![CISCO]

Cisco Firepower 4100/9300 FXOS Firepower Chassis Manager Configuration Guide, 2.10(1)

First Published: 2021-05-26

Last Modified: 2023-01-19

CONTENTS

CHAPTER 4   **User Management** **39**

CHAPTER 10   **Logical Devices**   **183**

Introduction to the Security Appliance

About the Firepower Security Appliance

The Cisco Firepower 4100/9300 chassis is a next-generation platform for network and content security solutions. The Firepower 4100/9300 chassis is part of the Cisco Application Centric Infrastructure (ACI) Security Solution and provides an agile, open, secure platform that is built for scalability, consistent control, and simplified management.

The Firepower 4100/9300 chassis provides the following features:

- Modular chassis-based security system—provides high performance, flexible input/output configurations, and scalability.

- Firepower Chassis Manager—graphical user interface provides streamlined, visual representation of current chassis status and simplified configuration of chassis features.

- Firepower eXtensible Operating System (FXOS) CLI—provides command-based interface for configuring features, monitoring chassis status, and accessing advanced troubleshooting features.

- FXOS REST API—allows users to programmatically configure and manage their chassis.

How the Logical Device Works with the Firepower 4100/9300

The Firepower 4100/9300 runs its own operating system on the supervisor called the Firepower eXtensible Operating System (FXOS). The on-the-box Firepower Chassis Manager provides simple, GUI-based management capabilities. You configure hardware interface settings, smart licensing (for the ASA), and other basic operating parameters on the supervisor using the Firepower Chassis Manager.

A logical device lets you run one application instance and also one optional decorator application to form a service chain. When you deploy the logical device, the supervisor downloads an application image of your choice and establishes a default configuration. You can then configure the security policy within the application operating system.

Logical devices cannot form a service chain with each other, and they cannot communicate over the backplane with each other. All traffic must exit the chassis on one interface and return on another interface to reach

another logical device. For container instances, you can share data interfaces; only in this case can multiple logical devices communicate over the backplane.

Supported Applications

You can deploy logical devices on your chassis using the following application types.

FTD

The FTD provides next-generation firewall services, including stateful firewalling, routing, VPN, Next-Generation Intrusion Prevention System (NGIPS), Application Visibility and Control (AVC), URL filtering, and malware defense.

You can manage the FTD using one of the following managers:

- FMC—A full-featured, multidevice manager on a separate server.

- FDM—A simplified, single device manager included on the device.

- CDO—A cloud-based, multidevice manager.

ASA

The ASA provides advanced stateful firewall and VPN concentrator functionality in one device. You can manage the ASA using one of the following managers:

- ASDM—A single device manager included on the device.

- CLI

- CDO—A cloud-based, multidevice manager.

- CSM—A multidevice manager on a separate server.

Radware DefensePro (Decorator)

You can install Radware DefensePro (vDP) to run in front of the ASA or the FTD as a decorator application. vDP is a KVM-based virtual platform that provides distributed denial-of-service (DDoS) detection and mitigation capabilities on the Firepower 4100/9300. Traffic from the network must first pass through the vDP before reaching the ASA or the FTD.

Firepower Chassis Manager Overview

The FXOS provides a web interface that makes it easy to configure platform settings and interfaces, provision devices, and monitor system status. The navigation bar at the top of the user interface provides access to the following:

- Overview—From the Overview page you can easily monitor the status of the chassis. For more information, see Monitoring the Chassis Status, on page 3.

- Interfaces—From the Interfaces page, you can view the status of the installed interfaces on the chassis, edit interface properties, enable or disable an interface, and create port channels. For more information, see Interface Management, on page 151.

- Logical Devices—From the Logical Devices page, you can create, edit, and delete logical devices. You can also view the current status of existing logical devices. For more information, see Logical Devices, on page 183.

- Security Modules/Security Engine—From the Security Modules/Security Engine page, you can view the status of and can perform various functions on a security module/engine, such as power cycling, reinitializing, acknowledging, and decommissioning. For more information, see Security Module/Engine Management, on page 271.

- Platform Settings—From the Platform Settings page, you can configure chassis settings for the following: date and time, SSH, SNMP, HTTPS, AAA, Syslog, and DNS. For more information, see Platform Settings, on page 101.

- System Settings—From the System menu, you can manage the following settings:

  - Licensing—From the Licensing page, you can configure Smart Call Home settings and register your chassis with the Licensing Authority. For more information, see License Management for the ASA, on page 17.

  - Updates—From the Updates page, you can upload Platform Bundle and Application images to the chassis. For more information, see Image Management, on page 55.

  - User Management—From the User Management page you can configure user settings and define user accounts for the Firepower 4100/9300 chassis. For more information, see User Management, on page 39.

Monitoring the Chassis Status

From the Overview page you can easily monitor the status of the Firepower 4100/9300 chassis. The Overview page provides the following elements:

- Device Information—The top of the Overview page contains the following information about the Firepower 4100/9300 chassis:

  - Chassis name—shows the name assigned to the chassis during initial configuration.

  - IP address—shows the management IP address assigned to the chassis during initial configuration.

  - Model—shows the Firepower 4100/9300 chassis model.

  - Version—shows the FXOS version running on the chassis.

  - Operational State—shows the operable status for the chassis.

  - Chassis uptime—shows the elapsed time since the system was last restarted.

  - Shutdown button—gracefully shuts down the Firepower 4100/9300 chassis (see Powering Off the Firepower 4100/9300 Chassis, on page 98).

Note You can power off/on a security module/engine from the Security Modules/Security Engine page (see Power-Cycling a Security Module/Engine, on page 274).

- Reboot button—gracefully shuts down the Firepower 4100/9300 chassis (see Rebooting the Firepower 4100/9300 Chassis, on page 98).

- Uptime Information Icon—hover over the icon to see uptime for the chassis and for any installed security module/engine.

- Visual Status Display—Below the Device Information section is a visual representation of the chassis that shows the components that are installed in the chassis and provides a general status for those components. You can hover over the ports that are shown in the Visual Status Display to get additional information such as interface name, speed, type, admin state, and operational state. For models with multiple security modules, you can hover over the security modules that are shown in the Visual Status Display to get additional information such as device name, template type, admin state, and operational state. If a logical device is installed on that security module, you can also see the management IP address, software version, and logical device mode.

- Detailed Status Information—Below the Visual Status Display is a table containing detailed status information for the chassis. The status information is broken up into five sections: Faults, Interfaces, Devices, License, and Inventory. You can see a summary for each of those sections above the table and you can see additional details for each of those sections by clicking on the summary area for the information you want to view.

  The system provides the following detailed status information for the chassis:

  - Faults—Lists the faults that have been generated in the system. The faults are sorted by severity: Critical, Major, Minor, Warning, and Info. For each fault that is listed, you can see the severity, a description of the fault, the cause, the number of occurrences, and the time of the most recent occurrence. You can also see whether the fault has been acknowledged or not.

    You can click on any of the faults to see additional details for the fault or to acknowledge the fault. To acknowledge multiple faults, click the check box next to each fault you want to acknowledge and then click **Acknowledge**. You can use the **Select All Faults** and **Cancel Selected Faults** buttons to quickly select or deselect multiple faults.

> **Note** Once the underlying cause of the fault has been addressed, the fault will automatically be cleared from the listing during the next polling interval. If a user is working on a resolution for a specific fault, they can acknowledge the fault to let other users know that the fault is currently being addressed.

  - Interfaces—Lists the interfaces installed in the system. The **All Interfaces** tab shows the interface name, operational status, administrative status, number of received bytes, and number of transmitted bytes. The **Hardware Bypass** tab shows only interface pairs that are supported for the Hardware Bypass feature on the FTD application. For each pair, the operational state is shown: disabled (Hardware Bypass is not configured for the pair), standby (Hardware Bypass is configured, but not currently active), and bypass (actively in Hardware Bypass).

  - Instances—Lists the logical devices configured in the system and provides the following details for each logical device (hover your cursor over the bar): device name, status, image version, management IP address, and number of cores. You can also view the Ingress VLAN Group Entry Utilisation and Switch Forwarding Path Entry Utilisation at the bottom of the page.

  - License—(For ASA logical devices) Shows whether smart licensing is enabled, provides the current registration status of your license, and shows license authorization information for the chassis.

- Inventory—Lists the components installed in the chassis and provides relevant details for those components, such as: component name, number of cores, installation location, operational status, operability, capacity, power, thermal, serial number, model number, part number, and vendor.

Note   If power redundancy is implemented, do not change any settings related to power redundancy in FXOS.

Getting Started

Task Flow

The following procedure shows the basic tasks that should be completed when configuring your Firepower 4100/9300 chassis.

Procedure

Step 1	Configure the Firepower 4100/9300 chassis hardware (see the *Cisco Firepower Security Appliance Hardware Installation Guide*).
Step 2	Complete the initial configuration (see Initial Configuration, on page 7).
Step 3	Log in to the Firepower Chassis Manager (see Log In or Out of the Firepower Chassis Manager, on page 14).
Step 4	Set the Date and Time (see Setting the Date and Time, on page 101).
Step 5	Configure a DNS server (see Configuring DNS Servers, on page 143).
Step 6	Register your product license (see License Management for the ASA, on page 17).
Step 7	Configure users (see User Management, on page 39).
Step 8	Perform software updates as required (see Image Management, on page 55).
Step 9	Configure additional platform settings (see Platform Settings, on page 101).
Step 10	Configure interfaces (see Interface Management, on page 151).
Step 11	Create logical devices (see Logical Devices, on page 183).

Initial Configuration

Before you can use Firepower Chassis Manager or the FXOS CLI to configure and manage your system, you must perform some initial configuration tasks. You can perform the initial configuration using the FXOS CLI

accessed through the console port or using SSH, HTTPS, or REST API accessed through the management port (this procedure is also referred to as low-touch provisioning).

Initial Configuration Using Console Port

The first time that you access the Firepower 4100/9300 chassis using the FXOS CLI, you will encounter a setup wizard that you can use to configure the system.

✎

Note   To repeat the initial setup, you need to erase any existing configuration using the following commands:

```
Firepower-chassis# connect local-mgmt
firepower-chassis(local-mgmt)# erase configuration
```

You must specify only one IPv4 address, gateway, and subnet mask, or only one IPv6 address, gateway, and network prefix for the single management port on the Firepower 4100/9300 chassis. You can configure either an IPv4 or an IPv6 address for the management port IP address.

Before you begin

1. Verify the following physical connections on the Firepower 4100/9300 chassis:

   - The console port is physically connected to a computer terminal or console server.

   - The 1 Gbps Ethernet management port is connected to an external hub, switch, or router.

   For more information, refer to the hardware installation guide.

2. Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:

   - 9600 baud

   - 8 data bits

   - No parity

   - 1 stop bit

3. Gather the following information for use with the setup script:

   - New admin password

   - Management IP address and subnet mask

   - Gateway IP address

   - Subnets from which you want to allow HTTPS and SSH access

   - Hostname and domain name

   - DNS server IP address

Procedure

Step 1   Power on the chassis.

Step 2   Connect to the serial console port using a terminal emulator.

The Firepower 4100/9300 includes an RS-232–to–RJ-45 serial console cable. You might need to use a third party serial-to-USB cable to make the connection. Use the following serial parameters:

- 9600 baud

- 8 data bits

- No parity

- 1 stop bit

Step 3   Complete the system configuration as prompted.

Note   You can optionally enter the debug menu at any time during initial configuration to debug any setup issues or abort configurations and reboot the system. To enter the debug menu, press Ctrl-C. To exit the debug menu, press Ctrl-D twice. Note that anything you type in the interim between pressing Ctrl-D the first time and pressing it a second time will run after the second time Ctrl-D is pressed.

Example:

```
            ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Only minimal configuration including IP connectivity to
the FXOS Supervisor is performed through these steps.

Type Ctrl-C at any time for more options or to abort configuration
and reboot system.
To back track or make modifications to already entered values,
complete input till end of section and answer no when prompted
to apply configuration.


You have chosen to setup a new Security Appliance.
     Continue? (yes/no): y

Enforce strong password? (yes/no) [y]: n

Enter the password for "admin": Farscape&32
Confirm the password for "admin": Farscape&32
Enter the system name:  firepower-9300

Supervisor Mgmt IP address : 10.80.6.12

Supervisor Mgmt IPv4 netmask : 255.255.255.0

IPv4 address of the default gateway : 10.80.6.1

The system cannot be accessed via SSH if SSH Mgmt Access is not configured.

Do you want to configure SSH Mgmt Access? (yes/no) [y]: y

SSH Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0
```

```
        SSH Mgmt Access IPv4 netmask: 255.0.0.0

        Firepower Chassis Manager cannot be accessed if HTTPS Mgmt Access is not configured.

        Do you want to configure HTTPS Mgmt Access? (yes/no) [y]: y

        HTTPS Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0

        HTTPS Mgmt Access IPv4 netmask: 255.0.0.0

        Configure the DNS Server IP address? (yes/no) [n]: y

          DNS IP address : 10.164.47.13

        Configure the default domain name? (yes/no) [n]: y

          Default domain name : cisco.com

        Following configurations will be applied:

          Switch Fabric=A
          System Name=firepower-9300
          Enforced Strong Password=no
          Supervisor Mgmt IP Address=10.89.5.14
          Supervisor Mgmt IP Netmask=255.255.255.192
          Default Gateway=10.89.5.1
          SSH Access Configured=yes
              SSH IP Address=10.0.0.0
              SSH IP Netmask=255.0.0.0
          HTTPS Access Configured=yes
              HTTPS IP Address=10.0.0.0
              HTTPS IP Netmask=255.0.0.0
          DNS Server=72.163.47.11
          Domain Name=cisco.com

        Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): y
        Applying configuration. Please wait...  Configuration file – Ok
.....

Cisco FPR Series Security Appliance
firepower-9300 login:  admin
Password: Farscape&32
Successful login attempts for user 'admin' : 1
Cisco Firepower Extensible Operating System (FX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2009-2019, Cisco Systems, Inc. All rights reserved.

[...]

firepower-chassis#
```

Low-Touch Provisioning Using Management Port

When your Firepower 4100/9300 chassis boots up, if it does not find the startup configuration, the device enters the Low-Touch Provisioning mode in which the device locates a Dynamic Host Control Protocol (DHCP) server and then bootstraps itself with its management interface IP address. You can then connect through the management interface to configure the system using SSH, HTTPS, or the FXOS REST API.

> **Note**   To repeat the initial setup, you need to erase any existing configuration using the following commands:
>
> ```
> Firepower-chassis# connect local-mgmt
> firepower-chassis(local-mgmt)# erase configuration
> ```

You must specify only one IPv4 address, gateway, and subnet mask, or only one IPv6 address, gateway, and network prefix for the single management port on the Firepower 4100/9300 chassis. You can configure either an IPv4 or an IPv6 address for the management port IP address.

Before you begin

Gather the following information for use with the setup script:

- New admin password
- Management IP address and subnet mask
- Gateway IP address
- Subnets from which you want to allow HTTPS and SSH access
- Hostname and domain name
- DNS server IP address

Procedure

Step 1   Configure your DHCP server to assign an IP address to management port of the Firepower 4100/9300 chassis.

The DHCP client request from the Firepower 4100/9300 chassis will contain the following:

- The management interface's MAC address.
- DHCP option 60 (vendor-class-identifier)—Set to "FPR9300" or "FPR4100".
- DHCP option 61 (dhcp-client-identifier)—Set to the Firepower 4100/9300 chassis serial number. This serial number can be found on a pull-out tab on the chassis.

Step 2   Power on the Firepower 4100/9300 chassis.
If the startup configuration is not found when the chassis boots up, the device enters the Low-Touch Provisioning mode.

Step 3   To configure your system using HTTPS:

a) Using a supported browser, enter the following URL in the address bar:

https://<ip_address>/api

where *<ip_address>* is the IP address of the management port on the Firepower 4100/9300 chassis that was assigned by your DHCP server.

> **Note**   For information on supported browsers, refer to the release notes for the version you are using (see http://www.cisco.com/c/en/us/support/security/firepower-9000-series/products-release-notes-list.html).

b) When prompted, log in with the username **install** and the password *<chassis_serial_number>*.

The *<chassis_serial_number>* can be obtained by inspecting a tag on the chassis.

c) Complete the system configuration as prompted.

- Strong password enforcement policy (for strong password guidelines, see User Accounts, on page 39) .

- Password for the admin account.

- System name

- Supervisor Management IPv4 address and subnet mask, or IPv6 address and prefix.

- Default gateway IPv4 or IPv6 address.

- Host/network address and netmask/prefix from which SSH access is allowed.

- Host/network address and netmask/prefix from which HTTPS access is allowed.

- DNS Server IPv4 or IPv6 address.

- Default domain name.

d) Click **Submit**.

Step 4 To configure your system using SSH:

a) Connect to the management port using the following command:

ssh install@*<ip_address>*

where *<ip_address>* is the IP address of the management port on the Firepower 4100/9300 chassis that was assigned by your DHCP server.

b) When prompted, log in with the password **Admin123**.

c) Complete the system configuration as prompted.

Note You can optionally enter the debug menu at any time during initial configuration to debug any setup issues or abort configurations and reboot the system. To enter the debug menu, press Ctrl-C. To exit the debug menu, press Ctrl-D twice. Note that anything you type in the interim between pressing Ctrl-D the first time and pressing it a second time will run after the second time Ctrl-D is pressed.

Example:

```
        ---- Basic System Configuration Dialog ----

This setup utility will guide you through the basic configuration of
the system. Only minimal configuration including IP connectivity to
the FXOS Supervisor is performed through these steps.

Type Ctrl-C at any time for more options or to abort configuration
and reboot system.
To back track or make modifications to already entered values,
complete input till end of section and answer no when prompted
to apply configuration.


You have chosen to setup a new Security Appliance.
```

```
        Continue? (yes/no): y

  Enforce strong password? (yes/no) [y]: n

  Enter the password for "admin": Farscape&32
  Confirm the password for "admin": Farscape&32
  Enter the system name:  firepower-9300

  Supervisor Mgmt IP address : 10.80.6.12

  Supervisor Mgmt IPv4 netmask : 255.255.255.0

  IPv4 address of the default gateway : 10.80.6.1

  The system cannot be accessed via SSH if SSH Mgmt Access is not configured.

  Do you want to configure SSH Mgmt Access? (yes/no) [y]: y

  SSH Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0

  SSH Mgmt Access IPv4 netmask: 255.0.0.0

  Firepower Chassis Manager cannot be accessed if HTTPS Mgmt Access is not configured.

  Do you want to configure HTTPS Mgmt Access? (yes/no) [y]: y

  HTTPS Mgmt Access host/network address (IPv4/IPv6): 10.0.0.0

  HTTPS Mgmt Access IPv4 netmask: 255.0.0.0

  Configure the DNS Server IP address? (yes/no) [n]: y

    DNS IP address : 10.164.47.13

  Configure the default domain name? (yes/no) [n]: y

    Default domain name : cisco.com

  Following configurations will be applied:

    Switch Fabric=A
    System Name=firepower-9300
    Enforced Strong Password=no
    Supervisor Mgmt IP Address=10.89.5.14
    Supervisor Mgmt IP Netmask=255.255.255.192
    Default Gateway=10.89.5.1
    SSH Access Configured=yes
        SSH IP Address=10.0.0.0
        SSH IP Netmask=255.0.0.0
    HTTPS Access Configured=yes
        HTTPS IP Address=10.0.0.0
        HTTPS IP Netmask=255.0.0.0
    DNS Server=72.163.47.11
    Domain Name=cisco.com

  Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): y
  Applying configuration. Please wait...  Configuration file - Ok
.....

Initial Setup complete, Terminating sessions
.Connection to <ip_address> closed.
```

Step 5   To configure your system using the FXOS REST API:

Use the following examples for configuring the system using the REST API. For more information, see https://developer.cisco.com/site/ssp/firepower/.

Note The attributes dns, domain_name, https_net, https_mask, ssh_net, and ssh_mask are optional. All other attributes are mandatory for REST API configuration.

```
IPv4 REST API example:

{
    "fxosBootstrap": {
        "dns": "1.1.1.1",
        "domain_name": "cisco.com",
        "mgmt_gw": "192.168.0.1",
        "mgmt_ip": "192.168.93.3",
        "mgmt_mask": "255.255.0.0",
        "password1": "admin123",
        "password2": "admin123",
        "strong_password": "yes",
        "system_name": "firepower-9300",
        "https_mask": "2",
        "https_net": "::",
        "ssh_mask": "0",
        "ssh_net": "::"
    }
}


IPV6 REST API example

{
    "fxosBootstrap": {
        "dns": "2001::3434:4343",
        "domain_name": "cisco.com",
        "https_mask": "2",
        "https_net": "::",
        "mgmt_gw": "2001::1",
        "mgmt_ip": "2001::2001",
        "mgmt_mask": "64",
        "password1": "admin123",
        "password2": "admin123",
        "ssh_mask": "0",
        "ssh_net": "::",
        "strong_password": "yes",
        "system_name": "firepower-9300"
    }
}
```

Log In or Out of the Firepower Chassis Manager

Before you can configure your Firepower 4100/9300 chassis using Firepower Chassis Manager, you must log in using a valid user account. For more information on user accounts, see User Management, on page 39.

You are automatically logged out of the system if a certain period of time passes without any activity. By default, the system will log you out after 10 minutes of inactivity. To configure this timeout setting, see Configuring the Session Timeout, on page 48. You can also configure an absolute timeout setting that will

log users out of the system after a certain period of time even if the session is active. To configure the absolute timeout setting, see Configuring the Absolute Session Timeout, on page 49.

For a list of all system changes that cause you to be automatically logged out of Firepower Chassis Manager, see System Changes that Cause Firepower Chassis Manager Sessions to be Closed, on page 81.

Note You can optionally configure your Firepower Chassis Manager to allow only a certain number of unsuccessful login attempts before the user is locked out of the system for a specified amount of time. For more information, see Set the Maximum Number of Login Attempts, on page 50.

Procedure

Step 1 To log in to the Firepower Chassis Manager:

a) Using a supported browser, enter the following URL in the address bar:

`https://<chassis_mgmt_ip_address>`

where *<chassis_mgmt_ip_address>* is the IP address or host name of the Firepower 4100/9300 chassis that you entered during initial configuration.

Note For information on supported browsers, refer to the release notes for the version you are using (see http://www.cisco.com/c/en/us/support/security/firepower-9000-series/products-release-notes-list.html).

b) Enter your username and password.

c) Click **Login**.

You are logged in and the Firepower Chassis Manager opens to show the Overview page.

Step 2 To log out of the Firepower Chassis Manager, point at your username in the navigation bar and then select **Logout**.

You are logged out of the Firepower Chassis Manager and are returned to the login screen.

Accessing the FXOS CLI

You can connect to the FXOS CLI using a terminal plugged into the console port. Verify that the console port parameters on the computer terminal (or console server) attached to the console port are as follows:

• 9600 baud

• 8 data bits

• No parity

• 1 stop bit

You can also connect to the FXOS CLI using SSH and Telnet. The FXOS supports up to eight simultaneous SSH connections. To connect with SSH, you need to know the hostname or IP address of the Firepower 4100/9300 chassis.

Use one of the following syntax examples to log in with SSH, Telnet, or Putty:

> ✎
>
> **Note** SSH log in is case-sensitive.

From a Linux terminal using SSH:

- **ssh ucs-***auth-domain**username*@{*UCSM-ip-address|UCMS-ipv6-address*}

  ssh ucs-example\\jsmith@192.0.20.11

  ssh ucs-example\\jsmith@2001::1

- **ssh -l ucs-***auth-domain**username* {*UCSM-ip-address| UCSM-ipv6-address| UCSM-host-name*}

  ssh -l ucs-example\\jsmith 192.0.20.11

  ssh -l ucs-example\\jsmith 2001::1

- **ssh** {*UCSM-ip-address | UCSM-ipv6-address | UCSM-host-name*} **-l ucs-***auth-domain**username*

  ssh 192.0.20.11 -l ucs-example\\jsmith

  ssh 2001::1 -l ucs-example\\jsmith

- **ssh ucs-***auth-domain**username*@{*UCSM-ip-address|UCSM-ipv6-address*}

  ssh ucs-ldap23\\jsmith@192.0.20.11

  ssh ucs-ldap23\\jsmith@2001::1

From a Linux terminal using Telnet:

> ✎
>
> **Note** Telnet is disabled by default. See Configuring Telnet, on page 108 for instructions on enabling Telnet.

- **telnet ucs-***UCSM-host-name* **ucs-***auth-domain**username*

  ```
  telnet ucs-qa-10
  login: ucs-ldap23\blradmin
  ```

- **telnet ucs-**{*UCSM-ip-address|UCSM-ipv6-address*}*ucs-auth-domain**username*

  ```
  telnet 10.106.19.12 2052
  ucs-qa-10-A login: ucs-ldap23\blradmin
  ```

From a Putty client:

- Login as: **ucs-***auth-domain**username*

  ```
  Login as: ucs-example\jsmith
  ```

  > ✎
  >
  > **Note** If the default authentication is set to local, and the console authentication is set to LDAP, you can log in to the fabric interconnect from a Putty client using **ucs-local\admin**, where admin is the name of the local account.

CHAPTER 3

License Management for the ASA

Cisco Smart Licensing is a flexible licensing model that provides you with an easier, faster, and more consistent way to purchase and manage software across the Cisco portfolio and across your organization. And it's secure—you control what users can access. With Smart Licensing you get:

- **Easy Activation:** Smart Licensing establishes a pool of software licenses that can be used across the entire organization—no more PAKs (Product Activation Keys).

- **Unified Management:** My Cisco Entitlements (MCE) provides a complete view into all of your Cisco products and services in an easy-to-use portal, so you always know what you have and what you are using.

- **License Flexibility:** Your software is not node-locked to your hardware, so you can easily use and transfer licenses as needed.

To use Smart Licensing, you must first set up a Smart Account on Cisco Software Central (software.cisco.com).

For a more detailed overview on Cisco Licensing, go to cisco.com/go/licensingguide

Note This section only applies to ASA logical devices on the Firepower 4100/9300 chassis. For more information on licensing for FTD logical devices, see the FMC Configuration Guide.

About Smart Software Licensing

This section describes how Smart Software Licensing works.

Note This section only applies to ASA logical devices on the Firepower 4100/9300 chassis. For more information on licensing for FTD logical devices, see the FMC Configuration Guide.

Smart Software Licensing for the ASA

For the ASA application on the Firepower 4100/9300 chassis, Smart Software Licensing configuration is split between the Firepower 4100/9300 chassis supervisor and the application.

- Firepower 4100/9300 chassis—Configure all Smart Software Licensing infrastructure in the supervisor, including parameters for communicating with the License Authority. The Firepower 4100/9300 chassis itself does not require any licenses to operate.

Note Inter-chassis clustering requires that you enable the same Smart Licensing method on each chassis in the cluster.

- ASA Application—Configure all license entitlements in the application.

Note Cisco Transport Gateway is not supported on Firepower 4100/9300 security appliances.

Smart Software Manager and Accounts

When you purchase 1 or more licenses for the device, you manage them in the Cisco Smart Software Manager:

https://software.cisco.com/#module/SmartLicensing

The Smart Software Manager lets you create a master account for your organization.

Note If you do not yet have an account, click the link to set up a new account. The Smart Software Manager lets you create a master account for your organization.

By default, your licenses are assigned to the *Default Virtual Account* under your master account. As the account administrator, you can optionally create additional virtual accounts; for example, you can create accounts for regions, departments, or subsidiaries. Multiple virtual accounts let you more easily manage large numbers of licenses and devices.

Offline Management

If your devices do not have Internet access, and cannot register with the License Authority, you can configure offline licensing.

Permanent License Reservation

If your devices cannot access the internet for security reasons, you can optionally request permanent licenses for each ASA. Permanent licenses do not require periodic access to the License Authority. Like PAK licenses, you will purchase a license and install the license key for the ASA. Unlike a PAK license, you obtain and manage the licenses with the Smart Software Manager. You can easily switch between regular smart licensing mode and permanent license reservation mode.

You can obtain a license that enables all features: Standard tier with maximum Security Contexts and the Carrier license. The license is managed on the Firepower 4100/9300 chassis, but you also need to request the entitlements in the ASA configuration so that the ASA allows their use.

Satellite Server

If your devices cannot access the internet for security reasons, you can optionally install a local Smart Software Manager satellite server as a virtual machine (VM). The satellite provides a subset of Smart Software Manager functionality, and allows you to provide essential licensing services for all your local devices. Only the satellite needs to connect periodically to the main License Authority to sync your license usage. You can sync on a schedule or you can sync manually.

Once you download and deploy the satellite application, you can perform the following functions without sending data to Cisco SSM using the Internet:

- Activate or register a license

- View your company's licenses

- Transfer licenses between company entities

For more information, see the Smart Software Manager satellite installation and configuration guides on Smart Account Manager satellite.

Licenses and Devices Managed per Virtual Account

Licenses and devices are managed per virtual account: only that virtual account's devices can use the licenses assigned to the account. If you need additional licenses, you can transfer an unused license from another virtual account. You can also transfer devices between virtual accounts.

Only the Firepower 4100/9300 chassis registers as a device, while the ASA applications in the chassis request their own licenses. For example, for a Firepower 9300 chassis with 3 security modules, the chassis counts as one device, but the modules use 3 separate licenses.

Evaluation License

The Firepower 4100/9300 chassis supports two types of evaluation license:

- Chassis-level evaluation mode—Before the Firepower 4100/9300 chassis registers with the Licensing Authority, it operates for 90 days (total usage) in evaluation mode. The ASA cannot request specific entitlements in this mode; only default entitlements are enabled. When this period ends, the Firepower 4100/9300 chassis becomes out-of-compliance.

- Entitlement-based evaluation mode—After the Firepower 4100/9300 chassis registers with the Licensing Authority, you can obtain time-based evaluation licenses that can be assigned to the ASA. In the ASA,

you request entitlements as usual. When the time-based license expires, you need to either renew the time-based license or obtain a permanent license.

Note You cannot receive an evaluation license for Strong Encryption (3DES/AES); only permanent licenses support this entitlement.

Smart Software Manager Communication

This section describes how your device communicates with the Smart Software Manager.

Device Registration and Tokens

For each virtual account, you can create a registration token. This token is valid for 30 days by default. Enter this token ID plus entitlement levels when you deploy each chassis, or when you register an existing chassis. You can create a new token if an existing token is expired.

At startup after deployment, or after you manually configure these parameters on an existing chassis, the chassis registers with the Cisco License Authority. When the chassis registers with the token, the License Authority issues an ID certificate for communication between the chassis and the License Authority. This certificate is valid for 1 year, although it will be renewed every 6 months.

Periodic Communication with the License Authority

The device communicates with the License Authority every 30 days. If you make changes in the Smart Software Manager, you can refresh the authorization on the device so the change takes place immediately. Or you can wait for the device to communicate as scheduled.

You can optionally configure an HTTP proxy.

The Firepower 4100/9300 chassis must have internet access either directly or through an HTTP proxy at least every 90 days. Normal license communication occurs every 30 days, but with the grace period, your device will operate for up to 90 days without calling home. After the grace period, you must contact the Licensing Authority, or you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected.

Note If your device is unable to communicate with the license authority for one year, the device will enter an unregistered state but will not lose any previously enabled strong encryption capabilities.

Out-of-Compliance State

The device can become out of compliance in the following situations:

- Over-utilization—When the device uses unavailable licenses.

- License expiration—When a time-based license expires.

- Lack of communication—When the device cannot reach the Licensing Authority for re-authorization.

To verify whether your account is in, or approaching, an Out-of-Compliance state, you must compare the entitlements currently in use by your Firepower 4100/9300 chassis against those in your Smart Account.

In an out-of-compliance state, you will not be able to make configuration changes to features requiring special licenses, but operation is otherwise unaffected. For example, existing contexts over the Standard license limit can continue to run, and you can modify their configuration, but you will not be able to add a *new* context.

Smart Call Home Infrastructure

By default, a Smart Call Home profile exists in the FXOS configuration that specifies the URL for the Licensing Authority. You cannot remove this profile. Note that the only configurable option for the License profile is the destination address URL for the License Authority. Unless directed by Cisco TAC, you should not change the License Authority URL.

Note Cisco Transport Gateway is not supported on Firepower 4100/9300 security appliances.

Cisco Success Network

Cisco Success Network is a user-enabled cloud service. When you enable Cisco Success Network, a secure connection is established between the Firepower 4100/9300 chassis and the Cisco cloud to stream usage information and statistics. Streaming telemetry provides a mechanism that selects data of interest from the ASA and transmits it in a structured format to remote management stations to do the following:

• Inform you of available unused features that can improve the effectiveness of the product in your network

• Inform you of additional technical support services and monitoring that might be available for your product

• Help Cisco improve our products

You enable Cisco Success Network when you register the Firepower 4100/9300 with the Cisco Smart Software Manager. See Register the Firepower 4100/9300 chassis with the License Authority, on page 33.

You can enroll in the Cisco Success Network only if all the following conditions are met:

• Smart Software License is registered.

• Smart License Satellite mode is disabled.

• Permanent License is disabled.

Once you enroll in the Cisco Success Network, the chassis establishes and maintains the secure connection at all times. You can turn off this connection at any time by disabling Cisco Success Network, which disconnects the device from the Cisco Success Network cloud.

You can view your current Cisco Success Network enrollment status on the **System** > **Licensing** > **Cisco Success Network** page, and you can change your enrollment status. See Change Cisco Success Network Enrollment, on page 34.

Cisco Success Network Telemetry Data

Cisco Success Network allows the chassis to stream configuration and operating state information once in every 24 hours to the Cisco Success Network cloud. Collected and monitored data include the following:

- **Enrolled device information**—Firepower 4100/9300 chassis model name, product identifier, serial number, UUID, system uptime, and Smart Licensing information. See Enrolled Device Data, on page 22.

- **Software information**—Type and version number for the software running on the Firepower 4100/9300 chassis. See Software Version Data, on page 23.

- **ASA device information**—Information about the ASA devices running on the security module/engine of the Firepower 4100/9300. Note that for the Firepower 4100 series, only the information about a single ASA device is included. ASA device information includes smart licenses in use for each device, device models, serial numbers, and software version. See ASA Device Data, on page 23.

  - **Performance information**—System uptime, CPU usage, memory usage, disk space usage, and bandwidth usage information of the ASA devices. See Performance Data, on page 23.

  - **Usage information**—Feature status, cluster, failover, and login information:

    - **Feature status**—List of enabled ASA features that you have configured or are enabled by default.

    - **Cluster information**—Includes cluster information if the ASA device is in clustered mode. If the ASA device is not in clustered mode, this information is not displayed. The cluster information includes the cluster group name of the ASA device, cluster interface mode, unit name, and state. For the other peer ASA devices in the same cluster, the information includes the name, state, and serial number.

    - **Failover information**—Includes failover information if the ASA is in failover mode. If the ASA is not in failover mode, this information is not displayed. The failover information includes the role and state of the ASA, and the role, state, and serial number of the peer ASA device.

    - **Login history**—User login frequency, login time, and date stamp for the most recent successful login on the ASA device. However, the login history does not include the user login name, credentials, or any other personal information.

  See Usage Data, on page 24 for more information.

Enrolled Device Data

Once you enroll the Firepower 4100/9300 chassis in Cisco Success Network, select telemetry data about the chassis is streamed to the Cisco cloud. The following table describes the collected and monitored data.

Table 1: Enrolled Device Telemetry Data

Data Point	Example Value
Device model	Cisco Firepower FP9300 Security Appliance
Serial number	GMX1135L01K
Smart license PIID	752107e9-e473-4916-8566-e26d0c4a5bd9
Smart license virtual account name	FXOS-general
System uptime	32115
UDI product identifier	FPR-C9300-AC

Software Version Data

Cisco Success Network collects software information that pertains to the chassis including type and software version. The following table describes the collected and monitored software information.

Table 2: Software Version Telemetry Data

Data Point	Example Value
Type	package_version
Version	2.7(1.52)

ASA Device Data

Cisco Success Network collects information about the ASA devices running on the security module/engine of the Firepower 4100/9300. The following table describes the collected and monitored information about ASA devices.

Table 3: ASA Device Telemetry Data

Data Point	Example Value
ASA device PID	FPR9K-SM-36
ASA device model	Cisco Adaptive Security Appliance
ASA device serial number	XDQ311841WA
Deployment type (native or container)	Native
Security context mode (single or multiple)	Single

ASA software version	{
	type: "asa_version",
	ersion: "9.13.1.5"
	}

Device manager version	{
	type: "device_mgr_version",
	version: "7.10.1"
	}

Activated smart licenses in use	{
	"type": "Strong encryption",
	"tag":
	"regid.2016-05.com.cisco.ASA-GEN-STRONG-ENCRYPTION,
	5.7_982308k4-74w2-5f38-64na-707q99g10cce",
	"count": 1
	}

Performance Data

Cisco Success Network collects the performance-specific information for the ASA devices. The information includes system uptime, CPU usage, memory usage, disk space usage, and bandwidth usage information.

- **CPU usage**—CPU usage information for the past five minutes

• **Memory usage**—Free, used, and total memory of the system

• **Disk usage**—Free, used, and total disk space information

• **System uptime**—System uptime information

• **Bandwidth usage**—System bandwidth usage; aggregated from all nameif-ed interfaces

This shows the statistics for received and transmitted packets (or bytes) per second since system up time.

The following table describes the collected and monitored information.

Table 4: Performance Telemetry Data

Data Point	Example Value

System CPU usage in past five minutes	{
	"fiveSecondsPercentage":0.2000000,
	"oneMinutePercentage": 0,
	"fiveMinutesPercentage": 0
	
	}

System memory usage	{
	"freeMemoryInBytes":225854966384,
	"usedMemoryInBytes": 17798281616,
	"totalMemoryInBytes":243653248000
	
	}

System disk usage	{
	"freeGB": 21.237285,
	"usedGB": 0.238805,
	"totalGB": 21.476090
	
	}

System uptime	99700000

System bandwidth usage	{
	"receivedPktsPerSec": 3,
	"receivedBytesPerSec": 212,
	"transmittedPktsPerSec": 3,
	"transmittedBytesPerSec": 399
	}

Usage Data

Cisco Success Network collects feature status, cluster, failover, and login information for the ASA devices running on the security module/engine of the chassis. The following table describes the collected and monitored data about ASA device usage.

Table 5: Usage Telemetry Data

Data Point	Example Value
Feature status	```[{ "name": "cluster", "status": "enabled" }, { "name": "webvpn", "status": "enabled" }, { "name": "logging-buffered", "status": "debugging" }]```
Cluster information	```{ "clusterGroupName": "asa-cluster", "interfaceMode": "spanned", "unitName": "unit-3-3", "unitState": "SLAVE", "otherMembers": { "items": [{ "memberName": "unit-2-1", "memberState": "MASTER", "memberSerialNum": "DAK391674E" }] } }```
Failover information	```{ myRole: "Primary", peerRole: "Secondary", myState: "active", peerState: "standby", peerSerialNum: "DAK39162B" }```
Login history	```{ "loginTimes": "1 times in last 1 days", "lastSuccessfulLogin": "12:25:36 PDT Mar 11 2019" }```

Telemetry Example File

Firepower 4100/9300 chassis aggregates the data received from all ASA devices that have telemetry enabled and are online with the chassis-specific information and additional fields before sending the data to Cisco cloud. If there are no applications with telemetry data, then telemetry is still sent to the Cisco cloud with the chassis information.

The following is an example of a Cisco Success Network telemetry file that includes the information sent to the Cisco cloud for two ASA devices on a Firepower 9300.

```
{
  "version": "1.0",
```

```
            "metadata": {
              "topic": "ASA.telemetry",
              "contentType": "application/json",
              "msgID": "2227"
            },
            "payload": {
              "recordType": "CST_ASA",
              "recordVersion": "1.0",
              "recordedAt": 1560868270055,
              "FXOS": {
                "FXOSdeviceInfo": {
                  "deviceModel": "Cisco Firepower FP9300 Security Appliance",
                  "serialNumber": "HNY4475P01K",
                  "smartLicenseProductInstanceIdentifier": "413509m0-f952-5822-7492-r62c0a5h4gf4",
                  "smartLicenseVirtualAccountName": "FXOS-general",
                  "systemUptime": 32115,
                  "udiProductIdentifier": "FPR-C9300-AC"
                },
                "versions": {
                  "items": [
                    {
                      "type": "package_version",
                      "version": "2.7(1.52)"
                    }
                  ]
                }
              },
              "asaDevices": {
                "items": [
                  {
                    "CPUUsage": {
                      "fiveMinutesPercentage": 0,
                      "fiveSecondsPercentage": 0,
                      "oneMinutePercentage": 0
                    },
                    "bandwidthUsage": {
                      "receivedBytesPerSec": 1,
                      "receivedPktsPerSec": 0,
                      "transmittedBytesPerSec": 1,
                      "transmittedPktsPerSec": 0
                    },
                    "deviceInfo": {
                      "deploymentType": "Native",
                      "deviceModel": "Cisco Adaptive Security Appliance",
                      "securityContextMode": "Single",
                      "serialNumber": "ADG2158508T",
                      "systemUptime": 31084,
                      "udiProductIdentifier": "FPR9K-SM-24"
                    },
                    "diskUsage": {
                      "freeGB": 19.781810760498047,
                      "totalGB": 20.0009765625,
                      "usedGB": 0.21916580200195312
                    },
                    "featureStatus": {
                      "items": [
                        {
                          "name": "aaa-proxy-limit",
                          "status": "enabled"
                        },
                        {
                          "name": "firewall_user_authentication",
                          "status": "enabled"
                        },
```

```
{
  "name": "IKEv2 fragmentation",
  "status": "enabled"
},
{
  "name": "inspection-dns",
  "status": "enabled"
},
{
  "name": "inspection-esmtp",
  "status": "enabled"
},
{
  "name": "inspection-ftp",
  "status": "enabled"
},
{
  "name": "inspection-hs232",
  "status": "enabled"
},
{
  "name": "inspection-netbios",
  "status": "enabled"
},
{
  "name": "inspection-rsh",
  "status": "enabled"
},
{
  "name": "inspection-rtsp",
  "status": "enabled"
},
{
  "name": "inspection-sip",
  "status": "enabled"
},
{
  "name": "inspection-skinny",
  "status": "enabled"
},
{
  "name": "inspection-snmp",
  "status": "enabled"
},
{
  "name": "inspection-sqlnet",
  "status": "enabled"
},
{
  "name": "inspection-sunrpc",
  "status": "enabled"
},
{
  "name": "inspection-tftp",
  "status": "enabled"
},
{
  "name": "inspection-xdmcp",
  "status": "enabled"
},
{
  "name": "management-mode",
  "status": "normal"
},
```

```
                          {
                            "name": "mobike",
                            "status": "enabled"
                          },
                          {
                            "name": "ntp",
                            "status": "enabled"
                          },
                          {
                            "name": "sctp-engine",
                            "status": "enabled"
                          },
                          {
                            "name": "smart-licensing",
                            "status": "enabled"
                          },
                          {
                            "name": "static-route",
                            "status": "enabled"
                          },
                          {
                            "name": "threat_detection_basic_threat",
                            "status": "enabled"
                          },
                          {
                            "name": "threat_detection_stat_access_list",
                            "status": "enabled"
                          }
                        ]
                      },
                      "licenseActivated": {
                        "items": []
                      },
                      "loginHistory": {
                        "lastSuccessfulLogin": "05:53:18 UTC Jun 18 2019",
                        "loginTimes": "1 times in last 1 days"
                      },
                      "memoryUsage": {
                        "freeMemoryInBytes": 226031548496,
                        "totalMemoryInBytes": 241583656960,
                        "usedMemoryInBytes": 15552108464
                      },
                      "versions": {
                        "items": [
                          {
                            "type": "asa_version",
                            "version": "9.13(1)248"
                          },
                          {
                            "type": "device_mgr_version",
                            "version": "7.13(1)31"
                          }
                        ]
                      }
                    },
                    {
                      "CPUUsage": {
                        "fiveMinutesPercentage": 0,
                        "fiveSecondsPercentage": 0,
                        "oneMinutePercentage": 0
                      },
                      "bandwidthUsage": {
                        "receivedBytesPerSec": 1,
                        "receivedPktsPerSec": 0,
```

```
            "transmittedBytesPerSec": 1,
            "transmittedPktsPerSec": 0
        },
        "deviceInfo": {
            "deploymentType": "Native",
            "deviceModel": "Cisco Adaptive Security Appliance",
            "securityContextMode": "Single",
            "serialNumber": "RFL21764S1D",
            "systemUptime": 31083,
            "udiProductIdentifier": "FPR9K-SM-24"
        },
        "diskUsage": {
            "freeGB": 19.781543731689453,
            "totalGB": 20.0009765625,
            "usedGB": 0.21943283081054688
        },
        "featureStatus": {
            "items": [
                {
                    "name": "aaa-proxy-limit",
                    "status": "enabled"
                },
                {
                    "name": "call-home",
                    "status": "enabled"
                },
                {
                    "name": "crypto-ca-trustpoint-id-usage-ssl-ipsec",
                    "status": "enabled"
                },
                {
                    "name": "firewall_user_authentication",
                    "status": "enabled"
                },
                {
                    "name": "IKEv2 fragmentation",
                    "status": "enabled"
                },
                {
                    "name": "inspection-dns",
                    "status": "enabled"
                },
                {
                    "name": "inspection-esmtp",
                    "status": "enabled"
                },
                {
                    "name": "inspection-ftp",
                    "status": "enabled"
                },
                {
                    "name": "inspection-hs232",
                    "status": "enabled"
                },
                {
                    "name": "inspection-netbios",
                    "status": "enabled"
                },
                {
                    "name": "inspection-rsh",
                    "status": "enabled"
                },
                {
                    "name": "inspection-rtsp",
```

```
                                "status": "enabled"
                            },
                            {
                              "name": "inspection-sip",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-skinny",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-snmp",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-sqlnet",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-sunrpc",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-tftp",
                              "status": "enabled"
                            },
                            {
                              "name": "inspection-xdmcp",
                              "status": "enabled"
                            },
                            {
                              "name": "management-mode",
                              "status": "normal"
                            },
                            {
                              "name": "mobike",
                              "status": "enabled"
                            },
                            {
                              "name": "ntp",
                              "status": "enabled"
                            },
                            {
                              "name": "sctp-engine",
                              "status": "enabled"
                            },
                            {
                              "name": "smart-licensing",
                              "status": "enabled"
                            },
                            {
                              "name": "static-route",
                              "status": "enabled"
                            },
                            {
                              "name": "threat_detection_basic_threat",
                              "status": "enabled"
                            },
                            {
                              "name": "threat_detection_stat_access_list",
                              "status": "enabled"
                            }
                          ]
                    },
```

```
              "licenseActivated": {
                "items": []
              },
              "loginHistory": {
                "lastSuccessfulLogin": "05:53:16 UTC Jun 18 2019",
                "loginTimes": "1 times in last 1 days"
              },
              "memoryUsage": {
                "freeMemoryInBytes": 226028740080,
                "totalMemoryInBytes": 241581195264,
                "usedMemoryInBytes": 15552455184
              },
              "versions": {
                "items": [
                  {
                    "type": "asa_version",
                    "version": "9.13(1)248"
                  },
                  {
                    "type": "device_mgr_version",
                    "version": "7.13(1)31"
                  }
                ]
              }
            }
          }
        ]
      }
    }
  }
}
```

Prerequisites for Smart Software Licensing

• Note that this chapter only applies to ASA logical devices on the Firepower 4100/9300 chassis. For more information on licensing for FTD logical devices, see the FMC Configuration Guide.

• Create a master account on the Cisco Smart Software Manager:

  https://software.cisco.com/#module/SmartLicensing

  If you do not yet have an account, click the link to set up a new account. The Smart Software Manager lets you create a master account for your organization.

• Purchase 1 or more licenses from the Cisco Commerce Workspace. On the home page, search for your platform in the **Find Products and Solutions** search field. Some licenses are free, but you still need to add them to your Smart Software Licensing account.

• Ensure internet access or HTTP proxy access from the chassis, so the chassis can contact the Licensing Authority.

• Configure a DNS server so the chassis can resolve the name of the Licensing Authority.

• Set the time for the chassis.

• Configure the Smart Software Licensing infrastructure on the Firepower 4100/9300 chassis before you configure the ASA licensing entitlements.

Guidelines for Smart Software Licensing

ASA Guidelines for Failover and Clustering

Each Firepower 4100/9300 chassis must be registered with the License Authority or satellite server. There is no extra cost for secondary units. For permanent license reservation, you must purchase separate licenses for each chassis.

Defaults for Smart Software Licensing

The Firepower 4100/9300 chassis default configuration includes a Smart Call Home profile called "SLProfile" that specifies the URL for the Licensing Authority.

