Name        : vx-2
        Node Health      : ok
        Node Address     : 10.1.1.2
        Node Role        : node
        Node Utilization : 96 %

        Node Name        : vx-3
        Node Health      : ok
        Node Address     : 10.1.1.3
        Node Role        : node
        Node Utilization : 83 %
```

The following example shows the status from the vx-4 node, which is not yet a member of a cluster.

```
vx-4 # show mvx cluster status

  MVX Cluster nodes status is not yet available.
```

Output Fields

Field	Description
MVX Cluster Status	
Cluster Name	The cluster name.
Cluster Health	The cluster health: • `ok`—The cluster is healthy. • `degraded`—One or more nodes are unhealthy or offline. • `critical`—All nodes are unhealthy or no broker is available.
Cluster Size	The number of nodes in the cluster.
Cluster Utilization	The current cluster utilization, as a percentage of the capacity of the whole cluster. This calculation is based on the cluster's submission queue size and the member nodes' VM load. **NOTE:** Any node in a `critical` health state is not included in this calculation.
MVX Cluster Nodes	
Node Name	The node name.
Node Health	The node health: • `ok`—The node is healthy. • `critical`—Some major functionality is not working (for example, there is no broker node) or the node is offline (for example, it is reloading after an upgrade).
Node Address	The IP address of the Virtual Execution appliance (node).
Node Role	The role of the node in the cluster: • `broker`—broker node • `node`—compute node
Node Utilization	The current node utilization, as a percentage of capacity. This calculation is based on the node's VM load.

Viewing Cluster Database Statistics Using the Virtual Execution CLI

Use the commands in this section to view cluster database statistics for each broker in an MVX cluster.

To view cluster database statistics:

1. Log in to the CLI of any broker (Virtual Execution appliance) in the cluster.
2. Enable the CLI enable mode:

   ```
   vx-hostname > enable
   ```

3. View the statistics:

   ```
   vx-hostname # show mvx node cdbmgr status
   ```

Examples

The following example from the vx-1 broker shows the statistics for it and the other broker in the cluster.

```
vx-1 # show mvx node cdbmgr status

  CdbMgr Node Stats:
    Ip Address                                : 10.11.121.12
    CDB Node ID                               : 7xxxxxxx-bxxx-4xxx-gxxx-2xxxxx
    Status                                    : Up and running
    Load (Data size)                          : 94.29 KB
    Owns (% of cluster data stored on node)   : 100.0%

    Ip Address                                : 10.11.121.18
    CDB Node ID                               : bxxxxxxx-cxxx-5xxx-axxx-2xxxxx
    Status                                    : Up and running
    Load (Data size)                          : 70.33 KB
    Owns (% of cluster data stored on node)   : 100.0%
```

The following example shows the output of the command on a compute node that is not a broker.

```
vx-2 # show mvx node cdbmgr status
% Broker functionality is disabled
```

Output Fields

Field	Description
Ip Address	The IP address of the broker.
CDB Node ID	The unique identifier for the database node.

Field	Description
Status	The current status of the broker: • `Up and running`—The broker is running normally. • `Up and leaving cluster`—The broker is running and is leaving the cluster. • `Up and joining cluster`—The broker is running and is joining the cluster. • `Node is down`—The broker is not running. • `Down and joining cluster`—The broker is not running and is joining the cluster. • `Down and leaving cluster`—The broker is not running and is leaving the cluster.
Load (Data size)	The total amount of data this broker is handling. The data includes submission data from sensors and analysis results from compute nodes.
Owns (% of cluster data stored on node)	The percentage of the overall cluster data stored on this broker. For example, `100.0%` means all submission and analysis results from all member brokers have been replicated to storage on this broker.

Viewing Node Status

The following topics describe how to view information about the cluster, nodes, and brokers using the Central Management Web UI and the Virtual Execution CLI.

Prerequisites

- Monitor, Operator, or Admin access

Viewing Node Status Using the Central Management Web UI

Use the **Appliances** > **Nodes** page in the Central Management Web UI to view a list of the nodes this Central Management appliance manages and their overall status. You can then click a node name for more detail.

The following table shows the node status fields.

Field	Description
Name	The name of the node.
Connection	The state of the connection between the node and the MVX cluster.
Health	The state of the health of the node.
Cluster	The cluster of which the node is a member.
IP	The IP address of the node's management interface.
Appliance ID	The appliance ID of the node.
OS	The Virtual Execution system image version running on the node.
Guest Images	The guest images version installed on the node.
Node Details	
Last Contact	The last time the Central Management appliance contacted the node to get its status and health check data.

Field	Description
Connection	The status of the connection between the Central Management appliance and the node.
Member Groups	The groups of which the node is a member, including the reserved system group.
EULA	Whether the terms of the FireEye End User License Agreement (EULA) were accepted when the appliance was first configured.

To view node status:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Nodes**.
3. Click the node name in the **Name** column to view additional details.

Viewing Node Status Using the Virtual Execution CLI

Use the commands in this section to view the status of brokers and compute nodes in an MVX cluster.

To view node status:

1. Log in to the Virtual Execution CLI.
2. Enable the CLI enable mode:

   ```
   vx-hostname > enable
   ```

3. To view the node status:

   ```
   vx-hostname # show mvx node status
   ```

4. To view detailed status information:

   ```
   vx-hostname # show mvx node status full
   ```

Examples

The following example shows the status of the vx-1 broker.

```
vx-1 # show mvx node status
MVX Cluster: Node Status

Broker Role:
    Enabled                    : yes
    Ready                      : yes
    SSH port                   : 22
    Submission Interface       : ether1
    Cluster Interface          : ether1
    Key Hash                   : f2:xx:xx:xx:xx:xx:...
    Compute Enabled            : yes
```

```
Health Infomration:
    Overall Status Ok          : yes
    Overall Status Desc        : healthy


Sensor information:
    Number of connected sensors :  2
    Sensor ID list:
        Sensor ID        : 002XXXXXXXX4
        Sensor Hostname  : nx-2
        Sensor Address   : 10.13.65.14

        Sensor ID        : 002XXXXXXXX5
        Sensor Hostname  : nx-1
        Sensor Address   : 172.17.74.50

Node information:
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1 (self)
    Broker ID               : 002XXXXXXXX7 (self)
    Broker Address          : 10.11.121.12 (self)
    Broker State            : N/A
    Failure Reason          : N/A
    Last Connection Attempt : N/A
    Connection Last Formed  : N/A
    Connection Last Broken  : N/A
```

The following example shows the status of the vx-2 compute node.

```
vx-2 # show mvx node status
MVX Cluster: Node Status

Broker Role:
    Enabled              : no
    Ready                : yes
    SSH port             : 22
    Submission Interface : ether1
    Cluster Interface    : ether1
    Key Hash             : f4:xx:xx:xx:xx:xx:...
    Compute Enabled      : yes

Health Information:
    Overall Status Ok    : yes
    Overall Status Desc  : healthy

Node information:
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1
    Broker ID               : 002XXXXXXXX7
    Broker Address          : 10.11.121.12
    Broker State            : Connected
    Failure Reason          : None
    Last Connection Attempt : 2019/08/07 22:37:18
    Connection Last Formed  : 2019/08/07 22:37:18
    Connection Last Broken  : 2019/08/07 22:32:56
```

The following example shows detailed information about the vx-1 broker.

```
vx-1 # show mvx node status full
MVX Cluster: Node Status

Broker Role:
    Enabled              : yes
    Ready                : yes
```

```
    SSH port             : 22
    Submission Interface : ether1
    Cluster Interface    : ether1
    Key Hash             : f2:xx:xx:xx:xx:xx:...

Health Information:
    Overall Status Ok    : yes
    Overall Status Desc  : healthy

Detailed Health Information:
    CCD ok                 : yes.
    MvxClient Ok           : yes. Healthy
    Guest Images Ok        : yes. Installed
    Notification Client Ok : yes. Healthy
    WSAPI Ok               : yes. Running
    Queuemgr Ok            : yes. Healthy

Sensor information:
    Number of connected sensors :  2
    Sensor ID list:
        Sensor ID        : 002XXXXXXXX4
        Sensor Hostname  : nx-2
        Sensor Address   : 10.13.65.14

        Sensor ID        : 002XXXXXXXX5
        Sensor Hostname  : nx-1
        Sensor Address   : 172.17.74.50

Node information:
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1 (self)
    Broker ID               : 002XXXXXXXX7 (self)
    Broker Address          : 10.11.121.12 (self)
    Broker State            : N/A
    Failure Reason          : N/A
    Last Connection Attempt : N/A
    Connection Last Formed  : N/A
    Connection Last Broken  : N/A
```

Output Fields

Field	Description
Broker Role	
Enabled	Whether the node is enabled as a broker.
Ready	Whether the node can be enabled as a broker.
SSH port	Port used for secure shell (SSH) communication between the broker and sensors, the broker and compute nodes, and the broker and the Central Management appliance.
Submission Interface	Name of the interface used for communication between the broker and the sensors and hybrid appliances.

Field	Description
Cluster Interface	Name of the interface used for communication between the broker and compute nodes.
Key Hash	SSH host key used to authenticate secure communication.
Compute Enabled	Whether the node has been added to a cluster.
Health Information	
Overall Status Ok	Whether the status of the node is acceptable.
Overall Status Desc	Description of the overall status, such as "healthy."
Detailed Health Information	
CCD Ok	Whether the cluster communication process the enrollment service uses is running on the broker nodes. If this value is `no`, this field includes a list of the broker nodes that are not connected and the reasons. Some reasons for failure follow: • `no.Unknown` • `no.No route to host` • `no.Connection refused` • `no.Timeout` See Sensor, Hybrid Appliance, or Compute Node Cannot Connect to Broker on page 183 for troubleshooting information.
MvxClient Ok	Whether the submission process is running and a description. For example: • `yes.Healthy` • `no.Broker not configured` • `no.Connection refused retrying !!!` • `no.Queues not configured`
Guest Images Ok	Whether the guest images on the compute node are installed.
Notification Client Ok	(*Broker only*) Whether the notification process is healthy and a description. For example: `no.Sensor is down`.

Field	Description
WSAPI Ok	(*Broker only*) Whether the Web services API process is enabled and running.
Queuemgr Ok	(*Broker only*) Whether the queue manager process is healthy and a description. For example: `no.Broker disabled`.
Sensor information (*Broker only*)	
Number of connected sensors	Number of enrolled sensors that are connected to the broker.
Sensor ID list	List of connected sensors.
Sensor ID	Appliance ID of the sensor.
Sensor Hostname	Hostname of the sensor.
Sensor Address	IP address of the sensor's management interface.
Node information	
Cluster Name	Name of the cluster of which this node is a member.
Broker Name	Hostname of the broker connected to this node.
Broker ID	Appliance ID of the broker.
Broker Address	IP address of the broker's management interface.
Broker State	Status of the broker, such as "connected." (This field does not apply to brokers.)
Failure Reason	Reason the node is not connected to the broker. (This field does not apply to brokers.)
Last Connection Attempt	Date and time the broker and node last tried to connect. (This field does not apply to brokers.)

Field	Description
Connection Last Formed	Date and time the broker and node last connected. (This field does not apply to brokers.)
Connection Last Broken	Date and time the broker and node last lost their connection. (This field does not apply to brokers.)
Other Failure Reason	Additional information about a failure condition.

Viewing Queue Status on a Broker Using the Virtual Execution CLI

Use the commands in this section to view the status of the MVX engine queue on a broker.

To view the queue status:

1. Log in to the broker (Virtual Execution) CLI.
2. Enable the CLI enable mode:

   ```
   vx-hostname > enable
   ```

3. View the status:

   ```
   vx-hostname # show mvx node queuemgr status
   ```

Examples

The following example shows the status of the MVX engine queue on the vx-1 broker (the only broker in the cluster).