Configure Regular Smart Software Licensing

To communicate with the Cisco License Authority, you can optionally configure an HTTP proxy. To register with the License Authority, you must enter the registration token ID on the Firepower 4100/9300 chassis that you obtained from your Smart Software License account.

Procedure

Step 1 (Optional) Configure the HTTP Proxy, on page 32.

Step 2 (Optional) Delete the Call Home URL, on page 33

Step 3 Register the Firepower 4100/9300 chassis with the License Authority, on page 33.

(Optional) Configure the HTTP Proxy

If your network uses an HTTP proxy for Internet access, you must configure the proxy address for Smart Software Licensing. This proxy is also used for Smart Call Home in general.

Note   HTTP proxy with authentication is not supported.

Procedure

Step 1 Choose **System** > **Licensing** > **Call Home**.

The Call Home page provides fields for configuring the destination address URL for the License Authority and for configuring an HTTP proxy.

Note   Unless directed by Cisco TAC, you should not change the License Authority URL.

Step 2   In the Server Enable drop-down list, select **on**.

Step 3   Enter the proxy IP address and port in the **Server URL** and **Server Port** fields. For example, enter port 443 for an HTTPS server.

Step 4   Click **Save**.

(Optional) Delete the Call Home URL

Use the following procedure to delete a previously configured Call Home URL.

Procedure

Step 1   Choose **System** > **Licensing** > **Call Home**.

Step 2   In the **Call home Configuration** area, select Delete.

Register the Firepower 4100/9300 chassis with the License Authority

When you register the Firepower 4100/9300 chassis, the License Authority issues an ID certificate for communication between the Firepower 4100/9300 chassis and the License Authority. It also assigns the Firepower 4100/9300 chassis to the appropriate virtual account. Normally, this procedure is a one-time instance. However, you might need to later re-register the Firepower 4100/9300 chassis if the ID certificate expires because of a communication problem, for example.

Procedure

Step 1   In the Smart Software Manager or the Smart Software Manager Satellite, request and copy a registration token for the virtual account to which you want to add this Firepower 4100/9300 chassis.

For more information on how to request a registration token using the Smart Software Manager Satellite, see the Cisco Smart Software Manager Satellite User Guide (https://www.cisco.com/c/en/us/buy/smart-accounts/software-manager-satellite.html).

Step 2   In Firepower Chassis Manager, choose **System** > **Licensing** > **Smart License**.

Step 3   Enter the registration token in the **Enter Product Instance Registration Token** field.

Step 4   (Optional) You can uncheck the **Enable Cisco Success Network** check box to disable the Cisco Success Network feature.

See Cisco Success Network, on page 21 for more information.

Step 5   Click **Register**.

The Firepower 4100/9300 chassis attempts to register with the License Authority.

To unregister the device, click **Unregister**.

Deregistering the Firepower 4100/9300 chassis removes the device from your account. All license entitlements and certificates on the device are removed. You might want to deregister to free up a license for a new Firepower 4100/9300 chassis. Alternatively, you can remove the device from the Smart Software Manager.

Change Cisco Success Network Enrollment

You enable Cisco Success Network when you register the Firepower 4100/9300 with the Cisco Smart Software Manager. After that, use the following procedure to view or change enrollment status.

Note Cisco Success Network does not work in evaluation mode.

Procedure

Step 1 Choose **System** > **Licensing** > **Cisco Success Network**.

Step 2 Under **Cisco Success Network Preferences**, read the information provided by Cisco, and click **Click here** to check out the sample data that will be sent to Cisco.

Step 3 Choose whether you want to **Enable Cisco Success Network**, and click **Save**.

Configure a Smart License Satellite Server for the Firepower 4100/9300 chassis

The following procedure shows how to configure the Firepower 4100/9300 chassis to use a Smart License satellite server.

Before you begin

- Complete all prerequisites listed in the Prerequisites for Smart Software Licensing, on page 31.

- Deploy and set up a Smart Software Satellite Server:

  Download the Smart License Satellite OVA file from Cisco.com and install and configure it on a VMwareESXi server. For more information, see the Smart Software Manager satellite Install Guide.

- Verify that the FQDN of the Smart Software Satellite Server can be resolved by your internal DNSserver.

- Verify whether the satellite trustpoint is already present:

  scope security

  show trustpoint

  Note that the trustpoint is added by default in FXOS version 2.4(1) and later. If the trustpoint is not present, you must add one manually using the following steps:

1. Go to http://www.cisco.com/security/pki/certs/clrca.cer and copy the entire body of the SSL certificate (from "-----BEGIN CERTIFICATE-----" to "-----END CERTIFICATE-----") into a place you can access during configuration.

2. Enter security mode:

   scope security

3. Create and name a trusted point:

   create trustpoint *trustpoint_name*

4. Specify certificate information for the trust point. Note: the certificate must be in Base64 encoded X.509 (CER) format.

   set certchain *certchain*

   For the *certchain* variable, paste the certificate text that you copied in step 1.

   If you do not specify certificate information in the command, you are prompted to enter a certificate or a list of trust points defining a certification path to the root certificate authority (CA). On the next line following your input, type **ENDOFBUF** to finish.

5. Commit the configuration:

   commit-buffer

Procedure

Step 1 Choose **System** > **Licensing** > **Call Home**.

Step 2 In the **Call home Configuration** area, replace the default URL in the **Address** field with the URL of your Smart Software Satellite Server using information that you gathered in the prerequisites for this procedure, using the following format: **https://**[*FQDN of Satellite server*]**/Transportgateway/services/DeviceRequestHandler**

Step 3 Register the Firepower 4100/9300 chassis with the License Authority, on page 33. Note that you must request and copy the registration token from the Smart License Manager satellite.

Configure Permanent License Reservation

You can assign a permanent license to your Firepower 4100/9300 chassis. This universal reservation allows you to use any entitlement for an unlimited count on your device.

Note Before you begin, you must purchase the permanent licenses so they are available in the Smart Software Manager. Not all accounts are approved for permanent license reservation. Make sure you have approval from Cisco for this feature before you attempt to configure it.

Install the Permanent License

The following procedure shows how to assign a permanent license to your Firepower 4100/9300 chassis.

Procedure

Step 1   Choose **System** > **Licensing** > **Permanent License**.

Step 2   Click **Generate** to generate a reservation request code. Copy the reservation request code to your clipboard.

Step 3   Go to the Smart Software Manager Inventory screen in the Cisco Smart Software Manager portal, and click the **Licenses** tab:

https://software.cisco.com/#SmartLicensing-Inventory

The **Licenses** tab displays all existing licenses related to your account, both regular and permanent.

Step 4   Click **License Reservation**, and paste the generated reservation request code into the box.

Step 5   Click **Reserve License**.

The Smart Software Manager generates an authorization code. You can download the code or copy it to the clipboard. At this point, the license is now in use according to the Smart Software Manager.

If you do not see the **License Reservation** button, then your account is not authorized for permanent license reservation. In this case, you should disable permanent license reservation and re-enter the regular smart license commands.

Step 6   In Firepower Chassis Manager, enter the generated authorization code into the **Authorization Code** text box.

Step 7   Click **Install**.

Once your Firepower 4100/9300 chassis is fully licensed with PLR, the Permanent License page displays your license status and offers the option to return your permanent license.

Step 8   Enable feature entitlements on the ASA logical device. See the ASA licensing chapter to enable entitlements.

(Optional) Return the Permanent License

If you no longer need a permanent license, you must officially return it to the Smart Software Manager using this procedure. If you do not follow all steps, the license stays in an in-use state and cannot be used elsewhere.

Procedure

Step 1   Choose **System** > **Licensing** > **Permanent License**.

Step 2   Click **Return** to generate a return code. Copy the return code to your clipboard.

The Firepower 4100/9300 chassis immediately becomes unlicensed and moves to the Evaluation state.

Step 3   Go to the Smart Software Manager Inventory screen, and click on the **Product Instances** tab:

https://software.cisco.com/#SmartLicensing-Inventory

Step 4   Search for your Firepower 4100/9300 chassis using its universal device identifier (UDI).

Step 5 Choose **Actions** > **Remove**, and paste the generated return code into the box.

Step 6 Click **Remove Product Instance**.

The permanent license is returned to the available pool.

Step 7 Reboot the system. For details on how to reboot your Firepower 4100/9300 chassis, see Rebooting the Firepower 4100/9300 Chassis, on page 98.

History for Smart Software Licensing

Feature Name	Platform Releases	Description
Cisco Success Network	2.7.1	Cisco Success Network is a user-enabled cloud service. When you enable Cisco Success Network, a secure connection is established between the Firepower 4100/9300 chassis and the Cisco cloud to stream usage information and statistics. Streaming telemetry provides a mechanism that selects data of interest from the ASA and transmits it in a structured format to remote management stations to do the following: • Inform you of available unused features that can improve the effectiveness of the product in your network • Inform you of additional technical support services and monitoring that might be available for your product • Help Cisco improve our products Once you enroll in the Cisco Success Network, the chassis establishes and maintains the secure connection at all times. You can turn off this connection at any time by disabling Cisco Success Network, which disconnects the device from the Cisco Success Network cloud. We introduced the following commands: **scope telemetry {enable | disable}** We introduced the following screens: **System** > **Licensing** > **Cisco Success Network**

Feature Name	Platform Releases	Description
Cisco Smart Software Licensing for the Firepower 4100/9300 chassis	1.1(1)	Smart Software Licensing lets you purchase and manage a pool of licenses. Smart licenses are not tied to a specific serial number. You can easily deploy or retire devices without having to manage each unit's license key. Smart Software Licensing also lets you see your license usage and needs at a glance. Smart Software Licensing configuration is split between the Firepower 4100/9300 chassis supervisor and the security module. We introduced the following screens: **System** > **Licensing** > **Call Home** **System** > **Licensing** > **Smart License**

C H A P T E R **4**

User Management

- User Accounts, on page 39
- Guidelines for Usernames, on page 40
- Guidelines for Passwords, on page 41
- Guidelines for Remote Authentication, on page 42
- User Roles, on page 44
- Password Profile for Locally Authenticated Users, on page 44
- Configuring User Settings, on page 45
- Configuring the Session Timeout, on page 48
- Configuring the Absolute Session Timeout, on page 49
- Set the Maximum Number of Login Attempts, on page 50
- Configure Minimum Password Length Check, on page 50
- Creating a Local User Account, on page 51
- Deleting a Local User Account, on page 52
- Activating or Deactivating a Local User Account, on page 53
- Clearing the Password History for a Locally Authenticated User, on page 53

User Accounts

User accounts are used to access the system. You can configure up to 48 local user accounts. Each user account must have a unique username and password.

Admin Account

The admin account is a default user account and cannot be modified or deleted. This account is the system administrator or superuser account and has full privileges. There is no default password assigned to the admin account; you must choose the password during the initial system setup.

The admin account is always active and does not expire. You cannot configure the admin account as inactive.

Locally Authenticated User Accounts

A locally authenticated user account is authenticated directly through the chassis and can be enabled or disabled by anyone with admin or AAA privileges. Once a local user account is disabled, the user cannot log in. Configuration details for disabled local user accounts are not deleted by the database. If you reenable a disabled

local user account, the account becomes active again with the existing configuration; however, the account password must be reset.

Remotely Authenticated User Accounts

A remotely authenticated user account is any user account that is authenticated through LDAP, RADIUS, or TACACS+. All remote users are initially assigned the **Read-Only** role by default.

If a user maintains a local user account and a remote user account simultaneously, the roles defined in the local user account override those maintained in the remote user account.

The fallback authentication method is to use the local database. This fallback method is not configurable.

See the following topics for more information on guidelines for remote authentication, and how to configure and delete remote authentication providers:

Expiration of User Accounts

You can configure user accounts to expire at a predefined time. When the expiration time is reached, the user account is disabled.

By default, user accounts do not expire.

After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. You can, however, configure the account with the latest expiration date available.

Guidelines for Usernames

The username is also used as the login ID for Firepower Chassis Manager and the FXOS CLI. When you assign login IDs to user accounts, consider the following guidelines and restrictions:

- The login ID can contain between 1 and 32 characters, including the following:

  - Any alphabetic character

  - Any digit

  - _ (underscore)

  - - (dash)

  - . (dot)

- The login ID must be unique.

- The login ID must start with an alphabetic character. It cannot start with a number or a special character, such as an underscore.

- The login ID is case-sensitive.

- You cannot create an all-numeric login ID.

- After you create a user account, you cannot change the login ID. You must delete the user account and create a new one.

Guidelines for Passwords

A password is required for each locally authenticated user account. A user with admin or AAA privileges can configure the system to perform a password strength check on user passwords. If the password strength check is enabled, each user must have a strong password.

We recommend that each user have a strong password. If you enable the password strength check for locally authenticated users, the FXOS rejects any password that does not meet the following requirements:

- Must contain a minimum of 8 characters and a maximum of 127 characters.

Note You can optionally configure a minimum password length of 15 characters on the system, to comply with Common Criteria requirements. For more information, see Configure Minimum Password Length Check, on page 50.

- Must include at least one uppercase alphabetic character.

- Must include at least one lowercase alphabetic character.

- Must include at least one non-alphanumeric (special) character.

- Must not contain a space.

- Must not contain a character that is repeated more than 3 times consecutively, such as aaabbb.

- Must not contain three consecutive numbers or letters in any order, such as passwordABC or password321.

- Must not be identical to the username or the reverse of the username.

- Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word.

- Must not contain the following symbols: $ (dollar sign), ? (question mark), and = (equals sign).

Note This restriction applies whether the password strength check is enabled or not.

- Must not be blank for local user and admin accounts.

Guidelines for Remote Authentication

If a system is configured for one of the supported remote authentication services, you must create a provider for that service to ensure that the Firepower 4100/9300 chassis can communicate with the system. The following guidelines impact user authorization:

User Accounts in Remote Authentication Services

User accounts can exist locally in the Firepower 4100/9300 chassis or in the remote authentication server.

You can view the temporary sessions for users who log in through remote authentication services from the Firepower Chassis Manager or the FXOS CLI.

User Roles in Remote Authentication Services

If you create user accounts in the remote authentication server, you must ensure that the accounts include the roles those users require for working in the Firepower 4100/9300 chassis and that the names of those roles match the names used in FXOS. Based on the role policy, a user might not be allowed to log in, or is granted only read-only privileges.

User Attributes in Remote Authentication Providers

For RADIUS and TACACS+ configurations, you must configure a user attribute for the Firepower 4100/9300 chassis in each remote authentication provider through which users log in to Firepower Chassis Manager or the FXOS CLI. This user attribute holds the roles and locales assigned to each user.

When a user logs in, FXOS does the following:

1. Queries the remote authentication service.

2. Validates the user.

3. If the user is validated, checks the roles and locales assigned to that user.

The following table contains a comparison of the user attribute requirements for the remote authentication providers supported by FXOS:

Authenication Provider	Custom Attribute	Schema Extension	Attribute ID Requirements
LDAP	Optional	You can choose to do one of the following: • Do not extend the LDAP schema and configure an existing, unused attribute that meets the requirements. • Extend the LDAP schema and create a custom attribute with a unique name, such as CiscoAVPair.	The Cisco LDAP implementation requires a unicode type attribute. If you choose to create the CiscoAVPair custom attribute, use the following attribute ID: 1.3.6.1.4.1.9.287247.1 A sample OID is provided in the following section.

Authentication Provider	Custom Attribute	Schema Extension	Attribute ID Requirements
RADIUS	Optional	You can choose to do one of the following: • Do not extend the RADIUS schema and use an existing, unused attribute that meets the requirements. • Extend the RADIUS schema and create a custom attribute with a unique name, such as cisco-avpair.	The vendor ID for the Cisco RADIUS implementation is 009 and the vendor ID for the attribute is 001. The following syntax example shows how to specify multiples user roles and locales if you choose to create the cisco-avpair attribute: `shell:roles="admin,aaa" shell:locales="L1,abc"`. Use a comma "," as the delimiter to separate multiple values.
TACACS+	Required	You must extend the schema and create a custom attribute with the name cisco-av-pair.	The cisco-av-pair name is the string that provides the attribute ID for the TACACS+ provider. The following syntax example shows how to specify multiples user roles and locales when you create the cisco-av-pair attribute: `cisco-av-pair=shell:roles="admin aaa" shell:locales*"L1 abc"`. Using an asterisk (*) in the cisco-av-pair attribute syntax flags the locale as optional, preventing authentication failures for other Cisco devices that use the same authorization profile. Use a space as the delimiter to separate multiple values.

Sample OID for LDAP User Attribute

The following is a sample OID for a custom CiscoAVPair attribute:

```
CN=CiscoAVPair,CN=Schema,
CN=Configuration,CN=X
objectClass: top
objectClass: attributeSchema
cn: CiscoAVPair
distinguishedName: CN=CiscoAVPair,CN=Schema,CN=Configuration,CN=X
instanceType: 0x4
uSNCreated: 26318654
attributeID: 1.3.6.1.4.1.9.287247.1
attributeSyntax: 2.5.5.12
isSingleValued: TRUE
showInAdvancedViewOnly: TRUE
adminDisplayName: CiscoAVPair
adminDescription: UCS User Authorization Field
oMSyntax: 64
```

```
lDAPDisplayName: CiscoAVPair
name: CiscoAVPair
objectCategory: CN=Attribute-Schema,CN=Schema,CN=Configuration,CN=X
```

User Roles

The system contains the following user roles:

Administrator

Complete read-and-write access to the entire system. The default admin account is assigned this role by default and it cannot be changed.

Read-Only

Read-only access to system configuration with no privileges to modify the system state.

Operations

Read-and-write access to NTP configuration, Smart Call Home configuration for Smart Licensing, and system logs, including syslog servers and faults. Read access to the rest of the system.

AAA Administrator

Read-and-write access to users, roles, and AAA configuration. Read access to the rest of the system.

Password Profile for Locally Authenticated Users

The password profile contains the password history and password change interval properties for all locally authenticated users. You cannot specify a different password profile for each locally authenticated user.

Password History Count

The password history count allows you to prevent locally authenticated users from reusing the same password over and over again. When this property is configured, the Firepower chassis stores passwords that were previously used by locally authenticated users up to a maximum of 15 passwords. The passwords are stored in reverse chronological order with the most recent password first to ensure that the only the oldest password can be reused when the history count threshold is reached.

A user must create and use the number of passwords configured in the password history count before being able to reuse one. For example, if you set the password history count to 8, a locally authenticated user cannot reuse the first password until after the ninth password has expired.

By default, the password history is set to 0. This value disables the history count and allows users to reuse previously passwords at any time.

If necessary, you can clear the password history count for a locally authenticated user and enable reuse of previous passwords.

Password Change Interval

The password change interval enables you to restrict the number of password changes a locally authenticated user can make within a given number of hours. The following table describes the two configuration options for the password change interval.

Interval Configuration	Description	Example
No password change allowed	This option does not allow passwords for locally authenticated users to be changed within a specified number of hours after a password change. You can specify a no change interval between 1 and 745 hours. By default, the no change interval is 24 hours.	For example, to prevent passwords from being changed within 48 hours after a locally authenticated user changes his or her password, set the following: • Change during interval to disable • No change interval to 48
Password changes allowed within change interval	This option specifies the maximum number of times that passwords for locally authenticated users can be changed within a pre-defined interval. You can specify a change interval between 1 and 745 hours and a maximum number of password changes between 0 and 10. By default, a locally authenticated user is permitted a maximum of 2 password changes within a 48 hour interval.	For example, to allow a password to be changed a maximum of once within 24 hours after a locally authenticated user changes his or her password, set the following: • Change during interval to enable • Change count to 1 • Change interval to 24

Configuring User Settings

Procedure

Step 1 Choose **System** > **User Management**.

Step 2 Click the **Settings** tab.

Step 3 Complete the following fields with the required information:

Note If **Default Authentication** and **Console Authentication** are both set to use the same remote authentication protocol (RADIUS, TACACS+, or LDAP), you cannot change certain aspects of that server's configuration (for example, deleting that server, or changing its order of assignment) without updating these user settings.

Name	Description
Default Authentication field	The default method by which a user is authenticated during remote login. This can be one of the following:
	• **Local**—The user account must be defined locally on the chassis.
	• **Radius**—The user account must be defined on the RADIUS server specified for the chassis.
	• **TACACS**—The user account must be defined on the TACACS+ server specified for the chassis.
	• **LDAP**—The user account must be defined on the LDAP/MS-AD server specified for the chassis.
	• **None**—If the user account is local to the chassis, no password is required when the user logs in remotely.
	Note All **Radius**, **TACACS**, and **LDAP** settings must be configured under Platform Settings. For more information, see About AAA, on page 131 in the Platform Settings chapter.
Console Authentication field	The method by which a user is authenticated when connecting to the FXOS CLI via the console port. This can be one of the following:
	• **Local**—The user account must be defined locally on the chassis.
	• **Radius**—The user account must be defined on the RADIUS server specified for the chassis.
	• **TACACS**—The user account must be defined on the TACACS+ server specified for the chassis.
	• **LDAP**—The user account must be defined on the LDAP/MS-AD server specified for the chassis.
	• **None**—If the user account is local to the chassis, no password is required when the user connects to the FXOS CLI using the console port.
Remote User Settings	
Remote User Role Policy	Controls what happens when a user attempts to log in and the remote authentication provider does not supply a user role with the authentication information:
	• **Assign Default Role**—The user is allowed to log in with a read-only user role.
	• **No-Login**—The user is not allowed to log in to the system, even if the username and password are correct.
Local User Settings	

Name	Description
Password Strength Check check box	If checked, all local user passwords must conform to the guidelines for a strong password (see Guidelines for Passwords, on page 41). The strong password check is enabled by default.
History Count field	The number of unique passwords a user must create before the user can reuse a previously used password. The history count is in reverse chronological order with the most recent password first to ensure that only the oldest password can be reused when the history count threshold is reached. This value can be anywhere from 0 to 15. You can set the **History Count** field to 0 to disable the history count and allow users to reuse previously used passwords at any time.
Change During Interval field	Controls when a locally authenticated user can change his or her password. This can be: • **Enable**—Locally authenticated users can change their passwords based on the settings for Change Interval and Change Count. • **Disable**—Locally authenticated users cannot change their passwords for the period of time specified for No Change Interval.
Change Interval field	The number of hours over which the number of password changes specified in the **Change Count** field are enforced. This value can be anywhere from 1 to 745 hours. For example, if this field is set to 48 and the **Change Count** field is set to 2, a locally authenticated user can make no more than 2 password changes within a 48 hour period.
Change Count field	The maximum number of times a locally authenticated user can change his or her password during the Change Interval. This value can be anywhere from 0 to 10.
No Change Interval field	The minimum number of hours that a locally authenticated user must wait before changing a newly created password. This value can be anywhere from 1 to 745 hours. This interval is ignored if the **Change During Interval** property is not set to **Disable**.
Passphrase Expiration Days field	Set the expiration between 1 and 9999 days. By default, expiration is disabled.
Passphrase Expiration Warning Period field	Set the number of days before expiration to warn the user about their password expiration at each login, between 0 and 9999. The default is 14 days.
Expiration Grace Period field	Set the number of days a user has to change their password after expiration, between 0 and 9999. The default is 3 days.

Name	Description
Password Reuse Interval field	Set the number of days before you can reuse a password, between 1 and 365. The default is 15 days. If you enable both the **History Count** and the **Password Reuse Interval**, then both requirements must be met. For example, if you set the history count to 3, and the reuse interval to 10 days, then you can change your password only after 10 days have passed, and you have changed your password 3 times.

Step 4   Click **Save**.

Configuring the Session Timeout

You can use the FXOS CLI to specify the amount of time that can pass without user activity before the Firepower 4100/9300 chassis closes user sessions. You can configure different settings for console sessions and for HTTPS, SSH, and Telnet sessions.

You can set a timeout value up to 3600 seconds (60 minutes). The default value is 600 seconds. To disable this setting, set the session timeout value to 0.

Procedure

Step 1   Enter security mode:

Firepower-chassis # **scope security**

Step 2   Enter default authorization security mode:

Firepower-chassis /security # **scope default-auth**

Step 3   Set the idle timeout for HTTPS, SSH, and Telnet sessions:

Firepower-chassis /security/default-auth # **set session-timeout** *seconds*

Step 4   (Optional) Set the idle timeout for console sessions:

Firepower-chassis /security/default-auth # **set con-session-timeout** *seconds*

Step 5   Commit the transaction to the system configuration:

Firepower-chassis /security/default-auth # **commit-buffer**

Step 6   (Optional) View the session and absolute session timeout settings:

Firepower-chassis /security/default-auth # **show detail**

Example:

```
Default authentication:
 Admin Realm: Local
 Operational Realm: Local
 Web session refresh period(in secs): 600
 Idle Session timeout (in secs) for web, ssh, telnet sessions: 600
 Absolute Session timeout (in secs) for web, ssh, telnet sessions: 3600
```

```
Serial Console Session timeout(in secs): 600
Serial Console Absolute Session timeout(in secs): 3600
Admin Authentication server group:
Operational Authentication server group:
Use of 2nd factor: No
```

Configuring the Absolute Session Timeout

The Firepower 4100/9300 chassis has an absolute session timeout setting that closes user sessions after the absolute session timeout period has passed, regardless of session use. This absolute timeout functionality is global across all forms of access including serial console, SSH, and HTTPS.

The absolute timeout value defaults to 3600 seconds (60 minutes) and can be changed using the FXOS CLI. To disable this setting, set the absolute session timeout value to 0.

Procedure

Step 1   Enter security mode:

Firepower-chassis # **scope security**

Step 2   Enter default authorization security mode:

Firepower-chassis /security # **scope default-auth**

Step 3   Set the absolute session timeout:

Firepower-chassis /security/default-auth # **set absolute-session-timeout** *seconds*

Step 4   Commit the transaction to the system configuration:

Firepower-chassis /security/default-auth # **commit-buffer**

Step 5   (Optional) View the session and absolute session timeout settings:

Firepower-chassis /security/default-auth # **show detail**

Example:

```
Default authentication:
 Admin Realm: Local
 Operational Realm: Local
 Web session refresh period(in secs): 600
 Idle Session timeout (in secs) for web, ssh, telnet sessions: 600
 Absolute Session timeout (in secs) for web, ssh, telnet sessions: 3600
 Serial Console Session timeout(in secs): 600
 Serial Console Absolute Session timeout(in secs): 3600
 Admin Authentication server group:
 Operational Authentication server group:
 Use of 2nd factor: No
```

Set the Maximum Number of Login Attempts

You can configure the maximum number of failed login attempts allowed before a user is locked out of the Firepower 4100/9300 chassis for a specified amount of time. If a user exceeds the set maximum number of login attempts, the user is locked out of the system. No notification appears indicating that the user is locked out. In this event, the user must wait the specified amount of time before attempting to log in.

Perform these steps to configure the maximum number of login attempts.

Note
- All types of user accounts (including admin) are locked out of the system after exceeding the maximum number of login attempts.

- The default maximum number of unsuccessful login attempts is 0. The default amount of time the user is locked out of the system after exceeding the maximum number of login attemps is 30 minutes (1800 seconds).

This option is one of a number offered for achieving Common Criteria certification compliance on your system. For more information, see Security Certifications Compliance, on page 65.

Procedure

Step 1 From the FXOS CLI, enter security mode:

scope security

Step 2 Set the maximum number of unsuccessful login attempts.

set max-login-attempts *num_attempts*

The *num_attempts* value is any integer from 0-10.

Step 3 Specify the amount of time (in seconds) the user should remain locked out of the system after reaching the maximum number of login attempts:

set user-account-unlock-time

unlock_time

Step 4 Commit the configuration:

commit-buffer

Configure Minimum Password Length Check

If you enable minimum password length check, you must create passwords with the specified minimum number of characters. For example, if the *min_length* option is set to 15, you must create passwords using 15

characters or more. This option is one of a number that allow for Common Criteria certification compliance on your system. For more information, see Security Certifications Compliance.

Perform these steps to configure the minimum password length check.

Procedure

Step 1   From the FXOS CLI, enter security mode:

scope security

Step 2   Specify the minimum password length:

set min-password-length *min_length*

Step 3   Commit the configuration:

commit-buffer

Creating a Local User Account

Procedure

Step 1   Choose **System** > **User Management**.

Step 2   Click the **Local Users** tab.

Step 3   Click **Add User** to open the **Add User** dialog box.

Step 4   Complete the following fields with the required information about the user:

Name	Description
User Name field	The account name that is used when logging into this account. This name must be unique and meet the guidelines and restrictions for user account names (see Guidelines for Usernames, on page 40). After you save the user, the login ID cannot be changed. You must delete the user account and create a new one.
First Name field	The first name of the user. This field can contain up to 32 characters.
Last Name field	The last name of the user. This field can contain up to 32 characters.
Email field	The email address for the user.
Phone Number field	The telephone number for the user.

Name	Description
Password field	The password associated with this account. If password strength check is enabled, a user's password must be strong and the FXOS rejects any password that does not meet the strength check requirements (see Guidelines for Passwords, on page 41). **Note** Passwords must not contain the following symbols: $ (dollar sign), ? (question mark), and = (equals sign). This restriction applies whether the password strength check is enabled or not.
Confirm Password field	The password a second time for confirmation purposes.
Account Status field	If the status is set to **Active**, a user can log into Firepower Chassis Manager and the FXOS CLI with this login ID and password.
User Role list	The role that represents the privileges you want to assign to the user account (see User Roles, on page 44). All users are assigned the Read-Only role by default and this role cannot be deselected. To assign multiple roles, hold down **Ctrl** and click the desired roles. **Note** When you delete a user role, current session IDs for the user are revoked, meaning all of the user's active sessions (both CLI and Web) are immediately terminated.
Account Expires check box	If checked, this account expires and cannot be used after the date specified in the **Expiration Date** field. **Note** After you configure a user account with an expiration date, you cannot reconfigure the account to not expire. You can, however, configure the account with the latest expiration date available.
Expiry Date field	The date on which the account expires. The date should be in the format yyyy-mm-dd. Click the calendar icon at the end of this field to view a calendar that you can use to select the expiration date.

Step 5 Click **Add**.

Deleting a Local User Account

Procedure

Step 1 Choose **System** > **User Management**.

Step 2 Click the **Local Users** tab.

Step 3 In the row for the user account that you want to delete, click **Delete**.

Step 4 In the **Confirm** dialog box, click **Yes**.

Activating or Deactivating a Local User Account

You must be a user with admin or AAA privileges to activate or deactivate a local user account.

Procedure

Step 1 Choose **System** > **User Management**.

Step 2 Click the **Local Users** tab.

Step 3 In the row for the user account that you want to activate or deactivate, click **Edit (pencil icon)**.

Step 4 In the **Edit User** dialog box, do one of the following:

- To activate a user account, click the **Active** radio button in the **Account Status** field. Note that when you reactivate a user account, the account password must be reset.

- To deactivate a user account, click the **Inactive** radio button in the **Account Status** field.

The admin user account is always set to active. It cannot be modified.

Step 5 Click **Save**.

Step 6 Commit the transaction to the system configuration:

Firepower-chassis /security/local-user # **commit-buffer**

Clearing the Password History for a Locally Authenticated User

Procedure

Step 1 Enter security mode:

Firepower-chassis # **scope security**

Step 2 Enter local user security mode for the specified user account:

Firepower-chassis /security # **scope local-user** *user-name*

Step 3 Clear the password history for the specified user account:

Firepower-chassis /security/local-user # **clear password-history**

Step 4 Commit the transaction to the system configuration:

Firepower-chassis /security/local-user # **commit-buffer**

Example

The following example clears the password history and commits the transaction:

```
Firepower-chassis # scope security
Firepower-chassis /security # scope local-user admin
Firepower-chassis /security/local-user # clear password-history
Firepower-chassis /security/local-user* # commit-buffer
Firepower-chassis /security/local-user #
```

Image Management

About Image Management

The Firepower 4100/9300 chassis uses two basic types of images:

Note   All images are digitally signed and validated through Secure Boot. Do not modify the image in any way or you will receive a validation error.

- Platform Bundle—The platform bundle is a collection of multiple independent images that operate on the Supervisor and security module/engine. The platform bundle is a FXOS software package.

- Application—Application images are the software images you want to deploy on the security module/engine of the Firepower 4100/9300 chassis. Application images are delivered as Cisco Secure Package files (CSP) and are stored on the supervisor until deployed to a security module/engine as part of logical device creation or in preparation for later logical device creation. You can have multiple different versions of the same application image type stored on the Supervisor.

Note   If you are upgrading both the Platform Bundle image and one or more Application images, you must upgrade the Platform Bundle first.

✎

Note   If you are installing an ASA application in the device, you can delete the images of the existing application FTD and vice versa. When you try to delete all the FTD images, at least one image deletion will be denied with an error message `Invalid operation as no default FTD/ASA APP will be left.` `Please select a new default FTD app.` In order to delete all the FTD images, you must leave the default image alone and delete the rest of the images and then finally delete the default image.

Downloading Images from Cisco.com

Download FXOS and application images from Cisco.com so you can upload them to the chassis.

Before you begin

You must have a Cisco.com account.

Procedure

Step 1   Using a web browser, navigate to http://www.cisco.com/go/firepower9300-software or http://www.cisco.com/go/firepower4100-software.
The software download page for the Firepower 4100/9300 chassis is opened in the browser.

Step 2   Find and then download the appropriate software image to your local computer.

Uploading an Image to the Security Appliance

You can upload FXOS and application images to the chassis.

Before you begin

Make sure the image you want to upload is available on your local computer.

Procedure

Step 1   Choose **System** > **Updates**.
The Available Updates page shows a list of the FXOS platform bundle images and application images that are available on the chassis.

Step 2   Click **Upload Image** to open the Upload Image dialog box.

Step 3   Click **Choose File** to navigate to and select the image that you want to upload.

Step 4   Click **Upload**.
The selected image is uploaded to the Firepower 4100/9300 chassis. While the image is uploading, the system displays a progress bar to indicate the percentage of the upload that has been completed.

Step 5   For certain software images you will be presented with an end-user license agreement after uploading the image. Follow the system prompts to accept the end-user license agreement.

Verifying the Integrity of an Image

The integrity of the image is automatically verified when a new image is added to the Firepower 4100/9300 chassis. If needed, you can use the following procedure to manually verify the integrity of an image.

Procedure

Step 1   Choose **System** > **Updates**.
The Available Updates page shows a list of the FXOS platform bundle images and application images that are available on the chassis.

Step 2   Click **Verify** (check mark icon) for the image you want to verify.

The system will verify the integrity of the image and display the staus in the Image Integrity field.

Upgrading the FXOS Platform Bundle

Before you begin

Download the platform bundle software image from Cisco.com (see Downloading Images from Cisco.com, on page 56) and then upload that image to the Firepower 4100/9300 chassis (see Uploading an Image to the Security Appliance, on page 56).

Note   The upgrade process typically takes between 20 and 30 minutes.

If you are upgrading a Firepower 9300 or 4100 Series security appliance that is running a standalone logical device or if you are upgrading a Firepower 9300 security appliance that is running an intra-chassis cluster, traffic will not traverse through the device while it is upgrading.

If you are upgrading Firepower 9300 or 4100 Series security appliance that is part of an inter-chassis cluster, traffic will not traverse through the device being upgraded while it is upgrading. However, the other devices in the cluster will continue to pass traffic.

Procedure

Step 1   Choose **System** > **Updates**.
The Available Updates page shows a list of the FXOS platform bundle images and application images that are available on the chassis.

Step 2   Click **Upgrade** for the FXOS platform bundle to which you want to upgrade.

The system will first verify the software package that you want to install. It will inform you of any incompatibility between currently installed applications and the specified FXOS platform software package. It will also warn you that any existing sessions will be terminated and that the system will need to be rebooted as part of the upgrade.

Step 3 Click **Yes** to confirm that you want to proceed with installation, or click **No** to cancel the installation.

The FXOS unpacks the bundle and upgrades/reloads the components.

Downloading a Logical Device Software Image to the Firepower 4100/9300 chassis

You can use FTP, HTTP/HTTPS, SCP, SFTP, or TFTP to copy the logical device software image to the Firepower 4100/9300 chassis.

Before you begin

Collect the following information that you will need to import a configuration file:

- IP address and authentication credentials for the server from which you are copying the image

- Fully qualified name of the software image file

Note   FXOS 2.8.1 and later versions support HTTP/HTTPS protocols for firmware and application image downloads.

Procedure

Step 1 Enter Security Services mode:

Firepower-chassis # **scope ssa**

Step 2 Enter Application Software mode:

Firepower-chassis /ssa # **scope app-software**

Step 3 Download the logical device software image:

Firepower-chassis /ssa/app-software # **download image** *URL*

Specify the URL for the file being imported using one of the following syntax:

- **ftp**://*username@hostname*/*path*

- **http**://*username@hostname*/*path*

- **https**://*username@hostname*/*path*

- **scp**://*username@hostname*/*path*

- **sftp**://*username@hostname*/*path*

- **tftp**://*hostname*:*port-num*/*path*

  Note   Do not use tftpdnld to install the image as it throws error.

Step 4   To monitor the download process:

Firepower-chassis /ssa/app-software # **show   download-task**

Step 5   To view the downloaded applications:

Firepower-chassis /ssa/app-software # **up**

Firepower-chassis /ssa # **show   app**

Step 6   To view details for a specific application:

Firepower-chassis /ssa # **scope   app**   *application_type   image_version*

Firepower-chassis /ssa/app # **show   expand**

Example

The following example copies an image using the SCP protocol:

```
Firepower-chassis # scope ssa
Firepower-chassis /ssa # scope app-software
Firepower-chassis /ssa/app-software # download image
scp://user@192.168.1.1/images/cisco-asa.9.4.1.65.csp
Firepower-chassis /ssa/app-software # show download-task

Downloads for Application Software:
    File Name                      Protocol   Server               Userid          State
    ----------------------------- ---------- -------------------- --------------- -----
    cisco-asa.9.4.1.65.csp         Scp        192.168.1.1          user            Downloaded

Firepower-chassis /ssa/app-software # up

Firepower-chassis /ssa # show app

Application:
    Name       Version    Description Author     Deploy Type CSP Type    Is Default App
    ---------- ---------- ----------- ---------- ----------- ----------- --------------
    asa        9.4.1.41   N/A                    Native      Application No
    asa        9.4.1.65   N/A                    Native      Application Yes

Firepower-chassis /ssa # scope app asa 9.4.1.65
Firepower-chassis /ssa/app # show expand

Application:
    Name: asa
    Version: 9.4.1.65
    Description: N/A
    Author:
    Deploy Type: Native
    CSP Type: Application
    Is Default App: Yes
```

```
        App Attribute Key for the Application:
            App Attribute Key Description
            ----------------- -----------
            cluster-role      This is the role of the blade in the cluster
            mgmt-ip           This is the IP for the management interface
            mgmt-url          This is the management URL for this application

        Net Mgmt Bootstrap Key for the Application:
            Bootstrap Key Key Data Type Is the Key Secret Description
            ------------- ------------- ----------------- -----------
            PASSWORD      String        Yes               The admin user password.

        Port Requirement for the Application:
            Port Type: Data
            Max Ports: 120
            Min Ports: 1

            Port Type: Mgmt
            Max Ports: 1
            Min Ports: 1

            Mgmt Port Sub Type for the Application:
                Management Sub Type
                ------------------
                Default

            Port Type: Cluster
            Max Ports: 1
            Min Ports: 0
Firepower-chassis /ssa/app #
```

Updating the Image Version for a Logical Device

Use this procedure to upgrade the ASA application image to a new version, or set the FTD application image to a new startup version that will be used in a disaster recovery scenario.

When you change the startup version on a FTD logical device using Firepower Chassis Manager or the FXOS CLI, the application does not immediately upgrade to the new version. The logical device startup version is the version that FTD reinstalls to in a disaster recovery scenario. After initial creation of a FTD logical device, you do not upgrade the FTD logical device using Firepower Chassis Manager or the FXOS CLI. To upgrade a FTD logical device, you must use FMC. See the System Release Notes for more information: http://www.cisco.com/c/en/us/support/security/defense-center/products-release-notes-list.html.

Also, note that any updates to the FTD logical device will not be reflected on the **Logical Devices** > **Edit** and **System** > **Updates** pages in Firepower Chassis Manager. On these pages, the version shown indicates the software version (CSP image) that was used to create the FTD logical device.

Note   When you set the startup version for FTD, startup version of the application gets updated. Hence, you must manually reinstall the application or reinitialize the blade to apply the selected version. This procedure is not the equivalent of upgrading or downgrading the FTD software, rather a complete reinstallation (reimage). Therefore, the application gets deleted and the existing configuration gets lost.

When you change the startup version on an ASA logical device, the ASA upgrades to that version and all configuration is restored. Use the following workflows to change the ASA startup version, depending on your configuration:

Note When you set the startup version for ASA, the application gets automatically restarted. This procedure is the equivalent of upgrading or downgrading the ASA software (existing configuration gets preserved).

ASA High Availability -

1. Change the logical device image version(s) on the standby unit.

2. Make the standby unit active.

3. Change the application version(s) on the other unit.

ASA Inter-Chassis Cluster -

1. Change the startup version on the data unit.

2. Make the data unit the control unit.

3. Change the startup version on the original control unit (now data).

Before you begin

Download the application image you want to use for the logical device from Cisco.com (see Downloading Images from Cisco.com, on page 56) and then upload that image to the Firepower 4100/9300 chassis (see Uploading an Image to the Security Appliance, on page 56).

If you are upgrading both the Platform Bundle image and one or more Application images, you must upgrade the Platform Bundle first.

Procedure

Step 1 Choose **Logical Devices** to open the Logical Devices page.
The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown instead.

Step 2 Click **Update Version** for the logical device that you want to update to open the **Update Image Version** dialog box.

Step 3 For the **New Version**, choose the software version.

Step 4 Click **OK**.

Firmware Upgrade

For information about upgrading the firmware on your Firepower 4100/9300 chassis, see the *Cisco Firepower 4100/9300 FXOS Firmware Upgrade Guide*.

Manually Downgrading to Version 2.0.1 or Lower

Follow these CLI steps to manually downgrade the CIMC image on a security module.

> ✎
>
> **Note**   This procedure is used specifically to downgrade to version 2.0.1 or lower, from version 2.1.1 or higher.

Before you begin

Ensure the application image you want to downgrade to has been downloaded to the Firepower 4100/9300 chassis (see Downloading Images from Cisco.com, on page 56 and Downloading a Logical Device Software Image to the Firepower 4100/9300 chassis, on page 58).

Procedure

Step 1   Disable image version comparison before downgrading the CIMC image.

Follow the steps in this example to clear the default platform image version:

Example:

```
firepower# scope org
firepower /org # scope fw-platform-pack default
firepower /org/fw-platform-pack # set platform-bundle-version ""
Warning: Set platform version to empty will result software/firmware incompatibility issue.
firepower /org/fw-platform-pack* # commit-buffer
firepower /org/fw-platform-pack #
```

Step 2   Downgrade the module image.

Follow the steps in this example to change the CIMC image:

Example:

```
firepower# scope server 1/1
firepower /chassis/server # scope cimc
firepower /chassis/server/cimc # update firmware <version_num>
firepower /chassis/server/cimc* # activate firmware <version_num>
firepower /chassis/server/cimc* # commit-buffer
firepower /chassis/server/cimc #
```

Repeat this step as necessary to update other modules.

Step 3   Install the new firmware bundle.

Follow the steps in this example to install the downgrade image:

Example:

```
firepower# scope firmware
firepower /firmware # scope auto-install
firepower /firmware/auto-install # install platform platform-vers <version_num>
The currently installed FXOS platform software package is <version_num>

WARNING: If you proceed with the upgrade, the system will reboot.
```

```
This operation upgrades firmware and software on Security Platform Components
Here is the checklist of things that are recommended before starting Auto-Install
(1) Review current critical/major faults
(2) Initiate a configuration backup
Do you want to proceed? (yes/no):
```

What to do next

You can use the **show fsm status expand** command in `firmware/auto-install` mode to monitor the installation process.

CHAPTER 6

Security Certifications Compliance

Security Certifications Compliance

United States federal government agencies are sometimes required to use only equipment and software complying with security standards established by the U.S. Department of Defense and global certification organizations. The Firepower 4100/9300 chassis supports compliance with several of these security certification standards.

See the following topics for steps to enable features that support compliance with these standards:

• Enable FIPS Mode

• Enable Common Criteria Mode

• Configure IPSec Secure Channel, on page 67

• Configure Static CRL for a Trustpoint, on page 73

• About the Certificate Revocation List Check, on page 74

• Configure CRL Periodic Download, on page 78

• Setting the Date and Time Using NTP, on page 102

• Set the LDAP Key Ring Certificate, on page 79

• Configure the IP Access List, on page 145

• Configure Minimum Password Length Check

• Set the Maximum Number of Login Attempts, on page 50

Note

Note that these topics discuss enabling certifications compliance on the Firepower 4100/9300 chassis only. Enabling certification compliance on the Firepower 4100/9300 chassis does not automatically propagate compliance to any of its attached logical devices.

Generate the SSH Host Key

Prior to FXOS release 2.0.1, the existing SSH host key created during initial setup of a device was hard-coded to 1024 bits. To comply with FIPS and Common Criteria certification, you must destroy this old host key and generate a new one. See Enable FIPS Mode or Enable Common Criteria Mode for more information.

Perform these steps to destroy the old SSH host key and generate a new certifications-compliant one.

Procedure

Step 1 From the FXOS CLI, enter services mode:

scope system

scope services

Step 2 Delete the SSH host key:

delete ssh-server host-key

Step 3 Commit the configuration:

commit-buffer

Step 4 Set the SSH host key size to 2048 bits:

set ssh-server host-key rsa 2048

Step 5 Commit the configuration:

commit-buffer

Step 6 Create a new SSH host key:

create ssh-server host-key

commit-buffer

Step 7 Confirm the new host key size:

show ssh-server host-key

Host Key Size: 2048

Configure IPSec Secure Channel

IPSec is a framework of open standards developed by the Internet Engineering Task Force (IETF). It creates secure, authenticated, and reliable communication over IP networks. The IPSec security service provides:

- Connectionless Integrity – Assurance the received traffic has not been modified.

- Data origin authentication – Assurance the traffic is sent by legitimate party.

- Confidentiality (encryption) – Assurance the user's traffic is not examined by non-authorized parties.

- Access control – Prevention of unauthorized use of a resource.

IPSec Secure Channel supports the following algorithms:

- Phase 1

```
aes128gcm16-prfsha384-prfsha512-prfsha256-prfsha1-ecp256-ecp384-ecp521-modp2048-modp3072-modp4096
aes128-aes192-aes256-sha256-sha384-sha1_160-sha1-sha512-prfsha384-prfsha512-prfsha256-prfsha1-ecp256-ecp384-ecp521
aes128-aes192-aes256-sha256-sha384-sha1_160-sha1-sha512-prfsha384-prfsha512-prfsha256-prfsha1-modp2048-modp3072-modp4096
```

- Phase 2

    - Only AES SHA based encryption algorithms are supported. (DES and MD5 are not supported)

    - Supported DH groups are 14,15,16,19,20, and 21.

Note   IPSec connections can only be initiated from FXOS. FXOS does not accept incoming IPSec connection requests.

IPsec tunnels are sets of SAs that FXOS establishes between peers. The SAs specify the protocols and algorithms to apply to sensitive data and also specify the keying material that the peers use. IPsec SAs control the actual transmission of user traffic. SAs are unidirectional, but are generally established in pairs (inbound and outbound).

IPSec on Chassis Manager has two modes:

Transport Mode
    IP Header, IPSec Header, TCP Header, Data

Tunnel Mode
    New IP Header, IPSec Header, Original IP Header, TCP Header, Data

IPSec's operation can be broken down into five main steps:

1. Traffic Selection – Interesting traffic which matches IPSec policy starts the IKE process. For example, traffic can be selected using src/dst host IP or subnet. Alternatively, user also can trigger IKE process through admin command.

2. IKE Phase 1 – authenticate IPSec peers and to setup a secure channel to enable IKE exchanges

3. IKE phase 2 – negotiate SAs to set up the IPSec tunnel. SA stands for Security Association, it is a relationship between IPSec end-points that describe what security services are used to protect data traffic.

4. Data transfer – Data packets are encrypted and encapsulated in IPSec header using parameters and keys stored in the SA

5. IPSec tunnel termination – IPSec SAs terminate through deletion or by timing out.

You can configure IPSec on your Firepower 4100/9300 chassis to provide end-to-end data encryption and authentication service on data packets going through the public network. This option is one of a number offered for achieving Common Criteria certification compliance on your system. For more information, see Security Certifications Compliance, on page 65.