```
vx-1 # show mvx node queuemgr status

  QueueMgr Queue Stats:
    Queue Name             : high
    Queue Size             : 1
    Running submissions    : 1

    Queue Name             : low
    Queue Size             : 0
    Running submissions    : 0

    Queue Name             : normal
    Queue Size             : 116
    Running submissions    : 111

    Queue Name             : urgent
    Queue Size             : 2
    Running submissions    : 2
```

```
  QueueMgr Cluster Node Status:
    Ip address    : 10.11.121.12
    Running       : true
```

The following example shows the status of the MVX queue engine from the vx-3 broker (one of two brokers in the cluster).

```
vx-3 # show mvx node queuemgr status

  QueueMgr Queue Stats:
    Queue Name             : high
    Queue Size             : 1
    Running submissions    : 1

    Queue Name             : low
    Queue Size             : 2
    Running submissions    : 2

    Queue Name             : normal
    Queue Size             : 129
    Running submissions    : 129

    Queue Name             : urgent
    Queue Size             : 2
    Running submissions    : 2

  QueueMgr Cluster Node Status:
    Ip address    : 10.11.121.15
    Running       : true

    Ip address    : 10.11.121.16
    Running       : true
```

The following example shows the output of the command on a compute node that is not a broker.

```
vx-4 # show mvx node queuemgr status
QueueMgr is disabled
```

Output Fields

Field	Description
QueueMgr Queue Stats	
Queue Name	Priority of the submissions in the named queue (high, low, normal, or urgent).
Queue Size	Number of submissions in the queue.
Running Submissions	Number of submissions being analyzed.
QueueMgr Cluster Node Status	
IP address	IP address of the broker.

Field	Description
Running	Whether the queue process is running on the broker.

CHAPTER 19: Viewing Sensor and Hybrid Appliance Status

You can view MVX cluster-related information about sensors and hybrid appliances using the Central Management Web UI and the sensor or hybrid appliance CLI.

Prerequisites

- Monitor, Operator, or Admin access to the Central Management appliance
- Monitor or Admin access to the sensor or hybrid appliance

Viewing Sensor and Hybrid Appliance Status Using the Web UI

Use the **Appliances** > **Sensors** page in the Central Management Web UI to view the enrolled sensors and hybrid appliances this Central Management appliance manages and their overall status.

To view the overall status of a sensor:

1. Log in to the Central Management Web UI.
2. Select **Appliances** > **Clusters**.
3. Click the name of the cluster.
4. Click **Enrolled Sensors**.
5. For additional details, and to edit or delete a sensor, click **Sensors Page** in the row for the sensor or hybrid appliance. (You can also use the following procedure to navigate to this page.)

To view detailed status information:

1. Select **Appliances** > **Sensors**.
2. Click the sensor or hybrid appliance name to view additional details.

Viewing Sensor and Hybrid Appliance Status Using the CLI

Use the commands in this section to view the status of sensors and hybrid appliances enrolled in an MVX cluster.

To view the sensor or hybrid appliance status:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI enable mode:

   ```
   appl-hostname > enable
   ```

3. View the sensor status:

   ```
   appl-hostname # show mvx status
   ```

Examples

The following example shows the status of a virtual sensor.

```
VNX-4 # show mvx status
MVX Mode Status:
   Sensor Config Enabled:   yes
   Current Operating Mode:  sensor
   Mode Reboot Required:    no
   Submission Interface:    ether1
   Modes Supported:         mvx sensor-only
   Virtual Model:           yes
   Virtual System:          yes
   WSAPI Current State:     running
```

The following example shows the status of a physical integrated Network Security appliance that was enabled as a sensor.

```
nx-06 # show mvx status
MVX Mode Status:
   Sensor Config Enabled:   yes
   Current Operating Mode:  sensor
   Mode Reboot Required:    no
   Submission Interface:    ether1
   Modes Supported:         mvx configurable
   Virtual Model:           no
   Virtual System:          no
   WSAPI Current State:     running
```


Output Fields

Field	Description
Sensor Config Enabled	Whether the appliance is enabled as a sensor.
Current Operating Mode	Current operating mode (`sensor`, `hybrid`, or `integrated`).
Mode Reboot Required	Whether you need to reload the appliance after changing the operating mode for the change to take effect.
Submission Interface	Name of the interface used for communication between the sensor and broker.
Modes Supported	Operating mode for this appliance model: • `mvx configurable`—The appliance has an MVX engine. It can operate as an integrated appliance, in which its own MVX engine performs the analysis. It can also be converted to sensor mode, in which it submits objects to an MVX cluster instead of its own MVX engine, or to hybrid mode, in which it submits objects to an MVX cluster when a predefined capacity threshold is reached. • `mvx sensor-only`—The appliance has no MVX engine, and must submit objects to an MVX cluster for analysis. • `mvx integrated-only`—The appliance cannot submit objects to an MVX cluster and must use its own MVX engine for analysis.
Virtual Model	Whether this is a virtual appliance model.
Virtual System	Whether this is a virtual software image.
WSAPI Current State	The state of the Web services API process (`running` or `not running`)
Detailed Health Information	
Submission Client (MVX_SC) Ok	Whether the submission client is running: • `yes. Healthy`—The submission client is running. • `****` —The submission client is not available.

CHAPTER 20: Viewing Enrollment Status

The following sections describe how to view enrollment status.

Viewing Enrollment Status Using the Web UI

Use the **Cluster Enrollment** column on the **Appliances** > **Sensors** page to view the enrollment status of managed sensors and hybrid appliances.

For example:

- The name of the cluster and broker is displayed if the sensor or hybrid appliance is enrolled.
- **Unenrolled** is displayed if the sensor or hybrid appliance is not currently enrolled with a cluster.
- **Local** is displayed if the appliance is an integrated appliance and is not in MVX sensor or hybrid mode. In this case, the appliance always submits objects to its own (local) MVX analysis engine instead of submitting to a cluster.

To view the enrollment status:

1. Log in to the Web UI of the Central Management that manages the sensors.
2. Select **Appliances** > **Sensors**.
3. View the status in the **Cluster Enrollment** column.

Viewing Enrollment Status Using the CLI

Use the commands in this section to view the enrollment status of brokers and sensors or hybrid appliances.

To view the enrollment status from a broker:

1. Log in to the broker (Virtual Execution) CLI.
2. Enable the CLI enable mode:

   ```
   vx-hostname > enable
   ```

3. View the enrollment status:

   ```
   vx-hostname # show mvx cluster enrollment status
   ```

To view the enrollment status from a sensor or hybrid appliance:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI enable mode:

   ```
   appl-hostname > enable
   ```

3. View the enrollment status:

   ```
   appl-hostname # show mvx cluster enrollment status
   ```

To view the enrollment status from the Central Management appliance:

1. Log in to the CLI of the Central Management appliance that manages the cluster.
2. Enable the CLI enable mode:

   ```
   cm-hostname > enable
   ```

3. View the enrollment status:

   ```
   cm-hostname # show cmc mvx cluster enrollment status
   ```

Examples

The following example shows the enrollment status for the nx-1 sensor.

```
nx-1 # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client:
    Status ok               : yes
    Status description      : enrolled
    Last checked at         : 2019/07/26 20:42:01

  Enrollment Service:
    Auto enabled            : yes
    Service address         : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster       : any (less loaded)
    Cloud enabled           : no
    Cloud License enabled   : no
    Connect on demand       : no

Broker Info:
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1
```

```
    Broker ID               : 002XXXXXXXX7
    Broker Address          : 10.11.121.12
    Broker State            : Connected
    Failure Reason          : None
    Last Connection Attempt : 2019/06/27 18:14:04
    Connection Last Formed  : 2019/06/27 18:14:04
    Connection Last Broken  :
```

The following example shows the enrollment status for the vx-1 broker.

```
vx-1 # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client:
    Status ok               : yes
    Status description      : Ok
    Last checked at         : 2019/07/26 20:42:01

  Enrollment Service:
    Auto enabled            : yes
    Service address         : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster       :

    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1 (self)
    Broker ID               : 002XXXXXXXX7 (self)
    Broker Address          : 10.11.121.12 (self)
    Broker State            : N/A
    Failure Reason          : N/A
    Last Connection Attempt : N/A
    Connection Last Formed  : N/A
    Connection Last Broken  : N/A
    Connect on demand       : no
```

The following example shows the enrollment status for the fx-1 hybrid appliance.

```
fx-1 # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client:
    Status ok               : yes
    Status description      : enrolled
    Last checked at         : 2019/07/26 20:42:01

  Enrollment Service:
    Auto enabled            : yes
    Service address         : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster       : any (less loaded)
    Cloud enabled           : no
    Cloud License enabled   : no
    Connect on demand       : no

Broker Info:
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1
    Broker ID               : 002XXXXXXXX7
    Broker Address          : 10.11.121.12
    Broker State            : Connected
    Failure Reason          : None
    Last Connection Attempt : 2019/06/27 18:14:04
    Connection Last Formed  : 2019/06/27 18:14:04
    Connection Last Broken  :
```

The following example shows that the nx-1 and nx-2 sensors are enrolled with Cluster-Acme and are connected to the vx-1 broker.

```
cm-5 (config) # show cmc mvx cluster enrollment status
SENSOR NAME         CLUSTER NAME     BROKER NAME      BROKER ADDRESS
===========         ============     ===========      ==============
nx-1                Cluster-Acme     vx-1             10.11.121.12
nx-2                Cluster-Acme     vx-1             10.11.121.12
```

Output Fields

Field	Description
Enrollment Client	
Status ok	Whether the enrollment client status is acceptable.
Status description	Description of the status (`enrolled` or `unenrolled`). If the appliance is not a sensor, `local` is the description, because the appliance uses its own MVX engine for analysis.
Last checked at	Date and time the enrollment client was last checked.
Enrollment Service	
Auto enabled	Whether automatic enrollment is enabled.
Service address	Enrollment service type and address. The Central Management appliance manages the enrollment service for the sensors, hybrid appliances, and MVX clusters it manages.
Preferred cluster	The name of a preferred cluster, if one was configured as described in Enrolling with a Preferred Cluster on page 68. If none was configured, the value of this field is `any (less loaded)` to indicate that the sensor was enrolled with the cluster with the most capacity. (This field does not apply to brokers.)
Cloud enabled	Whether the sensor is able to send submissions to the FireEye cloud MVX service for analysis. (This functionality is enabled by default when the `Cloud License enabled` field is `yes`).
Cloud License enabled	Whether the sensor has a valid and active CLOUD_MVX license.

Field	Description
Connect on demand	Whether the sensor or hybrid appliance connects to the enrollment service only when it has objects to submit for analysis. This field is `yes` for FireEye cloud MVX deployments. This field is `no` for on-premises MVX cluster deployments.
Cluster Name	Name of the cluster with which the sensor or hybrid appliance is enrolled.
Broker Name	Hostname of the broker with which the sensor or hybrid appliance is connected.
Broker ID	Appliance ID of the broker.
Broker Address	IP address of the broker.
Broker State	State of the connection between the sensor or hybrid appliance and the broker node or between two broker nodes. • **Connected**—The sensor or hybrid appliance and the broker node or the two broker nodes are connected. • **Disconnected**—The sensor or hybrid appliance and the broker node or the two broker nodes are disconnected.
Failure Reason	Reason the connection failed, if any. • **Authentication Failure**—The DTI credentials on the sensor or hybrid appliance are invalid. • **Authorization Failure**—The sensor or hybrid appliance is using an invalid or inactive license. • **Enrollment Service Unavailable**—The sensor or hybrid appliance cannot reach the enrollment service. Otherwise, the value of this field is `none`.
Last Connection Attempt	Date and time the sensor or hybrid appliance and broker or two brokers last tried to connect.
Connection Last Formed	Date and time the sensor or hybrid appliance and broker or the two brokers last connected.

Field	Description
Connection Last Broken	Date and time the sensor or hybrid appliance and broker or the two brokers last lost their connection.
Other Failure Reason	Other reasons the connection between the sensor or hybrid appliance and broker or between two brokers failed. Other reasons include `Enrollment Service Error`, `Params missing or invalid`, and `Server error`.

CHAPTER 21: Viewing Cluster Utilization

The Central Management appliance continuously gathers and reports relevant data about MVX cluster utilization and performance. Statistics are collected from each broker, but because the cluster utilization is synchronized between brokers, the utilization is reported for the cluster as a whole.

The following recommended levels of utilization are set by default.

- **Healthy**—The cluster utilization is up to 60% capacity.
- **Warning**—The cluster utilization is from 60% to 85% capacity.
- **Critical**—The cluster utilization is over 85% capacity.

NOTE: You can change these levels, as described in Changing Alert Levels on page 95.

The Central Management appliance raises alerts and displays messages when the cluster continuously or critically exceeds these limits. Exceeding these limits can cause submissions to be dropped, which can result in reduced malware detection efficacy.

You can also use the utilization data as a tool for future capacity planning.

Prerequisites

- Monitor, Operator, or Admin access to the Central Management appliance

Viewing Cluster Utilization Using the Web UI

The Central Management appliance continuously gathers and reports relevant data about the cluster utilization and performance. You can view statistics for the current day, past week, or past month.

On the Central Management Dashboard, you can view the utilization zone in which the cluster is operating. The Central Management appliance warns you if the utilization reaches certain thresholds.

To view cluster utilization from the Central Management appliance:

1. Log in to the Central Management Web UI.
2. Click the **Dashboard** tab, if it is not already displayed.
3. Select the cluster from the group menu.
4. Locate the **Cluster Utilization** pane at the bottom of the Dashboard.
5. View the data in the **Cluster Utilization** section.

Viewing Cluster Utilization Using the CLI

Use the commands in this section to monitor cluster utilization and view related configuration settings.

NOTE: You can also view cluster utilization information from the Virtual Execution CLI. See Viewing Cluster Status Using the Virtual Execution CLI on page 108.

To view utilization statistics:

1. Log in to the Central Management CLI.
2. Enable the CLI enable mode:

   ```
   cm-hostname > enable
   ```

3. To view daily statistics:

   ```
   cm-hostname # show cmc mvx cluster <cluster name> stats daily
   ```

4. To view hourly statistics:

   ```
   cm-hostname # show cmc mvx cluster <cluster name> stats hourly
   ```

To view utilization configuration settings:

1. Log in to the Central Management CLI.
2. Enable the enable CLI mode:

   ```
   cm-hostname > enable
   ```

3. View configuration settings:

   ```
   cm-hostname # show cmc mvx status cluster-sizing config
   ```


Examples

The following example shows daily statistics for Cluster-Acme.

```
cm-1 # show cmc mvx cluster Cluster-Acme stats daily

                      Util    Submission      Submission    Submission
 Time                   (%)    Incoming            Done       Dropped
=============================================================================
 2019/05/23 00:00:00    0.41         431             115           135
 2019/05/22 00:00:00    0.00           0               0             0
```

The following example shows hourly statistics for Cluster-Acme.

```
cm-1 # show cmc mvx cluster Cluster-Acme stats hourly

                      Util      Submission      Submission    Submission
 Time                   (%)      Incoming            Done       Dropped
=============================================================================
 2019/05/23 23:00:00    0.00           30              10            20
 2019/05/22 22:00:00    0.00           21              10            11
 2019/05/22 21:00:00    0.33          915             902            13
 2019/05/22 20:00:00    0.33           27               6            21
...
```

The following example shows utilization configuration data for the clusters managed by this Central Management appliance.

```
cm-1 # show cmc mvx status cluster-sizing config

MVX Cluster Sizing Configurations:
  Enabled:                          yes
  Utilization Warning Threshold:    80%
  Utilization Critical Threshold:   95%
```

Output Fields

Field	Description
Time	Time period covered by the reported data.
Util (%)	The cluster utilization shown as a percentage of total capacity.
Submission Incoming	Number of submissions the cluster received.
Submission Done	Number of submissions that completed analysis.
Submission Dropped	Number of submissions that were dropped because the cluster was oversubscribed.
Enabled	Whether cluster utilization and performance data is currently being collected and reported. This functionality is enabled by default. You can disable and re-enable this functionality as described in Disabling Cluster Utilization Statistics on page 96.

Field	Description
Utilization Warning Threshold	The percentage of total capacity at which warning alerts are raised. The default is 80%. You can change the percentage as described in Changing Alert Levels on page 95.
Utilization Critical Threshold	The percentage of total capacity at which critical alerts are raised. The default is 95%. You can change the percentage as described in Changing Alert Levels on page 95.

CHAPTER 22: Viewing Submission Statistics

The Central Management appliance continuously gathers and reports statistics about the submissions the cluster received from its enrolled sensors and hybrid appliances. Statistics are collected from each broker, but because the cluster utilization is synchronized between brokers, the statistics are reported for the cluster as a whole.

Submissions can be dropped and messages are displayed when the cluster continuously or critically exceeds certain levels. Dropped submissions can result in reduced malware detection efficacy. You can use the submission statistics as a tool for future capacity planning.

You can also view submission statistics and results from the broker CLI.

Prerequisites

- Monitor or Admin access to the broker (Virtual Execution appliance) to view detailed statistics and analysis results from the CLI
- Monitor, Operator, or Admin access to the Central Management appliance to view the Dashboard panel

Viewing Submission Statistics Using the Web UI

Use the **Submission Statistics** graph in the **Cluster Utilization** panel of the Central Management Dashboard to view cluster submission statistics.

The graph shows the following statistics:

- **All Submissions**—The number of submissions the MVX cluster received.
- **Submissions Done**—The number of submissions the cluster processed.

- **Submissions Dropped**—The number of submissions that were dropped because the cluster did not have the capacity to process them.

To view submission statistics:

1. Log in to the Central Management Web UI.
2. Select **Dashboard**, if the Dashboard is not already open.
3. Select the cluster in the **All Groups** menu at the top right side of the Dashboard.
4. Locate the **Submission Statistics** graph in the **Cluster Utilization** panel at the bottom of the Dashboard.
5. Click **Day**, **Week**, or **Month** to select the time period to cover in the graph.

Viewing Submission Statistics Using the CLI

Use the commands in this section in the broker CLI to view statistics and analysis results for the submissions the MVX cluster processed.

NOTE: This section describes the basic form of the commands. For a full list of commands and their usage, see the *CLI Command Reference*.

Viewing Submission Statistics

Use the following commands (or a variation of them) to view submission statistics:

- `show mvx submission`—Statistics about all submissions the cluster processed.
- `show mvx submission from` —Statistics about submissions the cluster processed during the specified time period.
- `show mvx submission <sensor-id>`—Statistics about submissions from the sensor with the specified ID.
- `show mvx submission since`—Statistics about submissions since the specified date and time.

Viewing Analysis Results

Use the following commands (or a variation of them) to view analysis results:

- `show mvx submission done`—Analysis results for all submissions the cluster processed.

- `show mvx submission limit`—Analysis results for the most recent specified number of submissions the cluster processed.
- `show mvx submission malicious`—Analysis results for malicious submissions the cluster processed.
- `show mvx submission md5sum`—Analysis results for submissions with the specified MD5 hash.
- `show mvx submission sha256`—Analysis results for submissions with the specified SHA-256 hash.
- `show mvx submission uuid`—Analysis results for submissions with the specified universally unique identifier (UUID).

Examples

The following example shows statistics for all submissions the cluster processed.