> **Note**
> - If you are using an IPSec secure channel in FIPS mode, the IPSec peer must support RFC 7427.
> - If you elect to configure enforcement of matching cryptographic key strength between IKE and SA connections (set sa-strength-enforcement to yes in the below procedure):
>
> | If SA enforcement is enabled | then when IKE negotiated key size is less then ESP negotiated key size, the connection fails. |
> |---|---|
> | | then when IKE negotiated key size is large or equal than ESP negotiated key size, SA enforcement check passes and the connection is successful. |
> | If SA enforcement is disabled | then SA enforcement check passes and the connection is successful. |

Perform these steps to configure an IPSec secure channel.

Procedure

Step 1   From the FXOS CLI, enter security mode:

scope security

Step 2   Create the keyring:

enter keyring ssp

! create certreq subject-name *subject-name* **ip** *ip*

Step 3   Enter the associated certificate request information:

enter certreq

Step 4   Set the country:

set country *country*

Step 5   Set the DNS:

set dns *dns*

Step 6   Set the email:

set e-mail *email*

Step 7   Set the IP information:

set ip *ip-address*

set ipv6 *ipv6*

Step 8   Set the locality:

set locality *locality*

Step 9   Set the organization name:

set org-name *org-name*

Step 10   Set the organization unit name:

set org-unit-name *org-unit-name*

Step 11   Set the password:

! set password

Step 12   Set the state:

set state *state*

Step 13   Set the subject name for the certreq:

set subject-name *subject-name*

Step 14   Exit:

exit

Step 15   Set the modulus:

set modulus *modulus*

Step 16   Set the regeneration for the certificate request:

set regenerate *{ yes | no }*

Step 17   Set the trustpoint:

set trustpoint interca

Step 18   Exit:

exit

Step 19   Enter the newly created trustpoint:

enter trustpoint interca

Step 20   Generate certificate signing request:

set certchain

Example:

```
-----BEGIN CERTIFICATE-----
MIIF3TCCA8WgAwIBAgIBADANBgkqhkiG9w0BAQsFADBwMQswCQYDVQQGEwJVUzEL
MAkGA1UECAwCQ0ExDDAKBgNVBAcMA1NKQzEOMAwGA1UECgwFQ2lzY28xDTALBgNV
```

BAsMBFNUQlUxCzAJBgNVBAMMAkNBMRowGAYJKoZIhvcNAQkBFgtzc3BAc3NwLm5l
dDAeFw0xNjEyMDgxOTMzNTJaFw0yNjEyMDYxOTMzNTJaMHAxCzAJBgNVBAYTAlVT
MQswCQYDVQQIDAJDQTEMMAoGA1UEBwwDU0pDMQ4wDAYDVQQKDAVDaXNjbzENMAsG
A1UECwwEU1RCVTELMAkGA1UEAwwCQ0ExGjAYBgkqhkiG9w0BCQEWC3NzcEBzc3Au
bmV0MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEA2ukWyMLQuLqTvhq7
zFb3Oz/iyDG/ui6mrLIYn8wE3E39XcXA1/x9IHCmxFKNJdD7EbsggfOuy0Bj+Y4s
+uZ1VapBXV/JrAie7bNn3ZYrI29yuyOrIqoi9k9gL/oRBzH18BwBwGHBOz3hGrSK
Yc2yhsq9y/6yI3nSuLZm6ybmUKjTa+B4YuhDTz4hl/I9x/J5nbGiab3vLDKss1nO
xP9+1+Lc690Vl8/mNPWdjCjDI+U/L9keYs/rbZdRSeXy9kMae42+4FlRHDJjPcSN
Yw1g/gcR2F7QUKRygKckJKXDX2QIiGYSctlSHj18O87o5s/pmQAWWRGkKpfDv3oH
cMPgI2T9rC0D8NNcgPXj9PFKfexoGNGwNTO85fK3kjgMOdWbdeMG3EihxEEOUPD0
Fdu0HrTM5lvwb+vr5wE9HsAiMJ8UuujmHqH5mlwyy3Me+cEDHo0hLeNs+AFrqEXQ
e9S+KZC/dq/9zOLpRsVqSfJsAuVl/QdPDbWShjflE/fP2Wj01PqXywQydzymVvgE
wEZaoFg+mlGJm0+q4RDvnpzEviOYNSAGmOkILh5HQ/eYDcxvd0qbORWb31H32ySl
Ila6UTT9+vnND1f838fxvNvr8nyGD2S/LVaxnZlO4jcSIvtdizbbT8u5B4VcLKIC
x0vkqjo6RvNZJ52sUaD9C3UodTUCAwEAAaOBgTB/MC8GA1UdHwQoMCYwJKAioCCG
Hmh0dHA6Ly8xOTIuMTY4LjUuQuMjkvcm9vdGNhLmNybDAdBgNVHQ4EFgQU7Jg01A74
jpx8U0APk76pVfYQQ5AwHwYDVR0jBBgwFoAU7Jg01A74jpx8U0APk76pVfYQQ5Aw
DAYDVR0TBAUwAwEB/zANBgkqhkiG9w0BAQsFAAOCAgEAvI8ky2jiXc4wPiMuxIfY
W7DRmszPUWQ7edor7yxuCqzHLVFFOwYRudsyXbv7INR3rJ/X1cRQj9+KidWWVxpo
pFahRhzYxVZ10DHKIzGTQS3jiHgrF3Z8ohWbL15L7PEDlrxMBoJvabPeQRgTmY/n
XZJ7qRYbypO3gUMCaCZ12raJc3/DIpBQ29yweCbUkc9qiHKA0IbnvAxoroHWmBld
94LrJCggfMQTuNJQszJiVVsYJfZ+utlDp2QwfdDv7B0JkwTBjdwRSfotEbc5R18n
BNXYHqxuoNMmqbS3KjCLXcH6xIN8t+UkfP89hvJt/fluJ+s/VJSVZWK4tAWvR7wl
QngCKRJW6FYpzeyNBctiJ07wO+Wt4e3KhIjJDYvA9hFixWcVGDf2r6QW5BYbgGOK
DkHb/gdr/bcdLBKN/PtSJ+prSrpBSaA6rJX8D9UmfhqqN/3f+sS1fM4qWORJc6G2
gAcg7AjEQ/0do512vAI8p8idOg/Wv1O17mavzLpcue05cwMCX9fKxKZZ/+7Pk19Y
ZrXS6uMn/CGnViptn0w+uJ1IRj1oulk+/ZyPtBvFHUkFRnhoWj5SMFyds2IaatyI
47N2ViaZBxhU3GICaH+3O+8rs9Kkz9tBZDSnEJVZA6yxaNCVP1bRUO20G3oRTmSx
8iLBjN+BXggxMmG8ssHisgw=
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIIFqDCCA5CgAwIBAgIBBDANBgkqhkiG9w0BAQsFADBwMQswCQYDVQQGEwJVUzEL
MAkGA1UECAwCQ0ExDDAKBgNVBAcMA1NKQzEOMAwGA1UECgwFQ2lzY28xDTALBgNV
BAsMBFNUQlUxCzAJBgNVBAMMAkNBMRowGAYJKoZIhvcNAQkBFgtzc3BAc3NwLm5l
dDAeFw0xNjEyMTUyMTM0NTRaFw0yNjEyMTMyMTM0NTRaMHwxCzAJBgNVBAYTAlVT
MQswCQYDVQQIDAJDQTEPMA0GA1UECgwGbV3c3RnMRAwDgYDVQQLDAduZXdzdGJ1
MRMwEQYDVQQDDAppbRlcm0xLWNhMSgwJgYJKoZIhvcNAQkBFhlpbRlcm0xLWNh
QGludGVybTEtY2EubmV0MIICIjANBgkqhkiG9w0BAQEFAAOCAg8AMIICCgKCAgEA
wLpNnyEx5I4P8uDoWKWF3IZsegjhLANsodxuAUmhmwKekd0OpZZxHMw1wSO4IBX5
4itJS0xyXFzPmeptG3OXvNqCcsT+4BXl3DoGgPMULccc4NesHeg2z8+q3SPA6uZh
iseWNvKfnUjixbQEBtcrWBiSKnZuOz1cpuBn34gtgeFFoCEXN+EZVpPESiancDVh
8pCPlipc/08ZJ3o9GW2j0eHJN84sguIEDL812ROejQvpmfqGUq11stkIIuh+wB+V
VRhUBVG7pV57I6DHeeRp6cDMLXaM3iMTelhdShyo5YUaRJMak/t8kCqhtGXfuLlI
E2AkxKXeeveR9n6cpQd5JiNzCT/t9IQL/T/CCqMICRXLFpLCS9o5S5O2B6QFgcTZ
yKR6hsmwe22wpK8QI7/5oWNXlolb96hHJ7RPbG7RXYqmcLiXY/d2j9/RuNoPJawI
hLkfhoIdPA28xlnfIB1azCmMmdPcBO6cbUQfCj5hSmk3StVQKgJCjaujz55TGGd1
GjnxDMX9twwz7Ee51895Xmtr24qqaCXJoW/dPhcIIXRdJPMsTJ4yPG0BieuRwd0p
i8w/rFwbHzv4C9Fthw1JrRxH1yeHJHrLlZgJ5txSaVUIgrgVCJaf6/jrRRWoRJwt
AzvnzYql2dZPCcEAYgP7JcaQpvdpuDgq++NgBtygiqECAwEAAaNBMD8wDAYDVR0T
BAUwAwEB/zAvBgNVHR8EKDAmMCSgIqAghh5odHRwOi8vMTkyLjE2OC40LjI5L2lu
dGVrbS5jmwwDQYJKoZIhvcNAQELBQADggIBAG/XujJh5G5UWo+cwTSitAezWbJA
h1dAiXZ/OYWZSxkFRliErKdupLqL0ThjnX/wRFfEXbrBQwm5kWAUUDr97D1Uz+2A
8LC5I8SWKXmyf0jUtsnEQbDZb33oVL7yXJk/A0SF0jihpPheMA+YRazalT9xj9KH
PE7nHCJMbb2ptrHUyvBrKSYrSeEqOpQU2+otnFyV3rS9aelgVjuaWyaWOc3lZ1Oi
CC2tJvY3NnM56j5iesxUCeY/SZ2/ECXN7RRBViLHmA3gFKmWf3xeNiKkxmJCxOaa
UWPC1x2V66I8DG9uUzlWyd79O2dy52aAphAHC6hqlzb6v+gw1Tld7UxaqVd8CD5W

ATjNs+ifkJS1h5ERxHjgcurZXOpR+NWpwF+UDzbMXxx+KAAXCI6ltCd8Pb3wOUC3
PKvwEXaIcCcxGx71eRLpWPZFyEoi4N2NGE9OXRjz0K/KERZgNhsIW3bQMjcw3aX6
OXskEuKGgsayctnWyxVqNnqvpuz06kqyubh4+ZgGKZ5LNEXYmGNz3oED1rUN636Tw
SjGAPHgeROzyTFDixCei6aROlGdP/Hwvb0/+uThIe89g8WZ0djTKFUM8uBO3f+II
/cbuyBO1+JrDMq8NkAjxKlJlp1c3WbfCue/qcwtcfUBYZ4i53a56UNF5Ef0rpy/8
B/+07Me/p2y9Luqa
-----END CERTIFICATE-----
ENDOFBUF

Step 21   Show the certificate signing request:

show certreq

Example:

Firepower-chassis# /security/keyring # show certreq
Certificate request subject name: SSP
Certificate request ip address: 192.168.0.111
Certificate request FI A ip address: 0.0.0.0
Certificate request FI B ip address: 0.0.0.0
Certificate request e-mail name:
Certificate request ipv6 address: ::
Certificate request FI A ipv6 address: ::
Certificate request FI B ipv6 address: ::
Certificate request country name: US
State, province or county (full name): CA
Locality name (eg, city): SJC
Organisation name (eg, company): Cisco
Organisational Unit Name (eg, section): Sec
DNS name (subject alternative name):
Request:
-----BEGIN CERTIFICATE REQUEST-----
MIICwTCCAakCAQAwVTELMAkGA1UEBhMCVVMxCzAJBgNVBAgMAkNBMQwwCgYDVQQH
DANTSkMxDjAMBgNVBAoMBUNpc2NvMQwwCwYDVQQLDARTEJVMQwwCgYDVQQDDANT
U1AwggEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQDq292Rq3t0laoxPbfE
p/lTKr6rxFhPqSSbtm6sXer//VZFiDTWODockDItuf4Kja215mIS0RyvEYVeRgAs
wbN459wm0BASd8xCjIhsuHDV7yHu539BnvRW6Q2o+gHeSRwckqjClK/tsIxsPkV0
6OduZYXk2bnsLWs6tNk3uzOIT2Q0FcZ1ET66C8fyyKWTrmvcZjDjkMm2nDFsPlX9
39TYPItDkJE3PocqyaCqmT4uobOuvQeLJh/efkBvwhb4BF8vwzRpHWTdjjU5YnR1
qiR4q7j1RmzVFxCDY3IVP/KDBoa5NyCLEUZECP5QCQFDzIRETZwVOKtxUVG0Nljd
K5TxAgMBAAGgJzAlBgkqhkiG9w0BCQ4xGDAWMBQGA1UdEQQNMAuCA1NTUIcEwKgA
rjANBgkqhkiG9w0BAQsFAAOCAQEARtRBoInxXkBYNlVeEoFCqKttu3+Hc7UdyoRM
2L2pjx5OHbQICC+8NRVRMYujTnp67BWuUZZl03dGP4/lbN6bC9P3CvkZdKUsJkN0
m1Ye9dgz7MO/KEcosarmoMl9WB8LlweVdt6ycSdJzs9shOxwT6TAZPwL7gq/1ShF
RJh6sq5W9p6E0SjYefK62E7MatRjDjS8DXoxj6gfn9DqK15iVpkK2QqT5rneSGj+
R+20TcUnT0h/S5K/bySEM/3U1gFxQCOzbzPuHkj28kXAVczmTxXEkJBFLVduWNo6
DT3u0xImiPR1sqW1jpMwbhC+ZGDtvgKjKHToagup9+8R9IMcBQ==
-----END CERTIFICATE REQUEST-----

Step 22   Enter IPSec mode:

scope ipsec

Step 23   Set the log verbose level:

set log-level *log_level*

Step 24   Create and enter an IPSec connection:

	enter connection *connection_name*
Step 25	Set IPSec mode to tunnel or transport:
	set mode *tunnel_or_transport*
Step 26	Set the local IP address:
	set local-addr *ip_address*
Step 27	Set the remote IP address:
	set remote-addr *ip_address*
Step 28	If using tunnel mode, set the remote subnet:
	set remote-subnet *ip/mask*
Step 29	(Optional) Set the remote identity:
	set remote-ike-ident *remote_identity_name*
Step 30	Set the keyring name:
	set keyring-name *name*
Step 31	(Optional) Set the keyring password:
	set keyring-passwd *passphrase*
Step 32	(Optional) Set the IKE-SA lifetime in minutes:
	set ike-rekey-time *minutes*
	The *minutes* value can be any integer between 60-1440, inclusive.
Step 33	(Optional) Set the Child SA lifetime in minutes (30-480):
	set esp-rekey-time *minutes*
	The *minutes* value can be any integer between 30-480, inclusive.
Step 34	(Optional) Set the number of retransmission sequences to perform during initial connect:
	set keyringtries *retry_number*
	The *retry_number* value can be any integer between 1-5, inclusive.
Step 35	(Optional) Enable or disable the certificate revocation list check:
	set revoke-policy *{ relaxed	strict }*
Step 36	Enable the connection:
	set admin-state enable
Step 37	Reload connections:
	reload-conns
	The system stops all connections and then reloads them. All connections will try to re-establish.
Step 38	(Optional) Add the existing trustpoint name to IPsec:

create authority *trustpoint_name*

Step 39 Configure the enforcement of matching cryptographic key strength between IKE and SA connections:

set sa-strength-enforcement *yes_or_no*

Configure Static CRL for a Trustpoint

Revoked certifications are kept in the Certification Revocation List (CRL). Client applications use the CRL to check the authentication of a server. Server applications utilize the CRL to grant or deny access requests from client applications which are no longer trusted.

You can configure your Firepower 4100/9300 chassis to validate peer certificates using Certification Revocation List (CRL) information. This option is one of a number offered for achieving Common Criteria certification compliance on your system. For more information, see Security Certifications Compliance, on page 65.

Perform these steps to validate peer certificates using CRL information.

Procedure

Step 1 From the FXOS CLI, enter security mode:

scope security

Step 2 Enter trustpoint mode:

scope trustpoint *trustname*

Step 3 Enter revoke mode:

scope revoke

Step 4 Download the CRL file(s):

import crl *protocol*://*user_id@CA_or_CRL_issuer_IP*/tmp/*DoDCA1CRL1.crl*

Note DER format static CRL is not supported in FXOS. You must convert the DER format CRL file to PEM format using the following command:

```
openssl crl -in filename.crl -inform DER -outform PEM -out crl.pem
```

Step 5 (Optional) Show the status of the import process of CRL information:

show import-task detail

Step 6 Set the certificate revocation method to CRL-only:

set certrevokemethod {crl}

About the Certificate Revocation List Check

You can configure your Certificate Revocation List (CRL) check mode to be either strict or relaxed in IPSec and secure LDAP connections.

FXOS harvests dynamic (non-static) CRL information from the CDP information of an X.509 certificate, which indicates dynamic CRL information. System administration downloads static CRL information manually, which indicates local CRL information in the FXOS system. FXOS processes dynamic CRL information against the current processing certificate in the certificate chain. The static CRL is applied to the whole peer certificate chain.

For steps to enable or disable certificate revocation checks for your secure LDAP and IPSec connections, see Configure IPSec Secure Channel, on page 67 and Creating an LDAP Provider, on page 134.

Note

- If the Certificate Revocation Check Mode is set to Strict, static CRL is only applicable when the peer certificate chain has a level of 1 or higher. (For example, when the peer certificate chain contains only the root CA certificate and the peer certificate signed by the root CA.)

- When configuring static CRL for IPSec, the Authority Key Identifier (authkey) field must be present in the imported CRL file. Otherwise, IPSec considers it invalid.

- Static CRL takes precedence over Dynamic CRL from the same issuer. When FXOS validates the peer certificate, if a valid (determined) static CRL of the same issuer exists, FXOS ignores the CDP in the peer certificate.

- Strict CRL checking is enabled by default in the following scenarios:
    - Newly created secure LDAP provider connections, IPSec connections, or Client Certificate entries
    - Newly deployed FXOS chassis managers (deployed with an initial starting version of FXOS 2.3.1.x or later)

The following tables describe the connection results, depending on your certificate revocation list check setting and certificate validation.

Table 6: Certificate Revocation Check Mode set to Strict without a local static CRL

Without local static CRL	LDAP Connection	IPSec Connection
Checking peer certificate chain	Full certificate chain is required	Full certificate chain is required
Checking CDP in peer certificate chain	Full certificate chain is required	Full certificate chain is required
CDP checking for Root CA certificate of the peer certificate chain	Yes	Not applicable
Any certificate validation failure in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message

Without local static CRL	LDAP Connection	IPSec Connection
Any certificate revoked in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
One CDP is missing in the peer certificate chain	Connection fails with syslog message	Peer certificate: connection fails with syslog message Intermediate CAs: connection succeeded
One CDP CRL is empty in the peer certificate chain with valid signature	Connection fails with syslog message	Connection succeeds
Any CDP in the peer certificate chain cannot be downloaded	Connection fails with syslog message	Peer certificate: Connection fails with syslog message Intermediate CA: connection succeeded
Certificate has CDP, but the CDP server is down	Connection fails with syslog message	Peer certificate: Connection fails with syslog message Intermediate CA: connection succeeded
Certificate has CDP, server is up, and CRL is on CDP, but the CRL has an invalid signature	Connection fails with syslog message	Peer certificate: Connection fails with syslog message Intermediate CA: connection succeeded

Table 7: Certificate Revocation Check Mode set to Strict with a local static CRL

With local static CRL	LDAP Connection	IPSec Connection
Checking peer certificate chain	Full certificate chain is required	Full certificate chain is required
Checking CDP in peer certificate chain	Full certificate chain is required	Full certificate chain is required
CDP checking for Root CA certificate of the peer certificate chain	Yes	Not applicable
Any certificate validation failure in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
Any certificate revoked in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
One CDP is missing in the peer certificate chain (Certificate Chain level is 1)	Connection succeeds	Connection succeeds

With local static CRL	LDAP Connection	IPSec Connection
One CDP CRL is empty in the peer certificate chain (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Any CDP in the peer certificate chain cannot be downloaded (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Certificate has CDP, but the CDP server is down (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Certificate has CDP, server is up, and CRL is on CDP, but the CRL has an invalid signature (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Peer Certificate Chain level is higher than 1	Connection fails with syslog message	If combined with CDP, connection succeeds If there is no CDP, connection fails with syslog message

Table 8: Certificate Revocation Check Mode set to Relaxed without a local static CRL

Without local static CRL	LDAP Connection	IPSec Connection
Checking peer certificate chain	Full certificate chain	Full certificate chain
Checking CDP in the peer certificate chain	Full certificate chain	Full certificate chain
CDP checking for Root CA certificate of the peer certificate chain	Yes	Not applicable
Any certificate validation failure in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
Any certificate revoked in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
One CDP is missing in the peer certificate chain	Connection succeeds	Connection succeeds
One CDP CRL is empty in the peer certificate chain	Connection succeeds	Connection succeeds
Any CDP in the peer certificate chain cannot be downloaded	Connection succeeds	Connection succeeds
Certificate has CDP, but the CDP server is down	Connection succeeds	Connection succeeds

Without local static CRL	LDAP Connection	IPSec Connection
Certificate has CDP, server is up, and CRL is on CDP, but the CRL has an invalid signature	Connection succeeds	Connection succeeds

Table 9: Certificate Revocation Check Mode set to Relaxed with a local static CRL

With local static CRL	LDAP Connection	IPSec Connection
Checking peer certificate chain	Full certificate chain	Full certificate chain
Checking CDP in the peer certificate chain	Full certificate chain	Full certificate chain
CDP checking for Root CA certificate of the peer certificate chain	Yes	Not applicable
Any certificate validation failure in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
Any certificate revoked in the peer certificate chain	Connection fails with syslog message	Connection fails with syslog message
One CDP is missing in the peer certificate chain (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
One CDP CRL is empty in the peer certificate chain (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Any CDP in the peer certificate chain cannot be downloaded (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Certificate has CDP, but the CDP server is down (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Certificate has CDP, server is up, and CRL is on CDP, but the CRL has an invalid signature (Certificate Chain level is 1)	Connection succeeds	Connection succeeds
Peer Certificate Chain level is higher than 1	Connection fails with syslog message	If combined with CDP, connection succeeds If there is no CDP, connection fails with syslog message

Configure CRL Periodic Download

You can configure your system to periodically download a (CRL) so that a new CRL is used every 1 to 24 hours to validate certificates.

You can use the following protocols and interfaces with this feature:

- FTP

- SCP

- SFTP

- TFTP

- USB

Note

- SCEP and OCSP are not supported.

- You can only configure one periodic download per CRL.

- One CRL is supported per trustpoint.

Note   You can only configure the period in one-hour intervals.

Perform these steps to configure CRL periodic download.

Before you begin

Ensure that you have already configured your Firepower 4100/9300 chassis to validate peer certificates using (CRL) information. For more information, see Configure Static CRL for a Trustpoint, on page 73.

Procedure

Step 1   From the FXOS CLI, enter security mode:

scope security

Step 2   Enter trustpoint mode:

scope trustpoint

Step 3   Enter revoke mode:

scope revoke

Step 4   Edit the revoke configuration:

sh config

Step 5 Set your preferred configuration:

Example:

```
set certrevokemethod crl
set crl-poll-filename rootCA.crl
set crl-poll-path /users/myname
set crl-poll-period 1
set crl-poll-port 0
set crl-poll-protocol scp
! set crl-poll-pwd
set crl-poll-server 182.23.33.113
set crl-poll-user myname
```

Step 6 Exit the configuration file:

exit

Step 7 (Optional) Test the new configuration by downloading a new CRL:

Example:

```
Firepower-chassis /security/trustpoint/revoke # sh import-task

Import task:
File Name Protocol Server          Port    Userid       State
--------- -------- --------------- ---------- -------------- -----
rootCA.crl Scp     182.23.33.113   0       myname       Downloading
```

Set the LDAP Key Ring Certificate

You can configure a secure LDAP client key ring certificate to support a TLS connection on your Firepower 4100/9300 chassis. This option is one of a number offered for achieving Common Criteria certification compliance on your system. For more information, see Security Certifications Compliance, on page 65.

✎

Note If Common Criteria mode is enabled, you must have SSL enabled, and you must use the server DNS information to create the key ring certificate.

If SSL is enabled for the LDAP server entry, key ring information is referenced and checked when forming a connection.

LDAP server information has to be DNS information in the CC mode for the secure LDAP connection (with SSL enabled).

Perform these steps to configure a secure LDAP client key ring certificate:.

Procedure

Step 1 From the FXOS CLI, enter security mode:

scope security

Step 2 Enter LDAP mode:

scope ldap

Step 3 Enter LDAP server mode:

enter server *{server_ip|server_dns}*

Step 4 Set the LDAP key ring:

set keyring *keyring_name*

Step 5 Commit the configuration:

commit-buffer

CHAPTER **7**

System Administration

System Changes that Cause Firepower Chassis Manager Sessions to be Closed

The following system changes can cause the system to automatically log you out of Firepower Chassis Manager:

- If you modify the system time by more than 10 minutes.

- If the system is rebooted or shut down using Firepower Chassis Manager or the FXOS CLI.

- If you upgrade the FXOS version on Firepower 4100/9300 chassis.

- If you enable or disable FIPS or Common Criteria mode.

Note In addition to the above changes, you are automatically logged out of the system if a certain period of time passes without any activity. By default, the system will log you out after 10 minutes of inactivity. To configure this timeout setting, see Configuring the Session Timeout, on page 48. You can also configure an absolute timeout setting that will log users out of the system after a certain period of time even if the session is active. To configure the absolute timeout setting, see Configuring the Absolute Session Timeout, on page 49.

Changing the Management IP Address

Before you begin

You can change the management IP address on the Firepower 4100/9300 chassis from the FXOS CLI.

Note After changing the management IP address, you will need to reestablish any connections to Firepower Chassis Manager or the FXOS CLI using the new address.

Procedure

Step 2 To configure an IPv4 management IP address:

a) Set the scope for fabric-interconnect a:

Firepower-chassis# **scope fabric-interconnect a**

b) To view the current management IP address, enter the following command:

Firepower-chassis /fabric-interconnect # **show**

c) Enter the following command to configure a new management IP address and gateway:

Firepower-chassis /fabric-interconnect # **set out-of-band ip** *ip_address* **netmask** *network_mask* **gw** *gateway_ip_address*

d) Commit the transaction to the system configuration:

Firepower-chassis /fabric-interconnect* # **commit-buffer**

Step 3 To configure an IPv6 management IP address:

a) Set the scope for fabric-interconnect a:

Firepower-chassis# **scope fabric-interconnect a**

b) Set the scope for management IPv6 configuration:

Firepower-chassis /fabric-interconnect # **scope ipv6-config**

c) To view the current management IPv6 address, enter the following command:

Firepower-chassis /fabric-interconnect/ipv6-config # **show ipv6-if**

d) Enter the following command to configure a new management IP address and gateway:

Firepower-chassis /fabric-interconnect/ipv6-config # **set out-of-band ipv6** *ipv6_address* **ipv6-prefix** *prefix_length* **ipv6-gw** *gateway_address*

Note Only IPv6 Global Unicast addresses are supported as the chassis's IPv6 management address.

e) Commit the transaction to the system configuration:

Firepower-chassis /fabric-interconnect/ipv6-config* # **commit-buffer**

Example

The following example configures an IPv4 management interface and gateway:

```
Firepower-chassis# scope fabric-interconnect a
Firepower-chassis /fabric-interconnect # show

Fabric Interconnect:
    ID   OOB IP Addr     OOB Gateway     OOB Netmask     OOB IPv6 Address OOB IPv6 Gateway
 Prefix Operability
    ---- --------------- --------------- --------------- ---------------- ----------------
 ------ -----------
    A    192.0.2.112     192.0.2.1       255.255.255.0   ::               ::
 64     Operable
Firepower-chassis /fabric-interconnect # set out-of-band ip 192.0.2.111 netmask 255.255.255.0
 gw 192.0.2.1
Warning: When committed, this change may disconnect the current CLI session
Firepower-chassis /fabric-interconnect* #commit-buffer
Firepower-chassis /fabric-interconnect #
```

The following example configures an IPv6 management interface and gateway:

```
Firepower-chassis# scope fabric-interconnect a
Firepower-chassis /fabric-interconnect # scope ipv6-config
Firepower-chassis /fabric-interconnect/ipv6-config # show ipv6-if

Management IPv6 Interface:
    IPv6 Address                        Prefix     IPv6 Gateway
    ----------------------------------- ---------- ------------
    2001::8998                          64         2001::1
Firepower-chassis /fabric-interconnect/ipv6-config # set out-of-band ipv6 2001::8999
ipv6-prefix 64 ipv6-gw 2001::1
Firepower-chassis /fabric-interconnect/ipv6-config* # commit-buffer
Firepower-chassis /fabric-interconnect/ipv6-config #
```

Changing the Application Management IP

You can change the management IP address on the application(s) attached to your Firepower 4100/9300 chassis from the FXOS CLI. To do so, you must first change the IP information at the FXOS platform level, then change the IP information at the application level.

Note   Changing the application management IP will result in a service interruption.

Procedure

Step 1   Connect to the FXOS CLI. (See Accessing the FXOS CLI, on page 15).

Step 2   Scope to the logical device:

scope ssa

scope logical-device *logical_device_name*

Step 3   Scope to the management bootstrap and configure the new management bootstrap parameters. Note that there are differences between deployments:

For standalone configuration of an ASA logical device:

a)   Enter the logical device management bootstrap:

scope mgmt-bootstrap *asa*

b)   Enter the IP mode for the slot:

scope *ipv4_or_6 slot_number* default

c)   (IPv4 only) Set the new IP address:

set ip *ipv4_address* **mask** *network_mask*

d)   (IPv6 only) Set the new IP address:

set ip *ipv6_address* **prefix-length** *prefix_length_number*

e)   Set the gateway address:

set gateway *gateway_ip_address*

f)   Commit the configuration:

commit-buffer

For a clustered configuration of ASA logical devices:

a)   Enter the cluster management bootstrap:

scope cluster-bootstrap *asa*

b)   (IPv4 only) Set the new virtual IP:

set virtual ipv4 *ip_address* **mask** *network_mask*

c)   (IPv6 only) Set the new virtual IP:

set virtual ipv6 *ipv6_address* **prefix-length** *prefix_length_number*

d)   Set the new IP pool:

set ip pool *start_ip end_ip*

e)   Set the gateway address:

set gateway *gateway_ip_address*

f)   Commit the configuration:

commit-buffer

For standalone and clustered configurations of FTD:

a) Enter the logical device management bootstrap:

scope mgmt-bootstrap *ftd*

b) Enter the IP mode for the slot:

scope *ipv4_or_6 slot_number* firepower

c) (IPv4 only) Set the new IP address:

set ip *ipv4_address* **mask** *network_mask*

d) (IPv6 only) Set the new IP address:

set ip *ipv6_address* **prefix-length** *prefix_length_number*

e) Set the gateway address:

set gateway *gateway_ip_address*

f) Commit the configuration:

commit-buffer

Note   For a clustered configuration, you must set the new IP address for each application attached to the
Firepower 4100/9300 chassis. If you have an inter-chassis cluster or a HA configuration, you must
repeat these steps for each application on both chassis.

Step 4   Clear the management bootstrap information for each application:

a) Scope to ssa mode:

scope ssa

b) Scope to the slot:

scope slot *slot_number*

c) Scope to the application instance:

scope app-instance *asa_or_ftd*

d) Clear the management bootstrap information:

clear-mgmt-bootstrap

e) Commit the configuration:

commit-buffer

Step 5   Disable the application:

disable

commit-buffer

Note   For a clustered configuration, you must clear and disable the management bootstrap information
for each application attached to the Firepower 4100/9300 chassis. If you have an inter-chassis cluster
or a HA configuration, you must repeat these steps for each application on both chassis.

Step 6   When the application is offline and the slot comes online again, re-enable the application.

a) Scope back to ssa mode:

>
> **scope ssa**

b) Scope to the slot:

   scope slot *slot_number*

c) Scope to the application instance:

   scope app-instance *asa_or_ftd*

d) Enable the application:

   enable

e) Commit the configuration:

   commit-buffer

Note For a clustered configuration, you must repeat these steps to re-enable each application attached to the Firepower 4100/9300 chassis. If you have an inter-chassis cluster or a HA configuration, you must repeat these steps for each application on both chassis.

Changing the Firepower 4100/9300 Chassis Name

You can change the name used for your Firepower 4100/9300 chassis from the FXOS CLI.

Procedure

Step 1 Connect to the FXOS CLI (see Accessing the FXOS CLI, on page 15).

Step 2 Enter the system mode:

Firepower-chassis-A# **scope system**

Step 3 To view the current name:

Firepower-chassis-A /system # **show**

Step 4 To configure a new name:

Firepower-chassis-A /system # **set name** *device_name*

Step 5 Commit the transaction to the system configuration:

Firepower-chassis-A /fabric-interconnect* # **commit-buffer**

Example

The following example changes the devices name:

```
Firepower-chassis-A# scope system
```

```
Firepower-chassis-A /system # set name New-name
Warning: System name modification changes FC zone name and redeploys them non-disruptively
Firepower-chassis-A /system* # commit-buffer
Firepower-chassis-A /system # show

Systems:
    Name        Mode         System IP Address System IPv6 Address
    ---------- ----------- ----------------- ------------------
    New-name    Stand Alone 192.168.100.10    ::
New-name-A /system #
```

Install a Trusted Identity Certificate

After initial configuration, a self-signed SSL certificate is generated for use with the Firepower 4100/9300 chassis web application. Because that certificate is self-signed, client browsers do not automatically trust it. The first time a new client browser accesses the Firepower 4100/9300 chassis web interface, the browser will throw an SSL warning, requiring the user to accept the certificate before accessing the Firepower 4100/9300 chassis. You can use the following procedure to generate a Certificate Signing Request (CSR) using the FXOS CLI and install the resulting identity certificate for use with the Firepower 4100/9300 chassis. This identity certificate allows a client browser to trust the connection, and bring up the web interface with no warnings.

Procedure

Step 1	Connect to the FXOS CLI. (See Accessing the FXOS CLI, on page 15).
Step 2	Enter the security module:
	scope security
Step 3	Create a keyring:
	create keyring *keyring_name*
Step 4	Set a modulus size for the private key:
	set modulus *size*
Step 5	Commit the configuration:
	commit-buffer
Step 6	Configure the CSR fields. The certificate can be generated with basic options (for example, a subject-name), and optionally more advanced options that allow information like locale and organization to be embedded in the certificate. Note that when you configure the CSR fields, the system prompts for a certificate password.
	create certreq subject-name *subject_name*
	password
	set country *country*
	set state *state*
	set locality *locality*
	set org-name *organization_name*

set org-unit-name *organization_unit_name*

set subject-name *subject_name*

Step 7   Commit the configuration:

commit-buffer

Step 8   Export the CSR to provide to your certificate authority. The certificate authority uses the CSR to create your identity certificate.

a) Show the full CSR:

show certreq

b) Copy the output starting with (and including) "-----BEGIN CERTIFICATE REQUEST-----", ending with (and including) "-----END CERTIFICATE REQUEST-----":

Example:

```
-----BEGIN CERTIFICATE REQUEST-----
MIIC6zCCAdMCAQAwdzELMAkGA1UEBhMCVVMxEzARBgNVBAgMCkNhbGlmb3JuaWEx
ETAPBgNVBAcMCFNhbiBKb3NlMRYwFAYDVQQKDA1DaXNjbyBTeXN0ZW1zMQwwCgYD
VQQLDANUQUMxGjAYBgNVBAMMEWZwNDEyMC50ZXN0Lmxvy2FsMIIBIjANBgkqhkiG
9w0BAQEFAAOCAQ8AMIIBCgKCAQEAs0ON5gagkfZ2fi4JVEANG+7YGgcHbnUt7LpV
yMChnKOPJjBwkUMNQA1mQsRQDcbJ232/sK0fMSnyqOL8JzC7itxeVEZRyz7/ax7W
GNveg/XP+zd03nt4GXM63FsrPcPmA7EwgqDSLoShtBEV10hhf4+Nw4pKCZ+eSSkS
JkTB1ZHaKV9bttYg3kf/UEUUgk/EyrVq3B+u2DsooPVq76mTm8BwYMqHbJEv4Pmu
RjWE88yEvVwH7JTEij9OvxbatjDjVSJHZBURtCanvyBvGuLP/Q/Nmv3Lo3G9ITbL
L5gIYZVatTxp6HTUezH2MIIzOavU6d1tB9rnyxgGth5dPV0dhQIDAQABoC8wLQYJ
KoZIhvcNAQkOMSAwHjAcBgNVHREEFTATghFmcDQxMjAudGVzdC5sb2NhbDANBgkq
hkiG9w0BAQsFAAOCAQEAZUfCbwx9vt5aVDcL+tATu5xFE3LA310ck6GjlNv6W/6r
jBNLxusYi1rZZcW+CgnvNs4ArqYGyNVBySOavJO/VvQ1KfyxxJ1OIkyx3RzEjgK0
zzyoyrG+EZXC5ShiraS8HuWvE2wFM2wwWNtHWtvcQy55+/hDPD2Bv8pQOC2Zng3I
kLfG1dxWf1xAxLzf5J+AuIQ0CM5HzM9Zm8zREoWT+xHtLSqAqg/aCuomN9/vEwyU
OYfoJMvAqC6AZyUnMfUfCoyuLpLwgkxB0gyaRdnea5RhiGjYQ2lDXYDjEXp7rCx9
+6bvDl1n70JCegHdCWtP75SaNyaBEPkO0365rTckbw==
-----END CERTIFICATE REQUEST-----
```

Step 9   Exit the certreq mode:

exit

Step 10   Exit the keyring mode:

exit

Step 11   Provide the CSR output to the Certificate Authority in accordance with the Certificate Authority's enrollment process. If the request is successful, the Certificate Authority sends back an identity certificate that has been digitally signed using the CA's private key.

Step 12   **Note**   All identity certificates must be in Base64 format to be imported into FXOS. If the identity certificate chain received from the Certificate Authority is in a different format, you must first convert it with an SSL tool such as OpenSSL.

Create a new trustpoint to hold the identity certificate chain.

create trustpoint *trustpoint_name*

Step 13   Enter the identity certificate chain you received from the Certificate Authority in step 11, following the instructions on screen.

Note	For a Certificate Authority that uses intermediate certificates, the root and intermediate certificates must be combined. In a text file, paste the root certificate at the top, followed by each intermediate certificate in the chain, including all BEGIN CERTIFICATE and END CERTIFICATE flags. Copy and paste that entire text block into the trustpoint.

set certchain

Example:

```
firepower /security/trustpoint* # set certchain
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Trustpoint Certificate Chain:
>-----BEGIN CERTIFICATE-----
>MIICDTCCAbOgAwIBAgIQYIutxPDPw6BOp3uKNgJHZDAKBggqhkjOPQQDAjBTMRUw
>EwYKCZImiZPyLGQBGRYFbG9jYWwxGDAWBgoJkiaJk/IsZAEZFghuYWF1c3RpbjEg
>MB4GA1UEAxMXbmFhdXN0aW4tTkFBVVNUSU4tUEMtQ0EwHhcNMTUwNzI4MTc1NjU2
>WhcNMjAwNzI4MTgwNjU2WjBTMRUwEwYKCZImiZPyLGQBGRYFbG9jYWwxGDAWBgoJ
>kiaJk/IsZAEZFghuYWF1c3RpbjEgMB4GA1UEAxMXbmFhdXN0aW4tTkFBVVNUSU4t
>UEMtQ0EwWTATBgcqhkjOPQIBBggqhkjOPQMBBwNCAASvEA27V1Enq1gMtLkvJ6rx
>GXRpXWIEyuiBM4eQRoqZKnkeJUkm1xmqlubaDHPJ5TMGfJQYszLBRJPq+mdrKcDl
>o2kwZzATBgkrBgEEAYI3FAIEBh4EAEMAQTAOBgNVHQ8BAf8EBAMCAYYwDwYDVR0T
>AQH/BAUwAwEB/zAdBgNVHQ4EFgQUyInbDHPrFwEEBcbxGSgQW7pOVIkwEAYJKwYB
>BAGCNxUBBAMCAQAwCgYIKoZIzj0EAwIDSAAwRQIhAP++QJTUmniB/AxPDDN63Lqy
>18odMDoFTkG4p3Tb/2yMAiAtMYhlsv1gCxsQVOw0xZVRugSdoOak6n7wCjTFX9jr
>RA==
>-----END CERTIFICATE-----
>ENDOFBUF
```

Step 14	Commit the configuration:	
	commit-buffer	
Step 15	Exit the trustpoint mode:	
	exit	
Step 16	Enter the keyring mode:	
	scope keyring *keyring_name*	
Step 17	Associate the trustpoint created in step 13 with the keyring that was created for the CSR:	
	set trustpoint *trustpoint_name*	
Step 18	Import the signed identity certificate for the server.	
	set cert	
Step 19	Paste the contents of the identity certificate provided by the Certificate authority:	

Example:

```
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Keyring certificate:
>-----BEGIN CERTIFICATE-----
>MIIE8DCCBJagAwIBAgITRQAAAArehlUWgiTzvgAAAAAACjAKBggqhkjOPQQDAjBT
>MRUwEwYKCZImiZPyLGQBGRYFbG9jYWwxGDAWBgoJkiaJk/IsZAEZFghuYWF1c3Rp
>bjEgMB4GA1UEAxMXbmFhdXN0aW4tTkFBVVNUSU4tUEMtQ0EwHhcNMTYwNDI4MTMw
>OTU0WhcNMTgwNDI4MTMwOTU0WjB3MQswCQYDVQQGEwJVUzETMBEGA1UECBMKQ2Fs
>aWZvcm5pYTERMA8GA1UEBxMIU2FuIEpvc2UxFjAUBgNVBAoTDDNpc2NvIFN5c3Rl
>bXMxDDAKBgNVBAsTA1RBQzEaMBgGA1UEAxMRZnA0MTIwLnRlc3QubG9jYWwwggEi
>MA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCzQ43mBqCR9nZ+LglUQA0b7tga
```

```
>BwdudS3sulXIwKGco48mMHCRQw1ADWZCxFANxsnbfb+wrR8xKfKo4vwnMLuK3F5U
>RlHLPv9rHtYY296D9c/7N3Tee3gZczrcWys9w+YDsTCCoNIuhKG0ERXXSGF/j43D
>ikoJn55JKRImRMHVkdopX1u21iDeR/9QRRSCT8TKtWrcH67YOyig9WrvqZObwHBg
>yodskS/g+a5GNYTzzIS9XAfslMSKP06/Ftq2MONVIkdkFRG0Jqe/IG8a4s/9D82a
>/cujcb0hNssvmAhhlVq1PGnodNR7MfYwgjM5q9Tp3W0H2ufLGAa2Hl09XR2FAgMB
>AAGjggJYMIICVDAcBgNVHREEFTATghFmcDQxMjAudGVzdC5sb2NhbDAdBgNVHQ4E
>FgQU/1WpstiEYExs8DlZWcuHZwPtu5QwHwYDVR0jBBgwFoAUyInbDHPrFwEEBcbx
>GSgQW7pOVIkwgdwGA1UdHwSB1DCB0TCBzqCBy6CByIaBxWxkYXA6Ly8vQ049bmFh
>dXN0aW4tTkFFBVVNUSU4tUEMtQ0EsQ049bmFhdXN0aW4tcGMsQ049Q0RQLENOPVB1
>YmxpYyUyMEtleSUyMFNlcnZpY2VzLENOPVNlcnZpY2VzLENOPUNvbmZpZ3VyYXRp
>b24sREM9bmFhdXN4sREM9bG9jYWw/Y2VydGlmaWNhdGVSZXZvY2F0aW9uTGlz
>dD9iYXNlP29iamVjdENsYXNzPWNSTERpc3RyaWJ1dGlvblBvaW50MIHMBggrBgEF
>BQcBAQSBvzCBvDCBuQYIKwYBBQUHMAKGgaxsZGFwOi8vL0NOPW5hYXVzdGluLU5B
>QVVTVElOLVBDLUNBLENOPUFJQSxDTj1QdWJsaWMlMjBLZXklMjBTZXJ2aWNlcyxD
>Tj1TZXJ2aWNlcyxDTj1Db25maWd1cmF0aW9uLERDPW5hYXVzdGluLERDPWxvY2Fs
>P2NBQ2VydGlmaWNhdGU/YmFzZT9vYmplY3RDbGFzcz1jZXJ0aWZpY2F0aW9uQXV0
>aG9yaXR5MCEGCSsGAQQBgjcUAgQUHhIAVwBlAGIAUwBlAHIAdgBlAHIwDgYDVR0P
>AQH/BAQDAgWgMBMGA1UdJQQMMAoGCCsGAQUFBwMBMBMAoGCCqGSM49BAMCA0gAMEUC
>IFew7NcJirEtFRvyxjkQ4/dVo2oI6CRB308WQbYHNUu/AiEA7UdObiSJBG/PBZjm
>sgoIK60akbjotOTvUdUd9b6K1Uw=
>-----END CERTIFICATE-----
>ENDOFBUF
```

Step 20   Exit the keyring mode:

exit

Step 21   Exit the security mode:

exit

Step 22   Enter the system mode:

scope system

Step 23   Enter the services mode:

scope services

Step 24   Configure the FXOS web service to use the new certificate:

set https keyring *keyring_name*

Step 25   Commit the configuration:

commit-buffer

Step 26   Display the keyring associated with the HTTPS server. It should reflect the keyring name created in step 3 of
this procedure. If the screen output displays the default keyring name, the HTTPS server has not yet been
updated to use the new certificate:

show https

Example:

```
fp4120 /system/services # show https
Name: https
    Admin State: Enabled
    Port: 443
    Operational port: 443
    Key Ring: firepower_cert
    Cipher suite mode: Medium Strength
```

```
                    Cipher suite: ALL:!ADH:!EXPORT40:!EXPORT56:!LOW:!RC4:!MD5:!IDEA:+HIGH:+MEDIUM:+EXP:+eNULL
```

Step 27   Display the contents of the imported certificate, and verify that the **Certificate Status** value displays as **Valid**:

scope security

show keyring *keyring_name* **detail**

Example:

```
fp4120 /security #  scope security
fp4120 /security #  show keyring firepower_cert detail
Keyring firepower_cert:
    RSA key modulus: Mod2048
    Trustpoint CA: firepower_chain
    Certificate status: Valid
    Certificate:
    Data:
        Version: 3 (0x2)
        Serial Number:
            45:00:00:00:0a:de:86:55:16:82:24:f3:be:00:00:00:00:00:0a
    Signature Algorithm: ecdsa-with-SHA256
        Issuer: DC=local, DC=naaustin, CN=naaustin-NAAUSTIN-PC-CA
        Validity
            Not Before: Apr 28 13:09:54 2016 GMT
            Not After : Apr 28 13:09:54 2018 GMT
        Subject: C=US, ST=California, L=San Jose, O=Cisco Systems, OU=TAC,
CN=fp4120.test.local
        Subject Public Key Info:
            Public Key Algorithm: rsaEncryption
                Public-Key: (2048 bit)
                Modulus:
                    00:b3:43:8d:e6:06:a0:91:f6:76:7e:2e:09:54:40:
                    0d:1b:ee:d8:1a:07:07:6e:75:2d:ec:ba:55:c8:c0:
                    a1:9c:a3:8f:26:30:70:91:43:0d:40:0d:66:42:c4:
                    50:0d:c6:c9:db:7d:bf:b0:ad:1f:31:29:f2:a8:e2:
                    fc:27:30:bb:8a:dc:5e:54:46:51:cb:3e:ff:6b:1e:
                    d6:18:db:de:83:f5:cf:fb:37:74:de:7b:78:19:73:
                    3a:dc:5b:2b:3d:c3:e6:03:b1:30:82:a0:d2:2e:84:
                    a1:b4:11:15:d7:48:61:7f:8f:8d:c3:8a:4a:09:9f:
                    9e:49:29:12:26:44:c1:d5:91:da:29:5f:5b:b6:d6:
                    20:de:47:ff:50:45:14:82:4f:c4:ca:b5:6a:dc:1f:
                    ae:d8:3b:28:a0:f5:6a:ef:a9:93:9b:c0:70:60:ca:
                    87:6c:91:2f:e0:f9:ae:46:35:84:f3:cc:84:bd:5c:
                    07:ec:94:c4:8a:3f:4e:bf:16:da:b6:30:e3:55:22:
                    47:64:15:11:b4:26:a7:bf:20:6f:1a:e2:cf:fd:0f:
                    cd:9a:fd:cb:a3:71:bd:21:36:cb:2f:98:08:61:95:
                    5a:b5:3c:69:e8:74:d4:7b:31:f6:30:82:33:39:ab:
                    d4:e9:dd:6d:07:da:e7:cb:18:06:b6:1e:5d:3d:5d:
                    1d:85
                Exponent: 65537 (0x10001)
        X509v3 extensions:
            X509v3 Subject Alternative Name:
                DNS:fp4120.test.local
            X509v3 Subject Key Identifier:
                FF:55:A9:B2:D8:84:60:4C:6C:F0:39:59:59:CB:87:67:03:ED:BB:94
            X509v3 Authority Key Identifier:
                keyid:C8:89:DB:0C:73:EB:17:01:04:05:C6:F1:19:28:10:5B:BA:4E:54:89
            X509v3 CRL Distribution Points:
                Full Name:
                  URI:ldap:///CN=naaustin-NAAUSTIN-PC-CA,CN=naaustin-pc,CN=CDP,
                    CN=Public%20Key%20Services,CN=Services,CN=Configuration,DC=naaustin,
                  DC=local?certificateRevocationList?base?objectClass=cRLDistributionPoint
```

```
                Authority Information Access:
                    CA Issuers - URI:ldap:///CN=naaustin-NAAUSTIN-PC-CA,CN=AIA,
                      CN=Public%20Key%20Services,CN=Services,CN=Configuration,DC=naaustin,
                      DC=local?cACertificate?base?objectClass=certificationAuthority
                1.3.6.1.4.1.311.20.2:
                    ...W.e.b.S.e.r.v.e.r
                X509v3 Key Usage: critical
                    Digital Signature, Key Encipherment
                X509v3 Extended Key Usage:
                    TLS Web Server Authentication
        Signature Algorithm: ecdsa-with-SHA256
            30:45:02:20:57:b0:ec:d7:09:8a:b1:2d:15:1b:f2:c6:39:10:
            e3:f7:55:a3:6a:a0:8:e8:24:41:df:4f:16:41:b6:07:35:4b:bf:
            02:21:00:ed:47:4e:6e:24:89:04:6f:cf:05:98:e6:b2:0a:08:
            2b:ad:1a:91:b8:e8:b4:e4:ef:51:d5:1d:f5:be:8a:d5:4c
-----BEGIN CERTIFICATE-----
MIIE8DCCBJagAwIBAgITRQAAAArehlUWgiTzvgAAAAACjAKBggqhkjOPQQDAjBT
MRUwEwYKCZImiZPyLGQBGRYFbG9jYWwxGDAWBgoJkiaJk/IsZAEZFghuYWF1c3Rp
bjEgMB4GA1UEAxMXbmFhdXN0aW4tTkFBVVNUSU4tUEMtQ0EwHhcNMTYwNDI4MTMw
OTU0WhcNMTgwNDI4MTMwOTU0WjB3MQswCQYDVQQGEwJVUzETMBEGA1UECBMKQ2Fs
aWZvcm5pYTERMA8GA1UEBxMIU2FuIEpvc2UxFjAUBgNVBAoTDUNpc2NvIFN5c3Rl
bXMxDDAKBgNVBAsTA1RBQzEaMBgGA1UEAxMRZnA0MTIwLnRlc3QubG9jYWwwggEi
MA0GCSqGSIb3DQEBAQUAA4IBDwAwggEKAoIBAQCzQ43mBqCR9nZ+LglUQA0b7tga
BwdudS3sulXIwKGco48mMHCRQw1ADWZCxFANxsnbfb+wrR8xKfKo4vwnMLuK3F5U
RlHLPv9rHtYY296D9c/7N3Tee3gZczrcWys9w+YDsTCCoNIuhKG0ERXXSGF/j43D
ikoJn55JKRImRMHVkdopX1u21iDeR/9QRRSCT8TKtWrcH67YOyig9WrvqZObwHBg
yodskS/g+a5GNYTzzIS9XAfslMSKP06/Ftq2MONVIkdkFRG0Jqe/IG8a4s/9D82a
/cujcb0hNssvmAhhlVq1PGnodNR7MfYwgjM5q9Tp3W0H2ufLGAa2Hl09XR2FAgMB
AAGjggJYMIICVDAcBgNVHREEFTATghFmcDQxMjAudGVzdC5sb2NhbDAdBgNVHQ4E
FgQU/1WpstiEYExs8DlZWcuHZwPtu5QwHwYDVR0jBBgwFoAUyInbDHPrFwEEBcbx
GSgQW7pOVIkwgdwGA1UdHwSB1DCB0TCBzqqCBy6CByIaBxWxkYXA6Ly8vQ049bmFh
dXN0aW4tTkFBVVNUSU4tUEMtQ0EsQ049bmFhdXN0aW4tcGMsQ049Q0RQLENOPVB1
YmxpYyUyMEtleSUyMFNlcnZpY2VzLENOPVNlcnZpY2VzLENOPUNvbmZpZ3VyYXRp
b24sREM9bmFhdXN0aW4sREM9bG9jYWw/Y2VydGlmaWNhdGVSZXZvY2F0aW9uTGlz
dD9iYXNlP29iamVjdENsYXNzPWNSTERpc3RyaWJ1dGlvblBvaW50MIHMBggrBgEF
BQcBAQSBvzCBvDCBuQYIKwYBBQUHMAKGgaxsZGFwOi8vL0NOPW5hYXVzdGluLU5B
QVVTVElOLVBDLUNBLENOPUFJQSxDTj1QdWJsaWMlMjBLZXklMjBTZXJ2aWNlcyxD
Tj1TZXJ2aWNlcyxDTj1Db25maWd1cmF0aW9uLERDPW5hYXVzdGluLERDPWxvY2Fs
P2NBQ2VydGlmaWNhdGU/YmFzZT9vYmplY3RDbGFzcz1jZXJ0aWZpY2F0aW9uQXV0
aG9yaXR5MCEGCSsGAQQBgjcUAgQUHhIAVwBlAGIAUwBlAHIAdgBlAHIwDgYDVR0P
AQH/BAQDAgWgMBMGA1UdJQQMMAoGCCsGAQUFBwMBMAoGCCqGSM49BAMCA0gAMEUC
IFew7NcJirEtFRvyxjkQ4/dVo2oI6CRB308WQbYHNUu/AiEA7UdObiSJBG/PBZjm
sgoIK60akbjotOTvUdUd9b6K1Uw=
-----END CERTIFICATE-----

    Zeroized: No
```

What to do next

To verify that the new trusted certificate is presented, navigate to the Firepower Chassis Manager by entering *https://<FQDN_or_IP>/* in the address bar of a web browser.

✎

Note Browsers also verify the subject-name of a certificate against the input in the address bar. If the certificate is issued to the fully qualified domain name, it must be accessed that way in the browser. If it is accessed via IP address, a different SSL error is thrown (Common Name Invalid) even if the trusted certificate is used.

Auto-Import Certificate Update

When the Cisco certificate server changes its identity certificate to leverage a different root CA, the connectivity for the Smart Licensing on 4100 or 9300s running the ASA devices gets broken. Because the licensing connectivity is handled by the supervisor instead of Lina on the application, the Smart Licensing function fails. For FXOS-based devices, the issue can be resolved using the auto-import feature without an upgrade to the FXOS software.

By default, the auto-import feature is disabled. You can use the following procedure to enable the auto-import feature using the FXOS CLI.

Before you begin

DNS server should be configured to reach the cisco certificate server.

Procedure

Step 1 Connect to the FXOS CLI.

Step 2 Enter the security module:

scope security

Step 3 Enable the auto-import feature.

enter tp-auto-import

Example:

```
FXOS# scope security
FXOS /security # enter tp-auto-import
 FXOS /security #
```

Step 4 Commit the configuration.

commit-buffer

Step 5 Verify the auto-import status

show detail

Example:

Successful auto-import:

```
FXOS /security/tp-auto-import #
FXOS /security/tp-auto-import # show detail
Trustpoints auto import source URL: http://www.cisco.com/security/pki/trs/ios_core.p7b
TrustPoints auto import scheduled time : 22:00
```

```
Last Importing Status : Success, Imported with 23 TrustPoint(s)
TrustPoints auto Import function : Enabled
FXOS /security/tp-auto-import #
```

Auto-import failure:

```
FXOS /security/tp-auto-import #
FXOS /security/tp-auto-import # show detail
Trustpoints auto import source URL: http://www.cisco.com/security/pki/trs/ios_core.p7b
TrustPoints auto import scheduled time : 22:00
Last Importing Status : Failure
TrustPoints auto Import function : Enabled
FXOS /security/tp-auto-import #
```

Step 6 Configure the tp-auto-import feature. Set the import-time-hour.

set import-time-hour *hour* **import-time-min** *minutes*

Example:

```
FXOS /security/tp-auto-import # set
  import-time-hour  Trustpoints auto import hour time
FXOS /security/tp-auto-import # set import-time-hour
  0-23  Import Time Hour
FXOS /security/tp-auto-import # set import-time-hour 7 import-time-min
  0-59  Import Time Min
FXOS /security/tp-auto-import # set import-time-hour 7 import-time-min 20
<CR>
FXOS /security/tp-auto-import # set import-time-hour 7 import-time-min 20
FXOS /security/tp-auto-import* # commit-buffer
FXOS /security/tp-auto-import #
```

Note The auto-import source URL is fixed and you must change the import time detail to minute per day. Import occurs everyday on the scheduled time of the day. If hours and minutes are not set then the certificate import occurs only once while enabling it. Certificates get downloaded as a bundle into the box under the path /opt/certstore which can only be accessed through secure-login option. Along with the bundle (ios_core.p7b), individual certificates (AutoTP1 to AutoTPn) get extracted automatically.

Step 7 After the auto-import configuration completion, enter show detail command.

show detail

Example:

```
FXOS  /security/tp-auto-import # show detail
Trustpoints auto import source URL: http://www.cisco.com/security/pki/trs/ios_core.p7b
TrustPoints auto import scheduled time : 07:20
Last Importing Status : Success, Imported with 23 TrustPoint(s)
TrustPoints auto Import function : Enabled
```

Note The maximum certificates that can be imported is 30. Each import re-iterates for 6 times if there is any connectivity issue to Cisco Certificate Server and then updates the last importing status in the show command.

Step 8 (Optional) To disable the auto-import feature, enter the delete auto-import command.

delete tp-auto-import

Example:

```
FXOS /security #
FXOS /security # delete tp-auto-import
FXOS /security* # commit-buffer
```

```
FXOS /security # show detail
security mode:
    Password Strength Check: No
    Minimum Password Length: 8
    Is configuration export key set: No
    Current Task:
FXOS /security # scope tp-auto-import
Error: Managed object does not exist
FXOS /security #
FXOS /security # enter tp-auto-import
FXOS /security/tp-auto-import* # show detail
FXOS /security/tp-auto-import* #
```

Note   If you disable the auto-import feature, certificates that are imported remain persistent till the time there is no change in the build. Certificates get removed if you disable the auto-import feature and then downgrade/upgrade the build.

Pre-Login Banner

With a pre-login banner, when a user logs into Firepower Chassis Manager, the system displays the banner text and the user must click **OK** on the message screen before the system prompts for the username and password. If a pre-login banner is not configured, the system goes directly to the username and password prompt.

When a user logs into the FXOS CLI, the system displays the banner text, if configured, before it prompts for the password.

Creating the Pre-Login Banner

Procedure

Step 1   Connect to the FXOS CLI (see Accessing the FXOS CLI, on page 15).

Step 2   Enter security mode:

Firepower-chassis# **scope security**

Step 3   Enter banner security mode:

Firepower-chassis /security # **scope banner**

Step 4   Enter the following command to create a pre-login banner:

Firepower-chassis /security/banner # **create pre-login-banner**

Step 5   Specify the message that FXOS should display to the user before they log into Firepower Chassis Manager or the FXOS CLI:

Firepower-chassis /security/banner/pre-login-banner* # **set message**

Launches a dialog for entering the pre-login banner message text.

Step 6   At the prompt, type a pre-login banner message. You can enter any standard ASCII character in this field. You can enter multiple lines of text with each line having up to 192 characters. Press **Enter** between lines.

On the line following your input, type **ENDOFBUF** and press **Enter** to finish.

Press **Ctrl** and **C** to cancel out of the set message dialog.

Step 7   Commit the transaction to the system configuration:

Firepower-chassis /security/banner/pre-login-banner* # **commit-buffer**

Example

The following example creates the pre-login banner:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope banner
Firepower-chassis /security/banner # create pre-login-banner
Firepower-chassis /security/banner/pre-login-banner* # set message
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Enter prelogin banner:
>Welcome to the Firepower Security Appliance
>**Unauthorized use is prohibited**
>ENDOFBUF
Firepower-chassis /security/banner/pre-login-banner* # commit-buffer
Firepower-chassis /security/banner/pre-login-banner #
```

Modifying the Pre-Login Banner

Procedure

Step 1   Connect to the FXOS CLI (see Accessing the FXOS CLI, on page 15).

Step 2   Enter security mode:

Firepower-chassis# **scope security**

Step 3   Enter banner security mode:

Firepower-chassis /security # **scope banner**

Step 4   Enter pre-login-banner banner security mode:

Firepower-chassis /security/banner # **scope pre-login-banner**

Step 5   Specify the message that FXOS should display to the user before they log into Firepower Chassis Manager or the FXOS CLI:

Firepower-chassis /security/banner/pre-login-banner # **set message**

Launches a dialog for entering the pre-login banner message text.

Step 6   At the prompt, type a pre-login banner message. You can enter any standard ASCII character in this field. You can enter multiple lines of text with each line having up to 192 characters. Press **Enter** between lines.

On the line following your input, type **ENDOFBUF** and press **Enter** to finish.

Press **Ctrl** and **C** to cancel out of the set message dialog.

Step 7   Commit the transaction to the system configuration:

Firepower-chassis /security/banner/pre-login-banner* # **commit-buffer**

Example

The following example modifies the pre-login banner:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope banner
Firepower-chassis /security/banner # scope pre-login-banner
Firepower-chassis /security/banner/pre-login-banner # set message
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Enter prelogin banner:
>Welcome to the Firepower Security Appliance
>**Unauthorized use is prohibited**
>ENDOFBUF
Firepower-chassis /security/banner/pre-login-banner* # commit-buffer
Firepower-chassis /security/banner/pre-login-banner #
```

Deleting the Pre-Login Banner

Procedure

Step 1   Connect to the FXOS CLI (see Accessing the FXOS CLI, on page 15).

Step 2   Enter security mode:

Firepower-chassis# **scope security**

Step 3   Enter banner security mode:

Firepower-chassis /security # **scope banner**

Step 4   Delete the pre-login banner from the system:

Firepower-chassis /security/banner # **delete pre-login-banner**

Step 5   Commit the transaction to the system configuration:

Firepower-chassis /security/banner* # **commit-buffer**

Example

The following example deletes the pre-login banner:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope banner
```

```
Firepower-chassis /security/banner # delete pre-login-banner
Firepower-chassis /security/banner* # commit-buffer
Firepower-chassis /security/banner #
```

Rebooting the Firepower 4100/9300 Chassis

Procedure

Step 1 Choose **Overview** to open the Overview page.

Step 2 Click **Reboot** next to the Chassis Uptime in the upper-right corner of the Overview page.

Step 3 Click **Yes** to verify that you want to power off the Firepower 4100/9300 chassis.
The system will gracefully shut down any logical devices configured on the system and then power down each security module/engine before finally powering down and then restarting the Firepower 4100/9300 chassis. This process takes approximately 15-20 minutes.

Powering Off the Firepower 4100/9300 Chassis

Procedure

Step 1 Choose **Overview** to open the Overview page.

Step 2 Click **Shutdown** next to the Chassis Uptime in the upper-right corner of the Overview page.

Step 3 Click **Yes** to verify that you want to power off the Firepower 4100/9300 chassis.
The system will gracefully shut down any logical devices configured on the system and then power down each security module/engine before finally powering down the Firepower 4100/9300 chassis.

Restoring the Factory Default Configuration

You can use the FXOS CLI to restore your Firepower 4100/9300 chassis to factory default configuration.

Note This process erases all user configuration from the chassis including any logical device configuration. After completing this procedure, you will need to reconfigure the system (see Initial Configuration, on page 7).

Procedure

Step 1   (Optional) The **erase configuration** command does not remove the Smart License configuration from the chassis. If you also want to remove the Smart License configuration, perform the following steps:

scope license

deregister

Deregistering the Firepower 4100/9300 chassis removes the device from your account. All license entitlements and certificates on the device are removed.

Step 2   Connect to the local-management shell:

connect local-mgmt

Step 3   Enter the following command to erase all user configuration from your Firepower 4100/9300 chassis and restore the chassis to its original factory default configuration:

erase configuration

The system prompts you to verify that you are sure you want to erase all user configuration.

Step 4   Confirm that you want to erase the configuration by entering **yes** at the command prompt.
The system will erase all user configuration from your Firepower 4100/9300 chassis and then reboot the system.

Securely Erasing System Components

You can use the FXOS CLI to erase and securely erase components of your appliance.

The **erase configuration** command removes all user-configuration information on the chassis, restoring it to its original factory-default configuration, as described in .

The **secure erase** command securely erases the specified appliance component. That is, data is not just deleted—the physical storage is "wiped" (completely erased). This is important when transferring or returning the appliance as hardware storage components do not retain residual data or stubs.

Note   The device reboots during secure erase, which means SSH connections are terminated. Therefore, we recommend performing secure erase over a serial console-port connection.

Procedure

Step 1   Connect to the local-management shell:

connect local-mgmt

Step 2   Enter one of the following **erase configuration** commands to securely erase the specified appliance component:

a) **erase configuration chassis**

The system warns you that all data and images will be lost and cannot be recovered, and asks you to confirm that you want to proceed. If you enter **y**, the entire chassis is securely erased; security modules are erased first, followed by the Supervisor.

Since all data and software on the device are erased, device recovery can be accomplished only from the ROM Monitor (ROMMON).

b) **erase configuration security_module** *module_id*

The system warns you that all data and images on the module will be lost and cannot be recovered, and asks you to confirm that you want to proceed. If you enter **y**, the module is erased.

Note The **decommission-secure** command produces essentially the same result as this command.

After a security module is erased, it remains down until acknowledged (similar to a module that is decommissioned).

c) **erase configuration supervisor**

The system warns you that all data and images will be lost and cannot be recovered, and asks you to confirm that you want to proceed. If you enter **y**, the Supervisor is securely erased.

Since all data and software on the Supervisor are erased, device recovery can be accomplished only from the ROM Monitor (ROMMON).

Platform Settings

Setting the Date and Time

Use the NTP page to configure the network time protocol (NTP) on the system, to set the date and time manually, or to view the current system time.

NTP settings are automatically synced between the Firepower 4100/9300 chassis and any logical devices installed on the chassis.

Note If you are deploying FTD on the Firepower 4100/9300 chassis, you must configure NTP on the Firepower 4100/9300 chassis so that Smart Licensing will work properly and to ensure proper timestamps on device registrations. You should use the same NTP server for both the Firepower 4100/9300 chassis and the FMC, but note that you cannot use FMC as the NTP server for the Firepower 4100/9300 chassis.

If you are using NTP, you can view the overall synchronization status on the **Current Time** tab, or you can view the synchronization status for each configured NTP server by looking at the Server Status field in the

NTP Server table on the **Time Synchronization** tab. If the system is unable to synchronize with a particular NTP server, you can hover over the information icon next to the Server Status for more information.

Viewing the Configured Date and Time

Procedure

Step 1   Choose **Platform Settings** > **NTP**.

Step 2   Click the **Current Time** tab.

The system shows the date, time, and time zone that are configured on the device.

If you are using NTP, you can also view the overall synchronization status on the **Current Time** tab. You can view the synchronization status for each configured NTP server by looking at the Server Status field in the **NTP Server** table on the **Time Synchronization** tab. If the system is unable to synchronize with a particular NTP server, you can hover over the information icon next to the Server Status for more information.

Setting the Time Zone

Procedure

Step 1   Choose **Platform Settings** > **NTP**.

Step 2   Click the **Current Time** tab.

Step 3   Choose the appropriate time zone for the chassis from the **Time Zone** drop-down list.

Setting the Date and Time Using NTP

NTP is used to implement a hierarchical system of servers that provide a precisely synchronized time among network systems. This kind of accuracy is required for time-sensitive operations, such as validating CRLs, which include a precise time stamp. You can configure up to four NTP servers.

Note

- FXOS uses NTP version 3.

- If the stratum value of an external NTP server is 13 or greater, the application instance cannot sync to the NTP server on the FXOS chassis. Each time a NTP client syncs to a NTP server, the stratum value increases by one.

  If you have set up your own NTP server, you can find its stratum value in the /etc/ntp.conf file on the server. If the NTP server has stratum value of 13 or greater you can either change the stratum value in the ntp.conf file and restart the server, or use a different NTP server (for example: pool.ntp.org).

Before you begin

If you use a hostname for the NTP server, you must configure a DNS server. See .

Procedure

Step 1   Choose **Platform Settings** > **NTP**.

The **Time Synchronization** tab is selected by default.

Step 2   Under **Set Time Source**, click **Use NTP Server**.

Step 3   (Optional) Check the **NTP Server Authentication: Enable** check box if you need to authenticate with the NTP server.

Click **Yes** to require an authentication key ID and value.

Only SHA1 is supported for NTP server authentication.

Step 4   Click **Add** to identify up to 4 NTP servers by IP address or hostname.

Step 5   (Optional) Enter the NTP server's **Authentication Key** ID and **Authentication Value**.

Obtain the key ID and value from the NTP server. For example, to generate the SHA1 key on NTP server Version 4.2.8p8 or later with OpenSSL installed, enter the **ntp-keygen -M** command, and then view the key ID and value in the ntp.keys file. The key is used to tell both the client and server which value to use when computing the message digest.

Step 6   Click **Save**.

You can view the synchronization status of each server by looking at the Server Status field in the **NTP Server** table. If the system is unable to synchronize with a particular NTP server, you can hover over the information icon next to the Server Status for more information.

Note   If you modify the system time by more than 10 minutes, the system will log you out and you will need to log in to the Firepower Chassis Manager again.

Deleting an NTP Server

Procedure

Step 1   Choose **Platform Settings** > **NTP**.

Step 2   Click the **Time Synchronization** tab.

Step 3   For each NTP server that you want to remove, click the **Delete** icon for that server in the **NTP Server** table.

Step 4   Click **Save**.

Setting the Date and Time Manually

This section describes how to set the date and time manually on the chassis. Note that after you manually set the chassis date and time, it could take some time for the change to be reflected in the installed logical device(s).

Procedure

Step 1	Choose **Platform Settings** > **NTP**.
Step 2	Click the **Time Synchronization** tab.
Step 3	Under **Set Time Source**, click **Set Time Manually**.
Step 4	Click the **Date** drop-down list to display a calendar and then set the date using the controls available in the calendar.
Step 5	Use the corresponding drop-down lists to specify the time as hours, minutes, and AM/PM.

> **Tip** You can click **Get System Time** to set the date and time to match what is configured on the system you are using to connect to the Firepower Chassis Manager.

Step 6	Click **Save**.

The chassis is configured with the date and time specified.

> **Note** If you modify the system time by more than 10 minutes, the system will log you out and you will need to log in to the Firepower Chassis Manager again.

Configuring SSH

The following procedure describes how to enable or disable SSH access to the chassis, how to enable the FXOS chassis as an SSH client, and how to configure the various algorithms used by SSH for encryption, key exchange, and message authentication for both the SSH server and SSH client.

SSH is enabled by default.

Procedure

Step 1	Choose **Platform Settings** > **SSH** > **SSH Server**.
Step 2	To enable SSH access to the chassis, check the **Enable SSH** check box. To disable SSH access, uncheck the **Enable SSH** check box.
Step 3	For the server **Encryption Algorithm**, check the check boxes for each allowed encryption algorithm.

Note	• The following encryption algorithms are not supported in Common Criteria mode:

> • 3des-cbc
>
> • chacha20-poly1305@openssh.com

• chacha20-poly1305@openssh.com is not supported in FIPS. If FIPS mode is enabled on the FXOS chassis, you cannot use chacha20-poly1305@openssh.com as an encryption algorithm.

• The following encryption algorithms are not enabled by default:

```
aes128-cbc
aes192-cbc
aes265-cbc
```

Step 4   For the server **Key Exchange Algorithm**, check the check boxes for each allowed Diffie-Hellman (DH) key exchange. The DH key exchange provides a shared secret that cannot be determined by either party alone. The key exchange is combined with a signature and the host key to provide host authentication. This key-exchange method provides explicit server authentication. For more information about using DH key-exchange methods, see RFC 4253.

Note	• The following key exchange algorithms are not supported in Common Criteria mode:

> • diffie-hellman-group14-sha256
>
> • curve25519-sha256
>
> • curve25519-sha256@libssh.org

• The following key exchange algorithms are not supported in FIPS mode:

> • curve25519-sha256
>
> • curve25519-sha256@libssh.org

Step 5   For the server **Mac Algorithm**, check the check boxes for each allowed integrity algorithm.

Step 6   For the server **Host Key**, enter the modulus size for the RSA key pairs.

The modulus value (in bits) is in multiples of 8 from 1024 to 2048. The larger the key modulus size you specify, the longer it takes to generate an RSA key pair. We recommend a value of 2048.

Step 7   For the server **Volume Rekey Limit**, set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.

Step 8   For the server **Time Rekey Limit**, set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 9   Click **Save**.

Step 10   Click the **SSH Client** tab to customize the FXOS chassis SSH client.

Step 11   For the **Strict Host Keycheck**, choose **enable**, **disable**, or **prompt** to control SSH host key checking.

• **enable**-The connection is rejected if the host key is not already in the FXOS known hosts file. You must manually add hosts at the FXOS CLI using the **enter ssh-host** command in the system/services scope.

• **prompt**-You are prompted to accept or reject the host key if it is not already stored on the chassis.

> • **disable**-(The default) The chassis accepts the host key automatically if it was not stored before.

Step 12   For the client **Encryption Algorithm**, check the check boxes for each allowed encryption algorithm.

> **Note**   • The following encryption algorithms are not supported in Common Criteria mode:
>
> > • 3des-cbc
> >
> > • chacha20-poly1305@openssh.com
>
> > If Common Criteria mode is enabled on the FXOS chassis, you cannot use 3des-cbc as an encryption algorithm.
>
> • chacha20-poly1305@openssh.com is not supported in FIPS. If FIPS mode is enabled on the FXOS chassis, you cannot use chacha20-poly1305@openssh.com as an encryption algorithm.
>
> • The following encryption algorithms are not enabled by default:

```
aes128-cbc
aes192-cbc
aes265-cbc
```

Step 13   For the client **Key Exchange Algorithm**, check the check boxes for each allowed Diffie-Hellman (DH) key exchange. The DH key exchange provides a shared secret that cannot be determined by either party alone. The key exchange is combined with a signature and the host key to provide host authentication. This key-exchange method provides explicit server authentication. For more information about using DH key-exchange methods, see RFC 4253.

> **Note**   • The following Key Exchange Algorithms are not supported in Common Criteria mode:
>
> > • diffie-hellman-group14-sha256
> >
> > • curve25519-sha256
> >
> > • curve25519-sha256@libssh.org
>
> • The following Key Exchange Algorithms are not supported in FIPS mode:
>
> > • curve25519-sha256
> >
> > • curve25519-sha256@libssh.org

Step 14   For the client **Mac Algorithm**, check the check boxes for each allowed integrity algorithm.

Step 15   For the client **Volume Rekey Limit**, set the amount of traffic in KB allowed over the connection before FXOS disconnects from the session.

Step 16   For the client **Time Rekey Limit**, set the minutes for how long an SSH session can be idle before FXOS disconnects the session.

Step 17   Click **Save**.

Configuring TLS

The Transport Layer Security (TLS) protocol provides privacy and data integrity between two communicating applications. You can use the FXOS CLI to configure the minimum TLS version allowed when the FXOS chassis communicates with external devices. Newer TLS versions provide more secure communications, older TLS versions allow for backward compatibility with older applications.

For example, if the minimum TLS version configured on your FXOS chassis is v1.1, and a client browser is configured to only run v1.0, then the client will not be able to open a connection with the FXOS Chassis Manager via HTTPS. As such, peer applications and LDAP servers must be configured appropriately.

This procedure shows how to configure and view the minimum version of TLS allowed for communication between FXOS chassis and an external device.

Note
• As of the FXOS 2.3(1) release, the default minimum TLS version for the FXOS chassis is v1.1.

Procedure

Step 1 Enter system mode:

Firepower-chassis# **scope system**

Step 2 View the TLS version options available to your system:

Firepower-chassis /system # **set services tls-ver**

Example:

```
Firepower-chassis /system #
Firepower-chassis /system # set services tls-ver
   v1_0  v1.0
   v1_1  v1.1
   v1_2  v1.2
```

Step 3 Set the minimum TLS version:

Firepower-chassis /system # **set services tls-ver** *version*

Example:

```
Firepower-chassis /system #
Firepower-chassis /system # set services tls-ver v1_2
```

Step 4 Commit the configuration:

Firepower-chassis /system # **commit-buffer**

Step 5 Show the minimum TLS version configured on your system:

Firepower-chassis /system # **scope services**

Firepower-chassis /system/services # **show**

Example:

```
Firepower-chassis /system/services # show
Name: ssh
    Admin State: Enabled
    Port: 22
Kex Algo: Diffie Hellman Group1 Sha1,Diffie Hellman Group14 Sha1
Mac Algo: Hmac Sha1,Hmac Sha1 96,Hmac Sha2 512,Hmac Sha2 256
Encrypt Algo: 3des Cbc,Aes256 Cbc,Aes128 Cbc,Aes192 Cbc,Aes256 Ctr,Aes128 Ctr,Ae
s192 Ctr
Auth Algo: Rsa
    Host Key Size: 2048
Volume: None Time: None
Name: telnet
    Admin State: Disabled
    Port: 23
Name: https
    Admin State: Enabled
    Port: 443
    Operational port: 443
    Key Ring: default
    Cipher suite mode: Medium Strength
    Cipher suite: ALL:!ADH:!EXPORT40:!EXPORT56:!LOW:!RC4:!MD5:!IDEA:+HIGH:+MEDIU
M:+EXP:+eNULL
    Https authentication type: Cert Auth
    Crl mode: Relaxed
TLS:
    TLS version: v1.2
```

Configuring Telnet

The following procedure describes how to enable or disable Telnet access to the chassis. Telnet is disabled by default.

Note   Telnet configuration is currently only available using the CLI.

Procedure

Step 1   Enter system mode:

Firepower-chassis # **scope system**

Step 2   Enter system services mode:

Firepower-chassis /system # **scope services**

Step 3   To configure Telnet access to the chassis, do one of the following:

   • To allow Telnet access to the chassis, enter the following command:

   Firepower-chassis /system/services # **enable telnet-server**

   • To disallow Telnet access to the chassis, enter the following command:

   Firepower-chassis /system/services # **disable telnet-server**

Step 4 Commit the transaction to the system configuration:

Firepower /system/services # **commit-buffer**

Example

The following example enables Telnet and commits the transaction:

```
Firepower-chassis# scope system
Firepower-chassis /system # scope services
Firepower-chassis /services # enable telnet-server
Firepower-chassis /services* # commit-buffer
Firepower-chassis /services #
```

Configuring SNMP

Use the SNMP page to configure the Simple Network Management Protocol (SNMP) on the chassis. See the following topics for more information:

About SNMP

The Simple Network Management Protocol (SNMP) is an application-layer protocol that provides a message format for communication between SNMP managers and agents. SNMP provides a standardized framework and a common language used for the monitoring and management of devices in a network.

The SNMP framework consists of three parts:

- An SNMP manager—The system used to control and monitor the activities of network devices using SNMP.

- An SNMP agent—The software component within the chassis that maintains the data for the chassis and reports the data, as needed, to the SNMP manager. The chassis includes the agent and a collection of MIBs. To enable the SNMP agent and create the relationship between the manager and agent, enable and configure SNMP in the Firepower Chassis Manager or the FXOS CLI.

- A managed information base (MIB)—The collection of managed objects on the SNMP agent.

The chassis supports SNMPv1, SNMPv2c and SNMPv3. Both SNMPv1 and SNMPv2c use a community-based form of security. SNMP is defined in the following:

- RFC 3410 (http://tools.ietf.org/html/rfc3410)

- RFC 3411 (http://tools.ietf.org/html/rfc3411)

- RFC 3412 (http://tools.ietf.org/html/rfc3412)

- RFC 3413 (http://tools.ietf.org/html/rfc3413)

- RFC 3414 (http://tools.ietf.org/html/rfc3414)

- RFC 3415 (http://tools.ietf.org/html/rfc3415)

- RFC 3416 (http://tools.ietf.org/html/rfc3416)

- RFC 3417 (http://tools.ietf.org/html/rfc3417)

- RFC 3418 (http://tools.ietf.org/html/rfc3418)

- RFC 3584 (http://tools.ietf.org/html/rfc3584)

Note Be aware that SNMP versions 1 and 2c have serious known security issues: they transmit all information without encryption, including the community string, which serves as the only form of authentication in these versions.

SNMP Notifications

A key feature of SNMP is the ability to generate notifications from an SNMP agent. These notifications do not require that requests be sent from the SNMP manager. Notifications can indicate improper user authentication, restarts, the closing of a connection, loss of connection to a neighbor router, or other significant events.

The chassis generates SNMP notifications as either traps or informs. Traps are less reliable than informs because the SNMP manager does not send any acknowledgment when it receives a trap, and the chassis cannot determine if the trap was received. An SNMP manager that receives an inform request acknowledges the message with an SNMP response protocol data unit (PDU). If the chassis does not receive the PDU, it can send the inform request again.

However, informs are available only with SNMPv2c, which is considered insecure, and is not recommended.

Note The ifindex order on the interface that uses SNMP does not change after you reboot the FXOS. However, the index number on the FXOS disk usage OID changes when you reboot the FXOS.

SNMP Security Levels and Privileges

SNMPv1, SNMPv2c, and SNMPv3 each represent a different security model. The security model combines with the selected security level to determine the security mechanism applied when the SNMP message is processed.

The security level determines the privileges required to view the message associated with an SNMP trap. The privilege level determines whether the message needs to be protected from disclosure or authenticated. The supported security level depends upon which security model is implemented. SNMP security levels support one or more of the following privileges:

- noAuthNoPriv—No authentication or encryption

- authNoPriv—Authentication but no encryption

- authPriv—Authentication and encryption

SNMPv3 provides for both security models and security levels. A security model is an authentication strategy that is set up for a user and the role in which the user resides. A security level is the permitted level of security

within a security model. A combination of a security model and a security level determines which security mechanism is employed when handling an SNMP packet.

Supported Combinations of SNMP Security Models and Levels

The following table identifies what the combinations of security models and levels mean.

Table 10: SNMP Security Models and Levels

Model	Level	Authentication	Encryption	What Happens
v1	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v2c	noAuthNoPriv	Community string	No	Uses a community string match for authentication.
v3	noAuthNoPriv	Username	No	Uses a username match for authentication. **Note** While you can configure it, FXOS does not support use of `noAuthNoPriv` with SNMP version 3.
v3	authNoPriv	HMAC-SHA	No	Provides authentication based on the HMAC Secure Hash Algorithm (SHA).
v3	authPriv	HMAC-SHA	DES	Provides authentication based on the HMAC-SHA algorithm. Provides Data Encryption Standard (DES) 56-bit encryption in addition to authentication based on the Cipher Block Chaining (CBC) DES (DES-56) standard.

SNMPv3 Security Features

SNMPv3 provides secure access to devices by a combination of authenticating and encrypting frames over the network. SNMPv3 authorizes management operations only by configured users and encrypts SNMP messages. The SNMPv3 User-Based Security Model (USM) refers to SNMP message-level security and offers the following services:

- Message integrity—Ensures that messages have not been altered or destroyed in an unauthorized manner and that data sequences have not been altered to an extent greater than can occur non-maliciously.

- Message origin authentication—Ensures that the claimed identity of the user on whose behalf received data was originated is confirmed.

- Message confidentiality and encryption—Ensures that information is not made available or disclosed to unauthorized individuals, entities, or processes.

SNMP Support

The chassis provides the following support for SNMP:

Support for MIBs

The chassis supports read-only access to MIBs.

For information about the specific MIBs available and where you can obtain them, see the Cisco FXOS MIB Reference Guide.

Authentication Protocol for SNMPv3 Users

The chassis supports the HMAC-SHA-96 (SHA) authentication protocol for SNMPv3 users.

AES Privacy Protocol for SNMPv3 Users

The chassis uses Advanced Encryption Standard (AES) as one of the privacy protocols for SNMPv3 message encryption and conforms with RFC 3826.

The privacy password, or priv option, offers a choice of DES or 128-bit AES encryption for SNMP security encryption. If you enable AES-128 configuration and include a privacy password for an SNMPv3 user, the chassis uses the privacy password to generate a 128-bit AES key. The AES privacy password can have a minimum of eight characters. If the passphrases are specified in clear text, you can specify a maximum of 64 characters.

Enabling SNMP and Configuring SNMP Properties

Procedure

Step 1   Choose **Platform Settings** > **SNMP**.

Step 2   In the **SNMP** area, complete the following fields:

Name	Description
Admin State check box	Whether SNMP is enabled or disabled. Enable this service only if your system includes integration with an SNMP server.
Port field	The port on which the chassis communicates with the SNMP host. You cannot change the default port.

Name	Description
Community/Username field	(Optional) The community string used for polling in SNMP v1 and v2. When you specify an SNMP community name, you are also automatically enabling SNMP versions 1 and 2c for polling requests from the SNMP remote manager. This field is not applicable to SNMP v3. Be aware that SNMP versions 1 and 2c have serious known security issues: they transmit all information without encryption, including the community string, which serves as the only form of authentication in these versions. Enter an alphanumeric string between 1 and 32 characters. Do not use @ (at sign), & (ampersand), ? (question mark) or an empty space. The default is **public**. If the **Community/Username** field is already set, the text to the right of the empty field reads **Set: Yes**. If the **Community/Username** field is not yet populated with a value, the text to the right of the empty field reads **Set: No**. **Note** You can use the CLI command **set snmp community** to delete an existing community string, thereby disabling SNMP versions 1 and 2c for polling requests from the SNMP remote manager.
System Administrator Name field	The contact person responsible for the SNMP implementation. Enter a string of up to 255 characters, such as an email address or a name and telephone number.
Location field	The location of the host on which the SNMP agent (server) runs. Enter an alphanumeric string up to 510 characters.

Step 3   Click **Save**.

What to do next

Create SNMP traps and users.

Creating an SNMP Trap

The following procedure describes how to create SNMP traps.

Note   You can define up to eight SNMP traps.

Procedure

Step 1	Choose **Platform Settings** > **SNMP**.
Step 2	In the **SNMP Traps** area, click **Add**.
Step 3	In the **Add SNMP Trap** dialog box, complete the following fields:

Name	Description
Host Name field	The hostname or IP address of the SNMP host to which the chassis should send the traps.
Community/Username field	Enter the SNMPv1/v2c community string, or the SNMPv3 user name, needed to permit access to the trap destination. This must be the same as the community or user name that is configured for the SNMP service.
	Enter an alphanumeric string between 1 and 32 characters. Do not use @ (at sign), \ (backslash), " (double quote), ? (question mark) or an empty space.
Port field	The port on which the chassis communicates with the SNMP host for the trap.
	Enter an integer between 1 and 65535.
Version field	The SNMP version and model used for the trap. This can be one of the following:
	• **V1**
	• **V2**
	• **V3**
	Note Be aware that SNMP versions 1 and 2c have serious known security issues: they transmit all information without encryption, including the community string, which serves as the only form of authentication in these versions.
Type field	Specify the type of trap to send:
	• **Traps**
	• **Informs** (only valid when **Version** is **V2**)
v3 Privilege field	If you selected **V3** for the version, specify the privilege level associated with the trap:
	• **Auth**—Authentication but no encryption.
	• **Noauth**—No authentication or encryption. Note that while you can select it, FXOS does not support this security level with SNMPv3.
	• **Priv**—Authentication and encryption.

Step 4   Click **OK** to close the **Add SNMP Trap** dialog box.

Step 5   Click **Save**.

Deleting an SNMP Trap

Procedure

Step 1   Choose **Platform Settings** > **SNMP**.

Step 2   In the **SNMP Traps** area, click the **Delete** icon in the row in the table that corresponds to the trap you want to delete.

Creating an SNMPv3 User

Procedure

Step 1   Choose **Platform Settings** > **SNMP**.

Step 2   In the **SNMP Users** area, click **Add**.

Step 3   In the **Add SNMP User** dialog box, complete the following fields:

Name	Description
Name field	The user name assigned to the SNMPv3 user.
	Enter up to 32 characters. The name must begin with a letter. Valid characters include letters, numbers, _ (underscore), . (period), @ (at sign), and - (hyphen).
Auth Type field	The authorization type: **SHA**.
Use AES-128 check box	If checked, this user uses AES-128 encryption.
	Note SNMPv3 does not support DES. If you leave the AES-128 box unchecked, no privacy encryption will be done and any configured privacy password will have no effect.

Name	Description
Password field	The password for this user. The FXOS rejects any password that does not meet the following requirements: • Must contain a minimum of 8 characters and a maximum of 80 characters. • Must contain only letters, numbers, and the following characters: ~`!@#%^&*()_-+{}[]|\\:;"'<,>./ • Must not contain the following symbols: $ (dollar sign), ? (question mark), or = (equals sign). • Must contain at least five different characters. • Must not contain too many consecutively incrementing or decrementing numbers or letters. For example, the string "12345" has four such characters, and the string "ZYXW" has three. If the total number of such characters exceeds a certain limit (typically more than around 4-6 such occurrences), the simplicity check will fail. **Note** The consecutively incrementing or decrementing character count is not reset when non-incrementing or decrementing characters are used in between. For example, abcd&!21 will fail the password check, but abcd&!25, will not.
Confirm Password field	The password again for confirmation purposes.

Name	Description
Privacy Password field	The privacy password for this user.
	The FXOS rejects any password that does not meet the following requirements:
	• Must contain a minimum of 8 characters and a maximum of 80 characters.
	• Must contain only letters, numbers, and the following characters:
	~`!@#%^&*()_-+{}[]|\:;"'<,>./
	• Must not contain the following symbols: $ (dollar sign), ? (question mark), or = (equals sign).
	• Must contain at least five different characters.
	• Must not contain too many consecutively incrementing or decrementing numbers or letters. For example, the string "12345" has four such characters, and the string "ZYXW" has three. If the total number of such characters exceeds a certain limit (typically more than around 4-6 such occurrences), the simplicity check will fail.
	Note The consecutively incrementing or decrementing character count is not reset when non-incrementing or decrementing characters are used in between. For example, abcd&!21 will fail the password check, but abcd&!25, will not.
Confirm Privacy Password field	The privacy password again for confirmation purposes.

Step 4   Click **OK** to close the **Add SNMP User** dialog box.

Step 5   Click **Save**.

Deleting an SNMPv3 User

Procedure

Step 1   Choose **Platform Settings** > **SNMP**.

Step 2   In the **SNMP Users** area, click the **Delete** icon in the row in the table that corresponds to the user you want to delete.

Configuring HTTPS

This section describes how to configure HTTPS on the Firepower 4100/9300 chassis.

Note You can change the HTTPS port using Firepower Chassis Manager or the FXOS CLI. All other HTTPS configuration can only be done using the FXOS CLI.

Certificates, Key Rings, and Trusted Points

HTTPS uses components of the Public Key Infrastructure (PKI) to establish secure communications between two devices, such as a client's browser and the Firepower 4100/9300 chassis.

Encryption Keys and Key Rings

Each PKI device holds a pair of asymmetric Rivest-Shamir-Adleman (RSA) encryption keys, one kept private and one made public, stored in an internal key ring. A message encrypted with either key can be decrypted with the other key. To send an encrypted message, the sender encrypts the message with the receiver's public key, and the receiver decrypts the message using its own private key. A sender can also prove its ownership of a public key by encrypting (also called 'signing') a known message with its own private key. If a receiver can successfully decrypt the message using the public key in question, the sender's possession of the corresponding private key is proven. Encryption keys can vary in length, with typical lengths from 512 bits to 2048 bits. In general, a longer key is more secure than a shorter key. FXOS provides a default key ring with an initial 2048-bit key pair, and allows you to create additional key rings.

The default key ring certificate must be manually regenerated if the cluster name changes or the certificate expires.

Certificates

To prepare for secure communications, two devices first exchange their digital certificates. A certificate is a file containing a device's public key along with signed information about the device's identity. To merely support encrypted communications, a device can generate its own key pair and its own self-signed certificate. When a remote user connects to a device that presents a self-signed certificate, the user has no easy method to verify the identity of the device, and the user's browser will initially display an authentication warning. By default, FXOS contains a built-in self-signed certificate containing the public key from the default key ring.

Trusted Points

To provide stronger authentication for FXOS, you can obtain and install a third-party certificate from a trusted source, or trusted point, that affirms the identity of your device. The third-party certificate is signed by the issuing trusted point, which can be a root certificate authority (CA) or an intermediate CA or trust anchor that is part of a trust chain that leads to a root CA. To obtain a new certificate, you must generate a certificate request through FXOS and submit the request to a trusted point.

Important The certificate must be in Base64 encoded X.509 (CER) format.

Creating a Key Ring

FXOS supports a maximum of 8 key rings, including the default key ring.

Procedure

Step 1 Enter security mode:

Firepower-chassis # **scope security**

Step 2 Create and name the key ring:

Firepower-chassis # **create keyring** *keyring-name*

Step 3 Set the SSL key length in bits:

Firepower-chassis # **set modulus** {**mod1024** | **mod1536** | **mod2048** | **mod512**}

Step 4 Commit the transaction:

Firepower-chassis # **commit-buffer**

Example

The following example creates a keyring with a key size of 1024 bits:

```
Firepower-chassis# scope security
Firepower-chassis /security # create keyring kr220
Firepower-chassis /security/keyring* # set modulus mod1024
Firepower-chassis /security/keyring* # commit-buffer
Firepower-chassis /security/keyring #
```

What to do next

Create a certificate request for this key ring.

Regenerating the Default Key Ring

The default key ring certificate must be manually regenerated if the cluster name changes or the certificate expires.

✎

Note The default keyring is only used by FCM on FXOS.

Procedure

Step 1 Enter security mode:

Firepower-chassis # **scope security**

Step 2 Enter key ring security mode for the default key ring:

Firepower-chassis /security # **scope keyring default**

Step 3 Regenerate the default key ring:

Firepower-chassis /security/keyring # **set regenerate yes**

Step 4 Commit the transaction:

Firepower-chassis # **commit-buffer**

Example

The following example regenerates the default key ring:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope keyring default
Firepower-chassis /security/keyring* # set regenerate yes
Firepower-chassis /security/keyring* # commit-buffer
Firepower-chassis /security/keyring #
```

Creating a Certificate Request for a Key Ring

Creating a Certificate Request for a Key Ring with Basic Options

Procedure

Step 1 Enter security mode:

Firepower-chassis # **scope security**

Step 2 Enter configuration mode for the key ring:

Firepower-chassis /security # **scope keyring** *keyring-name*

Step 3 Create a certificate request using the IPv4 or IPv6 address specified, or the name of the fabric interconnect. You are prompted to enter a password for the certificate request.