```
vx-1 # show mvx submission

Runtime Cluster Stats:
  Total queued:                 : 0
  Total running:                : 0
  Cluster Utilization           : 0%

MVX Submission Stats:
  Total urls                    : 613
  Total files                   : 1251
  Total submissions             : 1864
  Completed submissions         : 1864
  Malicious submissions count   : 10
```

The following example shows statistics for submissions the MVX cluster processed from July 1, 2017 to July 8, 2017.

```
vx-1 # show mvx submission from 2019/07/01 12:00:00 to 2019/07/08 12:00:00

Runtime Cluster Stats:
   Total queued                 : 0
   Total running                : 3
   Cluster Utilization          : 2%

MVX Submission Stats:
   Total urls
   Total files                  : 289
   Total submissions            : 1216
   Completed submissions        : 1505
   Malicious submissions count  : 8
```

The following example shows statistics for submissions the cluster processed in the last 5 1/2 days.

```
vx-1 # show mvx submission since 5 days 12 hours
Runtime Cluster Stats:
   Total queued                 : 0
   Total running                : 0
   Cluster Utilization          : 0%
```

```
MVX Submission Stats:
   Total urls                    : 747
   Total files                   : 1455
   Total submissions             : 2202
   Completed submissions         : 2201
   Malicious submissions count   : 18
```

The following example shows the analysis results for the two most recent submissions.

```
vx-1 # show mvx submission limit 2
  Sensor ID                      : 001XXX...
  UUID                           : 289XXXX-XXXX...
  Insert time                    : 2019-07-11T22:27:02.841901
  Start time                     : 2019-07-11T22:19:18.392246
  Complete time                  : 2019-07-11T22:17:49.425406
  Error Code                     : SUCCESS
  Sensor Sub ID                  : 2076
  Malicious                      : NO
  Riskware                       : NO
  Files Analyzed                 : 1
  Overall weight                 : 0

      Analysis Object Name       : SDFixCapcity.exe
      Start Time                 : 2019-07-11T22:17:49.425406
      SHA256                     : 1XXX...
      MD5SUM                     : 2XXX...
      File Type                  : exe
      Static Analysis Weight     : 80
      Dynamic Analysis Weight    : 0
      Child                      : NO

  Sensor ID                      : 002XXX...
  UUID                           : 364XXXX-XXXX...
  Insert time                    : 2019-07-11T22:26:47.814274
  Start time                     : 2019-07-11T22:17:38.337658
  Complete time                  : 2019-07-11T22:17:34.412521
  Error Code                     : STATIC_ANALYSIS_ONLY
  Sensor Sub ID                  : 2073
  Malicious                      : NO
  Riskware                       : NO
  Files Analyzed                 : 2
  Overall weight                 : 0

      Analysis Object Name       : home.aspx
      Start Time                 : 2019-07-11T22:17:34.412521
      SHA256                     : 3XXX...
      MD5SUM                     : 4XXX...
      File Type                  : gz
      Static Analysis Weight     : 0
      Dynamic Analysis Weight    : 0
      Child                      : NO

      Analysis Object Name       : file
      Start Time                 : 2019-07-11T22:17:34.412521
      SHA256                     : 5XXX...
      MD5SUM                     : 6XXX...
      File Type                  : htm
      Static Analysis Weight     : 0
      Dynamic Analysis Weight    : 0
      Child                      : YES
```


Output Fields

Field	Description
Total queued	Total number of submissions in the MVX engine queue waiting to be pulled by a compute node.
Total running	Total number of submissions that are currently running.
Cluster Utilization	Cluster utilization, displayed as a percentage of capacity.
Total files	Total number of files submitted.
Total submissions	Total number of submissions.
Completed submissions	Total number of submissions that completed analysis.
Malicious submissions count	Total number of submissions that were detected as malicious.
Sensor ID	Appliance ID of the sensor or hybrid appliance.
UUID	Unique universal identifier for the submission.
Insert Time	The date and time the submission was added to the MVX engine queue.
Start Time	Date and time the analysis began.
Complete Time	Date and time the analysis ended.

Field	Description
Error Code	Status of the analysis. Some common error codes follow: • `SUCCESS`—The submission was analyzed successfully. • `SUBMISSION_DUPLICATE`—The submission was not analyzed because it matches a submission that was analyzed within the last 24 hours. • `SUBMISSION_DISABLE`—The type of file that was submitted is disabled, so the file was not analyzed. (Security content determines whether a file type is enabled or disabled. You can use the `show guest-images file-association sort file-type` or `guest-images file-association sort os` command to view the status of each file type.) • `BLACKLIST`—The type of file that was submitted is on a blacklist, so it was not analyzed, but was processed according to policy (for example, it was blocked). • `WHITELIST`—The type of file that was submitted is on a whitelist, so it was not analyzed, but was processed according to policy (for example, it was allowed). • `UNKNOWN`—The type of file that was submitted is not recognized. • `NO_PROFILE_MATCH`—The type of file that was submitted is not associated with any guest images profiles installed on the MVX cluster. For example, an APK file is associated with the Android operating system on mobile devices, not with a Windows or OS X guest images profile installed on the cluster. • `TIMEOUT`—The submission was not analyzed within a specific period of time. If many submissions time out, it usually means the system is overloaded.
Sensor Sub ID	ID that the sensor assigned to the submission.
Malicious	Whether the submission was detected as malicious.
Riskware	Whether the submission was detected as riskware.
Files Analyzed	Number of files in the submission.
Overall weight	Weight that is assigned to the submission based on a set of rules and what the MVX engine detected during analysis.

Field	Description
Analysis Object Name	Name of the file that was analyzed.
Start Time	Date and time the analysis began.
SHA256	SHA-256 checksum of the file.
MD5SUM	MD5 checksum of the file.
File Type	Type of file that was analyzed.
Static Analysis Weight	Weight that is assigned to a static analysis job on a particular object.
Dynamic Analysis Weight	Weight that is assigned to a dynamic analysis job on a particular object.
Child	Whether the object is contained in another object, such as a PDF file in a ZIP file.

PART VI: Administration

CHAPTER 23: Upgrades

You can upgrade the system image and guest images for an entire cluster with a single action. The Central Management appliance orchestrates the upgrade to ensure that at least one broker and one compute node are running during the upgrade, so there is no interruption in service. (This assumes that you have more than one broker enabled.)

You can also upgrade a single node. This option allows you to upgrade a node that is connected to the Central Management appliance, but not currently part of a cluster.

CAUTION: Do not upgrade a single node that is currently part of a cluster. Upgrade the entire cluster from the Central Management appliance instead.

NOTE: In this document, a "full" upgrade means the operation upgrades both the system image and guest images.

The system image and guest images are cached on the Central Management appliance after the first download for any node. The Central Management appliance uses these images to update subsequent nodes, instead of downloading them again. This efficiency is realized whether you update the cluster or a single node.

The Central Management appliance provides warnings when nodes are not running the latest version of the Virtual Execution system image or the latest version of guest images. It also warns you when the nodes are not running the same version of guest images, which means you must upgrade one or more nodes to synchronize them.

The VX node hostname on the Central Management appliance must match the VX hostname on the Virtual Execution appliance. If they are different, the Central Management Web UI will prevent you from upgrading the cluster. You must change the hostnames to match before upgrading the cluster. When you create a cluster, ensure the Virtual Execution hostname on the Virtual Execution appliance is same as the node hostname in the Central Management Web UI.

Upgrade messages are displayed in various areas of the Central Management Web UI, such as on the Clusters, Nodes, and Summary pages.

You can also use the `show cmc mvx cluster detail` command output to determine whether upgrades are available:

```
cm-hostname # show cmc mvx cluster detail
MVX Cluster: Cluster-Acme
```

```
Health OK:            yes
...
Update Status:
  Latest OS version installed:    no
  GI update available:            no
...
```

You can perform the following upgrade operations from either the Central Management Web UI or CLI:

- Full upgrade (system image and guest images) with or without reboot
- System image upgrade with or without reboot
- Guest images upgrade (download and install)
- Guest images download
- Guest images install
- Cancel, suspend, and resume current operation
- Cancel, suspend, and resume guest-images upgrade

From the Central Management CLI, you can also delete downloaded guest images. For example, if the guest-images installation fails due to a bad download, you can delete the guest images to clean up the partial data. See Deleting a Guest-Images Download on page 161.

Each upgrade operation follows a specific sequence of tasks. Status messages keep you informed about the progress of the upgrade as it moves through the tasks.

NOTE: You can use the `show fenet update operations` command to view the tasks for each operation. See Configuring and Viewing Upgrade Settings on page 164.

Prerequisites

- Operator or Admin access

Performing Upgrades using the Web UI

Use the wizard to upgrade the cluster with the latest system image and guest images.

IMPORTANT: Click the appropriate icon to suspend, resume, or cancel the upgrade.

Upgrading a Cluster

To upgrade a cluster:

1. Select **Appliances** > **Clusters**.
2. Select **Upgrade** from the **Actions** menu. The Upgrade Cluster wizard opens.
3. On the Select updates screen, select the **System Image**, **Guest Image**, or both checkboxes.
4. On the System Image Upgrade screen:
    a. Select the system image version from the drop-down list.
    b. Select or clear the **Reboot** check box, depending on whether you want to automatically reload the appliance after the update is done.
    c. Review the information on the screen.
    d. (Optional) Click **View Nodes** to see the nodes that will be upgraded.
5. On the Guest Image Upgrade screen:
    a. Review the new guest-images version and other information on the screen.
    b. (Optional) Click **View Nodes** to see the nodes that will be upgraded.
6. On the Summary screen, review the information and then click **Finish**.
7. Monitor the upgrade progress. Progress messages such as the one shown below are displayed on the **Clusters** page and on the **Nodes** page, where you can see which node is currently being upgraded.

Performing Upgrades Using the CLI

The upgrade procedure depends on whether you are upgrading the system image, guest images, or both, and whether you are upgrading the cluster or an individual node.

CAUTION: Do not upgrade a single node that is currently part of a cluster. Upgrade the entire cluster from the Central Management appliance instead.

- Upgrading the System Image and Guest Images Using the CLI on the next page
- Upgrading the System Image Using the CLI on page 152
- Downloading and Installing Guest Images Using the CLI on page 155
- Downloading Guest Images Using the CLI on page 156
- Installing Guest Images Using the CLI on page 157
- Suspending, Resuming, or Canceling an Upgrade on page 158

- Deleting a Guest-Images Download on page 161
- Monitoring Upgrade Status Using the CLI on page 161

Upgrading the System Image and Guest Images Using the CLI

Use the commands in this section to upgrade both the system image and guest images on an MVX cluster or node using the Central Management CLI.

Upgrading a Cluster

The following procedure shows how to upgrade both the system image and guest images on all nodes in a cluster.

To upgrade both the system image and guest images on a cluster:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Upgrade the cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name>
   ```