Firepower-chassis /security/keyring # **create certreq** {**ip** [*ipv4-addr* | *ipv6-v6*] |**subject-name** *name*}

Step 4 Commit the transaction:

Firepower-chassis /security/keyring/certreq # **commit-buffer**

Step 5 Display the certificate request, which you can copy and send to a trust anchor or certificate authority:

Firepower-chassis /security/keyring # **show certreq**

Example

The following example creates and displays a certificate request with an IPv4 address for a key ring, with basic options:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope keyring kr220
Firepower-chassis /security/keyring # create certreq ip 192.168.200.123 subject-name sjc04
Certificate request password:
Confirm certificate request password:
Firepower-chassis /security/keyring* # commit-buffer
Firepower-chassis /security/keyring # show certreq
Certificate request subject name: sjc04
Certificate request ip address: 192.168.200.123
Certificate request e-mail name:
Certificate request country name:
State, province or county (full name):
Locality (eg, city):
Organization name (eg, company):
Organization Unit name (eg, section):
Request:
-----BEGIN CERTIFICATE REQUEST-----
MIIBfTCB5wIBADARMQ8wDQYDVQQDEwZzYW1jMDQwgZ8wDQYJKoZIhvcNAQEBBQAD
gY0AMIGJAoGBALpKn1t8qMZO4UGqILKFXQQc2c8b/vW2rnRF8OPhKbhghLA1YZ1F
JqcYEG5Yl1+vgohLBTd45s0GC8m4RTLJWHo4SwccAUXQ5Zngf45YtX1WsylwUWV4
0re/zgTk/WCd56RfOBvWR2Dtztu2pGA14sd761zLxt29K7R8mzj6CAUVAgMBAAGg
LTArBgkqhkiG9w0BCQ4xHjAcMBoGA1UdEQEB/wQQMA6CBnNhbWMwWMwNIcECsEiXjAN
BgkqhkiG9w0BAQQFAAOBgQCsxN0qUHYGFoQw56RwQueLTNPnrndqUwuZHUO03Teg
nhsyu4satpyiPqVV9viKZ+spvc6x5PWIcTWgHhH8BimOb/0OKuG8kwfIGGsEDlAv
TTYvUP+BZ9OFiPbRIA718S+V8ndXr1HejiQGxlDNqoN+odCXPc5kjoXD0lZTL09H
BA==
-----END CERTIFICATE REQUEST-----

Firepower-chassis /security/keyring #
```

What to do next

- Copy the text of the certificate request, including the BEGIN and END lines, and save it in a file. Send the file with the certificate request to a trust anchor or certificate authority to obtain a certificate for the key ring.

- Create a trusted point and set the certificate chain for the certificate of trust received from the trust anchor.

Creating a Certificate Request for a Key Ring with Advanced Options

Procedure

Step 1	Enter security mode:
	Firepower-chassis # **scope security**
Step 2	Enter configuration mode for the key ring:
	Firepower-chassis /security # **scope keyring** *keyring-name*
Step 3	Create a certificate request:

Firepower-chassis /security/keyring # **create certreq**

Step 4   Specify the country code of the country in which the company resides:

Firepower-chassis /security/keyring/certreq* # **set country** *country name*

Step 5   Specify the Domain Name Server (DNS) address associated with the request:

Firepower-chassis /security/keyring/certreq* # **set dns** *DNS Name*

Step 6   Specify the email address associated with the certificate request:

Firepower-chassis /security/keyring/certreq* # **set e-mail** *E-mail name*

Step 7   Specify the IP address of the Firepower 4100/9300 chassis:

Firepower-chassis /security/keyring/certreq* # **set ip** {*certificate request ip-address*|*certificate request ip6-address* }

Step 8   Specify the city or town in which the company requesting the certificate is headquartered:

Firepower-chassis /security/keyring/certreq* # **set locality** *locality name (eg, city)*

Step 9   Specify the organization requesting the certificate:

Firepower-chassis /security/keyring/certreq* # **set org-name** *organization name*

Step 10   Specify the organizational unit:

Firepower-chassis /security/keyring/certreq* # **set org-unit-name** *organizational unit name*

Step 11   Specify an optional password for the certificate request:

Firepower-chassis /security/keyring/certreq* # **set password** *certificate request password*

Step 12   Specify the state or province in which the company requesting the certificate is headquartered:

Firepower-chassis /security/keyring/certreq* # **set state** *state, province or county*

Step 13   Specify the fully qualified domain name of the Firepower 4100/9300 chassis:

Firepower-chassis /security/keyring/certreq* # **set subject-name** *certificate request name*

Step 14   Commit the transaction:

Firepower-chassis /security/keyring/certreq # **commit-buffer**

Step 15   Display the certificate request, which you can copy and send to a trust anchor or certificate authority:

Firepower-chassis /security/keyring # **show certreq**

Example

The following example creates and displays a certificate request with an IPv4 address for a key ring, with advanced options:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope keyring kr220
Firepower-chassis /security/keyring # create certreq
Firepower-chassis /security/keyring/certreq* # set "ip 192.168.200.123"
Firepower-chassis /security/keyring/certreq* #  set subject-name "sjc04"
Firepower-chassis /security/keyring/certreq* # set country "US"
Firepower-chassis /security/keyring/certreq* # set dns "bg1-samc-15A"
Firepower-chassis /security/keyring/certreq* # set email "test@cisco.com"
Firepower-chassis /security/keyring/certreq* # set locality "new york city"
Firepower-chassis /security/keyring/certreq* # set org-name "Cisco Systems"
Firepower-chassis /security/keyring/certreq* # set org-unit-name "Testing"
Firepower-chassis /security/keyring/certreq* # set state "new york"
Firepower-chassis /security/keyring/certreq* # commit-buffer
Firepower-chassis /security/keyring/certreq # show certreq
Certificate request subject name: sjc04
Certificate request ip address: 192.168.200.123
Certificate request e-mail name: test@cisco.com
Certificate request country name: US
State, province or county (full name): New York
Locality name (eg, city): new york city
Organization name (eg, company): Cisco
Organization Unit name (eg, section): Testing
Request:
-----BEGIN CERTIFICATE REQUEST-----
MIIBfTCB5wIBADARMQ8wDQYDVQQDEwZzYW1jMDQwgZ8wDQYJKoZIhvcNAQEBBQAD
gY0AMIGJAoGBALpKn1t8qMZO4UGqILKFXQQc2c8b/vW2rnRF8OPhKbhghLA1YZ1F
JqcYEG5Yl1+vgohLBTd45s0GC8m4RTLJWHo4SwccAUXQ5Zngf45YtX1WsylwUWV4
0re/zgTk/WCd56RfOBvWR2Dtztu2pGA14sd761zLxt29K7R8mzj6CAUVAgMBAAGg
LTArBgkqhkiG9w0BCQ4xHjAcMBoGA1UdEQEB/wQQMA6CBnNhbWMwNIcECsEiXjAN
BgkqhkiG9w0BAQQFAAOBgQCsxN0qUHYGFoQw56RwQueLTNPnrndqUwuZHUO03Teg
nhsyu4satpyiPqVV9viKZ+spvc6x5PWIcTWgHhH8BimOb/0OKuG8kwfIGGsEDlAv
TTYvUP+BZ9OFiPbRIA718S+V8ndXr1HejiQGxlDNqoN+odCXPc5kjoXD0lZTL09H
BA==
-----END CERTIFICATE REQUEST-----

Firepower-chassis /security/keyring/certreq #
```

What to do next

- Copy the text of the certificate request, including the BEGIN and END lines, and save it in a file. Send the file with the certificate request to a trust anchor or certificate authority to obtain a certificate for the key ring.

- Create a trusted point and set the certificate chain for the certificate of trust received from the trust anchor.

Creating a Trusted Point

Procedure

Step 1 Enter security mode:

Firepower-chassis # **scope security**

Step 2 Create a trusted point:

Firepower-chassis /security # **create trustpoint** *name*

Step 3 Specify certificate information for this trusted point:

Firepower-chassis /security/trustpoint # **set certchain** [*certchain*]

If you do not specify certificate information in the command, you are prompted to enter a certificate or a list of trustpoints defining a certification path to the root certificate authority (CA). On the next line following your input, type **ENDOFBUF** to finish.

Important The certificate must be in Base64 encoded X.509 (CER) format.

Step 4 Commit the transaction:

Firepower-chassis /security/trustpoint # **commit-buffer**

Example

The following example creates a trusted point and provides a certificate for the trusted point:

```
Firepower-chassis# scope security
Firepower-chassis /security # create trustpoint tPoint10
Firepower-chassis /security/trustpoint* # set certchain
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Trustpoint Certificate Chain:
> -----BEGIN CERTIFICATE-----
> MIIDMDCCApmgAwIBAgIBADANBgkqhkiG9w0BAQQFADB0MQswCQYDVQQGEwJVUzEL
> BxMMU2FuIEpvc2UsIENBMRUwEwYDVQQKEwxFeGFtcGxlIEluYy4xEzARBgNVBAsT
> ClRlc3QgR3JvdXAxGTAXBgNVBAMTEHRlc3QuZXhhbXBsZS5jb20xHzAdBgkqhkiG
> 9w0BCQEWEHVzZXJAZXhhbXBsZS5jb20wgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJ
> AoGBAMZw4nTepNIDhVzb0j7Z2Je4xAG56zmSHRMQeOGHemdh66u2/XAoLx7YCcYU
> ZgAMivyCsKgb/6CjQtsofvtrmC/eAehuK3/SINv7wd6Vv2pBt6ZpXgD4VBNKONDl
> GMbkPayVlQjbG4MD2dx2+H8EH3LMtdZrgKvPxTE+bF5wZVNAgMBAAGgJTAjBgkq
> hkiG9w0BCQcxFhMUQSBjaGFsbGVuZ2UgcGFzc3dvcmQwDQYJKoZIhvcNAQEFBQAD
> gYEAG61CaJoJaVMhzCl903O6Mg51zq1zXcz75+VFj2I6rH9asckCld3mkOVx5gJU
> Ptt5CVQpNgNLdvbDPSsXretysOhqHmp9+CLv8FDuy1CDYfuaLtvlWvfhevskV0j6
> jtcEMyZ+f7+3yh421ido3nO4MIGeBgNVHSMEgZYwgZOAFLlNjtcEMyZ+f7+3yh42
> 1ido3nO4oXikdjB0MQswCQYDVQQGEwJVUzELMAkGA1UECBMCQ0ExFDASBgNVBAcT
> C1NhbnRhIENsYXJhMRswGQYDVQQKExJOdW92YSBTeXN0ZW1zIEluYy4xFDASBgNV
> BAsTC0VuZ2luZWVyaW5nMQ8wDQYDVQQDEwZ0ZXN0Q0GCAQAwDAYDVR0TBAUwAwEB
> /zANBgkqhkiG9w0BAQQFAAOBgQAhWaRwXNR6B4g6Lsnr+fptHv+WVhB5fKqGQqXc
> wR4pYiO4z42/j9Ijenh75tCKMhW51az8copP1EBmOcyuhf5C6vasrenn1ddkkYt4
> PR0vxGc40whuiozBolesmsmjBbedUCwQgdFDWhDIZJwK5+N3x/kfa2EHU6id1avt
> 4YL5Jg==
> -----END CERTIFICATE-----
```

```
> ENDOFBUF
Firepower-chassis /security/trustpoint* # commit-buffer
Firepower-chassis /security/trustpoint #
```

What to do next

Obtain a key ring certificate from the trust anchor or certificate authority and import it into the key ring.

Importing a Certificate into a Key Ring

Before you begin

- Configure a trusted point that contains the certificate chain for the key ring certificate.

- Obtain a key ring certificate from a trust anchor or certificate authority.

✎

Note   If you change the certificate in a key ring that has already been configured on HTTPS, you must restart HTTPS in order for the new certificate to take effect. For more informaiton, see: Restarting HTTPS, on page 128.

Procedure

Step 1   Enter security mode:

Firepower-chassis # **scope security**

Step 2   Enter configuration mode for the key ring that will receive the certificate:

Firepower-chassis /security # **scope keyring** *keyring-name*

Step 3   Specify the trusted point for the trust anchor or certificate authority from which the key ring certificate was obtained:

Firepower-chassis /security/keyring # **set trustpoint** *name*

Step 4   Launch a dialog for entering and uploading the key ring certificate:

Firepower-chassis /security/keyring # **set cert**

At the prompt, paste the certificate text that you received from the trust anchor or certificate authority. On the next line following the certificate, type **ENDOFBUF** to complete the certificate input.

Important   The certificate must be in Base64 encoded X.509 (CER) format.

Step 5   Commit the transaction:

Firepower-chassis /security/keyring # **commit-buffer**

Example

The following example specifies the trust point and imports a certificate into a key ring:

```
Firepower-chassis# scope security
Firepower-chassis /security # scope keyring kr220
Firepower-chassis /security/keyring # set trustpoint tPoint10
Firepower-chassis /security/keyring* # set cert
Enter lines one at a time. Enter ENDOFBUF to finish. Press ^C to abort.
Keyring certificate:
> -----BEGIN CERTIFICATE-----
> MIIB/zCCAWgCAQAwgZkxCzAJBgNVBAYTAlVTMQswCQYDVQQIEwJDQTEVMBMGA1UE
> BxMMU2FuIEpvc2UsIENBMRUwEwYDVQQKEwxFeGFtcGxlIEluYy4xEzARBgNVBAsT
> ClRlc3QgR3JvdXAxGTAXBgNVBAMTEHRlc3QuZXhhbXBsZS5jb20xHzAdBgkqhkiG
> 9w0BCQEWEHZzZXJAZXhhbXBsZS5jb20wgZ8wDQYJKoZIhvcNAQEBBQADgY0AMIGJ
> AoGBAMZw4nTepNIDhVzb0j7Z2Je4xAG56zmSHRMQeOGHemdh66u2/XAoLx7YCcYU
> ZgAMivyCsKgb/6CjQtsofvtrmC/eAehuK3/SINv7wd6Vv2pBt6ZpXgD4VBNKONl
> GMbkPayVlQjbG4MD2dx2+H8EH3LMtdZrgKvPxPTE+bF5wZVNAgMBAAGgJTAjBgkq
> hkiG9w0BCQcxFhMUQSBjaGFsbGVuZ2UgcGFzc3dvcmQwDQYJKoZIhvcNAQEFBQAD
> gYEAG61CaJoJaVMhzCl903O6Mg51zq1zXcz75+VFj2I6rH9asckCld3mkOVx5gJU
> Ptt5CVQpNgNLdvbDPSsXretysOhqHmp9+CLv8FDuy1CDYfuaLtvlWvfhevskV0j6
> mK3Ku+YiORnv6DhxrOoqau8r/hyI/L43l7IPN1HhOi3oha4=
> -----END CERTIFICATE-----
> ENDOFBUF
Firepower-chassis /security/keyring* # commit-buffer
Firepower-chassis /security/keyring #
```

What to do next

Configure your HTTPS service with the key ring.

Configuring HTTPS

> ⚠️
>
> **Caution** After you complete the HTTPS configuration, including changing the port and key ring to be used by HTTPS, all current HTTP and HTTPS sessions are closed without warning as soon as you save or commit the transaction.

Procedure

Step 1   Enter system mode:

Firepower-chassis# **scope system**

Step 2   Enter system services mode:

Firepower-chassis /system # **scope services**

Step 3   Enable the HTTPS service:

Firepower-chassis /system/services # **enable https**

Step 4   (Optional) Specify the port to be used for the HTTPS connection:

Firepower-chassis /system/services # **set https port** *port-num*

Step 5   (Optional) Specify the name of the key ring you created for HTTPS:

Firepower-chassis /system/services # **set https keyring** *keyring-name*

Step 6   (Optional) Specify the level of Cipher Suite security used by the domain:

Firepower-chassis /system/services # **set https cipher-suite-mode** *cipher-suite-mode*

cipher-suite-mode can be one of the following keywords:

- **high-strength**

- **medium-strength**

- **low-strength**

- **custom**—Allows you to specify a user-defined Cipher Suite specification string.

Step 7   (Optional) If **cipher-suite-mode** is set to **custom**, specify a custom level of Cipher Suite security for the domain:

Firepower-chassis /system/services # **set https cipher-suite** *cipher-suite-spec-string*

cipher-suite-spec-string can contain up to 256 characters and must conform to the OpenSSL Cipher Suite specifications. You cannot use any spaces or special characters except ! (exclamation point), + (plus sign), - (hyphen), and : (colon). For details, see http://httpd.apache.org/docs/2.0/mod/mod_ssl.html#sslciphersuite.

For example, the medium strength specification string FXOS uses as the default is:
ALL:!ADH:!EXPORT56:!LOW:RC4+RSA:+HIGH:+MEDIUM:+EXP:+eNULL

Note   This option is ignored if **cipher-suite-mode** is set to anything other than **custom**.

Step 8   Commit the transaction to the system configuration:

Firepower-chassis /system/services # **commit-buffer**

Example

The following example enables HTTPS, sets the port number to 443, sets the key ring name to kring7984, sets the Cipher Suite security level to high, and commits the transaction:

```
Firepower-chassis# scope system
Firepower-chassis /system # scope services
Firepower-chassis /system/services # enable https
Firepower-chassis /system/services* # set https port 443
Warning: When committed, this closes all the web sessions.
Firepower-chassis /system/services* # set https keyring kring7984
Firepower-chassis /system/services* # set https cipher-suite-mode high
Firepower-chassis /system/services* # commit-buffer
Firepower-chassis /system/services #
```

Changing the HTTPS Port

The HTTPS service is enabled on port 443 by default. You cannot disable HTTPS, but you can change the port to use for HTTPS connections.

Procedure

Step 1   Choose **Platform Settings** > **HTTPS**.

Step 2   Enter the port to use for HTTPS connections in the **Port** field. Specify an integer between 1 and 65535. This service is enabled on port 443 by default.

Step 3   Click **Save**.

The chassis is configured with the HTTPS port specified.

After changing the HTTPS port, all current HTTPS sessions are closed. Users will need to log back in to the Firepower Chassis Manager using the new port as follows:

`https://<chassis_mgmt_ip_address>:<chassis_mgmt_port>`

where *<chassis_mgmt_ip_address>* is the IP address or host name of the chassis that you entered during initial configuration and *<chassis_mgmt_port>* is the HTTPS port you have just configured.

Restarting HTTPS

If you change the certificate in a key ring that has already been configured on HTTPS, you must restart HTTPS in order for the new certificate to take effect. Use the following procedure to reset HTTPS with an updated keyring.

Procedure

Step 1   Enter system mode:

Firepower-chassis# **scope system**

Step 2   Enter system services mode:

Firepower-chassis /system # **scope services**

Step 3   Set the HTTPS key ring back to its default value:

Firepower-chassis /system/services # **set https keyring default**

Step 4   Commit the transaction to the system configuration:

Firepower-chassis /system/services # **commit-buffer**

Step 5   Wait five seconds.

Step 6   Set HTTPS with the key ring you created:

Firepower-chassis /system/services # **set https keyring** *keyring-name*

Step 7   Commit the transaction to the system configuration:

Firepower-chassis /system/services # **commit-buffer**

Deleting a Key Ring

Procedure

Step 1   Enter security mode:

Firepower-chassis # **scope security**

Step 2   Delete the named key ring:

Firepower-chassis /security # **delete keyring** *name*

Step 3   Commits the transaction:

Firepower-chassis /security # **commit-buffer**

Example

The following example deletes a key ring:

```
Firepower-chassis# scope security
Firepower-chassis /security # delete keyring key10
Firepower-chassis /security* # commit-buffer
Firepower-chassis /security #
```

Deleting a Trusted Point

Before you begin

Ensure that the trusted point is not used by a key ring.

Procedure

Step 1   Enters security mode:

Firepower-chassis# **scope security**

Step 2   Delete the named trusted point:

Firepower-chassis /security # **delete trustpoint** *name*

Step 3   Commits the transaction:

Firepower-chassis /security # **commit-buffer**

Example

The following example deletes a trusted point:

```
Firepower-chassis# scope security
Firepower-chassis /security # delete trustpoint tPoint10
Firepower-chassis /security* # commit-buffer
Firepower-chassis /security #
```

Disabling HTTPS

Procedure

Step 1   Enter system mode:

Firepower-chassis# **scope system**

Step 2   Enter system services mode:

Firepower-chassis /system # **scope services**

Step 3   Disable the HTTPS service:

Firepower-chassis /system/services # **disable https**

Step 4   Commit the transaction to the system configuration:

Firepower-chassis /system/services # **commit-buffer**

Example

The following example disables HTTPS and commits the transaction:

```
Firepower-chassis# scope system
Firepower-chassis /system # scope services
Firepower-chassis /system/services # disable https
Firepower-chassis /system/services* # commit-buffer
Firepower-chassis /system/services #
```

Configuring AAA

This section describes authentication, authorization, and accounting. See the following topics for more information:

About AAA

Authentication, Authorization and Accounting (AAA) is a set of services for controlling access to network resources, enforcing policies, assessing usage, and providing the information necessary to bill for services. Authentication identifies the user. Authorization implements policies that determine which resources and services an authenticated user may access. Accounting keeps track of time and data resources that are used for billing and analysis. These processes are considered important for effective network management and security.

Authentication

Authentication provides a way to identify each user, typically by having the user enter a valid user name and valid password before access is granted. The AAA server compares the user's provided credentials with user credentials stored in a database. If the credentials are matched, the user is permitted access to the network. If the credentials do not match, authentication fails and network access is denied.

You can configure the Firepower 4100/9300 chassis to authenticate administrative connections to the chassis, including the following sessions:

- HTTPS

- SSH

- Serial console

Authorization

Authorization is the process of enforcing policies: determining what types of activities, resources, or services each user is permitted to access. After authentication, a user may be authorized for different types of access or activity.

Accounting

Accounting measures the resources a user consumes during access, which may include the amount of system time or the amount of data that a user has sent or received during a session. Accounting is carried out through the logging of session statistics and usage information, which is used for authorization control, billing, trend analysis, resource utilization, and capacity planning activities.

Interaction Between Authentication, Authorization, and Accounting

You can use authentication alone, or with authorization and accounting. Authorization always requires a user to be authenticated first. You can use accounting alone, or with authentication and authorization.

Supported Types of Authentication

FXOS supports the following types of user Authentication:

- **Remote** – The following network AAA services are supported:

    - LDAP

    - RADIUS

    - TACACS+

• **Local** – The chassis maintains a local database that you can populate with user profiles. You can use this local database instead of AAA servers to provide user authentication, authorization, and accounting.

User Roles

FXOS supports local and remote Authorization in the form of user-role assignment. The roles that can be assigned are:

• **Admin** – Complete read-and-write access to the entire system. The default admin account is assigned this role by default and it cannot be changed.

• **AAA Administrator** – Read-and-write access to users, roles, and AAA configuration. Read access to the rest of the system.

• **Operations** – Read-and-write access to NTP configuration, Smart Call Home configuration for Smart Licensing, and system logs, including syslog servers and faults. Read access to the rest of the system.

• **Read-Only** – Read-only access to system configuration with no privileges to modify the system state.

See User Management, on page 39 for more information about local users and role assignments.

Setting Up AAA

These steps provide a basic outline for setting up Authentication, Authorization and Accounting (AAA) on a Firepower 4100/9300 appliance.

1. Configure the desired type(s) of user authentication:

• **Local** – User definitions and local authentication are part of User Management, on page 39.

• **Remote** – Configuring remote AAA server access is part of Platform Settings, specifically:

• Configuring LDAP Providers, on page 133

• Configuring RADIUS Providers, on page 136

• Configuring TACACS+ Providers, on page 138

Note If you will be using remote AAA servers, be sure to enable and configure AAA services on the remote servers before configuring remote AAA server access on the chassis.

2. Specify the default authentication method—this also is part of User Management, on page 39.

Note If Default Authentication and Console Authentication are both set to use the same remote authentication protocol (RADIUS, TACACS+, or LDAP), you cannot change certain aspects of that server's configuration (for example, deleting that server, or changing its order of assignment) without updating these user settings.

Configuring LDAP Providers

Configuring Properties for LDAP Providers

The properties that you configure in this task are the default settings for all provider connections of this type. If an individual provider includes a setting for any of these properties, the FXOS uses that setting and ignores the default setting.

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with the FXOS. This account should be given a non-expiring password.

Procedure

Step 1 Choose **Platform Settings** > **AAA**.

Step 2 Click the **LDAP** tab.

Step 3 In the **Properties** area, complete the following fields:

Name	Description
Timeout field	The length of time in seconds the system will spend trying to contact the LDAP database before it times out.
	Enter an integer from 1 to 60 seconds. The default value is 30 seconds. This property is required.
Attribute field	An LDAP attribute that stores the values for the user roles and locales. This property is always a name-value pair. The system queries the user record for the value that matches this attribute.
	Note that the `shell:roles="admin,aaa"` attribute value is required when configuring properties for LDAP providers.
Base DN field	The specific distinguished name in the LDAP hierarchy where the server should begin a search when a remote user logs in and the system attempts to get the user's DN based on their user name. The length of the base DN can be a maximum of 255 characters minus the length of *cn=$userid*, where $userid identifies the remote user attempting to access the chassis using LDAP authentication.
	This property is required for LDAP providers. If you do not specify a base DN on this tab, then you must specify one for each LDAP provider that you define.
Filter field	Enter the filter attribute to use with your LDAP server, for example *cn=$userid* or *sAMAccountName=$userid*. The LDAP search is restricted to those user names that match the defined filter. The filter must include *$userid*.
	This property is required. If you do not specify a filter on this tab then you must specify one for each LDAP provider that you define.

Step 4 Click **Save**.

What to do next

Create an LDAP provider.

Creating an LDAP Provider

Follow these steps to define and configure a LDAP provider—that is, a specific remote server providing LDAP-based AAA services for this appliance.

Note The FXOS supports a maximum of 16 LDAP providers.

Before you begin

If you are using Active Directory as your LDAP server, create a user account in the Active Directory server to bind with the FXOS. This account should be given a non-expiring password.

Procedure

Step 1 Choose **Platform Settings** > **AAA**.

Step 2 Click the **LDAP** tab.

Step 3 For each LDAP provider that you want to add:

a) In the **LDAP Providers** area, click **Add**.

b) In the **Add LDAP Provider** dialog box, complete the following fields:

Name	Description
Hostname/FQDN (or IP Address) field	The hostname or IP address of the LDAP server. If SSL is enabled, this field must exactly match a Common Name (CN) in the security certificate of the LDAP database.
Order field	The order in which the FXOS uses this provider to authenticate users. Enter an integer between 1 and 16, or enter **lowest-available** or **0** (zero) if you want the FXOS to assign the next available order based on the other providers defined in Firepower Chassis Manager or the FXOS CLI.
Bind DN field	The distinguished name (DN) for an LDAP database account that has read and search permissions for all objects under the base DN. The maximum supported string length is 255 ASCII characters.

Name	Description
Base DN field	The specific distinguished name in the LDAP hierarchy where the server should begin a search when a remote user logs in and the system attempts to get the user's DN based on their user name. The length of the base DN can be set to a maximum of 255 characters minus the length of CN=$userid, where $userid identifies the remote user attempting to access Firepower Chassis Manager or the FXOS CLI using LDAP authentication. This value is required unless a default base DN has been set on the **LDAP** tab.
Port field	The port through which Firepower Chassis Manager or the FXOS CLI communicates with the LDAP database. The standard port number is 389.
Enable SSL check box	If checked, encryption is required for communications with the LDAP database. If unchecked, authentication information will be sent as clear text. LDAP uses STARTTLS. This allows encrypted communication using port 389. **Note** STARTTLS operation requires the CA cert of the LDAP provider to be installed on the FXOS certificate chain.
Filter field	Enter the filter attribute to use with your LDAP server, for example *cn=$userid* or *sAMAccountName=$userid*. The LDAP search is restricted to those user names that match the defined filter. The filter must include *$userid*. This value is required unless a default filter has been set on the **LDAP** tab.
Attribute field	An LDAP attribute that stores the values for the user roles and locales. This property is always a name-value pair. The system queries the user record for the value that matches this attribute name. This value is required unless a default attribute has been set on the **LDAP** tab.
Key field	The password for the LDAP database account specified in the **Bind DN** field. You can enter any standard ASCII characters except for space, § (section sign), ? (question mark), or = (equal sign).
Confirm Key field	The LDAP database password repeated for confirmation.
Timeout field	The length of time in seconds the system will spend trying to contact the LDAP database before it times out. Enter an integer from 1 to 60 seconds, or enter 0 (zero) to use the global timeout value specified on the **LDAP** tab. The default is 30 seconds.

Name	Description
Vendor field	This selection identifies the vendor that is providing the LDAP provider or server details: • If the LDAP provider is Microsoft Active Directory, select **MS AD**. • If the LDAP provider is not Microsoft Active Directory, select **Open LDAP**. The default is **Open LDAP**.

c) Click **OK** to close the **Add LDAP Provider** dialog box.

Step 4   Click **Save**.

Step 5   (Optional) Enable the certification revocation list check:

Firepower-chassis /security/ldap/server # **set revoke-policy** *{strict | relaxed}*

Note   This configuration only takes effect if the SSL connection is enabled.

Deleting an LDAP Provider

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **LDAP** tab.

Step 3   In the **LDAP Providers** area, click the **Delete** icon in the row in the table that corresponds to the LDAP Provider you want to delete.

Configuring RADIUS Providers

Configuring Properties for RADIUS Providers

The properties that you configure in this task are the default settings for all provider connections of this type. If an individual provider includes a setting for any of these properties, the FXOS uses that setting and ignores this default setting.

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **RADIUS** tab.

Step 3   In the **Properties** area, complete the following fields:

Name	Description
Timeout field	The length of time in seconds the system will spend trying to contact the RADIUS database before it times out. Enter an integer from 1 to 60 seconds. The default value is 5 seconds. This property is required.
Retries field	The number of times to retry the connection before the request is considered to have failed.

Step 4   Click **Save**.

What to do next

Create a RADIUS provider.

Creating a RADIUS Provider

Follow these steps to define and configure a RADIUS provider—that is, a specific remote server providing RADIUS-based AAA services for this appliance.

Note   The FXOS supports a maximum of 16 RADIUS providers.

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **RADIUS** tab.

Step 3   For each RADIUS provider that you want to add:

a) In the **RADIUS Providers** area, click **Add**.

b) In the **Add RADIUS Provider** dialog box, complete the following fields:

Name	Description
Hostname/FQDN (or IP Address) field	The hostname or IP address of the RADIUS server.
Order field	The order in which the FXOS uses this provider to authenticate users. Enter an integer between 1 and 16, or enter **lowest-available** or **0** (zero) if you want the FXOS to assign the next available order based on the other providers defined in Firepower Chassis Manager or the FXOS CLI.
Key field	The SSL encryption key for the database. You can enter any standard ASCII characters except for space, § (section sign), ? (question mark), or = (equal sign).

Name	Description
Confirm Key field	The SSL encryption key repeated for confirmation.
Authorization Port field	The port through which Firepower Chassis Manager or the FXOS CLI communicates with the RADIUS database. The valid range is 1 to 65535. The standard port number is 1700.
Timeout field	The length of time in seconds the system will spend trying to contact the RADIUS database before it times out. Enter an integer from 1 to 60 seconds, or enter 0 (zero) to use the global timeout value specified on the **RADIUS** tab. The default is 5 seconds.
Retries field	The number of times to retry the connection before the request is considered to have failed. If desired, enter an integer between 0 and 5. If you do not specify a value, Firepower Chassis Manager uses the value specified on the **RADIUS** tab.

c) Click **OK** to close the **Add RADIUS Provider** dialog box.

Step 4 Click **Save**.

Deleting a RADIUS Provider

Procedure

Step 1 Choose **Platform Settings** > **AAA**.

Step 2 Click the **RADIUS** tab.

Step 3 In the **RADIUS Providers** area, click the **Delete** icon in the row in the table that corresponds to the RADIUS Provider you want to delete.

Configuring TACACS+ Providers

Configuring Properties for TACACS+ Providers

The properties that you configure in this task are default settings for all provider connections of this type. If an individual provider configuration includes a setting for any of these properties, the FXOS uses that setting and ignores this default setting.

Note The FXOS chassis does not support command accounting for the Terminal Access Controller Access-Control System Plus (TACACS+) protocol.

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **TACACS** tab.

Step 3   In the **Properties** area, complete the following fields:

Name	**Description**
Timeout field	The length of time in seconds the system will spend trying to contact the TACACS+ database before it times out.
	Enter an integer from 1 to 60 seconds. The default value is 5 seconds.
	This property is required.

Step 4   Click **Save**.

What to do next

Create a TACACS+ provider.

Creating a TACACS+ Provider

Follow these steps to define and configure a TACACS+ provider—that is, a specific remote server providing TACACS-based AAA services for this appliance.

✎

Note   The FXOS supports a maximum of 16 TACACS+ providers.

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **TACACS** tab.

Step 3   For each TACACS+ provider that you want to add:

    a) In the **TACACS Providers** area, click **Add**.

    b) In the **Add TACACS Provider** dialog box, complete the following fields:

Name	**Description**
Hostname/FQDN (or IP Address) field	The hostname or IP address of the TACACS+ server.
Order field	The order in which the FXOS uses this provider to authenticate users.
	Enter an integer between 1 and 16, or enter **lowest-available** or **0** (zero) if you want the FXOS to assign the next available order based on the other providers defined in Firepower Chassis Manager or the FXOS CLI.

Name	Description
Key field	The SSL encryption key for the database. You can enter any standard ASCII characters except for space, § (section sign), ? (question mark), or = (equal sign).
Confirm Key field	The SSL encryption key repeated for confirmation.
Port field	The port through which Firepower Chassis Manager or the FXOS CLI communicates with this TACACS+ server. Enter an integer between 1 and 65535. The default port is 49.
Timeout field	The length of time in seconds the system will spend trying to contact the TACACS+ database before it times out. Enter an integer from 1 to 60 seconds, or enter 0 (zero) to use the global timeout value specified on the **TACACS+** tab. The default is 5 seconds.

    c) Click **OK** to close the **Add TACACS Provider** dialog box.

Step 4   Click **Save**.

Deleting a TACACS+ Provider

Procedure

Step 1   Choose **Platform Settings** > **AAA**.

Step 2   Click the **TACACS** tab.

Step 3   In the **TACACS Providers** area, click the **Delete** icon in the row in the table that corresponds to the TACACS+ Provider you want to delete.

Configuring Syslog

System logging is a method of collecting messages from devices to a server running a syslog daemon. Logging to a central syslog server helps in aggregation of logs and alerts. A syslog service accepts messages and stores them in files, or prints them according to a simple configuration file. This form of logging provides protected long-term storage for logs. Logs are useful both in routine troubleshooting and in incident handling.

Procedure

Step 1   Choose **Platform Settings** > **Syslog**.

Step 2   Configure Local Destinations:

    a) Click the **Local Destinations** tab.

b) On the **Local Destinations** tab, complete the following fields:

Name	Description
Console Section	
Admin State field	Whether the chassis displays syslog messages on the console.
	Check the **Enable** check box if you want to have syslog messages displayed on the console as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the console.
Level field	If you checked the **Enable** check box for **Console - Admin State**, select the lowest message level that you want displayed on the console. The chassis displays that level and above on the console. This can be one of the following:
	• **Emergencies**
	• **Alerts**
	• **Critical**
Monitor Section	
Admin State field	Whether the chassis displays syslog messages on the monitor.
	Check the **Enable** check box if you want to have syslog messages displayed on the monitor as well as added to the log. If the **Enable** check box is unchecked, syslog messages are added to the log but are not displayed on the monitor.
Level drop-down list	If you checked the **Enable** check box for **Monitor - Admin State**, select the lowest message level that you want displayed on the monitor. The system displays that level and above on the monitor. This can be one of the following:
	• **Emergencies**
	• **Alerts**
	• **Critical**
	• **Errors**
	• **Warnings**
	• **Notifications**
	• **Information**
	• **Debugging**

c) Click **Save**.

Step 3   Configure Remote Destinations:

a) Click the **Remote Destinations** tab.

b) On the **Remote Destinations** tab, complete the following fields for up to three external logs that can store messages generated by the chassis:

By sending syslog messages to a remote destination, you can archive messages according to the available disk space on the external syslog server, and manipulate logging data after it is saved. For example, you could specify actions to be executed when certain types of syslog messages are logged, extract data from the log and save the records to another file for reporting, or track statistics using a site-specific script.

Name	Description
Admin State field	Check the **Enable** check box if you want to have syslog messages stored in a remote log file.
Level drop-down list	Select the lowest message level that you want the system to store. The system stores that level and above in the remote file. This can be one of the following: • **Emergencies** • **Alerts** • **Critical** • **Errors** • **Warnings** • **Notifications** • **Information** • **Debugging**
Hostname/IP Address field	The hostname or IP address on which the remote log file resides. **Note** You must configure a DNS server if you use a hostname rather than an IP address.
Facility drop-down list	Choose a system log facility for syslog servers to use as a basis to file messages. This can be one of the following: • **Local0** • **Local1** • **Local2** • **Local3** • **Local4** • **Local5** • **Local6** • **Local7**

c) Click **Save**.

Step 4 Configure Local Sources:

a) Click the **Local Sources** tab.

b) On the **Local Sources** tab, complete the following fields:

Name	Description
Faults Admin State field	Whether system fault logging is enabled or not. If the **Enable** check box is checked, the chassis logs all system faults.
Audits Admin State field	Whether audit logging is enabled or not. If the **Enable** check box is checked, the chassis logs all audit log events.
Events Admin State field	Whether system event logging is enabled or not. If the **Enable** check box is checked, the chassis logs all system events.

c) Click **Save**.

Configuring DNS Servers

You need to specify a DNS server if the system requires resolution of host names to IP addresses. For example, you cannot use a name such as www.cisco.com when you are configuring a setting on the chassis if you do not configure a DNS server. You would need to use the IP address of the server, which can be either an IPv4 or an IPv6 address. You can configure up to four DNS servers.

Note When you configure multiple DNS servers, the system searches for the servers only in any random order. If a local management command requires DNS server lookup, it can only search for three DNS servers in random order.

Procedure

Step 1 Choose **Platform Settings** > **DNS**.

Step 2 Check the **Enable DNS Server** check box.

Step 3 For each DNS server that you want to add, up to a maximum of four, enter the IP address of the DNS server in the **DNS Server** field and click **Add**.

Step 4 Click **Save**.

Enable FIPS Mode

Perform these steps to enable FIPS mode on your Firepower 4100/9300 chassis.

Procedure

Step 1   Log into the Firepower 4100/9300 chassis as an admin user.

Step 2   Choose **Platform Settings** to open the Platform Settings window.

Step 3   Choose **FIPS/CC mode** to open the FIPS and Common Criteria window.

Step 4   Check the **Enable** checkbox for FIPS.

Step 5   Click **Save** to save the configuration.

Step 6   Follow the prompt to reboot the system.

When the FIPS Mode is enabled, it limits the key sizes and the algorithms allowed. The MIO uses CiscoSSL and the FIPS Object Module (FOM) for its cryptographic needs. It makes FIPS validation easier compared to ASA's proprietary cryptographic library implementation and HW acceleration.

What to do next

Prior to FXOS release 2.0.1, the existing SSH host key created during first-time setup of a device was hard coded to 1024 bits. To comply with FIPS and Common Criteria certification requirements, you must destroy this old host key and generate a new one using the procedure detailed in Generate the SSH Host Key. If you do not perform these additional steps, you will not be able to connect to the Supervisor using SSH after the device has rebooted with FIPS mode enabled. If you performed initial setup using FXOS 2.0.1 or later, you do not have to generate a new host key.

Enable Common Criteria Mode

Perform these steps to enable Common Criteria mode on your Firepower 4100/9300 chassis.

Procedure

Step 1   Log into the Firepower 4100/9300 chassis as an admin user.

Step 2   Choose **Platform Settings** to open the Platform Settings window.

Step 3   Choose **FIPS/CC mode** to open the FIPS and Common Criteria window.

Step 4   Check the **Enable** checkbox for Common Criteria.

Step 5   Click **Save** to save the configuration.

Step 6   Follow the prompt to reboot the system.

Common Criteria is an international standard for computer security. CC focuses on certificates, auditing, logging, passwords, TLS, SSH, etc. It essentially assumes FIPS compliance. Similar to FIPS, Cisco contracts with NIST accredited lab vendors to perform testing and submission to NIAP.

When the CC Mode is enabled, it limits the list of algorithms, cipher suites, and features that are needed to be supported. The MIO is evaluated against the Network Device Collaborative Protection Profile (NDcPP). CiscoSSL can only enforce part of the requirements most of which are covered in the CC compliance guide.

What to do next

Prior to FXOS release 2.0.1, the existing SSH host key created during first-time setup of a device was hard coded to 1024 bits. To comply with FIPS and Common Criteria certification requirements, you must destroy this old host key and generate a new one using the procedure detailed in Generate the SSH Host Key. If you do not perform these additional steps, you will not be able to connect to the Supervisor using SSH after the device has rebooted with Common Criteria mode enabled. If you performed initial setup using FXOS 2.0.1 or later, you do not have to generate a new host key.

Configure the IP Access List

By default, the Firepower 4100/9300 chassis denies all access to the local web server. You must configure your IP Access List with a list of allowed services for each of your IP blocks.

The IP Access List supports the following protocols:

- HTTPS

- SNMP

- SSH

For each block of IP addresses (v4 or v6), up to 100 different subnets can be configured for each service. A subnet of 0 and a prefix of 0 allows unrestricted access to a service.

Procedure

Step 1   Log into the Firepower 4100/9300 chassis as an admin user.

Step 2   Choose **Platform Settings** to open the Platform Settings page.

Step 3   Select **Access List** to open the Access List area.

Step 4   In this area, you can view, add, and delete the IPv4 and IPv6 addresses listed in your IP Access List.

To add an IPv4 block, you must enter a valid IPv4 IP address, a prefix [0-32] length, and select a protocol.

To add an IPv6 block, you must enter a valid IPv6 IP address, a prefix [0-128] length, and select a protocol

Add a MAC Pool Prefix and View MAC Addresses for Container Instance Interfaces

The FXOS chassis automatically generates MAC addresses for container instance interfaces, and guarantees that a shared interface in each instance uses a unique MAC address. The FXOS chassis generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or a system-defined prefix, and *zz.zzzz* is an internal counter generated by the chassis. The system-defined prefix matches the lower 2 bytes of the first MAC address in the burned-in

MAC address pool that is programmed into the IDPROM. Use **connect fxos**, then **show module** to view the MAC address pool. For example, if the range of MAC addresses shown for module 1 is b0aa.772f.f0b0 to b0aa.772f.f0bf, then the system prefix will be f0b0.

See Automatic MAC Addresses for Container Instance Interfaces, on page 190 for more information.

This procedure describes how to view the MAC addresses and how to optionally define the prefix used in generation.

> **Note** If you change the MAC address prefix after you deploy logical devices, you may experience traffic interruption.

Procedure

Step 1 Choose **Platform Settings** > **MAC Pool**.

This page shows generated MAC addresses along with the container instance and interface using the MAC address.

Step 2 (Optional) Add a MAC address prefix used in generating the MAC addresses.

a) Click **Add Prefix**.

The **Set the Prefix for the MAC Pool** dialogue box appears.

a) Enter a decimal value between 1 and 65535. This prefix is converted to a four-digit hexadecimal number, and used as part of the MAC address.

For an example of how the prefix is used, if you set a prefix of 77, then the chassis converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the chassis native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

b) Click **OK**.

New MAC addresses using the prefix are generated and assigned. The current prefix and the resulting hex value display above the table.

Add a Resource Profile for Container Instances

To specify resource usage per container instance, create one or more resource profiles. When you deploy the logical device/application instance, you specify the resource profile that you want to use. The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

• The minimum number of cores is 6.

> **Note** Instances with a smaller number of cores might experience relatively higher CPU utilization than those with larger numbers of cores. Instances with a smaller number of cores are more sensitive to traffic load changes. If you experience traffic drops, try assigning more cores.

- You can assign cores as an even number (6, 8, 10, 12, 14 etc.) up to the maximum.
- The maximum number of cores available depends on the security module/chassis model; see Requirements and Prerequisites for Container Instances, on page 198.

The chassis includes a default resource profile called "Default-Small," which includes the minimum number of cores. You can change the definition of this profile, and even delete it if it is not in use. Note that this profile is created when the chassis reloads and no other profile exists on the system.

You cannot change the resource profile settings if it is currently in use. You must disable any instances that use it, then change the resource profile, and finally reenable the instance. If you resize instances in an established High Availability pair or cluster, then you should make all members the same size as soon as possible.

If you change the resource profile settings after you add the FTD instance to the FMC, then update the inventory for each unit on the FMC **Devices** > **Device Management** > **Device** > **System** > **Inventory** dialog box.

Procedure

Step 1 Choose **Platform Settings** > **Resource Profiles** , and click **Add**.

The **Add Resource Profile** dialog box appears.

Step 2 Set the following paramters.

- **Name**—Sets the name of the profile between 1 and 64 characters. Note that you cannot change the name of this profile after you add it.
- **Description**—Sets the description of the profile up to 510 characters.
- **Number of Cores**—Sets the number of cores for the profile, between 6 and the maximum, depending on your chassis, as an even number.

Step 3 Click **OK**.

Configure a Network Control Policy

To permit the discovery of non-Cisco devices, FXOS supports the *Link Layer Discovery Protocol (LLDP)*, a vendor-neutral device discovery protocol that is defined in the IEEE 802.1ab standard. LLDP allows network devices to advertise information about themselves to other devices on the network. This protocol runs over the data-link layer, which allows two systems running different network layer protocols to learn about each other.

LLDP is a one-way protocol that transmits information about the capabilities and current status of a device and its interfaces. LLDP devices use the protocol to solicit information only from other LLDP devices.

To enable this functionality on your FXOS chassis, you can configure a network control policy, which specifies LLDP transmission and receiving behavior. Once a network control policy is created, it needs to be assigned to an interface. You can enable LLDP on any front interface, including fixed ports, EPM ports, port channels, and break out ports.

Note
- LLDP is not configurable on dedicated management ports.

- Internal backplane ports that connect to the blade have LLDP enabled by default, with no option to disable. All other ports have LLDP disabled by default.

Procedure

Step 1 Choose **Platform Settings** > **Network Control Policy**.

Step 2 Click **Add**.

Step 3 In the Network Control Policy dialog box, edit the following fields:

Name	Description
Name field	A unique name for the Network Control Policy.
LLDP recieve checkbox	Enables FXOS to receive LLDP packets.
LLDP transmit checkbox	Enables FXOS to transmit LLDP packets.
Description field	A description for the Network Control Policy.

Step 4 Click **Save**. After creating the Network Control Policy, you must assign it to an interface. For steps to edit and configure an interface with a Network Control Policy, see Configure a Physical Interface, on page 169.

Configure the Chassis URL

You can specify a management URL so that you can easily open Firepower Chassis Manager for an FTD instance directly from FMC. If you do not specify a chassis management URL, the chassis name is used instead.

If you change the chassis URL settings after you add the FTD instance to the FMC, then update the inventory for each unit on the **Devices** > **Device Management** > **Device** > **System** > **Inventory** dialog box.

Procedure

Step 1 Choose **Platform Settings** > **Chassis URL**.

Step 2 Set the following parameters.

- **Chassis Name**—Sets the name of the chassis between 1 and 60 characters.

- **Chassis URL**—Sets the URL that FMC should use to connect to an FTD instance within Firepower Chassis Manager. The URL must start with https://. If you do not specify a chassis management URL, the chassis name is used instead.

Step 3 Click **Update**.

CHAPTER **9**

Interface Management

About Interfaces

The Firepower 4100/9300 chassis supports physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces. EtherChannel interfaces can include up to 16 member interfaces of the same type.

Chassis Management Interface

The chassis management interface is used for management of the FXOS Chassis by SSH or Firepower Chassis Manager. This interface appears at the top of the **Interfaces** tab as **MGMT**, and you can only enable or disable this interface on the **Interfaces** tab. This interface is separate from the mgmt-type interface that you assign to the logical devices for application management.

To configure parameters for this interface, you must configure them from the CLI. See also Changing the Management IP Address, on page 82. To view information about this interface in the FXOS CLI, connect to local management and show the management port:

Firepower # **connect local-mgmt**

Firepower(local-mgmt) # **show mgmt-port**

Note that the chassis management interface remains up even if the physical cable or SFP module are unplugged, or if the **mgmt-port shut** command is performed.

> **Note** The chassis management interface does not support jumbo frames.

Interface Types

Physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces can be one of the following types:

- Data—Use for regular data. Data interfaces cannot be shared between logical devices, and logical devices cannot communicate over the backplane to other logical devices. For traffic on Data interfaces, all traffic must exit the chassis on one interface and return on another interface to reach another logical device.

- Data-sharing—Use for regular data. Only supported with container instances, these data interfaces can be shared by one or more logical devices/container instances (FTD-using-FMC only). Each container instance can communicate over the backplane with all other instances that share this interface. Shared interfaces can affect the number of container instances you can deploy. Shared interfaces are not supported for bridge group member interfaces (in transparent mode or routed mode), inline sets, passive interfaces, clusters, or failover links.

- Mgmt—Use to manage application instances. These interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface. You can only assign one management interface per logical device. Depending on your application and manager, you can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management.

> **Note** Mgmt interface change will cause reboot of the logical device, for example one change mgmt from e1/1 to e1/2 will cause the logical device to reboot to apply the new management.

- Eventing—Use as a secondary management interface for FTD-using-FMC devices. To use this interface, you must configure its IP address and other parameters at the FTD CLI. For example, you can separate management traffic from events (such as web events). See the management center configuration guide for more information. Eventing interfaces can be shared by one or more logical devices to access external hosts; logical devices cannot communicate over this interface with other logical devices that share the interface. If you later configure a data interface for management, you cannot use a separate eventing interface.

> **Note** A virtual Ethernet interface is allocated when each application instance is installed. If the application does not use an eventing interface, then the virtual interface will be in an admin down state.
>
> ```
> Firepower # show interface Vethernet775
> Firepower # Vethernet775 is down (Administratively down)
> Bound Interface is Ethernet1/10
> Port description is server 1/1, VNIC ext-mgmt-nic5
> ```

- Cluster—Use as the cluster control link for a clustered logical device. By default, the cluster control link is automatically created on Port-channel 48. The Cluster type is only supported on EtherChannel interfaces. For multi-instance clustering, you cannot share a Cluster-type interface across devices. You can add VLAN subinterfaces to the Cluster EtherChannel to provide separate cluster control links per cluster. If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster. The FDM and CDO does not support clustering.

| | Note | This chapter discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the FTD application. See FXOS Interfaces vs. Application Interfaces, on page 153 for more information. |

See the following table for interface type support for the FTD and ASA applications in standalone and cluster deployments.

Table 11: Interface Type Support

Application		Data	Data: Subinterface	Data-Sharing	Data-Sharing: Subinterface	Mgmt	Eventing	Cluster (EtherChannel only)	Cluster: Subinterface
FTD	Standalone Native Instance	Yes	—	—	—	Yes	Yes	—	—
	Standalone Container Instance	Yes	Yes	Yes	Yes	Yes	Yes	—	—
	Cluster Native Instance	Yes (EtherChannel only for inter-chassis cluster)	—	—	—	Yes	Yes	Yes	—
	Cluster Container Instance	Yes (EtherChannel only for inter-chassis cluster)	—	—	—	Yes	Yes	Yes	Yes
ASA	Standalone Native Instance	Yes	—	—	—	Yes	—	Yes	—
	Cluster Native Instance	Yes (EtherChannel only for inter-chassis cluster)	—	—	—	Yes	—	Yes	—

FXOS Interfaces vs. Application Interfaces

The Firepower 4100/9300 manages the basic Ethernet settings of physical interfaces, VLAN subinterfaces for container instances, and EtherChannel (port-channel) interfaces. Within the application, you configure

higher level settings. For example, you can only create EtherChannels in FXOS; but you can assign an IP address to the EtherChannel within the application.

The following sections describe the interaction between FXOS and the application for interfaces.

VLAN Subinterfaces

For all logical devices, you can create VLAN subinterfaces within the application.

For container instances in standalone mode only, you can *also* create VLAN subinterfaces in FXOS. Multi-instance clusters do not support subinterfaces in FXOS except on the Cluster-type interface. Application-defined subinterfaces are not subject to the FXOS limit. Choosing in which operating system to create subinterfaces depends on your network deployment and personal preference. For example, to share a subinterface, you must create the subinterface in FXOS. Another scenario that favors FXOS subinterfaces comprises allocating separate subinterface groups on a single interface to multiple instances. For example, you want to use Port-channel1 with VLAN 2–11 on instance A, VLAN 12–21 on instance B, and VLAN 22–31 on instance C. If you create these subinterfaces within the application, then you would have to share the parent interface in FXOS, which may not be desirable. See the following illustration that shows the three ways you can accomplish this scenario:

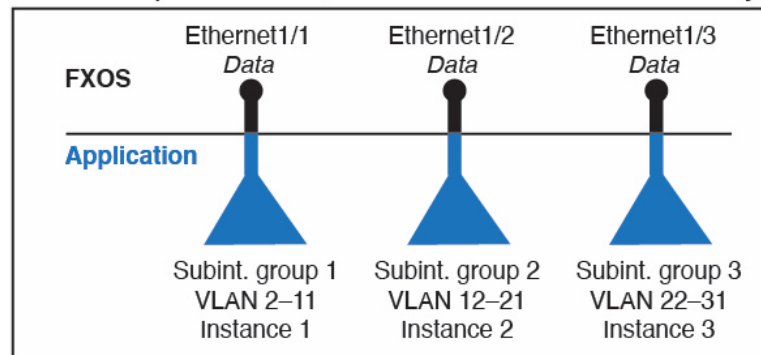*Figure 1: VLANs in FXOS vs. the Application for Container Instances*
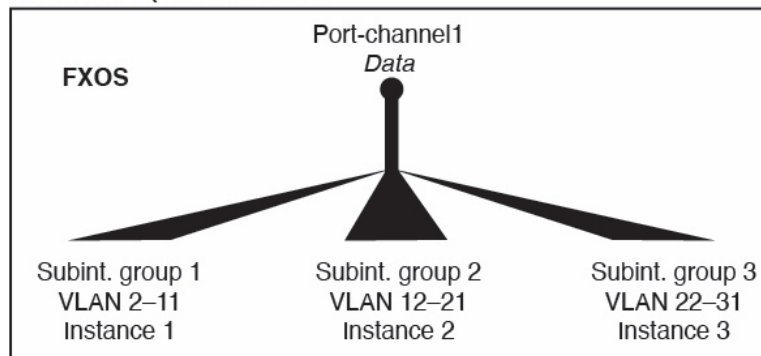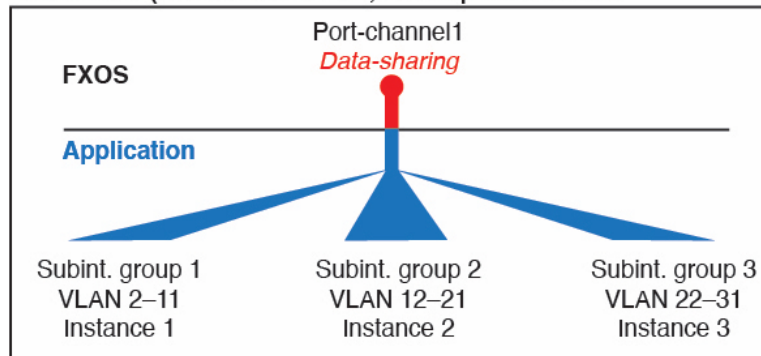
Independent Interface States in the Chassis and in the Application

You can administratively enable and disable interfaces in both the chassis and in the application. For an interface to be operational, the interface must be enabled in both operating systems. Because the interface state is controlled independently, you may have a mismatch between the chassis and application.

The default state of an interface within the application depends on the type of interface. For example, the physical interface or EtherChannel is disabled by default within the application, but a subinterface is enabled by default.

Hardware Bypass Pairs

For the FTD, certain interface modules on the Firepower 9300 and 4100 series let you enable the Hardware Bypass feature. Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures.

The Hardware Bypass feature is configured within the FTD application. You do not need to use these interfaces as Hardware Bypass pairs; they can be used as regular interfaces for both the ASA and the FTD applications. Note that Hardware Bypass-capable interfaces cannot be configured for breakout ports. If you want to use the Hardware Bypass feature, do not configure the ports as EtherChannels; otherwise, you can include these interfaces as EtherChannel members in regular interface mode.

When Hardware Bypass is enabled on an inline pair, switch bypass is attempted first. If the bypass configuration fails due a switch error, physical bypass is enabled.

> ✎
>
> **Note** Hardware Bypass (FTW) is not supported on FTD installed in service-chaining with third-party applications, such as VDP/Radware.

The FTD supports Hardware Bypass for interface pairs on specific network modules on the following models:

- Firepower 9300

- Firepower 4100 series

The supported Hardware Bypass network modules for these models include:

- Firepower 6-port 1G SX FTW Network Module single-wide (FPR-NM-6X1SX-F)

- Firepower 6-port 10G SR FTW Network Module single-wide (FPR-NM-6X10SR-F)

- Firepower 6-port 10G LR FTW Network Module single-wide (FPR-NM-6X10LR-F)

- Firepower 2-port 40G SR FTW Network Module single-wide (FPR-NM-2X40G-F)

- Firepower 8-port 1G Copper FTW Network Module single-wide (FPR-NM-8X1G-F)

Hardware Bypass can only use the following port pairs:

- 1 & 2

- 3 & 4

- 5 & 6

- 7 & 8

Jumbo Frame Support

The Firepower 4100/9300 chassis has support for jumbo frames enabled by default. To enable jumbo frame support on a specific logical device installed on the Firepower 4100/9300 chassis, you will need to configure the appropriate MTU settings for the interfaces on the logical device.

The maximum MTU that is supported for the application on the Firepower 4100/9300 chassis is 9184.

> ✏️
>
> **Note**   The chassis management interface does not support jumbo frames.

Shared Interface Scalability

Instances can share data-sharing type interfaces. This capability lets you conserve physical interface usage as well as support flexible networking deployments. When you share an interface, the chassis uses unique MAC addresses to forward traffic to the correct instance. However, shared interfaces can cause the forwarding table to grow large due to the need for a full mesh topology within the chassis (every instance must be able to communicate with every other instance that is sharing the same interface). Therefore, there are limits to how many interfaces you can share.

In addition to the forwarding table, the chassis maintains a VLAN group table for VLAN subinterface forwarding. You can create up to 500 VLAN subinterfaces.

See the following limits for shared interface allocation:
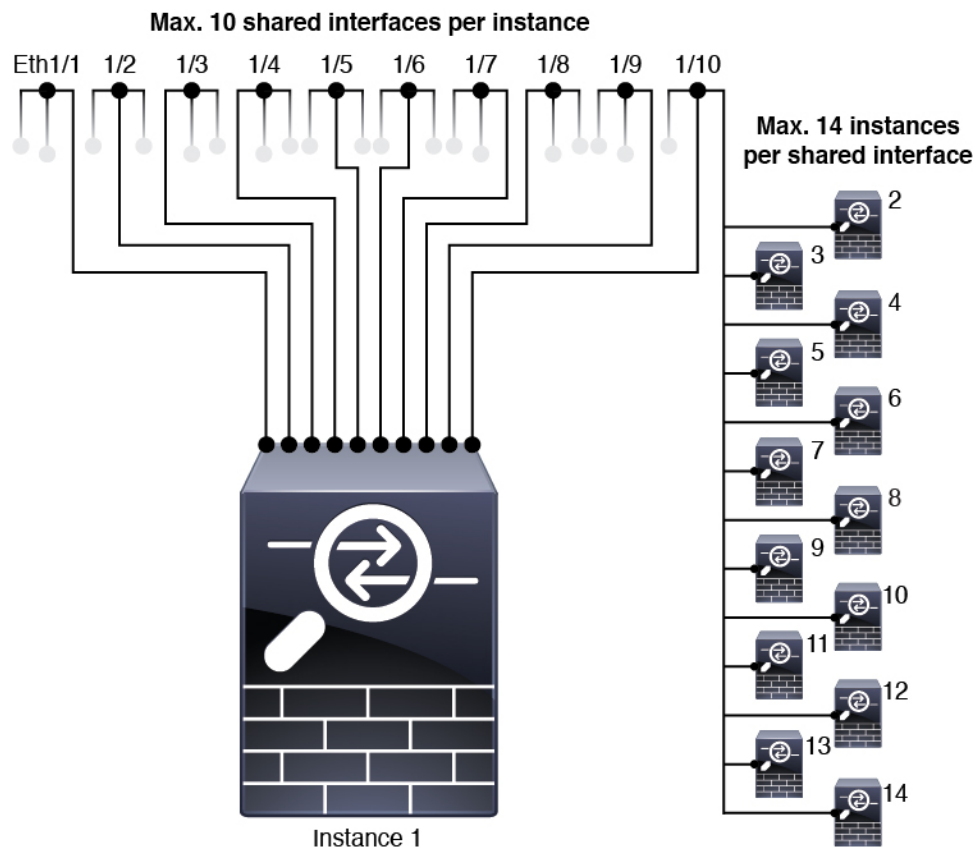

Shared Interface Best Practices

For optimal scalability of the forwarding table, share as few interfaces as possible. Instead, you can create up to 500 VLAN subinterfaces on one or more physical interfaces and then divide the VLANs among the container instances.

When sharing interfaces, follow these practices in the order of most scalable to least scalable:

1. Best—Share subinterfaces under a single parent, and use the same set of subinterfaces with the same group of instances.

   For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel: Port-Channel1.2, 3, and 4 instead of Port-Channel2, Port-Channel3, and Port-Channel4. When you share subinterfaces from a single parent, the VLAN group table provides better scaling of the forwarding table than when sharing physical/EtherChannel interfaces or subinterfaces across parents.

   Figure 2: Best: Shared Subinterface Group on One Parent
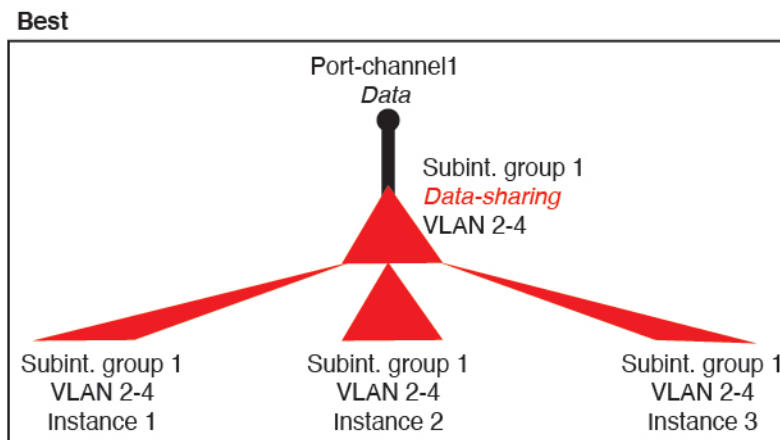
   If you do not share the same set of subinterfaces with a group of instances, your configuration can cause more resource usage (more VLAN groups). For example, share Port-Channel1.2, 3, and 4 with instances 1, 2, and 3 (one VLAN group) instead of sharing Port-Channel1.2 and 3 with instances 1 and 2, while sharing Port-Channel1.3 and 4 with instance 3 (two VLAN groups).

   Figure 3: Good: Sharing Multiple Subinterface Groups on One Parent
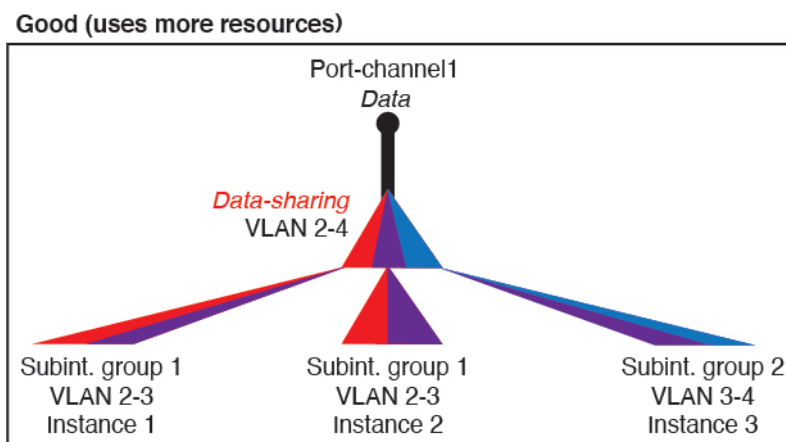
2. Fair—Share subinterfaces across parents.

   For example, share Port-Channel1.2, Port-Channel2.3, and Port-Channel3.4 instead of Port-Channel2, Port-Channel4, and Port-Channel4. Although this usage is not as efficient as only sharing subinterfaces on the same parent, it still takes advantage of VLAN groups.

Figure 4: Fair: Shared Subinterfaces on Separate Parents
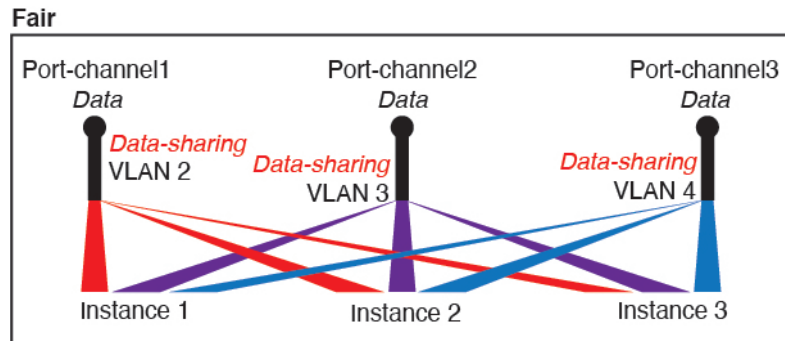
3. Worst—Share individual parent interfaces (physical or EtherChannel).

This method uses the most forwarding table entries.

Figure 5: Worst: Shared Parent Interfaces
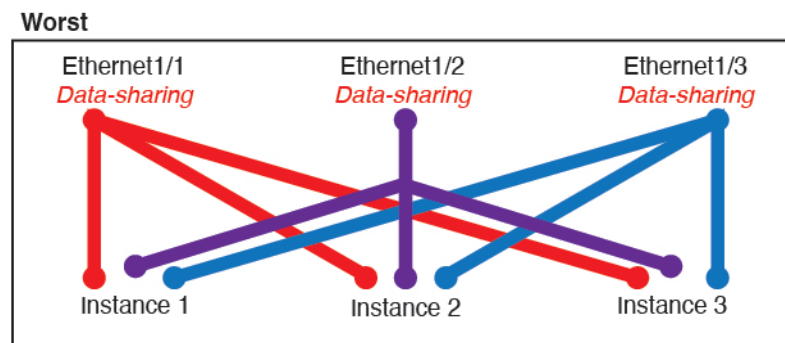
Shared Interface Usage Examples

See the following tables for examples of interface sharing and scalability. The below scenarios assume use of one physical/EtherChannel interface for management shared across all instances, and another physical or EtherChannel interface with dedicated subinterfaces for use with High Availability.

Firepower 9300 with Three SM-44s

The following table applies to three SM-44 security modules on a 9300 using only physical interfaces or EtherChannels. Without subinterfaces, the maximum number of interfaces are limited. Moreover, sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

Each SM-44 module can support up to 14 instances. Instances are split between modules as necessary to stay within limits.

Table 12: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with Three SM-44s

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 8 • 8 • 8 • 8	**0**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	16%
30: • 15 • 15	**0**	**2:** • Instance 1 • Instance 2	14%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
33: • 11 (1 ea.) • 11 (1 ea.) • 11 (1 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
33: • *11 (1 ea.)* • *11 (1 ea.)* • *12 (1 ea.)*	*3:* • *1* • *1* • *1*	*34:* • *Instance 1-Instance 11* • *Instance 12-Instance 22* • *Instance 23-Instance 34*	*102%* DISALLOWED
30: • 30 (1 ea.)	**1**	**6:** • Instance 1-Instance 6	25%
30: • 10 (5 ea.) • 10 (5 ea.) • 10 (5 ea.)	**3:** • 1 • 1 • 1	**6:** • Instance 1-Instance2 • Instance 2-Instance 4 • Instance 5-Instance 6	23%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
30: • 30 (6 ea.)	**2**	**5:** • Instance 1-Instance 5	28%
30: • 12 (6 ea.) • 18 (6 ea.)	**4:** • 2 • 2	**5:** • Instance 1-Instance2 • Instance 2-Instance 5	26%
24: • 6 • 6 • 6 • 6	**7**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	44%
24: • 12 (6 ea.) • 12 (6 ea.)	**14:** • 7 • 7	**4:** • Instance 1-Instance2 • Instance 2-Instance 4	41%

The following table applies to three SM-44 security modules on a 9300 using subinterfaces on a single parent physical interface. For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel. Sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

Each SM-44 module can support up to 14 instances. Instances are split between modules as necessary to stay within limits.

Table 13: Subinterfaces on One Parent and Instances on a Firepower 9300 with Three SM-44s

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
168: • 168 (4 ea.)	**0**	**42:** • Instance 1-Instance 42	33%
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	27%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
33: • 11 (1 ea.) • 11 (1 ea.) • 11 (1 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
165: • 55 (5 ea.) • 55 (5 ea.) • 55 (5 ea.)	**3:** • 1 • 1 • 1	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**2**	**14:** • Instance 1-Instance 14	46%
165: • 55 (5 ea.) • 55 (5 ea.) • 55 (5 ea.)	**6:** • 2 • 2 • 2	**33:** • Instance 1-Instance 11 • Instance 12-Instance 22 • Instance 23-Instance 33	98%
70: • 70 (5 ea.)	**10**	**14:** • Instance 1-Instance 14	46%
165: • *55 (5 ea.)* • *55 (5 ea.)* • *55 (5 ea.)*	*30:* • *10* • *10* • *10*	*33:* • *Instance 1-Instance 11* • *Instance 12-Instance 22* • *Instance 23-Instance 33*	*102%* DISALLOWED

Firepower 9300 with One SM-44

The following table applies to the Firepower 9300 with one SM-44 using only physical interfaces or EtherChannels. Without subinterfaces, the maximum number of interfaces are limited. Moreover, sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

The Firepower 9300 with one SM-44 can support up to 14 instances.

Table 14: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with One SM-44

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 8 • 8 • 8 • 8	**0**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	16%
30: • 15 • 15	**0**	**2:** • Instance 1 • Instance 2	14%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
14: • 7 (1 ea.) • 7 (1 ea.)	**2:** • **1** • **1**	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
32: • 8 • 8 • 8 • 8	**1**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	21%
32: • 16 (8 ea.) • 16 (8 ea.)	**2**	**4:** • Instance 1-Instance 2 • Instance 3-Instance 4	20%
32: • 8 • 8 • 8 • 8	**2**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	25%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 16 (8 ea.) • 16 (8 ea.)	**4:** • 2 • 2	**4:** • Instance 1-Instance 2 • Instance 3-Instance 4	24%
24: • 8 • 8 • 8	**8**	**3:** • Instance 1 • Instance 2 • Instance 3	37%
10: • 10 (2 ea.)	**10**	**5:** • Instance 1-Instance 5	69%
10: • 6 (2 ea.) • 4 (2 ea.)	**20:** • 10 • 10	**5:** • Instance 1-Instance 3 • Instance 4-Instance 5	59%
14: • *12 (2 ea.)*	*10*	*7:* • *Instance 1-Instance 7*	*109%* DISALLOWED

The following table applies to the Firepower 9300 with one SM-44 using subinterfaces on a single parent physical interface. For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel. Sharing multiple physical interfaces uses more forwarding table resources than sharing multiple subinterfaces.

The Firepower 9300 with one SM-44 can support up to 14 instances.

Table 15: Subinterfaces on One Parent and Instances on a Firepower 9300 with One SM-44

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
112: • 112 (8 ea.)	**0**	**14:** • Instance 1-Instance 14	17%
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	17%
14: • 14 (1 ea.)	**1**	**14:** • Instance 1-Instance 14	46%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
14: • 7 (1 ea.) • 7 (1 ea.)	**2:** • 1 • 1	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
112: • 112 (8 ea.)	**1**	**14:** • Instance 1-Instance 14	46%
112: • 56 (8 ea.) • 56 (8 ea.)	**2:** • 1 • 1	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
112: • 112 (8 ea.)	**2**	**14:** • Instance 1-Instance 14	46%
112: • 56 (8 ea.) • 56 (8 ea.)	**4:** • 2 • 2	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%
140: • 140 (10 ea.)	**10**	**14:** • Instance 1-Instance 14	46%
140: • 70 (10 ea.) • 70 (10 ea.)	**20:** • 10 • 10	**14:** • Instance 1-Instance 7 • Instance 8-Instance 14	37%

Viewing Shared Interface Resources

To view forwarding table and VLAN group usage, see the **Instances** > **Interface Forwarding Utilization** area. For example:
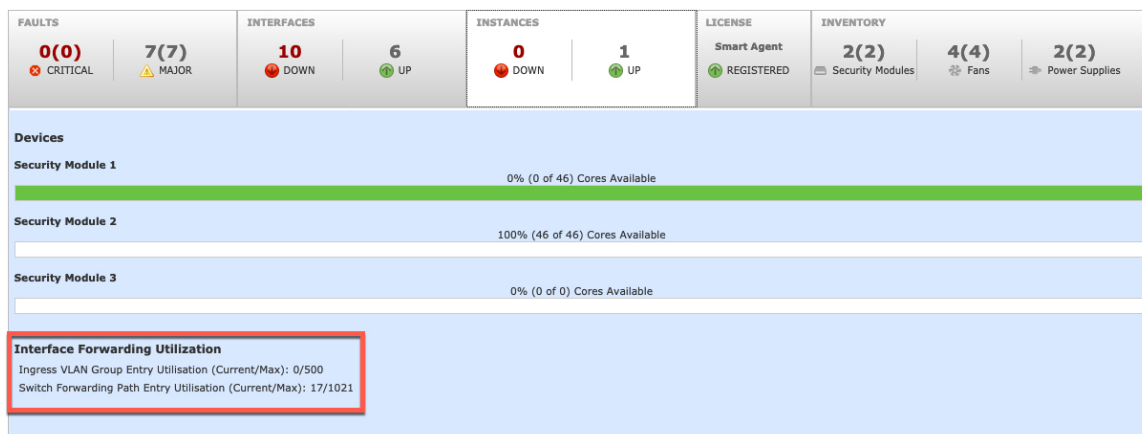
Inline Set Link State Propagation for the FTD

An inline set acts like a bump on the wire, and binds two interfaces together to slot into an existing network. This function allows the system to be installed in any network environment without the configuration of adjacent network devices. Inline interfaces receive all traffic unconditionally, but all traffic received on these interfaces is retransmitted out of an inline set unless explicitly dropped.

When you configure an inline set in the FTD application and enable link state propagation, the FTD sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down. When the downed interface comes back up, the second interface automatically comes back up, also. In other words, if the link state of one interface changes, the chassis senses the change and updates the link state of the other interface to match it. Note that the chassis requires up to 4 seconds to propagate link state changes. Link state propagation is especially useful in resilient network environments where routers are configured to reroute traffic automatically around network devices that are in a failure state.

Guidelines and Limitations for Interfaces

VLAN Subinterfaces

- This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the FTD application. See FXOS Interfaces vs. Application Interfaces, on page 153 for more information.

- Subinterfaces (and the parent interfaces) can only be assigned to container instances.

Note If you assign a parent interface to a container instance, it only passes untagged (non-VLAN) traffic. Do not assign the parent interface unless you intend to pass untagged traffic. For Cluster type interfaces, the parent interface cannot be used.

- Subinterfaces are supported on Data or Data-sharing type interfaces, as well as Cluster type interfaces. If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.

- For multi-instance clustering, FXOS subinterfaces are not supported on Data interfaces. However, subinterfaces are supported for the cluster control link, so you can use either a dedicated EtherChannel or a subinterface of an EtherChannel for the cluster control link. Note that *application*-defined subinterfaces are supported for Data interfaces.

- You can create up to 500 VLAN IDs.

- See the following limitations within the logical device application; keep these limitations in mind when planning your interface allocation.

  - You cannot use subinterfaces for an FTD inline set or as a passive interface.

  - If you use a subinterface for the failover link, then all subinterfaces on that parent, and the parent itself, are restricted for use as failover links. You cannot use some subinterfaces as failover links, and some as regular data interfaces.

Data-sharing Interfaces

- You cannot use a data-sharing interface with a native instance.

- Maximum 14 instances per shared interface. For example, you can allocate Ethernet1/1 to Instance1 through Instance14.

  Maximum 10 shared interfaces per instance. For example, you can allocate Ethernet1/1.1 through Ethernet1/1.10 to Instance1.
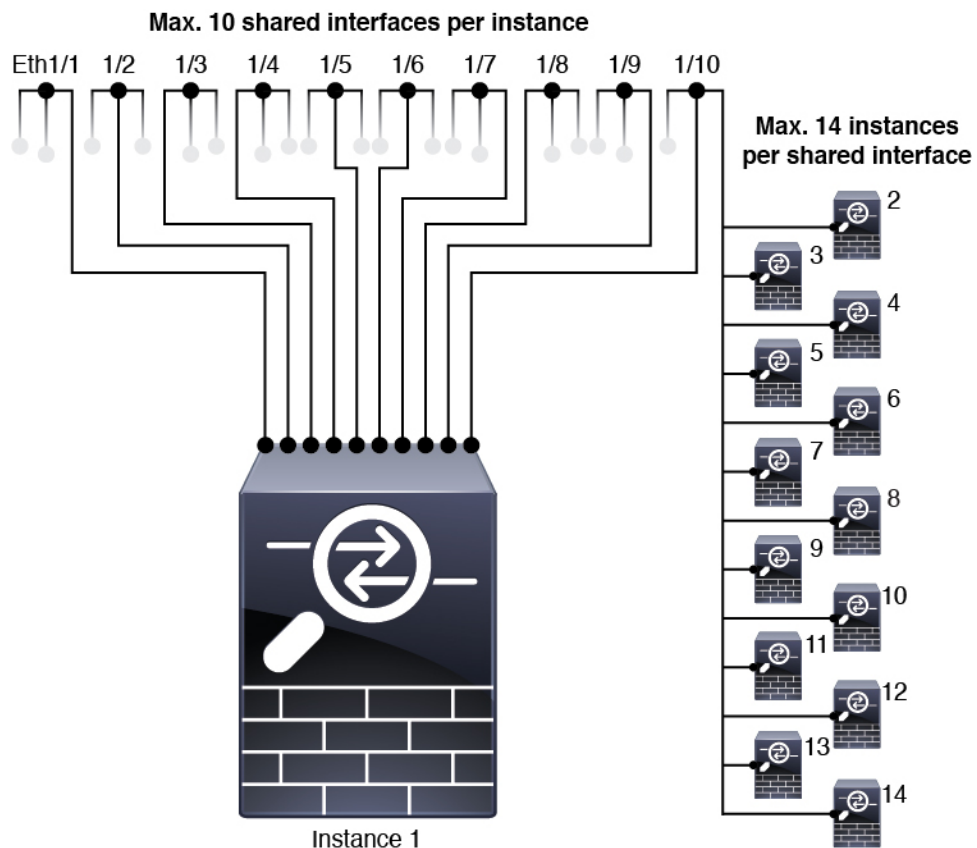
- You cannot use a data-sharing interface in a cluster.

- See the following limitations within the logical device application; keep these limitations in mind when planning your interface allocation.

    - You cannot use a data-sharing interface with a transparent firewall mode device.

    - You cannot use a data-sharing interface with FTD inline sets or passive interfaces.

    - You cannot use a data-sharing interface for the failover link.

Inline Sets for FTD

- Supported for physical interfaces (both regular and breakout ports) and EtherChannels. Subinterfaces are not supported.

- Link state propagation is supported.

Hardware Bypass

- Supported for the FTD; you can use them as regular interfaces for the ASA.

- The FTD only supports Hardware Bypass with inline sets.

- Hardware Bypass-capable interfaces cannot be configured for breakout ports.

- You cannot include Hardware Bypass interfaces in an EtherChannel and use them for Hardware Bypass; you can use them as regular interfaces in an EtherChannel.

- Hardware Bypass is not supported with High Availability.

Default MAC Addresses

For native instances:

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.

- EtherChannels—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.

For container instances:

- MAC addresses for all interfaces are taken from a MAC address pool. For subinterfaces, if you decide to manually configure MAC addresses, make sure you use unique MAC addresses for all subinterfaces on the same parent interface to ensure proper classification. See Automatic MAC Addresses for Container Instance Interfaces, on page 190.

Configure Interfaces

By default, physical interfaces are disabled. You can enable interfaces, add EtherChannels, add VLAN subinterfaces, edit interface properties, and configure breakout ports.

✎

Note

Enable or Disable an Interface

You can change the **Admin State** of each interface to be enabled or disabled. By default, physical interfaces are disabled. For VLAN subinterfaces, the admin state is inherited from the parent interface.

Procedure

Step 1   Choose **Interfaces** to open the Interfaces page.

The Interfaces page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2   To enable the interface, click the disabled **Slider disabled** ( ) so that it changes to the enabled **Slider enabled** ( ).

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from gray to green.

Step 3   To disable the interface, click the enabled **Slider enabled** ( ) so that it changes to the disabled **Slider disabled** ( ).

Click **Yes** to confirm the change. The corresponding interface in the visual representation changes from green to gray.

Configure a Physical Interface

You can physically enable and disable interfaces, as well as set the interface speed and duplex. To use an interface, it must be physically enabled in FXOS and logically enabled in the application.

✎

Note   For QSFPH40G-CUxM, auto-negotiation is always enabled by default and you cannot disable it.

Before you begin

- Interfaces that are already a member of an EtherChannel cannot be modified individually. Be sure to configure settings before you add it to the EtherChannel.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Edit** in the row for the interface you want to edit to open the **Edit Interface** dialog box.

Step 3 To enable the interface, check the **Enable** check box. To disable the interface, uncheck the **Enable** check box.

Step 4 Choose the interface **Type**:

- **Data**

- **Data-sharing**—For container instances only.

- **Mgmt**

- **Firepower-eventing**—For FTD only.

- **Cluster**—Do not choose the **Cluster** type; by default, the cluster control link is automatically created on Port-channel 48.

Step 5 (Optional) Choose the speed of the interface from the **Speed** drop-down list.

Step 6 (Optional) If your interface supports **Auto Negotiation**, click the **Yes** or **No** radio button.

Step 7 (Optional) Choose the duplex of the interface from the **Duplex** drop-down list.

Step 8 (Optional) Choose a previously-configured **Network Control Policy**.

Step 9 (Optional) Explicitly configure **Debounce Time (ms)**. Enter a value between 0-15000 milli-seconds.

Step 10 Click **OK**.

Add an EtherChannel (Port Channel)

An EtherChannel (also known as a port channel) can include up to 16 member interfaces of the same media type and capacity, and must be set to the same speed and duplex. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface. The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical Data or Data-sharing interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.

> • On—The EtherChannel is always on, and LACP is not used. An "on" EtherChannel can only establish a connection with another "on" EtherChannel.

> ✎
>
> **Note** It may take up to three minutes for an EtherChannel to come up to an operational state if you change its mode from On to Active or from Active to On.

Non-data interfaces only support active mode.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. "On" mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

When the Firepower 4100/9300 chassis creates an EtherChannel, the EtherChannel stays in a **Suspended** state for Active LACP mode or a **Down** state for On LACP mode until you assign it to a logical device, even if the physical link is up. The EtherChannel will be brought out of this **Suspended** state in the following situations:

> • The EtherChannel is added as a data or management interface for a standalone logical device

> • The EtherChannel is added as a management interface or cluster control link for a logical device that is part of a cluster

> • The EtherChannel is added as a data interface for a logical device that is part of a cluster and at least one unit has joined the cluster

Note that the EtherChannel does not come up until you assign it to a logical device. If the EtherChannel is removed from the logical device or the logical device is deleted, the EtherChannel will revert to a **Suspended** or **Down** state.

Procedure

Step 1 Choose **Interfaces** to open the Interfaces page.

The **All Interfaces** page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2 Click **Add Port Channel** above the interfaces table to open the **Add Port Channel** dialog box.

Step 3 Enter an ID for the port channel in the **Port Channel ID** field. Valid values are between 1 and 47.

Port-channel 48 is reserved for the cluster control link when you deploy a clustered logical device. If you do not want to use Port-channel 48 for the cluster control link, you can delete it and configure a Cluster type EtherChannel with a different ID. You can add multiple Cluster type EtherChannels and add VLAN subinterfaces for use with multi-instance clustering. For intra-chassis clustering, do not assign any interfaces to the Cluster EtherChannel.

Step 4 To enable the port channel, check the **Enable** check box. To disable the port channel, uncheck the **Enable** check box.

Step 5 Choose the interface **Type**:

> • **Data**

- **Data-sharing**—For container instances only.

- **Mgmt**

- **Firepower-eventing**—For FTD only.

- **Cluster**

Step 6   Set the required **Admin Speed** for the member interfaces from the drop-down list.

If you add a member interface that is not at the specified speed, it will not successfully join the port channel.

Step 7   For Data or Data-sharing interfaces, choose the LACP port-channel **Mode**, **Active** or **On**.

For non-Data or non-Data-sharing interfaces, the mode is always active.

Step 8   Set the required **Admin Duplex** for the member interfaces, **Full Duplex** or **Half Duplex**.

If you add a member interface that is configured with the specified duplex, it will not successfully join the port channel.

Step 9   To add an interface to the port channel, select the interface in the **Available Interface** list and click **Add Interface** to move the interface to the Member ID list.

You can add up to 16 member interfaces of the same media type and capacity. The member interfaces must be set to the same speed and duplex, and must match the speed and duplex that you configured for this port channel. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface.

Tip   You can add multiple interfaces at one time. To select multiple individual interfaces, click on the desired interfaces while holding down the **Ctrl** key. To select a range of interfaces, select the first interface in the range, and then, while holding down the **Shift** key, click to select the last interface in the range.

Step 10   To remove an interface from the port channel, click the **Delete** button to the right of the interface in the Member ID list.

Step 11   Click **OK**.

Add a VLAN Subinterface for Container Instances

You can add up to 500 subinterfaces to your chassis.

For multi-instance clustering, you can only add subinterfaces to the Cluster-type interface; subinterfaces on data interfaces are not supported.

VLAN IDs per interface must be unique, and within a container instance, VLAN IDs must be unique across all assigned interfaces. You can reuse VLAN IDs on *separate* interfaces as long as they are assigned to different container instances. However, each subinterface still counts towards the limit even though it uses the same ID.

This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the FTD application.

Procedure

Step 1   Choose **Interfaces** to open the **All Interfaces** tab.

The **All Interfaces** tab shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

Step 2   Click **Add New** > **Subinterface** to open the **Add Subinterface** dialog box.

Step 3   Choose the interface **Type**:

   • **Data**

   • **Data-sharing**

   • **Cluster**—If you add subinterfaces to a Cluster interface, you cannot use that interface for a native cluster.

For Data and Data-sharing interfaces: The type is independent of the parent interface type; you can have a Data-sharing parent and a Data subinterface, for example.

Step 4   Choose the parent **Interface** from the drop-down list.

You cannot add a subinterface to a physical interface that is currently allocated to a logical device. If other subinterfaces of the parent are allocated, you can add a new subinterface as long as the parent interface itself is not allocated.

Step 5   Enter a **Subinterface ID**, between 1 and 4294967295.

This ID will be appended to the parent interface ID as *interface_id*.*subinterface_id*. For example, if you add a subinterface to Ethernet1/1 with the ID of 100, then the subinterface ID will be: Ethernet1/1.100. This ID is not the same as the VLAN ID, although you can set them to match for convenience.

Step 6   Set the **VLAN ID** between 1 and 4095.

Step 7   Click **OK**.

Expand the parent interface to view all subinterfaces under it.

Configure Breakout Cables

The following procedure shows how to configure breakout cables for use with the Firepower 4100/9300 chassis. You can use a breakout cable to provide four 10 Gbps ports in place of a single 40 Gbps port.

Before you begin

Hardware Bypass-capable interfaces cannot be configured for breakout ports.

Procedure

Step 1   Choose **Interfaces** to open the Interfaces page.

The Interfaces page shows a visual representation of the currently installed interfaces at the top of the page and provides a listing of the installed interfaces in the table below.

The interfaces that are capable of supporting breakout cables but are not currently configured as such are indicated by a Breakout Port icon in the row for that interface. For interfaces that have already been configured as using a breakout cable, the individual breakout interfaces are listed separately (for example, Ethernet 2/1/1, 2/1/2, 2/1/3, and 2/1/4).

Step 2 To convert a 40 Gbps interface into four 10 Gbps interfaces:

a) Click the **Breakout Port** icon for the interface that you want to convert.

The Breakout Port Creation dialog box opens asking you to confirm that you want to proceed and warning you that the chassis will be rebooted.

b) Click **Yes** to confirm.

The chassis reboots and the specified interface is converted into four 10 Gbps interfaces.

Step 3 To convert the four 10 Gbps breakout interfaces back into a single 40 Gbps interface:

a) Click **Delete** for any of the breakout interfaces.

A confirmation dialog box opens asking you to confirm that you want to proceed and warning you that all four breakout interfaces will be deleted and that the chassis will be rebooted.

b) Click **Yes** to confirm.

The chassis reboots and the specified interfaces are converted into a single 40 Gbps interface.

Monitoring Interfaces

From the Interfaces page of the Firepower Chassis Manager, you can view the status of the installed interfaces on the chassis, edit interface properties, enable or disable an interface, and create port channels.

The Interfaces page is made up of two sections:

* The upper section shows a visual representation of the interfaces that are installed in the chassis. You can hover over any of the interfaces to get additional information about the interface.

  The interfaces are color coded to indicate their current status:

    * Green—The interface is installed and enabled.

    * Dark Grey—The interface is installed but disabled.

    * Red—There is a problem with the operational state of the interface.

    * Light Grey—The interface is not installed.

Note Interfaces that act as ports in port channels do not appear in this list.

* The lower section contains two tabs: **All Interfaces** and **Hardware Bypass**. On the **All Interfaces** tab: For each interface, you can enable or disable the interface. You can also click **Edit** to edit the properties of an interface, such as speed and interface type. For **Hardware Bypass**, see Hardware Bypass Pairs, on page 156.

Note   The port-channel 48 cluster type interface shows the **Operation State** as **failed** if it does not include any member interfaces. For intra-chassis clustering, this EtherChannel does not require any member interfaces, and you can ignore this Operational State.

Troubleshooting Interfaces

Error: The Switch Forwarding Path has 1076 entries and exceeds the limit of 1024. If you are adding an interface, reduce the number of shared interfaces assigned to logical devices, reduce the number of logical devices sharing interfaces, or use non-shared subinterfaces instead. If you are deleting a subinterface, you are seeing this message because the remaining configuration is no longer optimized to fit within the Switch Forwarding Path table. See the FXOS configuration guide for troubleshooting information about the deletion use case. Use 'show detail' under scope 'fabric-interconnect' to view the current Switch Forwarding Path Entry Count.

If you see this error when trying to delete a shared subinterface from a logical device, it is because your new configuration is not following this guideline for shared subinterfaces: use the same set of subinterfaces with the same group of logical devices. If you delete a shared subinterface from one logical device, you can end up with more VLAN groups and therefore less efficient usage of the forwarding table. To work around this situation, you need to add and delete shared subinterfaces simultaneously using the CLI so that you maintain the same set of subinterfaces for the same group of logical devices.

See the following scenarios for more information. These scenarios start with the following interfaces and logical devices:

- Shared subinterface set on the same parent: Port-Channel1.100 (VLAN 100), Port-Channel1.200 (VLAN 200), Port-Channel1.300 (VLAN 300)

- Logical device group: LD1, LD2, LD3, and LD4

Scenario 1: Remove a subinterface from one logical device, but leave it assigned to other logical devices

Do not remove the subinterface. Instead, just disable it in the application configuration. If you have to remove the subinterface, you will need to reduce the number of shared interfaces in general to continue to fit in the forwarding table.

Scenario 2: Remove all subinterfaces in the set from one logical device

Remove all subinterfaces in the set from the logical device at the CLI, and then save the configuration so that the removal is simultaneous.

1. View the VLAN groups for reference. In the following output, group 1 includes VLAN 100, 200, and 300, representing the 3 shared subinterfaces.

```
firepower# connect fxos
[...]
firepower(fxos)# show ingress-vlan-groups
ID   Class ID  Status          INTF        Vlan Status
1    1         configured
                                           100  present
                                           200  present
```

```
                                                     300   present
   2048 512        configured
                                                     0     present
   2049 511        configured
                                                     0     present
   firepower(fxos)# exit
   firepower#
```

2. View the shared subinterfaces assigned to the logical device you want to change.

```
   firepower# scope ssa
   firepower /ssa # scope logical-device LD1
   firepower /ssa/logical-device # show external-port-link

   External-Port Link:
      Name                          Port or Port Channel Name Port Type          App Name
      Description
      ---------------------------- ------------------------ ------------------ ----------
      -----------
      Ethernet14_ftd                Ethernet1/4                Mgmt               ftd
      PC1.100_ftd                   Port-channel1.100          Data Sharing       ftd
      PC1.200_ftd                   Port-channel1.200          Data Sharing       ftd
      PC1.300_ftd                   Port-channel1.300          Data Sharing       ftd
```

3. Remove the subinterfaces from the logical device, and then save the configuration.

```
   firepower /ssa/logical-device # delete external-port-link PC1.100_ftd
   firepower /ssa/logical-device* # delete external-port-link PC1.200_ftd
   firepower /ssa/logical-device* # delete external-port-link PC1.300_ftd
   firepower /ssa/logical-device* # commit-buffer
   firepower /ssa/logical-device #
```

If you had committed the configuration in the middle, you would have ended up with 2 VLAN groups, which could have generated the switch forwarding path error and prevented you from saving the configuration.

Scenario 3: Remove a subinterface from all logical devices in the group

Remove the subinterface from all logical devices in the group at the CLI, and then save the configuration so that the removal is simultaneous. For example:

1. View the VLAN groups for reference. In the following output, group 1 includes VLAN 100, 200, and 300, representing the 3 shared subinterfaces.

```
   firepower# connect fxos
   [...]
   firepower(fxos)# show ingress-vlan-groups
   ID   Class ID Status          INTF        Vlan Status
   1    1        configured
                                             100  present
                                             200  present
                                             300  present
   2048 512       configured
                                             0    present
   2049 511       configured
                                             0    present
```

2. View the interfaces assigned to each logical device, and note the shared subinterfaces in common. If they are on the same parent interface, they will belong to one VLAN group, and should match the **show ingress-vlan-groups** list. In Firepower Chassis Manager, you can hover over each shared subinterface to see which instances it is allocated to.

Figure 6: Instances per shared interface

At the CLI, you can view characteristics of all logical devices, including the allocated interfaces.

```
firepower# scope ssa
firepower /ssa # show logical-device expand

Logical Device:
    Name: LD1
    Description:
    Slot ID: 1
    Mode: Standalone
    Oper State: Ok
    Template Name: ftd

    External-Port Link:
        Name: Ethernet14_ftd
        Port or Port Channel Name: Ethernet1/4
        Port Type: Mgmt
        App Name: ftd
        Description:

        Name: PC1.100_ftd
        Port or Port Channel Name: Port-channel1.100
        Port Type: Data Sharing
        App Name: ftd
        Description:

        Name: PC1.200_ftd
        Port or Port Channel Name: Port-channel1.200
        Port Type: Data Sharing
        App Name: ftd
        Description:

        System MAC address:
            Mac Address
            -----------
            A2:F0:B0:00:00:25

        Name: PC1.300_ftd
        Port or Port Channel Name: Port-channel1.300
        Port Type: Data Sharing
```

```
            App Name: ftd
            Description:

[...]

        Name: LD2
        Description:
        Slot ID: 1
        Mode: Standalone
        Oper State: Ok
        Template Name: ftd

        External-Port Link:
            Name: Ethernet14_ftd
            Port or Port Channel Name: Ethernet1/4
            Port Type: Mgmt
            App Name: ftd
            Description:

            Name: PC1.100_ftd
            Port or Port Channel Name: Port-channel1.100
            Port Type: Data Sharing
            App Name: ftd
            Description:

            Name: PC1.200_ftd
            Port or Port Channel Name: Port-channel1.200
            Port Type: Data Sharing
            App Name: ftd
            Description:

            System MAC address:
                Mac Address
                -----------
                A2:F0:B0:00:00:28

            Name: PC1.300_ftd
            Port or Port Channel Name: Port-channel1.300
            Port Type: Data Sharing
            App Name: ftd
            Description:

[...]

        Name: LD3
        Description:
        Slot ID: 1
        Mode: Standalone
        Oper State: Ok
        Template Name: ftd

        External-Port Link:
            Name: Ethernet14_ftd
            Port or Port Channel Name: Ethernet1/4
            Port Type: Mgmt
            App Name: ftd
            Description:

            Name: PC1.100_ftd
            Port or Port Channel Name: Port-channel1.100
            Port Type: Data Sharing
            App Name: ftd
            Description:
```

```
            Name: PC1.200_ftd
            Port or Port Channel Name: Port-channel1.200
            Port Type: Data Sharing
            App Name: ftd
            Description:

            System MAC address:
                Mac Address
                -----------
                A2:F0:B0:00:00:2B

            Name: PC1.300_ftd
            Port or Port Channel Name: Port-channel1.300
            Port Type: Data Sharing
            App Name: ftd
            Description:

    [...]

        Name: LD4
        Description:
        Slot ID: 1
        Mode: Standalone
        Oper State: Ok
        Template Name: ftd

        External-Port Link:
            Name: Ethernet14_ftd
            Port or Port Channel Name: Ethernet1/4
            Port Type: Mgmt
            App Name: ftd
            Description:

            Name: PC1.100_ftd
            Port or Port Channel Name: Port-channel1.100
            Port Type: Data Sharing
            App Name: ftd
            Description:

            Name: PC1.200_ftd
            Port or Port Channel Name: Port-channel1.200
            Port Type: Data Sharing
            App Name: ftd
            Description:

            System MAC address:
                Mac Address
                -----------
                A2:F0:B0:00:00:2E

            Name: PC1.300_ftd
            Port or Port Channel Name: Port-channel1.300
            Port Type: Data Sharing
            App Name: ftd
            Description:

    [...]
```

3. Remove the subinterface from each logical device, and then save the configuration.

```
firepower /ssa # scope logical device LD1
firepower /ssa/logical-device # delete external-port-link PC1.300_ftd
firepower /ssa/logical-device* # exit
```

```
firepower /ssa* # scope logical-device LD2
firepower /ssa/logical-device* # delete external-port-link PC1.300_ftd
firepower /ssa/logical-device* # exit
firepower /ssa* # scope logical-device LD3
firepower /ssa/logical-device* # delete external-port-link PC1.300_ftd
firepower /ssa/logical-device* # exit
firepower /ssa* # scope logical-device LD4
firepower /ssa/logical-device* # delete external-port-link PC1.300_ftd
firepower /ssa/logical-device* # commit-buffer
firepower /ssa/logical-device #
```

If you had committed the configuration in the middle, you would have ended up with 2 VLAN groups, which could have generated the switch forwarding path error and prevented you from saving the configuration.