   where `<cluster name>` is the name of the cluster.

4. Monitor the status:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```

To upgrade to a specific version of the system and upgrade guest images on a cluster:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Upgrade the cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> version
   <version>
   ```

4. Monitor the status:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```


Upgrading a Node

The following procedure shows how to upgrade both the system image and guest images on a specific node.

To upgrade both the system image and guest images on a node:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Upgrade the node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name>
   ```

   where `<appliance name>` is the name of the Virtual Execution appliance.

4. Monitor the status:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance
   name>
   ```

Example

The following example shows a full upgrade on Cluster-Acme. In the example, the system image upgrade is complete on both nodes, and guest images are being downloaded to vx-2. The guest images download task is Task (07/10). This task will start on vx-1 after it is completed on vx-2.

```
cm-1 (config) # fenet update cluster Cluster-Acme
cluster update for Cluster-Acme:
success, update started
  Run 'show fenet update status cluster Cluster-Acme' for status
cm-1 (config) # show fenet update status cluster Cluster-Acme
Cluster Update Status:
  Cluster:                    Cluster-Acme
     Status:                  in-progress
     Current operation:       image-gi-update
     Current task:            gi-download
     Percent done:            35.07 %
     Start time:              2019/07/15 20:23:23.480
     End time:                ********

     Node:                    vx-2
        Status:               in-progress
        Percent done:         56.14 %
        Task (01/10):         image-check
           Status:            complete
           Percent done:      100.00 %
        Task (02/10):         gi-check
           Status:            complete
           Percent done:      100.00 %
        Task (03/10):         image-fetch
           Status:            complete
           Percent:           100.00 %
        Task (04/10):         image-install
           Status:            complete
```

```
      Percent done:       100.00 %
   Task (05/10):          image-rename
      Status:             complete
      Percent done:       100.00 %
   Task (06/10):          image-boot-next
      Status:             complete
      Percent done:       100.00 %
   Task (07/10):          gi-download
      Status:             in-progress
      Percent done:       56.19 %
Node:                     vx-1
   Status:                in-progress
   Percent done:          14.00 %
   Task (01/10):          image-check
      Status:             complete
      Percent done:       100.00 %
   Task (02/10):          gi-check
      Status:             complete
      Percent done:       100.00 %
   Task (03/10):          image-fetch
      Status:             complete
      Percent done:       100.00 %
   Task (04/10):          image-install
      Status:             complete
      Percent done:       100.00 %
   Task (05/10):          image-rename
      Status:             complete
      Percent done:       100.00 %
   Task (06/10):          image-boot-next
      Status:             complete
      Percent done:       100.00 %
```

Upgrading the System Image Using the CLI

Use the commands in this section to upgrade the system image on an MVX cluster or an individual node using the Central Management CLI.

NOTE: Use the `show fenet metadata status` command to see a list of the available versions of the system image.

Upgrading the System Image on a Cluster

The following procedure shows how to upgrade the Virtual Execution system image on all nodes in a cluster.

Install the system image:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. To upgrade to the latest version of the system image and then reboot the system:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> system-image
   ```

   or

   ```
   cm-hostname (config) # fenet update cluster <cluster name> system-image
   reboot
   ```

   where `<cluster name>` is the name of the cluster.

4. To upgrade to the latest version of the system image without rebooting the system:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> system-image
   no-reboot
   ```

5. To upgrade to a specific version of the system image:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> system-image
   version <version>
   ```

   where `<cluster name>` is the name of the cluster, and `<version>` is the system image version number.

6. Check the progress:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```

Upgrading the System Image on a Node

The following procedure shows how to upgrade the Virtual Execution system image on a specific node.

Install the system image:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. To install the latest version of the system image and then reboot the system:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> system
   image
   ```

   or

   ```
   cm-hostname (config) # fenet update appliance <appliance name> system
   image reboot
   ```

   where `<appliance name>` is the name of the Virtual Execution appliance.

4. To install the latest version of the system image without rebooting the system:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> system-
   image no-reboot
   ```

5. To install a specific version of the system image:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> system-image version <version>
   ```

   where `<appliance name>` is the name of the Virtual Execution appliance, and `<version>` is the system image version number.

6. To upgrade to a specific version of the system image and upgrade the guest images:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> version <version>
   ```

7. Check the progress:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance name>
   ```

Example

The following example shows a system image upgrade on Cluster-Acme. In the example, the image-check and image-fetch tasks are complete on both nodes, and the image-install task is in progress on vx-2.

```
cm-1 (config) # fenet update cluster Cluster-Acme  system-image
cluster update for Cluster-Acme:
update started: success
  Run 'show fenet update status cluster Cluster-Acme' for status
cm-1 (config) # show fenet update status cluster Cluster-Acme
Cluster Update Status:
  Cluster:                    Cluster-Acme
     Status:                  in-progress
     Current operation:       image-update
     Current task:            image-install
     Percent done:            19.28 %
     Start time:              2019/07/18 18:58:33.168
     End time:                ********

     Node:                    vx-2
        Status:               in-progress
        Percent done:         23.56 %
        Task (01/07):         image-check
           Status:            complete
           Percent done:      100.00 %
        Task (02/07):         image-fetch
           Status:            complete
           Percent:           100.00 %
        Task (03/07):         image-install
           Status:            in-progress
           Percent done:      42.82 %
     Node:                    vx-1
        Status:               in-progress
        Percent done:         15.00 %
        Task (01/07):         image-check
           Status:            complete
           Percent done:      100.00 %
        Task (02/07):         image-fetch
           Status:            complete
           Percent done:      100.00 %
```


Downloading and Installing Guest Images Using the CLI

Use the commands in this section to upgrade guest images on an MVX cluster or an individual node using the Central Management CLI.

NOTE: You can cancel, suspend, and resume a guest-images download. For details, see Suspending, Resuming, or Canceling an Upgrade on page 158.

Downloading and Installing Guest Images on a Cluster

The following procedure shows how to download and install guest images on all nodes in a cluster.

Download and install guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

2. Start the upgrade:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image
   ```

3. Check the progress:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```

Downloading and Installing Guest Images on a Node

The following procedure shows how to download and install guest images on a specific node.

Download and install guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

2. Start the upgrade:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-
   image
   ```

3. Check the progress:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance
   name>
   ```

Example

The following example downloads and installs guest images on the vx-1 node.

```
cm-1 (config) # fenet update appliance vx-1 guest-image
appliance update for vx-1
success, update started
  Run 'show fenet update status appliance vx-1' for status
cm-1 (config) # show fenet update status appliance vx-1
Appliance Update Status:
   Appliance:                      vx-1
      Status:                      complete
      Current operation:           gi-update
      Current task:                gi-install
      Percent done:                100.00 %
      Start time:                  2019/07/07 21:18:35.455
      End time:                    2019/07/07 22:18:49.335

      Node:                        vx-1
        Status:                    complete
        Percent done:              100.00 %
        Task (01/03):              gi-check
           Status:                 complete
           Percent done:           100.00 %
        Task (02/03):              gi-download
           Status:                 complete
           Percent done:           100.00 %
        Task (03/03):              gi-install
           Status:                 complete
           Percent done:           100.00 %
```

Downloading Guest Images Using the CLI

Use the commands in this section to download guest images to an MVX cluster or an individual node using the Central Management CLI.

NOTE: You can cancel, suspend, and resume a guest-images download. For details, see Suspending, Resuming, or Canceling an Upgrade on page 158.

Downloading Guest Images on a Cluster

The following procedure shows how to download guest images to all nodes in a cluster.

Download the guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

2. Start the upgrade:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image
   download
   ```

3. Check the progress:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```


Downloading Guest Images to a Node

The following procedure describes how to download guest images to a specific node.

Download the guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

2. Download the guest images:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-image download
   ```

3. Check the progress:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance name>
   ```

Example

The following example downloads guest images to all nodes of Cluster-Acme.

```
cm-1 (config) # fenet update cluster Cluster-Acme guest-image download
cluster update for Cluster-Acme
success, update started
  Run 'show fenet update status cluster Cluster-Acme' for status
```

Installing Guest Images Using the CLI

Use the commands in this section to install downloaded guest images on an MVX cluster or an individual node using the Central Management CLI.

Installing Guest Images on a Cluster

The following procedure shows how to install guest images on all nodes in a cluster.

Install guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname # configure terminal
   ```

2. Ensure that guest images were downloaded:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```

3. Install the guest images:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image install
   ```

4. Check the progress:

   ```
   cm-hostname (config) # show fenet update status cluster <cluster name>
   ```

Installing Guest Images on a Node

The following procedure shows how to install guest images on a specific node.

Install guest images:

1. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

2. Ensure that guest images were downloaded:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance
   name>
   ```

3. Install the guest images:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-
   image install
   ```

4. Check the progress:

   ```
   cm-hostname (config) # show fenet update status appliance <appliance
   name>
   ```

Example

The following example installs downloaded guest images on the vx-2 node.