Scenario 4: Add a subinterface to one or more logical devices

Add the subinterface to *all* logical devices in the group at the CLI, and then save the configuration so that the addition is simultaneous.

1. Add the subinterface to each logical device, and then save the configuration.

```
firepower# scope ssa
firepower /ssa # scope logical-device LD1
firepower /ssa/logical-device # create external-port-link PC1.400_ftd Port-channel1.400
 ftd
firepower /ssa/logical-device/external-port-link* # exit
firepower /ssa/logical-device* # exit
firepower /ssa # scope logical-device LD2
firepower /ssa/logical-device # create external-port-link PC1.400_ftd Port-channel1.400
 ftd
firepower /ssa/logical-device/external-port-link* # exit
firepower /ssa/logical-device* # exit
firepower /ssa # scope logical-device LD3
firepower /ssa/logical-device # create external-port-link PC1.400_ftd Port-channel1.400
 ftd
firepower /ssa/logical-device/external-port-link* # exit
firepower /ssa/logical-device* # exit
firepower /ssa # scope logical-device LD4
firepower /ssa/logical-device # create external-port-link PC1.400_ftd Port-channel1.400
 ftd
firepower /ssa/logical-device/external-port-link* # commit-buffer
firepower /ssa/logical-device/external-port-link #
```

If you had committed the configuration in the middle, you would have ended up with 2 VLAN groups, which could have generated the switch forwarding path error and prevented you from saving the configuration.

2. You can check that the Port-channel1.400 VLAN ID was added to VLAN group 1.

```
firepower /ssa/logical-device/external-port-link # connect fxos
[...]
firepower(fxos)# show ingress-vlan-groups
ID   Class ID Status        INTF      Vlan Status
1    1        configured
                                      200  present
                                      100  present
                                      300  present
                                      400  present
2048 512      configured
                                      0    present
```

```
2049 511       configured
                                      0     present
firepower(fxos)# exit
firepower /ssa/logical-device/external-port-link #
```

History for Interfaces

Feature Name	Platform Releases	Feature Information
Synchronization between the FTD operational link state and the physical link state	2.9.1	The chassis can now synchronize the FTD operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The FTD application interface admin state is not considered. Without synchronization from FTD, data interfaces can be in an Up state physically before the FTD application has completely come online, for example, or can stay Up for a period of time after you initiate an FTD shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the FTD before the FTD can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS. **Note** This feature is not supported for clustering, container instances, or an FTD with a Radware vDP decorator. It is also not supported for the ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices > Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled**, **show interface expand detail**
Support for VLAN subinterfaces on a Cluster type interface (multi-instance use only)	2.8.1	For use with multi-instance clusters, you can now create VLAN subinterfaces on cluster type interfaces. Because each cluster requires a unique cluster control link, VLAN subinterfaces provide a simple method to fulfill this requirement. You can alternatively assign a dedicated EtherChannel per cluster. Multiple Cluster type interfaces are now allowed. New/Modified screens: **Interfaces** > **All Interfaces** > **Add New** drop-down menu **> Subinterface > Type** field
Support for 500 VLANs, without contingencies	2.7.1	Previously, the device supported between 250 and 500 VLANs, depending on the number of parent interfaces and other deployment decisions. You can now use 500 VLANs in all cases.
VLAN subinterfaces for use with container instances	2.4.1	To provide flexible physical interface use, you can create VLAN subinterfaces in FXOS and also share interfaces between multiple instances. **Note** Requires FTD Version 6.3 or later. New/Modified screens: **Interfaces** > **All Interfaces** > **Add New** drop-down menu **> Subinterface** New/Modified FMC screens: **Devices** > **Device Management** > **Edit** icon **> Interfaces** tab

Feature Name	Platform Releases	Feature Information
Data-sharing interfaces for container instances	2.4.1	To provide flexible physical interface use, you can share interfaces between multiple instances. **Note** Requires FTD Version 6.3 or later. New/Modified screens: **Interfaces** > **All Interfaces** > **Type**
Support for data EtherChannels in On mode	2.4.1	You can now set data and data-sharing EtherChannels to either Active LACP mode or to On mode. Other types of EtherChannels only support Active mode. New/Modified screens: **Interfaces** > **All Interfaces** > **Edit Port Channel** > **Mode**
Support for EtherChannels in FTD inline sets	2.1.1	You can now use EtherChannels in a FTD inline set.
Inline set link state propagation support for the FTD	2.0.1	When you configure an inline set in the FTD application and enable link state propagation, the FTD sends inline set membership to the FXOS chassis. Link state propagation means that the chassis automatically brings down the second interface in the inline interface pair when one of the interfaces in an inline set goes down.
Support for Hardware bypass network modules for the FTD	2.0.1	Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. New/Modified FMC screens: **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface**
Firepower-eventing type interface for FTD	1.1.4	You can specify an interface as firepower-eventing for use with the FTD. This interface is a secondary management interface for FTD devices. To use this interface, you must configure its IP address and other parameters at the FTD CLI. For example, you can separate management traffic from events (such as web events). See the "Management Interfaces" section in the FMC configuration guide *System Configuration* chapter. New/Modified Firepower Chassis Manager screens: **Interfaces** > **All Interfaces** > **Type**

Logical Devices

About Logical Devices

A logical device lets you run one application instance (either ASA or FTD) and also one optional decorator application (Radware DefensePro) to form a service chain.

When you add a logical device, you also define the application instance type and version, assign interfaces, and configure bootstrap settings that are pushed to the application configuration.

Note For the Firepower 9300, you can install different application types (ASA and FTD) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

Standalone and Clustered Logical Devices

You can add the following logical device types:

- Standalone—A standalone logical device operates as a standalone unit or as a unit in a High Availability pair.

- Cluster—A clustered logical device lets you group multiple units together, providing all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. Multiple module devices, like the Firepower 9300, support intra-chassis clustering. For the Firepower 9300, all three modules must participate in the cluster, for both native and container instances. The FDM does not support clustering.

Logical Device Application Instances: Container and Native

Application instances run in the following deployment types:

- Native instance—A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance.

- Container instance—A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances. Multi-instance capability is only supported for the FTD using FMC; it is not supported for the ASA or the FTD using FDM.

Note Multi-instance capability is similar to ASA multiple context mode, although the implementation is different. Multiple context mode partitions a single application instance, while multi-instance capability allows independent container instances. Container instances allow hard resource separation, separate configuration management, separate reloads, separate software updates, and full FTD feature support. Multiple context mode, due to shared resources, supports more contexts on a given platform. Multiple context mode is not available on the FTD.

For the Firepower 9300, you can use a native instance on some modules, and container instances on the other module(s).

Container Instance Interfaces

To provide flexible physical interface use for container instances, you can create VLAN subinterfaces in FXOS and also share interfaces (VLAN or physical) between multiple instances. Native instances cannot use VLAN subinterfaces or shared interfaces. A multi-instance cluster cannot use VLAN subinterfaces or shared interfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel. See Shared Interface Scalability, on page 157 and Add a VLAN Subinterface for Container Instances, on page 172.

Note This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the FTD application. See FXOS Interfaces vs. Application Interfaces, on page 153 for more information.

How the Chassis Classifies Packets

Each packet that enters the chassis must be classified, so that the chassis can determine to which instance to send a packet.

- Unique Interfaces—If only one instance is associated with the ingress interface, the chassis classifies the packet into that instance. For bridge group member interfaces (in transparent mode or routed mode), inline sets, or passive interfaces, this method is used to classify packets at all times.

- Unique MAC Addresses—The chassis automatically generates unique MAC addresses for all interfaces, including shared interfaces. If multiple instances share an interface, then the classifier uses unique MAC addresses assigned to the interface in each instance. An upstream router cannot route directly to an instance without unique MAC addresses. You can also set the MAC addresses manually when you configure each interface within the application.

Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each instance.

Classification Examples

Packet Classification with a Shared Interface Using MAC Addresses

The following figure shows multiple instances sharing an outside interface. The classifier assigns the packet to Instance C because Instance C includes the MAC address to which the router sends the packet.

Figure 7: Packet Classification with a Shared Interface Using MAC Addresses
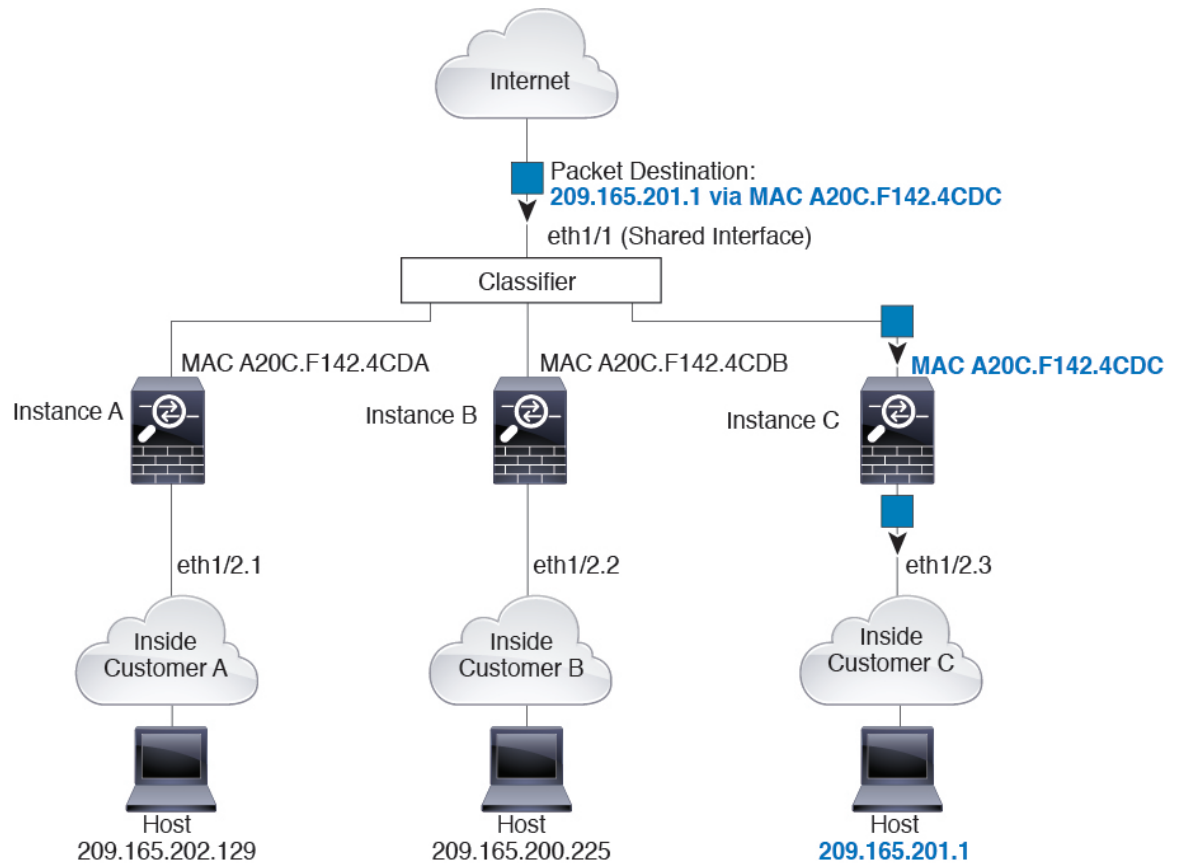
Incoming Traffic from Inside Networks

Note that all new incoming traffic must be classified, even from inside networks. The following figure shows a host on the Instance C inside network accessing the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 8: Incoming Traffic from Inside Networks
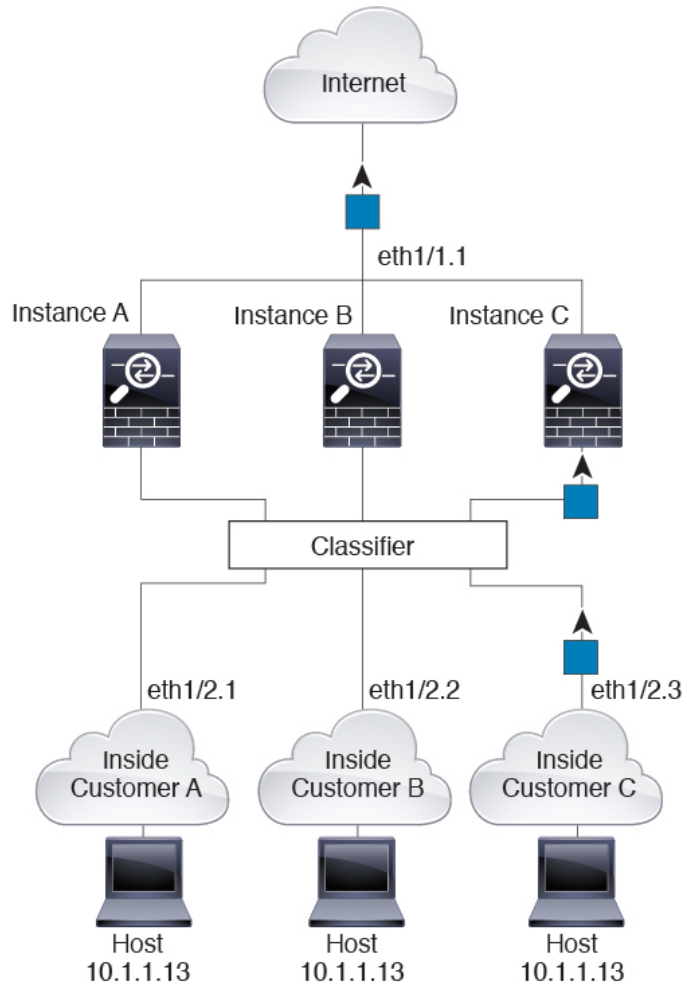
Transparent Firewall Instances

For transparent firewalls, you must use unique interfaces. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/2.3, which is assigned to Instance C.

Figure 9: Transparent Firewall Instances
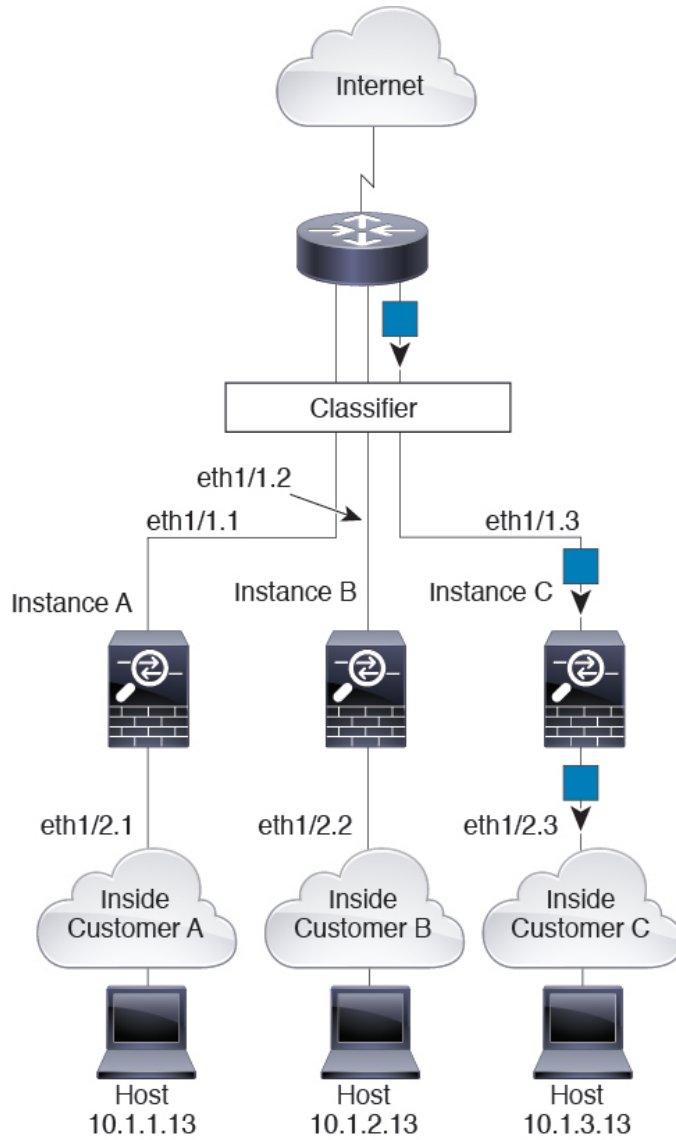
Inline Sets

For inline sets, you must use unique interfaces and they must be physical interfaces or EtherChannels. The following figure shows a packet destined to a host on the Instance C inside network from the internet. The classifier assigns the packet to Instance C because the ingress interface is Ethernet 1/5, which is assigned to Instance C.

Figure 10: Inline Sets
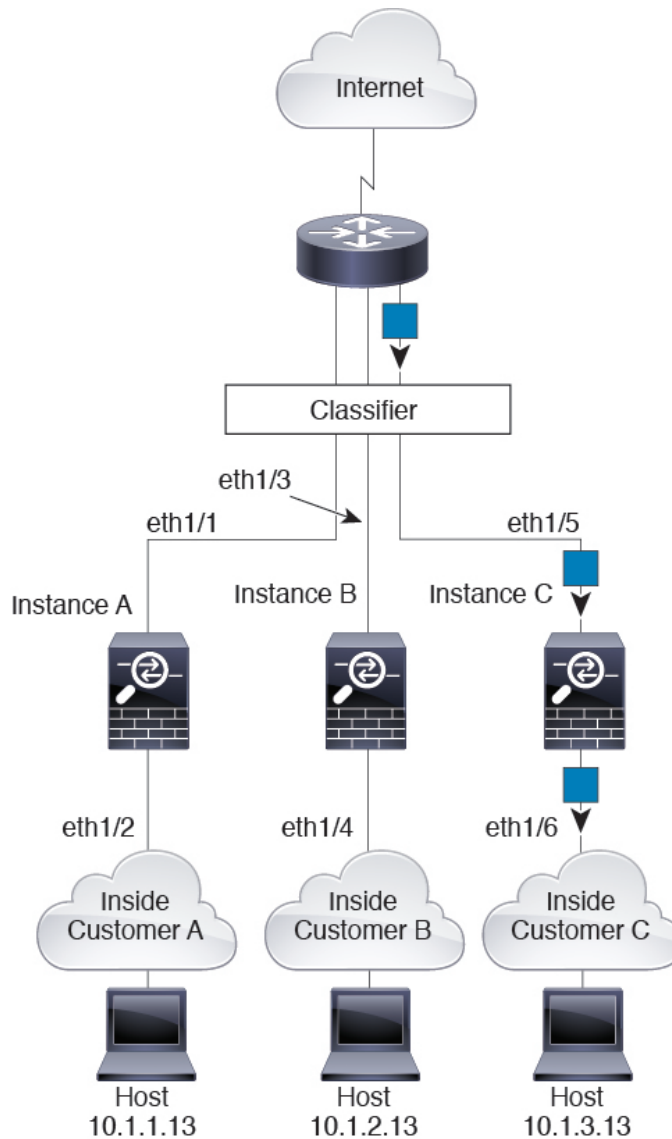
Cascading Container Instances

Placing an instance directly in front of another instance is called *cascading instances*; the outside interface of one instance is the same interface as the inside interface of another instance. You might want to cascade instances if you want to simplify the configuration of some instances by configuring shared parameters in the top instance.

The following figure shows a gateway instance with two instances behind the gateway.

Figure 11: Cascading nstances
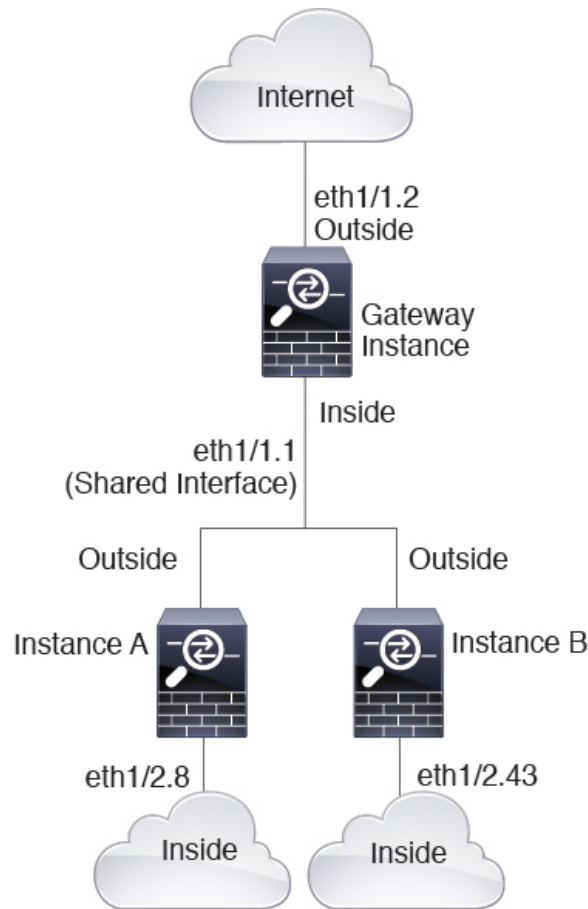
Typical Multi-Instance Deployment

The following example includes three container instances in routed firewall mode. They include the following interfaces:

- Management—All instances use the Port-Channel1 interface (management type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same management network.

- Inside—Each instance uses a subinterface on Port-Channel2 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.

- Outside—All instances use the Port-Channel3 interface (data-sharing type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Within each application, the interface uses a unique IP address on the same outside network.

- Failover—Each instance uses a subinterface on Port-Channel4 (data type). This EtherChannel includes two 10 Gigibit Ethernet interfaces. Each subinterface is on a separate network.
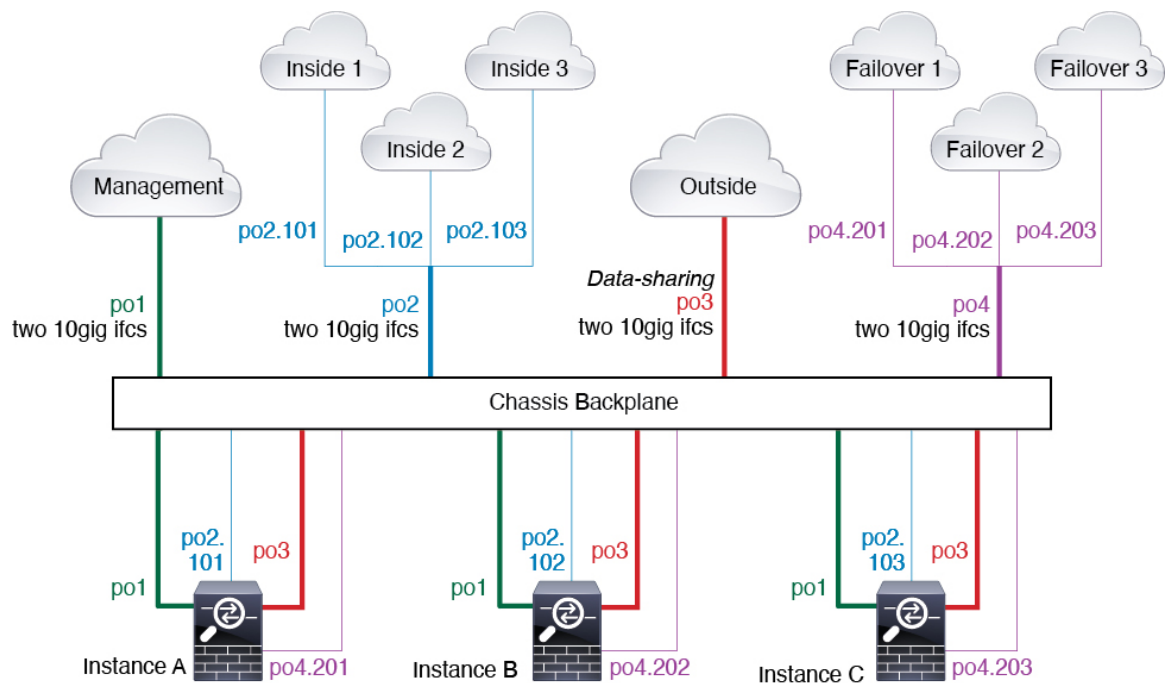
Automatic MAC Addresses for Container Instance Interfaces

The chassis automatically generates MAC addresses for instance interfaces, and guarantees that a shared interface in each instance uses a unique MAC address.

If you manually assign a MAC address to a shared interface within the instance, then the manually-assigned MAC address is used. If you later remove the manual MAC address, the autogenerated address is used. In the rare circumstance that the generated MAC address conflicts with another private MAC address in your network, we suggest that you manually set the MAC address for the interface within the instance.

Because autogenerated addresses start with A2, you should not start manual MAC addresses with A2 due to the risk of overlapping addresses.

The chassis generates the MAC address using the following format:

A2*xx.yyzz.zzzz*

Where *xx.yy* is a user-defined prefix or a system-defined prefix, and *zz.zzzz* is an internal counter generated by the chassis. The system-defined prefix matches the lower 2 bytes of the first MAC address in the burned-in MAC address pool that is programmed into the IDPROM. Use **connect fxos**, then **show module** to view the MAC address pool. For example, if the range of MAC addresses shown for module 1 is b0aa.772f.f0b0 to b0aa.772f.f0bf, then the system prefix will be f0b0.

The user-defined prefix is an integer that is converted into hexadecimal. For an example of how the user-defined prefix is used, if you set a prefix of 77, then the chassis converts 77 into the hexadecimal value 004D (*yyxx*). When used in the MAC address, the prefix is reversed (*xxyy*) to match the chassis native form:

A2**4D.00***zz.zzzz*

For a prefix of 1009 (03F1), the MAC address is:

A2**F1.03***zz.zzzz*

Container Instance Resource Management

To specify resource usage per container instance, create one or more resource profiles in FXOS. When you deploy the logical device/application instance, you specify the resource profile that you want to use. The resource profile sets the number of CPU cores; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance. To view the available resources per model, see Requirements and Prerequisites for Container Instances, on page 198. To add a resource profile, see Add a Resource Profile for Container Instances, on page 146.

Performance Scaling Factor for Multi-Instance Capability

The maximum throughput (connections, VPN sessions, and TLS proxy sessions) for a platform is calculated for a native instance's use of memory and CPU (and this value is shown in **show resource usage**). If you use multiple instances, then you need to calculate the throughput based on the percentage of CPU cores that you assign to the instance. For example, if you use a container instance with 50% of the cores, then you should initially calculate 50% of the throughput. Moreover, the throughput available to a container instance may be less than that available to a native instance.

For detailed instructions on calculating the throughput for instances, see https://www.cisco.com/c/en/us/products/collateral/security/firewalls/white-paper-c11-744750.html.

Container Instances and High Availability

You can use High Availability using a container instance on 2 separate chassis; for example, if you have 2 chassis, each with 10 instances, you can create 10 High Availability pairs. Note that High Availability is not configured in FXOS; configure each High Availability pair in the application manager.

For detailed requirements, see Requirements and Prerequisites for High Availability, on page 197 and Add a High Availability Pair, on page 217.

Container Instances and Clustering

You can create a cluster of container instances using one container instance per security module/engine. See Requirements and Prerequisites for Clustering, on page 193 for detailed requirements.

Requirements and Prerequisites for Logical Devices

See the following sections for requirements and prerequisites.

Requirements and Prerequisites for Hardware and Software Combinations

The Firepower 4100/9300 supports multiple models, security modules, application types, and high availability and scalability features. See the following requirements for allowed combinations.

Firepower 9300 Requirements

The Firepower 9300 includes 3 security module slots and multiple types of security modules. See the following requirements:

- Security Module Types—You can install modules of different types in the Firepower 9300. For example, you can install the SM-48 as module 1, SM-40 as module 2, and SM-56 as module 3.

- Native and Container instances—When you install a container instance on a security module, that module can only support other container instances. A native instance uses all of the resources for a module, so you can only install a single native instance on a module. You can use native instances on some modules, and container instances on the other module. For example, you can install a native instance on module 1 and module 2, but container instances on module 3.

- Native instance Clustering—All security modules in the cluster, whether it is intra-chassis or inter-chassis, must be the same type. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots. For example, you can install 2 SM-40s in chassis 1, and 3 SM-40s in chassis 2. You cannot use clustering if you install 1 SM-48 and 2 SM-40s in the same chassis.

- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.
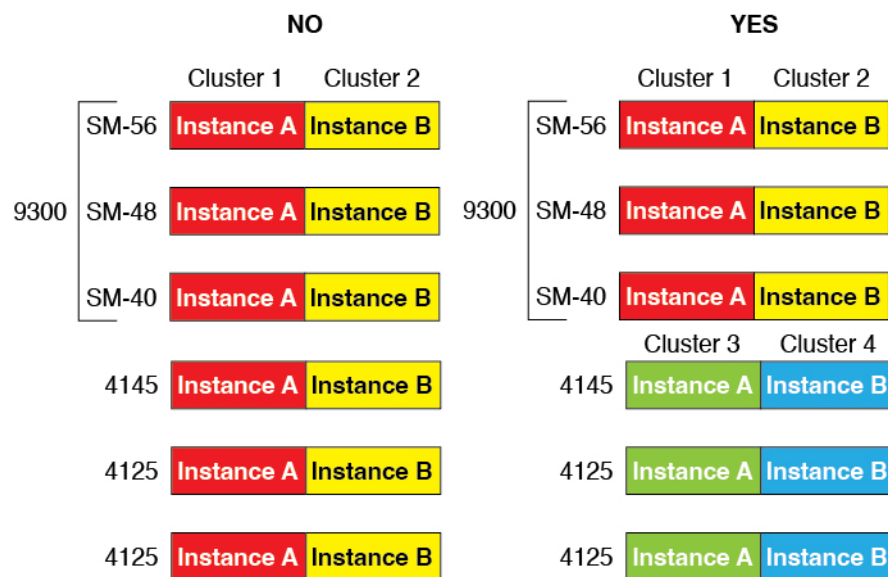
- High Availability—High Availability is only supported between same-type modules on the Firepower 9300. However, the two chassis can include mixed modules. For example, each chassis has an SM-40, SM-48, and SM-56. You can create High Availability pairs between the SM-40 modules, between the SM-48 modules, and between the SM-56 modules.

- ASA and FTD application types—You can install different application types on separate modules in the chassis. For example, you can install ASA on module 1 and module 2, and FTD on module 3.

- ASA or FTD versions—You can run different versions of an application instance type on separate modules, or as separate container instances on the same module. For example, you can install the FTD 6.3 on module 1, FTD 6.4 on module 2, and FTD 6.5 on module 3.

Firepower 4100 Requirements

The Firepower 4100 comes in multiple models. See the following requirements:

- Native and Container instances—When you install a container instance on a Firepower 4100, that device can only support other container instances. A native instance uses all of the resources for a device, so you can only install a single native instance on the device.

- Native instance Clustering—All chassis in the cluster must be the same model.

- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 4145 and a 4125. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.

NO   **YES**

	Cluster 1	Cluster 2			Cluster 1	Cluster 2
SM-56	Instance A	Instance B		SM-56	Instance A	Instance B
SM-48	Instance A	Instance B		SM-48	Instance A	Instance B
SM-40	Instance A	Instance B		SM-40	Instance A	Instance B

9300 (SM-56, SM-48, SM-40)

Cluster 3 Cluster 4

4145	Instance A	Instance B		4145	Instance A	Instance B
4125	Instance A	Instance B		4125	Instance A	Instance B
4125	Instance A	Instance B		4125	Instance A	Instance B
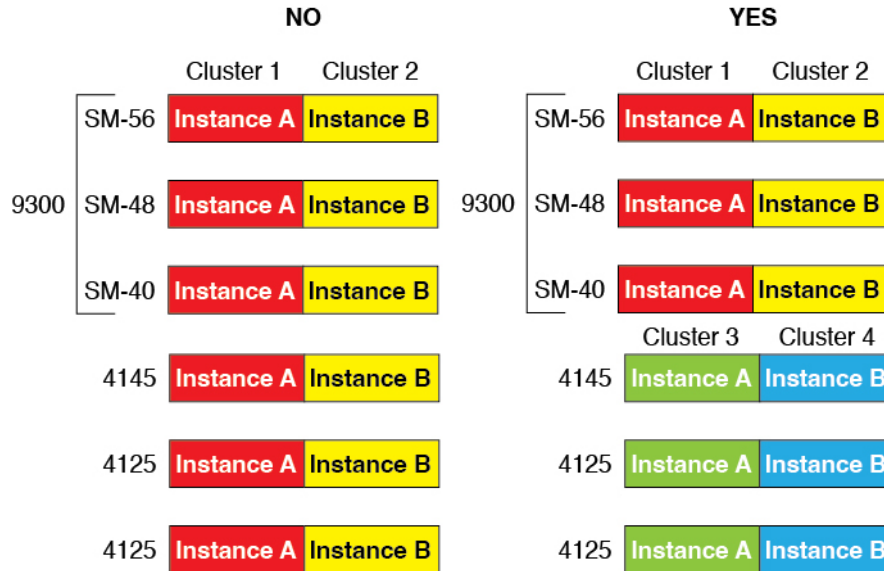
- High Availability—High Availability is only supported between same-type models.

- ASA and FTD application types—The Firepower 4100 can only run a single application type.

- The FTD container instance versions—You can run different versions of FTD as separate container instances on the same module.

Requirements and Prerequisites for Clustering

Cluster Model Support

- ASA on the Firepower 9300—Maximum 16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. Note that all modules in a chassis must belong to the cluster. Supported for intra-chassis, inter-chassis, and inter-site clustering.

- ASA on the Firepower 4100 series—Maximum 16 chassis. Supported for inter-chassis and inter-site clustering.

- FTD on the Firepower 9300 using FMC—Maximum 6 modules. For example, you can use 2 modules in 3 chassis, or 3 modules in 2 chassis, or any combination that provides a maximum of 6 modules. Note that all modules in a chassis must belong to the cluster. Supported for intra-chassis and inter-chassis clustering.

- FTD on the Firepower 4100 series using FMC—Maximum 6 chassis. Supported for inter-chassis clustering.

- Radware DefensePro—Supported for intra-chassis clustering with the ASA.

- Radware DefensePro—Supported for intra-chassis clustering with the FTD. Not supported for multi-instance clustering.

Clustering Hardware and Software Requirements

All chassis in a cluster:

- Native instance clustering—For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.

- Container instance clustering—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.
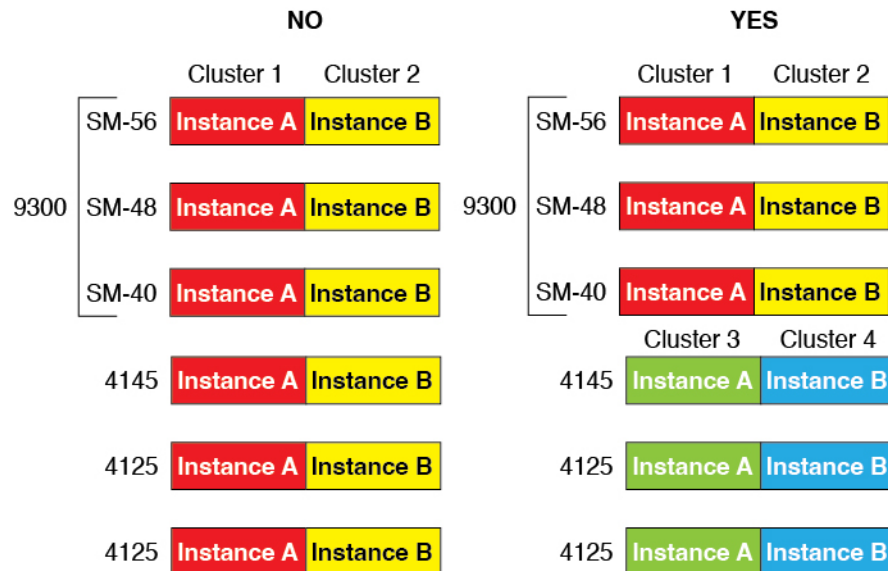
- Must run the identical FXOS software except at the time of an image upgrade.

- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in inter-chassis clustering. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.

- Must use the same NTP server. For FTD, the FMC must also use the same NTP server. Do not set the time manually.

- ASA: Each FXOS chassis must be registered with the License Authority or satellite server. There is no extra cost for data nodes. For permanent license reservation, you must purchase separate licenses for each chassis. For FTD, all licensing is handled by the FMC.

Multi-Instance Clustering Requirements

- No intra-security-module/engine clustering—For a given cluster, you can only use a single container instance per security module/engine. You cannot add 2 container instances to the same cluster if they are running on the same module.
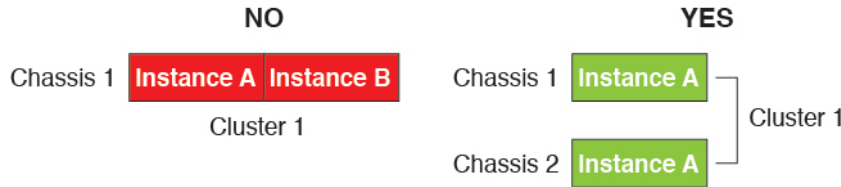
- Mix and match clusters and standalone instances—Not all container instances on a security module/engine need to belong to a cluster. You can use some instances as standalone or High Availability nodes. You can also create multiple clusters using separate instances on the same security module/engine.
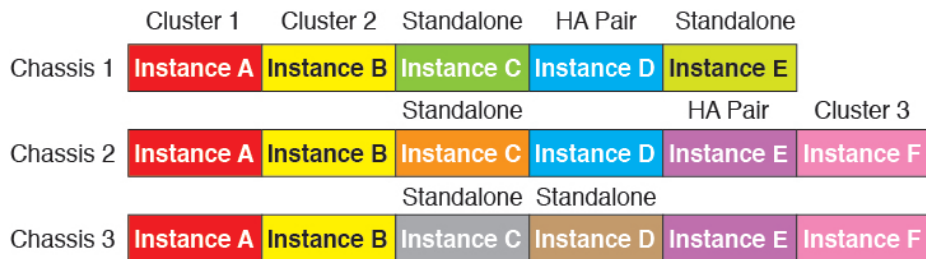
- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.

- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.

- Dedicated cluster control link—For inter-chassis clustering, each cluster needs a dedicated cluster control link. For example, each cluster can use a separate subinterface on the same cluster-type EtherChannel, or use separate EtherChannels.

Cluster 1   Cluster 2   Cluster 3

Chassis 1   Instance A   Instance B   Instance C

Po48.100: CCL1   Po48.101: CCL2   Po48.102: CCL3

Port-channel48

Switch

Port-channel48

Po48.101: CCL2

Po48.100: CCL1   Po48.102: CCL3

Chassis 2   Instance A   Instance B   Instance C

Cluster 1   Cluster 2   Cluster 3
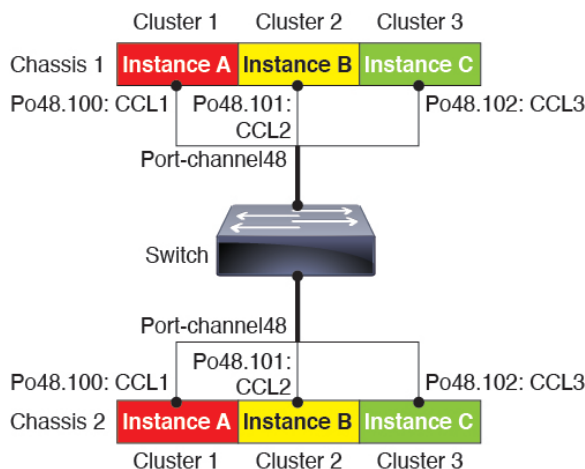
- No shared interfaces—Shared-type interfaces are not supported with clustering. However, the same Management and Eventing interfaces can used by multiple clusters.

- No subinterfaces—A multi-instance cluster cannot use FXOS-defined VLAN subinterfaces. An exception is made for the cluster control link, which can use a subinterface of the Cluster EtherChannel.

- Mix chassis models—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.

NO

Cluster 1   Cluster 2

SM-56   Instance A   Instance B

9300   SM-48   Instance A   Instance B

SM-40   Instance A   Instance B

4145   Instance A   Instance B

4125   Instance A   Instance B

4125   Instance A   Instance B

YES

Cluster 1   Cluster 2

SM-56   Instance A   Instance B

9300   SM-48   Instance A   Instance B

SM-40   Instance A   Instance B

Cluster 3   Cluster 4

4145   Instance A   Instance B

4125   Instance A   Instance B

4125   Instance A   Instance B
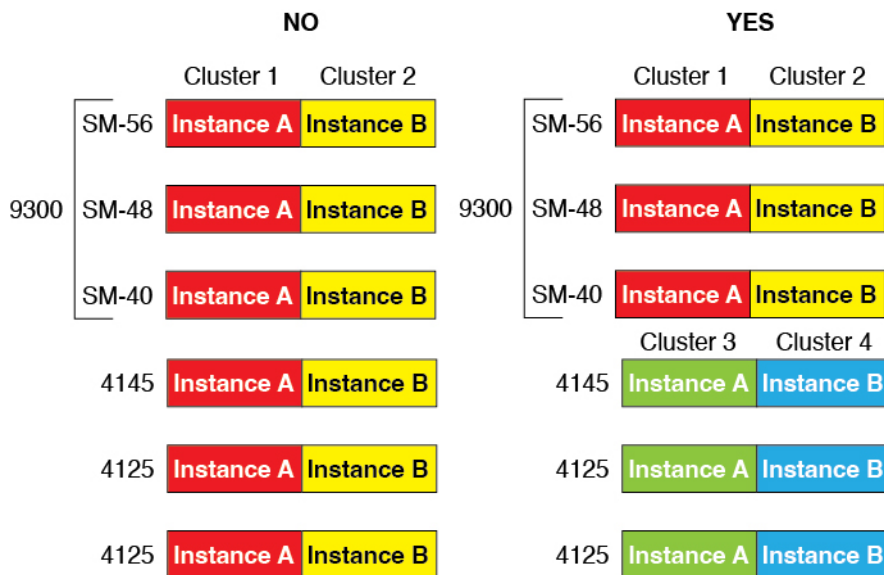
- Maximum 6 nodes—You can use up to six container instances in a cluster.

Switch Requirements for Inter-Chassis Clustering

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.

• For supported switch characteristics, see Cisco FXOS Compatibility.

Sizing the Data Center Interconnect for Inter-Site Clustering

You should reserve bandwidth on the data center interconnect (DCI) for cluster control link traffic equivalent to the following calculation:

$$\frac{\text{\# of cluster members per site}}{2} \times \text{cluster control link size per member}$$

If the number of members differs at each site, use the larger number for your calculation. The minimum bandwidth for the DCI should not be less than the size of the cluster control link for one member.

For example:

- For 4 members at 2 sites:
    - 4 cluster members total
    - 2 members at each site
    - 5 Gbps cluster control link per member

  Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).

- For 6 members at 3 sites, the size increases:
    - 6 cluster members total
    - 3 members at site 1, 2 members at site 2, and 1 member at site 3
    - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).

- For 2 members at 2 sites:
    - 2 cluster members total
    - 1 member at each site
    - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 10 Gbps (1/2 x 10 Gbps = 5 Gbps; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).

Requirements and Prerequisites for High Availability

- The two units in a High Availability Failover configuration must:
    - Be on a separate chassis; intra-chassis High Availability for the Firepower 9300 is not supported.
    - Be the same model.
    - Have the same interfaces assigned to the High Availability logical devices.

- Have the same number and types of interfaces. All interfaces must be preconfigured in FXOS identically before you enable High Availability.

- High Availability is only supported between same-type modules on the Firepower 9300; but the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.

- For container instances, each unit must use the same resource profile attributes.

- For other High Availability system requirements, see the application configuration guide chapter for High Availability.

Requirements and Prerequisites for Container Instances

Supported Application Types

- The FTD using FMC

Maximum Container Instances and Resources per Model

For each container instance, you can specify the number of CPU cores to assign to the instance. RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

Table 16: Maximum Container Instances and Resources per Model

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 4110	3	22	53 GB	125.6 GB
Firepower 4112	3	22	78 GB	308 GB
Firepower 4115	7	46	162 GB	308 GB
Firepower 4120	3	46	101 GB	125.6 GB
Firepower 4125	10	62	162 GB	644 GB
Firepower 4140	7	70	222 GB	311.8 GB
Firepower 4145	14	86	344 GB	608 GB
Firepower 4150	7	86	222 GB	311.8 GB
Firepower 9300 SM-24 security module	7	46	226 GB	656.4 GB
Firepower 9300 SM-36 security module	11	70	222 GB	640.4 GB
Firepower 9300 SM-40 security module	13	78	334 GB	1359 GB

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 9300 SM-44 security module	14	86	218 GB	628.4 GB
Firepower 9300 SM-48 security module	15	94	334 GB	1341 GB
Firepower 9300 SM-56 security module	18	110	334 GB	1314 GB

FMC Requirements

For all instances on a Firepower 4100 chassis or Firepower 9300 module, you must use the same FMC due to the licensing implementation.

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the FTD and ASA.

High Availability

- Configure high availability within the application configuration.

- You can use any data interfaces as the failover and state links. Data-sharing interfaces are not supported.

Multi-InstanceandContext Mode

- Multiple context mode is only supported on the ASA.

- Enable multiple context mode in the ASA after you deploy.

- Multi-instance capability with container instances is only available for the FTD using FMC.

- For FTD container instances, a single FMC must manage all instances on a security module/engine.

- You can enable TLS crypto acceleration on up to 16 container instances.

- For FTD container instances, the following features are not supported:

  - Radware DefensePro link decorator

  - FMC UCAPL/CC mode

  - Flow offload to hardware

Clustering Guidelines and Limitations

Switches for Inter-Chassis Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead.

- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.

- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.

- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.

- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.

- Some switches do not support dynamic port priority with LACP (active and standby links). You can disable dynamic port priority to provide better compatibility with Spanned EtherChannels.

- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.

- Port-channel bundling downtime should not exceed the configured keepalive interval.

- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

  router(config)# **port-channel** *id* **hash-distribution fixed**

  Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

- Firepower 4100/9300 clusters support LACP graceful convergence. So you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.

- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default. Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

EtherChannels for Inter-Chassis Clustering

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.