```
cm-1 (config) # fenet update appliance vx-2 guest-image install
appliance update for vx-2
success, update started
  Run 'show fenet update status appliance vx-2' for status
```

Suspending, Resuming, or Canceling an Upgrade

This section describes operations you can perform during upgrades.

Suspending an Upgrade

The following procedure shows how to suspend an upgrade on a cluster or node.

To suspend an upgrade:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. To suspend the upgrade on a cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> suspend
   ```

4. To suspend the upgrade on a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> suspend
   ```

Resuming a Suspended Upgrade

The following procedure shows how to resume an upgrade that was suspended.

To resume a upgrade:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. To resume an upgrade on cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> resume
   ```

4. To resume an upgrade on a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> resume
   ```

To resume a guest-images upgrade:

1. To resume a guest-images upgrade on a cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image
   resume
   ```

2. To resume a guest-images upgrade on a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-
   image resume
   ```

Canceling an Upgrade

The following procedure shows how to cancel an upgrade operation.

To cancel an upgrade:

1. Log into the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. To cancel an upgrade on a cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> cancel
   ```

4. To cancel an upgrade on a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> cancel
   ```

To cancel a guest-images upgrade:

1. To cancel a guest-images upgrade on a cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image
   cancel
   ```

2. To cancel a guest-images upgrade on a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-
   image cancel
   ```

Example

The following example suspends the upgrade on the vx-2 node and then resumes it. In this example, the guest-images download operation was in progress.

```
cm-1 (config) # fenet update appliance vx-2 suspend
appliance suspend for vx-2:
success, operation initiated
  Run 'show fenet update status appliance vx-2' for status
cm-1 (config) # show fenet update status appliance vx-2
Appliance Update Status:
  Appliance:                vx-2
    Status:                 suspend
    Current operation:      gi-download
    Current task:           gi-download
    Percent done:           18.66 %
    Start time:             2019/07/15 22:44:52.071
    End time:               ********

    Node:                   vx-2
      Status:               in-progress
      Percent done:         18.66 %
      Task (01/02):         gi-check
        Status:             complete
        Percent done:       100 %
      Task (02/02):         gi-download
        Status:             in-progress
        Percent done:       14.38 %

cm-1 (config) # fenet update appliance vx-2 resume
appliance resume for vx-2:
success, operation initiated
  Run 'show fenet update status appliance vx-2' for status
cm-1 (config) # show fenet update status appliance vx-2
Appliance Update Status:
  Appliance:                vx-2
    Status:                 in-progress
    Current operation:      gi-download
    Current task:           gi-download
    Percent done:           19.58 %
    Start time:             2019/07/15 22:44:52.071
    End time:               ********

    Node:                   vx-2
      Status:               in-progress
      Percent done:         19.58 %
      Task (01/02):         gi-check
        Status:             complete
        Percent done:       100 %
      Task (02/02):         gi-download
        Status:             in-progress
        Percent done:       15.35 %
```


Deleting a Guest-Images Download

Use the commands in this section to delete downloaded guest images. For example, if an installation fails due to a bad download, you can delete the guest images to clean up the partial data.

To delete guest images:

1. Log in to the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   ```

3. To delete guest images from all nodes in a cluster:

   ```
   cm-hostname (config) # fenet update cluster <cluster name> guest-image delete
   ```

4. To delete guest images from a node:

   ```
   cm-hostname (config) # fenet update appliance <appliance name> guest-image delete
   ```

Example

The following example deletes the guest images from Cluster-Acme.

```
cm-hostname (config) # fenet update cluster Cluster-Acme guest-image delete
cluster update for Cluster-Acme:
update started: success
  Run 'show fenet update status cluster Cluster-Acme' for status
```

The following example deletes the guest images from Cluster-mssp_01.

```
cm-hostname (config) # fenet update cluster Cluster-mssp_01 guest-image delete
cluster update for Cluster-mssp_01:
update started: success
  Run 'show fenet update status cluster Cluster-mssp_01' for status
```

Monitoring Upgrade Status Using the CLI

Use the commands in this section to monitor the status of an upgrade.

Monitoring Cluster Upgrades

The following procedure shows how to monitor the status of a cluster upgrade.

To monitor a cluster upgrade:

1. Log in to the Central Management CLI.
2. To view standard status information:

   ```
   cm-hostname > show fenet update status cluster <cluster name>
   ```

   where `<cluster name>` is the name of the cluster.
3. To view brief status information:

   ```
   cm-hostname > show fenet update status cluster <cluster name> brief
   ```
4. To view detailed status information:

   ```
   cm-hostname > show fenet update status cluster <cluster name> detail
   ```

Monitoring Node Upgrades

The following procedure shows how to monitor the status of a node upgrade.

To monitor a node upgrade:

1. Log in to the Central Management CLI.
2. To view standard status information:

   ```
   cm-hostname > show fenet update status appliance <appliance name>
   ```

   where `<appliance name>` is the name of the node (the name displayed in the `show cmc appliances` CLI command output).
3. To view brief status information:

   ```
   cm-hostname > show fenet update status appliance <appliance name> brief
   ```
4. To view verbose status information:

   ```
   cm-hostname > show fenet update status appliance <appliance name>
   detail
   ```

Examples

The following example shows standard status information for the vx-1 node.

```
cm-1 # show fenet update status appliance vx-1
Appliance Update Status:
  Appliance:                    vx-1
    Status:                     complete
    Current operation:          gi-update
    Current task:               gi-install
    Percent done:               100 %
    Start time:                 2019/07/07 21:18:35:455
    End time:                   2019/07/07 22:18:49.355

    Node:                       vx-1
      Status:                   complete
      Percent done:             100 %
      Task (01/03):             gi-check
```

```
         Status:                 complete
         Percent done:           100 %
      Task (02/03):              gi-download
         Status:                 complete
         Percent done:           100 %
      Task (03/03):              gi-install
         Status:                 complete
         Percent done:           100 %
```

The following example shows brief information for the Cluster-Acme cluster.

```
cm-1 # show fenet update status cluster Cluster-Acme brief
Cluster Update Status:
Name                    Operation             Percent   Status
----                    ---------             -------   ------
Cluster-Acme            image-gi-update       39.04     in-progress

   Node                 Task
   ----                 ----
   vx-2                 07/10                 64.08     in-progress
   vx-1                 07/10                 14.00     in-progress
```

The following example shows detailed information for the Cluster-Acme cluster.

NOTE: The following two fields report the status as a percentage:

- **Percent Done**—The percentage of the update operation as a whole that was completed on the node or the cluster.
- **Percent Complete**—The percentage of the total number of tasks that were completed on the node or the cluster.

```
cm-1 # show fenet update status cluster Cluster-Acme detail
Cluster Update Status:
  Cluster:                    Cluster-Acme
    Status:                   in-progress
    Current operation:        image-gi-update
    Current task:             gi-download
    Percent done:             39.28 %
    Percent complete:         60.00 %
    Current num nodes:        2
    Total num nodes:          2
    Version:
    Last updated at:          2019/07/14 19:48:54:821
    Last updated op:          image-gi-update
    Start time:               2019/07/15 20:23:23.480
    End time:                 ********

    Node:                     vx-2
      Status:                 in-progress
      Percent done:           64.57 %
      Percent complete:       60.00 %
      Task (01/10):           image-check
        Status:               complete
        Percent done:         100 %
        Retry count:          0
        Return code:          0
        Start time:           2019/07/15 20:23:23.486
        End time:             2019/07/15 20:23:33:659
        Return message:
Operation initiated in the background.
  Run 'show fenet image status' for status
```

```
      Task (02/10):              gi-check
        Status:                  complete
        Percent done:            100.00 %
        Retry count:             0
        Return code:             0
        Start time:              2019/07/15 20:23:33.661
        End time:                2019/07/15 20:23:35.732
        Return message:
Downloading server manifest.

       Task (03/10):             image-fetch
       ...
```

Configuring and Viewing Upgrade Settings

You can change the following settings for each task:

- **Timeout**: Number of seconds before a task times out.
- **Max Retry**: Number of times the system will retry a task that failed.
- **Parallel Exec**: Whether the task starts on all nodes in the cluster at the same time (parallel execution), or starts on each node one at a time.

You can also view the sequence of tasks that are performed in each update operation.

Upgrade Task Settings

The following table shows the default settings for each task.

Task	Description	Timeout (Seconds)	Max Retry	Parallel Exec
gi-check	Check for newer version of guest images.	60	2	no
gi-download	Download guest images.	86400	2	no
gi-install	Install guest images.	600	2	no
image-boot-next	Boot the system from the next partition when the node reboots.	300	2	no
image-check	Check for newer version of system image.	60	2	yes
image-fetch	Download the system image.	600	2	no
image-install	Install the system image.	600	2	no

Task	Description	Timeout (Seconds)	Max Retry	Parallel Exec
image-prep-reboot	Prepare the system for reboot.	600	2	no
image-reboot	Reboot the node.	900	2	no

Prerequisites

- Monitor, Operator, or Admin access to view settings
- Admin access to change settings

Viewing Upgrade Settings Using the CLI

Use the commands in this section to view upgrade settings.

To view the upgrade settings:

1. Log in to the Central Management CLI.
2. View the settings:

   ```
   cm-hostname > show fenet update config
   ```

To view the update operations:

1. Log in to the Central Management CLI.
2. View the operations:

   ```
   cm-hostname > show fenet update operations
   ```

Examples

The following example shows the default upgrade settings.

```
cm-1 > show fenet update config
Update Config:
  Task:                gi-check
    Timeout:           60
    Max retry:         2
    Parallel exec:     no
  Task:                gi-download
    Timeout:           86400
    Max retry:         2
    Parallel exec:     no
  Task:                gi-install
    Timeout:           600
    Max retry:         2
    Parallel exec:     no
  Task:                image-boot-next
    Timeout:           300
```

```
    Max retry:          2
    Parallel exec:      no
  Task:                 image-check
    Timeout:            60
    Max retry:          2
    Parallel exec:      yes
  Task:                 image-fetch
    Timeout:            600
    Max retry:          2
    Parallel exec:      no
  Task:                 image-install
    Timeout:            600
    Max retry:          2
    Parallel exec:      no
  Task:                 image-prep-reboot
    Timeout:            600
    Max retry:          2
    Parallel exec:      no
  Task:                 image-reboot
    Timeout:            900
    Max retry:          2
    Parallel exec:      no
```

The following example shows two upgrade operations and the tasks included in them.

```
cm-hostname > show fenet update operations
Update Operations:
  Operation id:         1
  Operation name:       image-update
  Operation cli:        fenet update cluster <name> system-image
    Task:               image-check
    Task:               image-fetch
    Task:               image-install
    Task:               image-rename
    Task:               image-boot-next
    Task:               image-prep-reboot
    Task:               image-reboot
  Operation id:         2
  Operation name:       gi-check
  Operation cli:        fenet update cluster <name> guest-image
    Task:               gi-check
    Task:               gi-download
    Task:               gi-install
  ...
```

Configuring Upgrade Settings Using the CLI

Use the commands in this section to configure settings for upgrade operations.

To change the timeout setting:

1. Log into the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Change the setting:

   ```
   cm-hostname (config) # fenet update config task <task> timeout
   <seconds>
   ```

   where `<task>` is the name of the task shown in Upgrade Task Settings on page 164, and `<seconds>` is the number of seconds (1–86400) before the task times out.

4. Verify your change:

   ```
   cm-hostname (config) # show fenet update config
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

To change the max retry setting:

1. Log into the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. Change the setting:

   ```
   cm-hostname (config) # fenet update config task <task> retry <number>
   ```

   where `<task>` is the name of the task shown in Upgrade Task Settings on page 164, and `<number>` is the number of times (1–5) a failed task is tried again.

4. Verify your change:

   ```
   cm-hostname (config) # show fenet update config
   ```

5. Save your change:

   ```
   cm-hostname (config) # write memory
   ```

To change whether the task is executed on nodes in parallel:

1. Log into the Central Management CLI.
2. Enable the CLI configuration mode:

   ```
   cm-hostname > enable
   cm-hostname # configure terminal
   ```

3. To enable parallel execution:

   ```
   cm-hostname (config) # fenet update config task <task> parallel-
   execution
   ```

   where `<task>` is the name of the task shown in Upgrade Task Settings on page 164.

4. To disable parallel execution:

   ```
   cm-hostname (config) # no fenet update config task <task> parallel-
   execution
   ```

5. Verify your change:

   ```
   cm-hostname (config) # show fenet update config
   ```

6. Save your change:

```
cm-hostname (config) # write memory
```

To reset default settings:

1. Log into the Central Management CLI.
2. Enable the CLI configuration mode:

```
cm-hostname > enable
cm-hostname # configure terminal
```

3. To reset the timeout setting:

```
cm-hostname (config) # no fenet update config task <task> timeout
```

where <task> is the task name (see Upgrade Task Settings on page 164).

4. To reset the maximum retry setting:

```
cm-hostname (config) # no fenet update config task <task> retry
```

5. To reset the parallel execution setting:

```
cm-hostname (config) # no fenet update config task <task> parallel-
execution
```

6. Verify your changes:

```
cm-hostname (config) # show fenet update config
```

7. Save your changes:

```
cm-hostname (config) # write memory
```

Examples

The following example changes the image-check timeout to 45 seconds.

```
cm-1 (config) # fenet update config task image-check timeout 45
cm-1 (config) # show fenet update config
Update Config:
  ...
  Task:              image-check
    Timeout:         45
    Max retry:       2
    Parallel exec:   yes
  ...
```

The following example changes the number of retries for the gi-install task to 3.

```
cm-1 (config) # fenet update config task gi-install retry 3
cm-1 (config) # show fenet update config
Update Config:
  ...
  Task:              gi-check
    Timeout:         600
    Max retry:       3
    Parallel exec:   no
  ...
```

The following example enables parallel execution for the image-boot-next task.


```
cm-1 (config) # fenet update config task image-boot-next parallel-execution
cm-1 (config) # show fenet update config
Update Config:
  ...
  Task:              image-boot-next
    Timeout:         300
    Max retry:       2
    Parallel exec:   yes
    ...
```

The following example restores the default setting for the image-fetch task.

```
cm-1 (config) # no fenet update config task image-fetch retry
```

CHAPTER 24: Working with Notifications and Logs

This section contains the following topics:

Sending SNMP Traps

The MVX cluster components send Simple Network Management Protocol (SNMP) data to convey abnormal conditions to administrative computers that monitor and control them. The administrative computers are called *SNMP managers*.

If SNMP is configured on the component, by default the component pushes the following events (known as *traps*) to the SNMP manager.

Event	Description	CM	NX	EX	FX	VX
dupe-appliance-detected	The token server received a license token request from a virtual appliance with the same appliance ID and UUID as another virtual appliance.	✓	✓	✓	✓	✓
mvx-cluster-state-changed	The health status of an MVX cluster changed.	✓				

Event	Description	CM	NX	EX	FX	VX
mvx-cluster-util-threshold-exceeded	The cluster utilization reached the warning or critical level.	✓				
token-server-reachable	Communication between a virtual appliance and the token server was restored.	✓	✓	✓	✓	✓
token-server-unreachable	The virtual appliance cannot communicate with the token server.	✓	✓	✓	✓	✓
token-state-change	The state of the token changed to "active" or "inactive."	✓	✓	✓	✓	✓
mvx-enrollment-failure	There is a failure in the communication between a sensor or node and the enrollment service.		✓	✓	✓	✓

NOTE: For details about configuring SNMP and enabling and disabling events, see the *System Administration Guide* or *Administration Guide* for the appliance.

Configuring Email Event Notifications

If the MVX cluster components and are configured to send system email notifications, by default they send notifications about events related to MVX cluster status and enrollment and the licensing of virtual sensors.

The following events generate an email notification that is sent to recipients who are configured to receive notifications for "fail" level events.

- dupe-appliance-detected
- token-server-unreachable

The following events generate an email notification that is sent to recipients who are configured to receive notifications for "info" level events.

- mvx-cluster-state-changed
- mvx-cluster-util-threshold-exceeded
- token-server-reachable
- token-change-state
- mvx-enrollment-failure

For information about configuring recipients and determining the notifications they receive, see the *System Administration Guide* or *Administration Guide* for the appliance.

Configuring Alert Notifications

Sensors and hybrid appliances support the standard appliance alert notification methods. Configure alert notifications to be sent from the managing Central Management appliance, not from the individual sensors or hybrid appliances. For details, see the "Centralized Notifications" information in the *Central Management Administration Guide*.

Viewing Local Log Files

You can view local log files from the managing Central Management appliance or from individual sensors or hybrid appliances and nodes. The log files contain information about all system activity. MVX cluster-specific log messages are only a part of the log file. You can filter the log files for only these messages.

When you filter log files, the entries that match or do not match a specified regular expression are displayed. You can filter the output from the current log file, an archived log file, or all log files. You can also specify which new entries from a continuous log should be displayed. When you filter all log files, the output is presented in a single view. The entries are displayed in forward chronological order, based on the last modification date.

CAUTION: When you filter all log files, the results are not displayed until all entries in all log files are processed. A progress bar indicates the status of the processing. Because the volume can be high with this option, be careful with your choice of regular expression.

For information about managing local logs, see the log management sections in the *System Administration Guide* or *Administration Guide* for the appliance and the *CLI Command Reference*.

Prerequisites

- Auditor, Operator, or Admin access

Filtering Log Output Using the Web UI

Use the **Log Management** page of the sensor, hybrid appliance, or Central Management appliance to generate log files and download them. You can then search the log files for

MVX cluster-specific messages, as described in this section.

Generate a log file:

1. Log in to the appliance Web UI.
2. Click the **About** tab.
3. Click **Log Manager**.
4. Click **Selected logs**. All categories are selected by default.
5. Select or clear checkboxes to specify the categories you want to include in the logs.
6. If a drop-down list is present, select the time period the log should cover. The default is **Today**. The other options are **Past week**, **Past 2 weeks**, and **Past month**.
7. Clear the **Password-protect generated log archive** checkbox.

   

   IMPORTANT: If this checkbox is selected, you will be unable to open the files.

8. Click **Create**. A status message is displayed while the log is being created.

Download a log file to your local file system:

1. Locate the log file archive in the **Log Archives** section.
2. Click the icon in the **Action** column, and then select **Download**.

   The log archive is downloaded to your local file system. The archive name begins with the hostname of the appliance.

Filter the log output:

1. Locate the file and open it in a text editor or program of your choice.
2. Search for the keywords of interest (for example, *enrollment* or *token*).

Filtering Log Output Using the CLI

Use the commands in this section to use regular expressions to filter the log display to include only MVX cluster-specific information.

To filter the active log:

1. Enable the CLI enable mode:

   ```
   hostname > enable
   ```

2. To display only the lines that match a regular expression:

   ```
   hostname # show log matching <regular expression>
   ```

3. To display only the lines that do not match a regular expression:

   ```
   hostname # show log not matching <regular expression>
   ```

To filter an archived log:

1. Enable the CLI enable mode:

   ```
   hostname > enable
   ```

2. List the logs:

   ```
   hostname # show log files
   ```

3. To display only the lines in an archived log that match a regular expression:

   ```
   hostname # show log files <number> matching <regular expression>
   ```

   where `<number>` is the number of the file in the list of log files.

4. To display only the lines in an archived log that do not match a regular expression:

   ```
   hostname # show log files <number> not matching <regular expression>
   ```

To filter the new lines as they are added in real time:

1. Enable the CLI enable mode:

   ```
   hostname > enable
   ```

2. To start printing only the new lines in the active log that match a regular expression:

   ```
   hostname # show log continuous matching <regular expression>
   ```

3. To start printing only the new lines in the active log that do not match a regular expression

   ```
   hostname # show log continuous not matching <regular expression>
   ```

To filter all logs:

1. Enable the CLI enable mode:

   ```
   hostname > enable
   ```

2. To view only the lines in all logs that match a regular expression:

   ```
   hostname # show log files all matching <regular expression>
   ```

3. To view only the lines in all logs that do not match a regular expression:

   ```
   hostname # show log files all not matching <regular expression>
   ```


Sending Log Messages to a Remote Syslog Server

The Central Management appliance, sensors, hybrid appliances, and nodes can send log messages to one or more remote syslog servers. You can specify the format and level of detail, the delivery frequency, and the event severity level.

For details about configuring the syslog server and the logging settings, see the *User Guide* or *Administration Guide* for the appliance.

CHAPTER 25: Enrollment Maintenance Tasks

Use the following procedures when you need to perform maintenance or troubleshooting activities on sensor or node enrollment.

Prerequisites

- Operator or Admin access

Unsubscribing and Re-Enrolling a Sensor or Hybrid Appliance

Use the commands in this section if you need to unsubscribe a sensor or hybrid appliance and then re-enroll it.

NOTE: You can also use the **Select** > **Un-Enroll And Delete** option on the **Sensors** page of the managing Central Management appliance to unsubscribe a sensor or hybrid appliance when you delete it from the Central Management appliance.

To unsubscribe a sensor or hybrid appliance:

1. Unsubscribe the sensor or hybrid appliance:

   ```
   appl-hostname (config) # mvx cluster unenroll now
   ```

2. Verify your change:

   ```
   appl-hostname (config) # show mvx cluster enrollment status
   ```

3. Save your change:

```
appl-hostname (config) # write memory
```

NOTE: If you unsubscribe a sensor or hybrid appliance, it will not be automatically enrolled again, even though the status shows that automatic enrollment is enabled. Use the `mvx cluster enroll now` command to re-enroll a sensor or hybrid appliance that was unsubscribed using the `mvx cluster unenroll now` command.

To re-enroll a sensor or hybrid appliance:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI configuration mode:

```
appl-hostname > enable
appl-hostname # configure terminal
```

3. Re-enroll the sensor or hybrid appliance:

```
appl-hostname (config) # mvx cluster enroll now
```

4. Verify your change:

```
appl-hostname (config) # show mvx cluster enrollment status
```

5. Save your change:

```
appl-hostname (config) # write memory
```

Examples

The following example unsubscribes the nx-6 sensor.

```
nx-6 (config) # mvx cluster unenroll now
Operation initiated in the background.
  Run 'show mvx cluster enrollment status'
nx-6 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client :
    Status ok              : no
    Status description     : unenrolled
    Last checked at        : 2019/08/21 00:39:03

  Enrollment Service :
    Auto enabled           : yes
    Service address        : DTI (DTIUser@10.11.121.13 : singleport)
    Preferred cluster      : any (less loaded)
    Cloud enabled          : yes
    Cloud License enabled  : yes
```

After you reload the sensor, the status is as follows:

```
nx-6 (config) # show mvx cluster enrollment status

Enrollment Client :
    Status ok              : no
```

```
    Status description     : Enrollment Service Unavailable
    Last checked at        : 2019/08/21 00:48:30

  Enrollment Service :
    Auto enabled           : yes
    Service address        : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster      : any (less loaded)
```

The following example sends a request to enroll the nx-05 sensor.

```
nx-05 (config) # mvx cluster enroll now
Operation initiated in the background.
  Run 'show mvx cluster enrollment status'
nx-05 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client :
     Status ok               : yes
     Status description      : enrolled
     Last checked at:        :2019/08/18 20:55:41
  ...
```

Restoring Automatic Sensor Enrollment

The enrollment service on a sensor or hybrid appliance must be enabled so it can submit objects to the cluster, and is enabled by default. Use the procedures in this section if you need to restore automatic enrollment after it is disabled.

NOTE: The `no mvx cluster enrollment-service client enable` command disables automatic enrollment. FireEye recommends that you use this command only with guidance from FireEye Technical Support.

To restore automatic enrollment:

1. Log in to the sensor or hybrid appliance CLI.
2. Enable the CLI configuration mode:

   ```
   appl-hostname > enable
   appl-hostname # configure terminal
   ```

3. Restore automatic enrollment:

   ```
   appl-hostname (config) # mvx cluster enrollment-service client enable
   ```

4. Verify your change:

   ```
   appl-hostname (config) # show mvx cluster enrollment status
   ```

5. Save your change:

   ```
   appl-hostname (config) # write memory
   ```

Example

The following example restores automatic enrollment on the nx-1 sensor.

```
nx-1 (config) # mvx cluster enrollment-service client enable
nx-1 (config) # show mvx cluster enrollment status

MVX Cluster Enrollment Status

  Enrollment Client :
    Status ok               : yes
    Status description      : enrolled
    Last checked at         : 2019/08/25 18:05:25

  Enrollment Service :
    Auto enabled            : yes
    Service address         : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster       : any (less loaded)
    Cloud enabled           : no
    Cloud License enabled   : no
    Connect on demand       : no

  Broker Info :
    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1
    Broker ID               : 002XXXXXXXXX
    Broker Address          : 10.11.121.12
    Broker State            : Connected
    Failure Reason          : None
    Last Connection Attempt : 2019/08/21 00:15:46
    Connection Last Formed  : 2019/08/21 00:15:47
    Connection Last Broken  :
```

Restoring Automatic Node Enrollment

The enrollment service on a node must be enabled so it can receive submissions from the sensor or hybrid appliance, and is enabled by default. Use the procedures in this section to re-enable the enrollment service after it is disabled.

NOTE: The `no mvx cluster enrollment-service client enable` command disables automatic enrollment. FireEye recommends that you use this command only with guidance from FireEye Technical Support.

To enable the enrollment service:

1. Log in to the Virtual Execution CLI.
2. Enable the CLI configuration mode:

   ```
   vx-hostname > enable
   vx-hostname # configure terminal
   ```

3. Enable the enrollment service:

   ```
   vx-hostname (config) # mvx cluster enrollment-service client enable
   ```

4. Verify your change:

   ```
   vx-hostname (config) # show mvx cluster enrollment status
   ```

5. Save your change:

   ```
   vx-hostname (config) # write memory
   ```

Example

The following example re-enables the enrollment service on the vx-2 node.

```
vx-2 (config) # mvx cluster enrollment-service client enable
vx-2 (config) # show mvx cluster enrollment status

  Enrollment Client :
    Status ok               : yes
    Status description      : Ok
    Last checked at         : 2019/07/21 00:34:01

  Enrollment Service :
    Auto enabled            : yes
    Service address         : CMS (DTIUser@10.11.121.13 : singleport)
    Preferred cluster       :

    Cluster Name            : Cluster-Acme
    Broker Name             : vx-1 (self)
    Broker ID               : 002XXXXXXXXX (self)
    Broker Address          : 10.11.121.12 (self)
    Broker State            : N/A
    Failure Reason          : N/A
    Last Connection Attempt : N/A
    Connection Last Formed  : N/A
    Connection Last Broken  : N/A
```


CHAPTER 26: Troubleshooting

This section shows how to handle problems you might encounter in your MVX Smart Grid deployment.

Sensor, Hybrid Appliance, or Compute Node Cannot Connect to Broker

Problem

A sensor, hybrid appliance, or compute node is unable to establish a connection with a broker.

Reason

The following are probable reasons for the problem:

- The sensor, hybrid appliance, or broker is unable to communicate with the enrollment service.
- The broker IP address is not configured correctly.
- The broker IP address does not match the IP address configured for the submission interface or cluster interface on the broker.
- Network settings are not configured correctly.
- Network deployment is preventing connectivity.
- A reason identified in the `show mvx node status full` command output. For a list of error codes, see CCD Ok on page 117.

Action

To correct the problem:

- Verify that broker information is displayed in the `show mvx cluster enrollment status` command output. If no broker information is displayed, troubleshoot the enrollment service.

- Verify that the broker node IP address is correct. If the IP address is incorrect, troubleshoot the Central Management connection or the enrollment service.
- If the sensor or hybrid appliance cannot connect to the broker node, verify that the IP address configured for the broker node matches the IP address configured for the submission interface on the broker node.
- If a compute node cannot connect to a broker node, verify that the broker node IP address matches the IP address configured for the cluster interface on the broker node.
- Try to ping the broker IP address:

  If a compute node or another broker node cannot reach the broker node:

  ```
  ping -I <cluster interface> <broker address>
  ```

  If a sensor or hybrid appliance cannot reach the broker node:

  ```
  ping -I <submission interface> <broker address>
  ```

  where

  - `<cluster interface>` is the name of the broker cluster interface.
  - `<submission interface>` is the name of the broker submission interface.
  - `<broker address>` is the IP address of the broker node.

  If the interface cannot be reached, fix network connectivity and network settings.
- If the interface can be reached, but the connection cannot be established (because the connection timed out or was refused), make sure that TCP port 22 is not blocked by a firewall.
- If the connection still cannot be established, the problem could be authentication failure or internal error. Examine the system logs and contact FireEye Technical Support if you need assistance.

Cluster is Unhealthy or Malformed

Problem

The nodes cannot authenticate with each other, there is a problem with a cluster process or component, or a sensor or broker is not connected.

Reason

The following are probable reasons for the problem:

- The authentication key hash is not the same on all nodes.
- The broker role is disabled.

- Cluster processes are unhealthy or not running.
- Guest images are not installed on the nodes.

Action

To correct the problem:

- Use the `show mvx node status full` command on a node to determine the specific issue and view a list of connected sensors, hybrid appliances, and brokers.
- If the `Key Hash` value is not the same on all nodes, remove the nodes from the cluster and then create the cluster again.
- If the node should be a broker, but the broker role is disabled, use the `cmc mvx cluster <cluster name> broker <node name> enable` command in the Central Management CLI.
- If guest images are not installed, install them.
- If the Web services API (WSAPI) process is not running, use the `wsapi enable` command in the Virtual Execution (node) CLI.

Technical Support

For technical support, contact FireEye through the Support portal:

https://csportal.fireeye.com

Documentation

Documentation for all FireEye products is available on the FireEye Documentation Portal (login required):

https://docs.fireeye.com/

FireEye, Inc. | 601 McCarthy Blvd. | Milpitas, CA | 1.408.321.6300 | 1.877.FIREEYE | www.fireeye.com