  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.
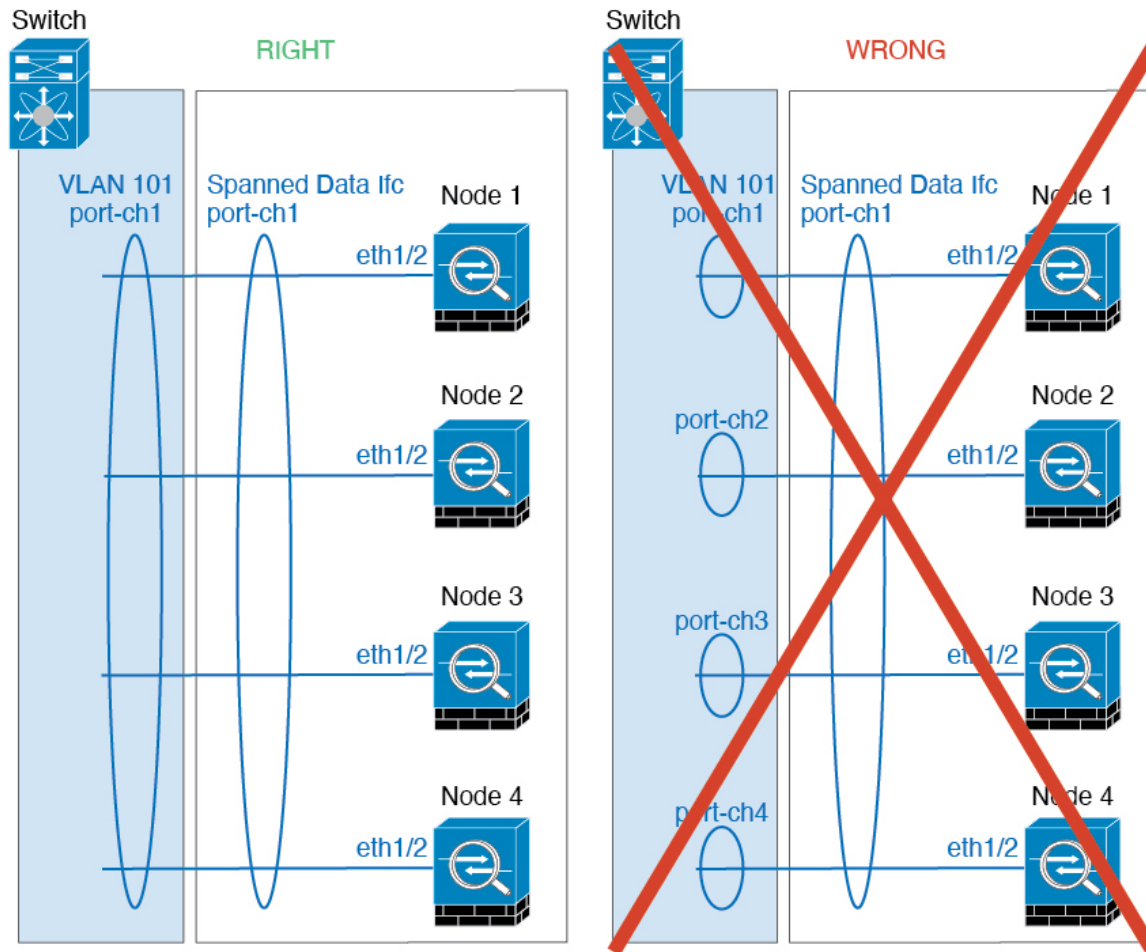
  - Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.
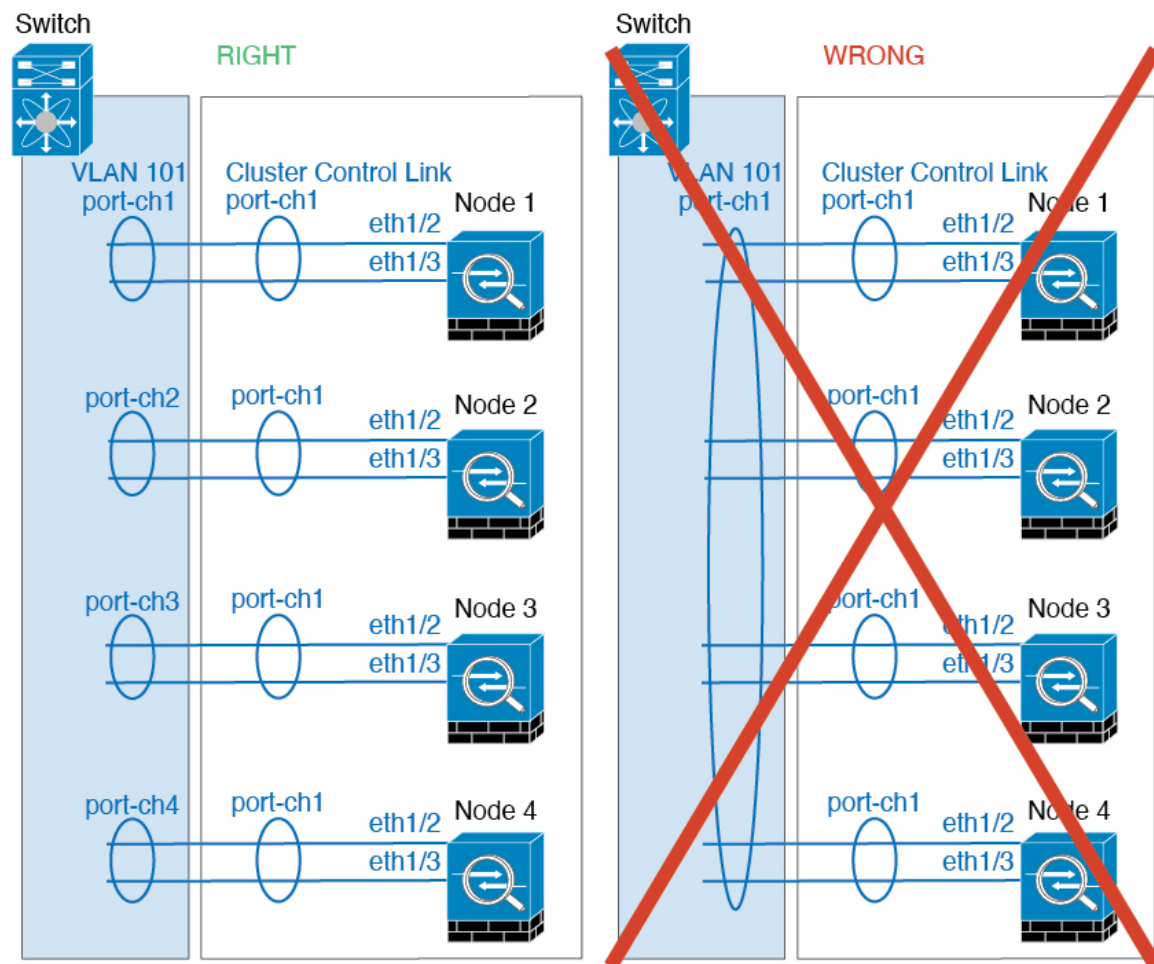
Inter-Site Clustering

See the following guidelines for inter-site clustering:

- The cluster control link latency must be less than 20 ms round-trip time (RTT).

- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.

- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.

- The does not encrypt forwarded data traffic on the cluster control link because it is a dedicated link, even when used on a Data Center Interconnect (DCI). If you use Overlay Transport Virtualization (OTV), or are otherwise extending the cluster control link outside of the local administrative domain, you can configure encryption on your border routers such as 802.1AE MacSec over OTV.

- The cluster implementation does not differentiate between members at multiple sites for incoming connections; therefore, connection roles for a given connection may span across sites. This is expected behavior. However, if you enable director localization, the local director role is always chosen from the same site as the connection owner (according to site ID). Also, the local director chooses a new owner

at the same site if the original owner fails (Note: if the traffic is asymmetric across sites, and there is continuous traffic from the remote site after the original owner fails, then a node from the remote site might become the new owner if it receives a data packet within the re-hosting window.).

- For director localization, the following traffic types do not support localization: NAT or PAT traffic; SCTP-inspected traffic; Fragmentation owner query.

- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.

- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.

- For transparent mode, if the cluster is connected to an HSRP router, you must add the router HSRP MAC address as a static MAC address table entry on the . When adjacent routers use HSRP, traffic destined to the HSRP IP address will be sent to the HSRP MAC Address, but return traffic will be sourced from the MAC address of a particular router's interface in the HSRP pair. Therefore, the MAC address table is typically only updated when the ARP table entry for the HSRP IP address expires, and the sends an ARP request and receives a reply. Because the 's ARP table entries expire after 14400 seconds by default, but the MAC address table entry expires after 300 seconds by default, a static MAC address entry is required to avoid MAC address table expiration traffic drops.

- For routed mode using Spanned EtherChannel, configure site-specific MAC addresses. Extend the data VLANs across the sites using OTV, or something similar. You need to create filters to prevent traffic that is destined to the global MAC address from being sent over the DCI to the other site. If the cluster becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's cluster nodes. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

Additional Guidelines

- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.

- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.

- We recommend connecting EtherChannels to a VSS, vPC, StackWise, or StackWise Virtual for redundancy.

• Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.

• For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults

• The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.

• The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.

• The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.

• Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Add a Standalone Logical Device

Standalone logical devices can be used alone or as high availability units. For more information about high availability usage, see Add a High Availability Pair, on page 217.

Add a Standalone ASA

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

You can deploy a routed or transparent firewall mode ASA from the Firepower 4100/9300 chassis.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

Before you begin

• Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

Note For the Firepower 9300, you can install different application types (ASA and FTD) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

• Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).

• Gather the following information:

• Interface IDs for this device

• Management interface IP address and network mask

• Gateway IP address

Procedure

Step 1   Choose **Logical Devices**.

Step 2   Click **Add** > **Standalone**, and set the following parameters:

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b) For the **Template**, choose **Cisco: Adaptive Security Appliance**.

c) Choose the **Image Version**.

d) Click **OK**.

You see the Provisioning - *device name* window.

Step 3   Expand the **Data Ports** area, and click each port that you want to assign to the device.

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces on the ASA, including setting the IP addresses.

Step 4   Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5   On the **General Information** page, complete the following:

a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.

b) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

c) Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

d) Configure the **Management IP** address.

Set a unique IP address for this interface.

      e) Enter a **Network Mask** or **Prefix Length**.

      f) Enter a **Network Gateway** address.

Step 6   Click the **Settings** tab.

Step 7   Choose the **Firewall Mode**: **Routed** or **Transparent**.

      In routed mode, the ASA is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

      The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

Step 8   Enter and confirm a **Password** for the admin user and for the enable password.

      The pre-configured ASA admin user/password and enable password is useful for password recovery; if you have FXOS access, you can reset the admin user password/enable password if you forget it.

Step 9   Click **OK** to close the configuration dialog box.

Step 10   Click **Save**.

      The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.
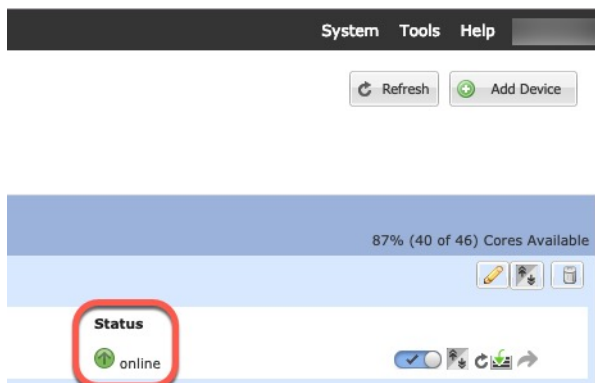
Step 11   See the ASA configuration guide to start configuring your security policy.

Add a Standalone FTD for the FMC

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

You can use native instances on some modules, and container instances on the other module(s).

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

  Note For the Firepower 9300, you can install different application types (ASA and FTD) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).

- You can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management. See the **configure network management-data-interface** command in the FTD command reference for more information.

- You must also configure at least one Data type interface. Optionally, you can also create a firepower-eventing interface to carry all event traffic (such as web events). See Interface Types, on page 152 for more information.

- For container instances, if you do not want to use the default profile, add a resource profile according to Add a Resource Profile for Container Instances, on page 146.

- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the **Reinitialize icon**. An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance. See Reinitializing a Security Module/Engine, on page 274 for more information.

- Gather the following information:

    - Interface IDs for this device

    - Management interface IP address and network mask

    - Gateway IP address

    - FMC IP address and/or NAT ID of your choosing

    - DNS server IP address

    - FTD hostname and domain name

Procedure

Step 1   Choose **Logical Devices**.

Step 2   Click **Add** > **Standalone**, and set the following parameters:

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.

c) Choose the **Image Version**.

d) Choose the **Instance Type**: **Container** or **Native**.

A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.

e) Click **OK**.

You see the Provisioning - *device name* window.

Step 3   Expand the **Data Ports** area, and click each interface that you want to assign to the device.
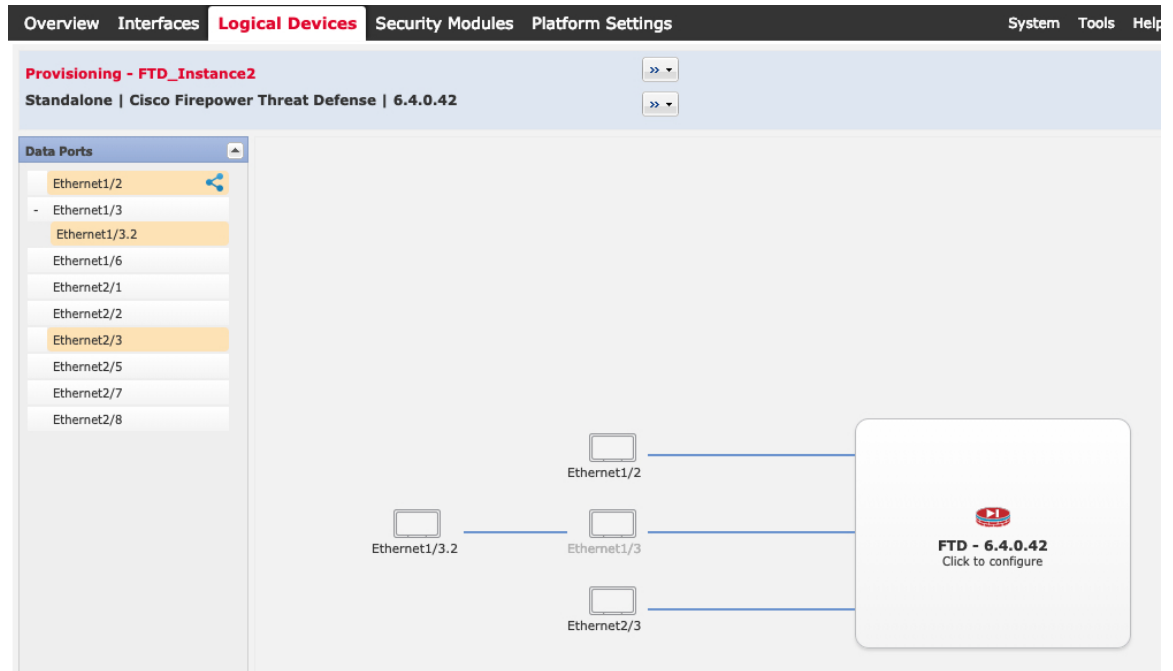
You can only assign data and data-sharing interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in FMC, including setting the IP addresses.

You can only assign up to 10 data-sharing interfaces to a container instance. Also, each data-sharing interface can be assigned to at most 14 container instances. A data-sharing interface is indicated by the sharing icon (⦉).

Hardware Bypass-capable ports are shown with the following icon: ⦿. For certain interface modules, you can enable the Hardware Bypass feature for Inline Set interfaces only (see the FMC configuration guide). Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. If you do not assign both interfaces in a Hardware Bypass pair, you see a warning message to make sure your assignment is intentional. You do not need to use the Hardware Bypass feature, so you can assign single interfaces if you prefer.

Step 4   Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5   On the **General Information** page, complete the following:

a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.

b) For a container instance, specify the **Resource Profile**.

If you later assign a different resource profile, then the instance will reload, which can take approximately 5 minutes. Note that for established High Availability pairs, if you assign a different-sized resource profile, be sure to make all members the same size as soon as possible.

c) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

d) Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

e) Configure the **Management IP** address.

Set a unique IP address for this interface.

f) Enter a **Network Mask** or **Prefix Length**.

g) Enter a **Network Gateway** address.

Step 6   On the **Settings** tab, complete the following:

a) For a native instance, in the **Management type of application instance** drop-down list, choose **FMC**.

Native instances also support FDM as a manager. After you deploy the logical device, you cannot change the manager type.

b) Enter the **Firepower Management Center IP** of the managing FMC. If you do not know the FMC IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.

c) For a container instance, **Permit Expert mode from FTD SSH sessions**: **Yes** or **No**. Expert Mode provides FTD shell access for advanced troubleshooting.

If you choose **Yes** for this option, then users who access the container instance directly from an SSH sesssion can enter Expert Mode. If you choose **No**, then only users who access the container instance from the FXOS CLI can enter Expert Mode. We recommend choosing **No** to increase isolation between instances.

Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the FTD CLI.

d) Enter the **Search Domains** as a comma-separated list.

e) Choose the **Firewall Mode**: **Transparent** or **Routed**.

In routed mode, the FTD is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

f) Enter the **DNS Servers** as a comma-separated list.

The FTD uses DNS if you specify a hostname for the FMC, for example.

g) Enter the **Fully Qualified Hostname** for the FTD.

h) Enter a **Registration Key** to be shared between the FMC and the device during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the FTD.

i) Enter a **Password** for the FTD admin user for CLI access.

j) Choose the **Eventing Interface** on which events should be sent. If not specified, the management interface will be used.

This interface must be defined as a Firepower-eventing interface.

k) For a container instance, set the **Hardware Crypto** as **Enabled** or **Disabled**.

This setting enables TLS crypto acceleration in hardware, and improves performance for certain types of traffic. This feature is enabled by default. You can enable TLS crypto acceleration for up to 16 instances per security module. This feature is always enabled for native instances. To view the percentage of hardware crypto resources allocated to this instance, enter the **show hw-crypto** command.

Step 7 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.
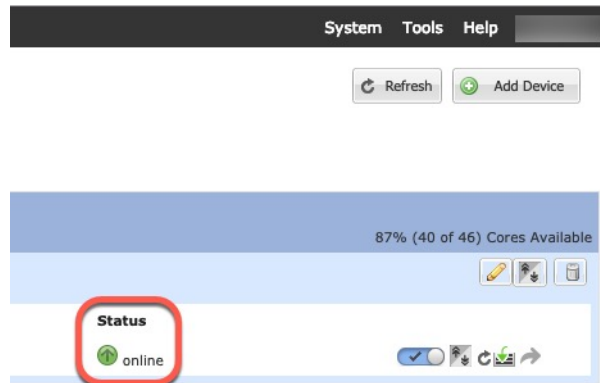
Step 10 See the FMC configuration guide to add the FTD as a managed device and start configuring your security policy.

Add a Standalone FTD for the FDM

You can use the FDM with a native instance. Container instances are not supported. Standalone logical devices work either alone or in a High Availability pair.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

Note For the Firepower 9300, you can install different application types (ASA and FTD) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).

- You must also configure at least one Data type interface.

- Gather the following information:

  - Interface IDs for this device

  - Management interface IP address and network mask

  - Gateway IP address

  - DNS server IP address

  - FTD hostname and domain name

Procedure

Step 1   Choose **Logical Devices**.

Step 2   Click **Add** > **Standalone**, and set the following parameters:

Add Standalone

Device Name:	FTD_1
Template:	Cisco Firepower Threat Defense
Image Version:	6.5.0.1159
Instance Type:	Native

OK   Cancel

a)   Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b)   For the **Template**, choose **Cisco Firepower Threat Defense**.

c)   Choose the **Image Version**.

d)   Choose the **Instance Type**: **Native**.

Container instances are not supported with the FDM.

e)   Click **OK**.

You see the Provisioning - *device name* window.

Step 3   Expand the **Data Ports** area, and click each interface that you want to assign to the device.

Overview   Interfaces   **Logical Devices**   Security Modules   Platform Settings   System   To

Provisioning - FTD-FDM2   »▾

Standalone | Cisco Firepower Threat Defense | 6.5.0.1159

Data Ports

Ethernet1/8
Ethernet2/1
Ethernet2/2
Ethernet2/3
Ethernet2/4
Ethernet2/5
Ethernet2/6
Ethernet2/7
Port-channel3

Ethernet1/8

Ethernet2/1

Port-channel3

Ethernet2/2

FTD - 6.5.0.1159
Click to configure

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in the FDM, including setting the IP addresses.

Step 4   Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5   On the **General Information** page, complete the following:

a)   (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.

b)   Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

c)   Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.

d)   Configure the **Management IP** address.

Set a unique IP address for this interface.

e)   Enter a **Network Mask** or **Prefix Length**.

f)   Enter a **Network Gateway** address.

Step 6   On the **Settings** tab, complete the following:

a) In the **Management type of application instance** drop-down list, choose **LOCALLY_MANAGED**.

Native instances also support the Firepower Management Center as a manager. If you change the manager after you deploy the logical device, then your configuration is erased and the device is reinitialized.

b) Enter the **Search Domains** as a comma-separated list.

c) The **Firewall Mode** only supports **Routed** mode.

d) Enter the **DNS Servers** as a comma-separated list.

e) Enter the **Fully Qualified Hostname** for the FTD.

f) Enter a **Password** for the FTD admin user for CLI access.

Step 7   On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 8   Click **OK** to close the configuration dialog box.

Step 9   Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.
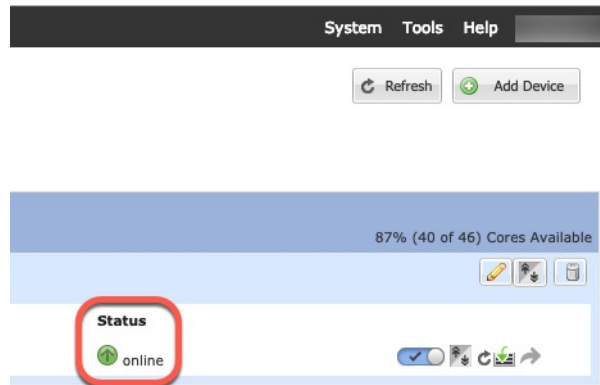
Step 10   See the FDM configuration guide to start configuring your security policy.

Add a High Availability Pair

FTD or ASA High Availability (also known as failover) is configured within the application, not in FXOS. However, to prepare your chassis for high availability, see the following steps.

Before you begin

See Requirements and Prerequisites for High Availability, on page 197.

Procedure

Step 1   Allocate the same interfaces to each logical device.

Step 2   Allocate 1 or 2 data interfaces for the failover and state link(s).

These interfaces exchange high availability traffic between the 2 chassis. We recommend that you use a 10 GB data interface for a combined failover and state link. If you have available interfaces, you can use separate failover and state links; the state link requires the most bandwidth. You cannot use the management-type interface for the failover or state link. We recommend that you use a switch between the chassis, with no other device on the same network segment as the failover interfaces.

For container instances, data-sharing interfaces are not supported for the failover link. We recommend that you create subinterfaces on a parent interface or EtherChannel, and assign a subinterface for each instance to use as a failover link. Note that you must use all subinterfaces on the same parent as failover links. You cannot use one subinterface as a failover link and then use other subinterfaces (or the parent interface) as regular data interfaces.

Step 3   Enable High Availability on the logical devices.

Step 4   If you need to make interface changes after you enable High Availability, perform the changes on the standby unit first, and then perform the changes on the active unit.

Note   For the ASA, if you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

Add a Cluster

Clustering lets you group multiple devices together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. The Firepower 9300, which includes multiple modules, supports intra-chassis clustering where you group all modules within a single chassis into a cluster. You can also use inter-chassis clustering, where multiple chassis are grouped together; inter-chasis clustering is the only option for single module devices like the Firepower 4100 series.

About Clustering on the Firepower 4100/9300 Chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- For native instance clustering: Creates a *cluster-control link* (by default, port-channel 48) for unit-to-unit communication.

  For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

  For intra-chassis clustering (Firepower 9300 only), this link utilizes the Firepower 9300 backplane for cluster communications.

  For inter-chassis clustering, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

  When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings. Some parts of the bootstrap configuration may be user-configurable within the application if you want to customize your clustering environment.

- Assigns data interfaces to the cluster as *Spanned* interfaces.

  For intra-chassis clustering, spanned interfaces are not limited to EtherChannels, like it is for inter-chassis clustering.The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For inter-chassis clustering, you must use Spanned EtherChannels for all data interfaces.

  Note   Individual interfaces are not supported, with the exception of a management interface.

- Assigns a management interface to all units in the cluster.

Primary and Secondary Unit Roles

One member of the cluster is the primary unit. The primary unit is determined automatically. All other members are secondary units.

You must perform all configuration on the primary unit only; the configuration is then replicated to the secondary units.

Cluster Control Link

For native instance clustering: The cluster control link is automatically created using the Port-channel 48 interface.

For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

For intra-chassis clustering, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications for intra-chassis clustering. For inter-chassis clustering, you must add one or more interfaces to the EtherChannel.

For a 2-member inter-chassis cluster, do not directly connect the cluster control link from one chassis to the other chassis. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Cluster control link traffic includes both control and data traffic.

Size the Cluster Control Link for Inter-Chassis Clustering

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster control link can handle the worst-case scenarios.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features. For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.

- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.

Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy for Inter-Chassis Clustering

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the

EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.
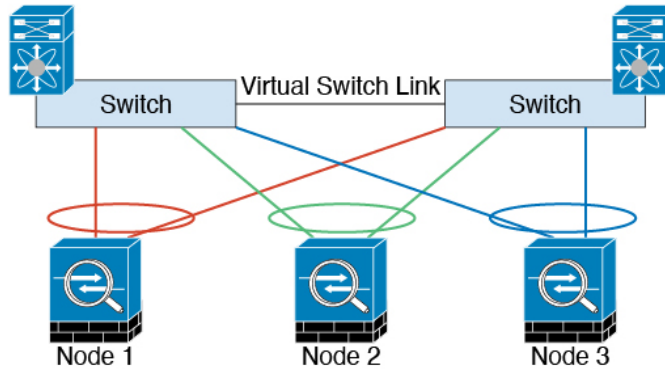
Cluster Control Link Reliability for Inter-Chassis Clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Cluster Control Link Network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id.slot_id*. For multi-instance clusters, which typically use different VLAN subinterfaces of the same EtherChannel, the same IP address can be used for different clusters because of VLAN separation. You can customize this IP address when you deploy the cluster. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed. For inter-site traffic, Cisco recommends using Overlay Transport Virtualization (OTV).

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

You must assign a Management type interface to the cluster. This interface is a special *individual* interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit.

For the ASA, the Main cluster IP address is a fixed address for the cluster that always belongs to the current primary unit. You must configure a range of addresses so that each unit, including the current primary unit, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a primary unit changes, the Main cluster IP address moves to the new primary unit, so management of the cluster continues seamlessly. The Local IP address is used for routing, and is also useful for troubleshooting. For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current primary unit. To manage an individual member, you can connect to

the Local IP address. For outbound management traffic such as TFTP or syslog, each unit, including the primary unit, uses the Local IP address to connect to the server.

For the FTD, assign a management IP address to each unit on the same network. Use these IP addresses when you add each unit to the FMC.

Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface. The EtherChannel inherently provides load balancing as part of basic operation.

For multi-instance clusters, each cluster requires dedicated data EtherChannels; you cannot use shared interfaces or VLAN subinterfaces.
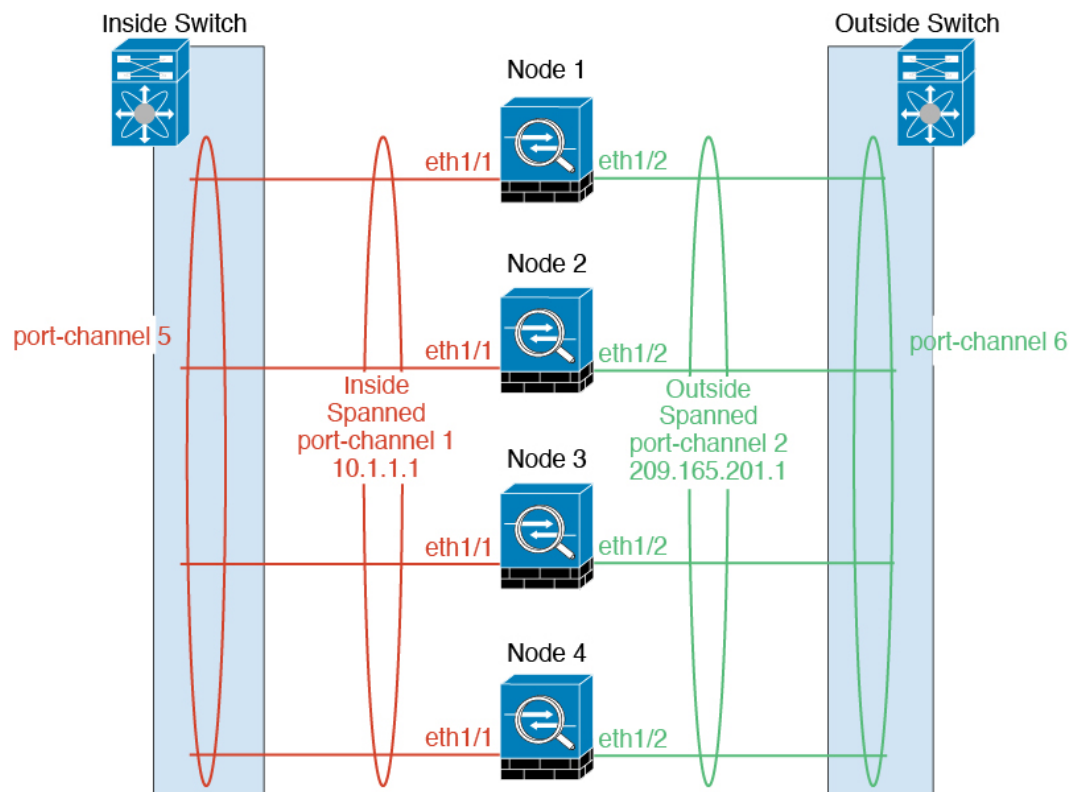
Inter-Site Clustering

For inter-site installations, you can take advantage of clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID.

Site IDs work with site-specific MAC addresses and IP addresses. Packets egressing the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. This feature prevents the switches from learning the same global MAC address from both

sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address. Site-specific MAC addresses and IP address are supported for routed mode using Spanned EtherChannels only.

Site IDs are also used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

See the following sections for more information about inter-site clustering:

- Sizing the Data Center Interconnect—Requirements and Prerequisites for Clustering, on page 193
- Inter-Site Guidelines—Clustering Guidelines and Limitations, on page 200
- Inter-Site Examples—Examples for Inter-Site Clustering, on page 262

Add an ASA Cluster

You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering. For inter-chassis clustering, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment

Create an ASA Cluster

Set the scope to the image version.

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For inter-chassis clustering, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- Gather the following information:
    - Management interface ID, IP address, and network mask
    - Gateway IP address

Procedure

Step 1   Configure interfaces.

Step 2   Choose **Logical Devices**.

Step 3   Click **Add** > **Cluster**, and set the following parameters:

a) Choose **I want to:** > **Create New Cluster**

b) Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

c) For the **Template**, choose **Cisco Adaptive Security Appliance**.

d) Choose the **Image Version**.

e) For the **Instance Type**, only the **Native** type is supported.

f) Click **OK**.

You see the Provisioning - *device name* window.

Step 4   Choose the interfaces you want to assign to this cluster.

All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

Step 5   Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6   On the **Cluster Information** page, complete the following.

a) For inter-chassis clustering, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

This field only appears if you added a member interface to cluster control link Port-Channel 48.

b) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8.

c) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.

The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

d) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.

The name must be an ASCII string from 1 to 38 characters.

e) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

f) (Optional) Set the **CCL Subnet IP** as *a*.*b*.**0.0**.

By default, the cluster control link uses the 127.2.0.0/16 network. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. In this case, specify any /16 network address on a unique network for the cluster, except for loopback (127.0.0.0/8), multicast (224.0.0.0/4), and internal (169.254.0.0/16) addresses. If you set the value to 0.0.0.0, then the default network is used.

The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: *a.b.chassis_id.slot_id*.

g) Choose the **Address Type** for the management interface.

This information is used to configure a management interface in the ASA configuration. Set the following information:

- **Management IP Pool**—Configure a pool of Local IP addresses, one of which will be assigned to each cluster unit for the interface, by entering the starting and ending addresses separated by a hyphen.

  Include at least as many addresses as there are units in the cluster. Note that for the Firepower 9300, you must include 3 addresses per chassis, even if you do not have all module slots filled. If you plan to expand the cluster, include additional addresses. The Virtual IP address (known as the Main cluster IP address) that belongs to the current control unit is *not* a part of this pool; be sure to reserve an IP address on the same network for the Main cluster IP address. You can use IPv4 and/or IPv6 addresses.

- **Network Mask** or **Prefix Length**

- **Network Gateway**

- **Virtual IP address**—Set the management IP address of the current control unit. This IP address must be on the same network as the cluster pool addresses, but not be part of the pool.

Step 7 On the **Settings** page, complte the following.

a) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the FTD is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

b) Enter and confirm a **Password** for the admin user and for the enable password.

The pre-configured ASA admin user is useful for password recovery; if you have FXOS access, you can reset the admin user password if you forget it.

Step 8   Click **OK** to close the configuration dialog box.

Step 9   Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for intra-chassis clusteringstart configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
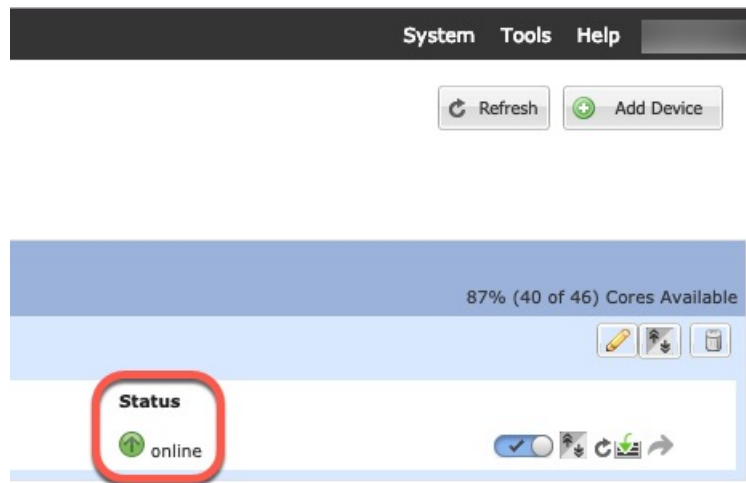
Step 10   For inter-chassis clustering, add the next chassis to the cluster:

a) On the first chassis of the Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

b) Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.

c) Choose **I want to: > Join an Existing Cluster**.

d) Click **OK**.

e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.

- **Site ID**—Enter the correct site ID.

- **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
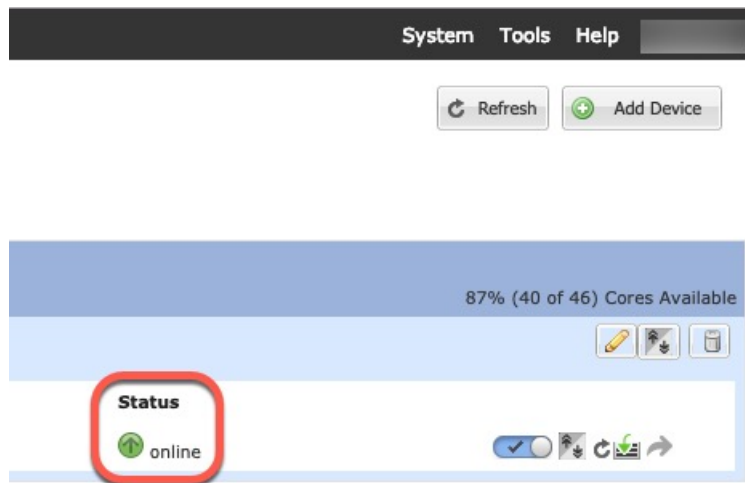
Step 11   Connect to the control unit ASA to customize your clustering configuration.

Add More Cluster Members

Add or replace the ASA cluster member.

Note   This procedure only applies to adding or replacing a *chassis*; if you are adding or replacing a module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- Make sure your existing cluster has enough IP addresses in the management IP address pool for this new member. If not, you need to edit the existing cluster bootstrap configuration on each chassis before you add this new member. This change causes a restart of the logical device.

- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

- For multiple context mode, enable multiple context mode in the ASA application on the first cluster member; additional cluster members will inherit the multiple context mode configuration automatically.

Procedure

Step 1   On an existing cluster the Firepower Chassis Manager, choose **Logical Devices** to open the **Logical Devices** page.

Step 2   Click the Show Configuration icon ( ) at the top right; copy the displayed cluster configuration.

Step 3   Connect to the Firepower Chassis Manager on the new chassis, and click **Add** > **Cluster**.

Step 4 Choose **I want to:** > **Join Existing Cluster**

Step 5 For the **Device Name**, provide a name for the logical device.

Step 6 Click **OK**.

Step 7 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

Step 8 Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.

- **Site ID**—Enter the correct site ID.

- **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
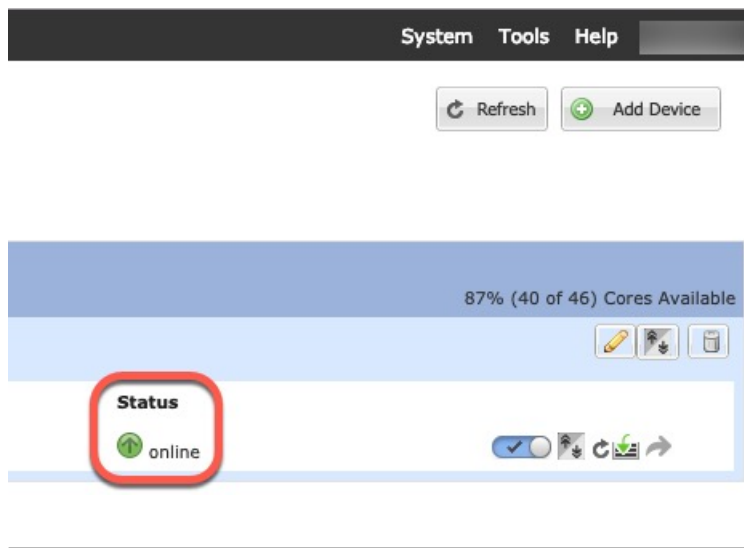
Add a FTD Cluster

In native mode: You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering.

In multi-instance mode: You can add one or more clusters on a single Firepower 9300 chassis as intra-chassis clusters (you must include an instance on each module), or add one or more clusters on multiple chassis for inter-chassis clustering.

For inter-chassis clustering, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment

Create a FTD Cluster

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For inter-chassis clustering, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.

- For container instances, if you do not want to use the default profile, add a resource profile according to Add a Resource Profile for Container Instances, on page 146.

- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the Reinitialize icon (⟳). An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance. See Reinitializing a Security Module/Engine, on page 274 for more information.

- Gather the following information:

    - Management interface ID, IP addresses, and network mask

    - Gateway IP address

    - FMC IP address and/or NAT ID of your choosing

    - DNS server IP address

    - FTD hostname and domain name

Procedure

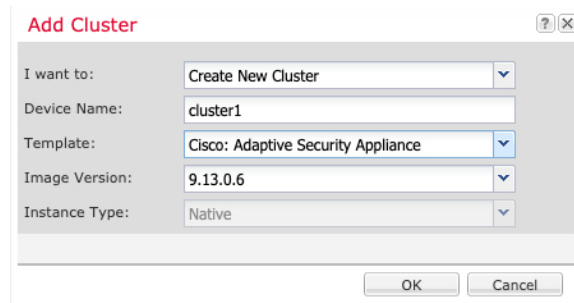**Step 1**   Configure interfaces.

Step 2   Choose **Logical Devices**.

Step 3   Click **Add** > **Cluster**, and set the following parameters:

Figure 12: Native Cluster

Figure 13: Multi-Instance Cluster

a) Choose **I want to:** > **Create New Cluster**

b) Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

c) For the **Template**, choose **Cisco Firepower Threat Defense**.

d) Choose the **Image Version**.

e) For the **Instance Type**, choose either **Native** or **Container**.

A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.

f) (Container Instance only) For the **Resource Type**, choose one of the resource profiles from the drop-down list.

For the Firepower 9300, this profile will be applied to each instance on each security module. You can set different profiles per security module later in this procedure; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model. We recommend choosing the correct profile before you create the cluster. If you need to create a new profile, cancel out of the cluster creation, and add one using .

g) Click **OK**.

You see the Provisioning - *device name* window.

Step 4   Choose the interfaces you want to assign to this cluster.
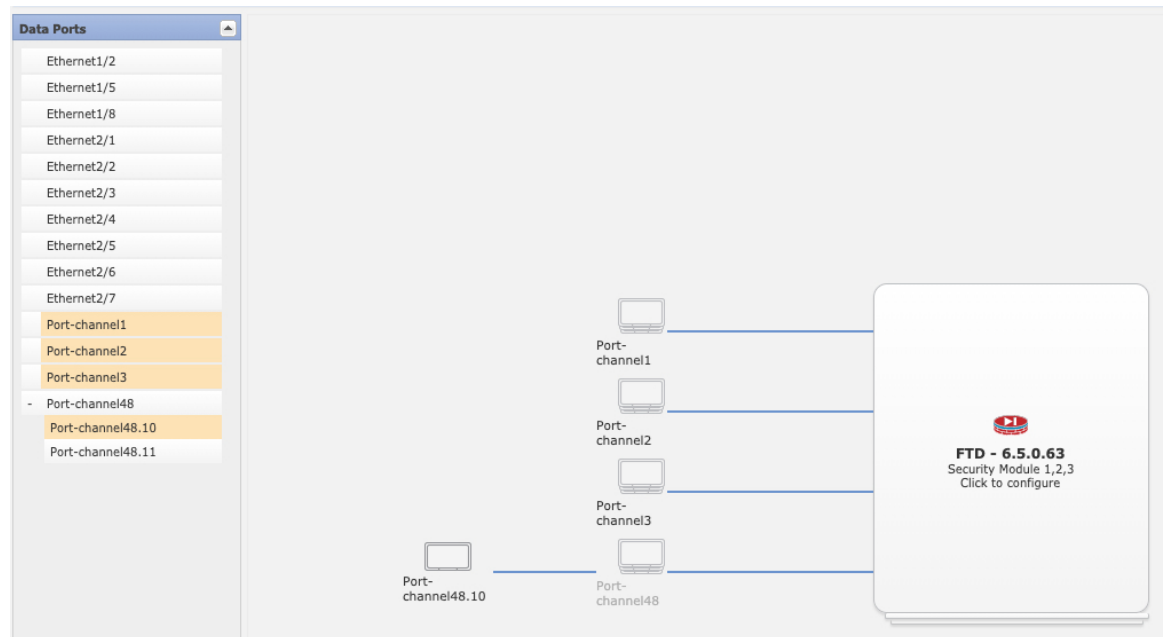
For native mode clustering: All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

For multi-instance clustering: Choose each data interface you want to assign to the cluster, and also choose the Cluster type port-channel or port-channel subinterface.

Step 5   Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Figure 14: Native Cluster

Figure 15: Multi-Instance Cluster

a) (Container Instance for the Firepower 9300 only) In the **Security Module (SM) and Resource Profile Selection** area, you can set a different resource profile per module; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model.

b) For inter-chassis clustering, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

This field only appears if you added a member interface to cluster control link Port-Channel 48.

c) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8. FlexConfig feature. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.

d) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.

The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.

e) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.

The name must be an ASCII string from 1 to 38 characters.

f) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

If you assign a Hardware Bypass-capable interface as the Management interface, you see a warning message to make sure your assignment is intentional.

g) (Optional) Set the **CCL Subnet IP** as *a*.*b*.**0.0**.

By default, the cluster control link uses the 127.2.0.0/16 network. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. In this case, specify any /16 network address on a unique network for the cluster, except for loopback (127.0.0.0/8), multicast (224.0.0.0/4), and internal (169.254.0.0/16) addresses. If you set the value to 0.0.0.0, then the default network is used.

The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: *a.b.chassis_id.slot_id*.

Step 7 On the **Settings** page, complete the following.

a) In the **Registration Key** field, enter the key to be shared between the FMC and the cluster members during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the FTD.

b) Enter a **Password** for the FTD admin user for CLI access.

c) In the **Firepower Management Center IP** field, enter the IP address of the managing FMC. If you do not know the FMC IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.

d) (Optional) For a container instance, **Permit Expert mode from FTD SSH sessions**: **Yes** or **No**. Expert Mode provides FTD shell access for advanced troubleshooting.

If you choose **Yes** for this option, then users who access the container instance directly from an SSH sesssion can enter Expert Mode. If you choose **No**, then only users who access the container instance from the FXOS CLI can enter Expert Mode. We recommend choosing **No** to increase isolation between instances.

Use Expert Mode only if a documented procedure tells you it is required, or if the Cisco Technical Assistance Center asks you to use it. To enter this mode, use the **expert** command in the FTD CLI.

e) (Optional) In the **Search Domains** field, enter a comma-separated list of search domains for the management network.

f) (Optional) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the FTD is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

g) (Optional) In the **DNS Servers** field, enter a comma-separated list of DNS servers.

The FTD uses DNS if you specify a hostname for the FMC, for example.

h) (Optional) In the **Firepower Management Center NAT ID** field, enter a passphrase that you will also enter on the FMC when you add the cluster as a new device.

Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. You can specify any text string as the NAT ID, from 1 to 37 characters. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

i) (Optional) In the **Fully Qualified Hostname** field, enter a fully qualified name for the FTD device.

Valid characters are the letters from a to z, the digits from 0 to 9, the dot (.), and the hyphen (-); maximum number of characters is 253.

j) (Optional) From the **Eventing Interface** drop-down list, choose the interface on which events should be sent. If not specified, the management interface will be used.

To specify a separate interface to use for events, you must configure an interface as a *firepower-eventing* interface. If you assign a Hardware Bypass-capable interface as the Eventing interface, you see a warning message to make sure your assignment is intentional.

Step 8   On the **Interface Information** page, configure a management IP address for each security module in the cluster. Select the type of address from the **Address Type** drop-down list and then complete the following for each security module.

Note   You must set the IP address for all 3 module slots in a chassis, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

a) In the **Management IP** field, configure an IP address.

Specify a unique IP address on the same network for each module.

b) Enter a **Network Mask** or **Prefix Length**.

c) Enter a **Network Gateway** address.

Step 9 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 10 Click **OK** to close the configuration dialog box.

Step 11 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for intra-chassis clusteringstart configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
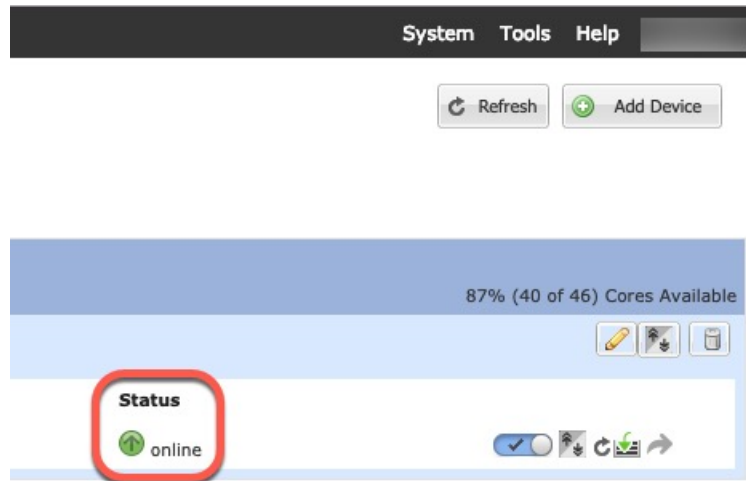
Step 12 For inter-chassis clustering, add the next chassis to the cluster:

a) On the first chassis of the Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

b) Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.

c) Choose **I want to: > Join an Existing Cluster**.

d) Click **OK**.

e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.

- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.

- **Cluster Key**—(Not prefilled) Enter the same cluster key.

- **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

Click **OK**.

g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
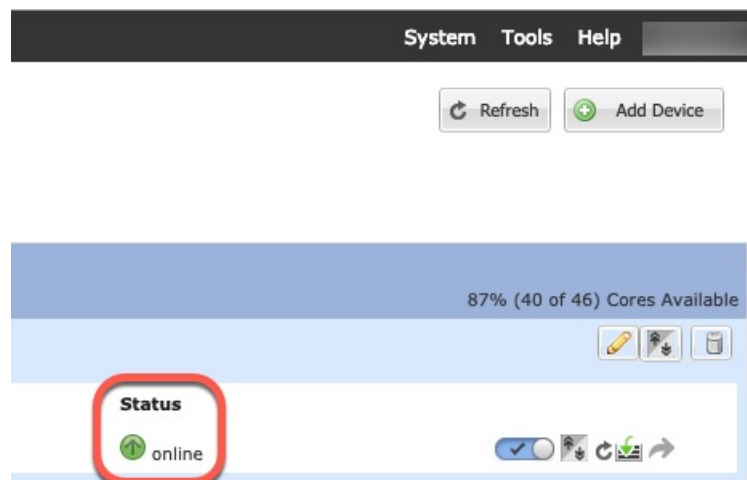
Step 13   Add the control unit to the FMC using the management IP address.

All cluster units must be in a successfully-formed cluster on FXOS prior to adding them to FMC.

The FMC then automatically detects the data units.

Add More Cluster Nodes

Add or replace the FTD cluster node in an existing cluster. When you add a new cluster node in FXOS, the FMC adds the node automatically.

Note   The FXOS steps in this procedure only apply to adding a new *chassis*; if you are adding a new module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- In the case of a replacement, you must delete the old cluster node from the FMC. When you replace it with a new node, it is considered to be a new device on the FMC.

- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

Procedure

Step 1   If you previously upgraded the FTD image using the FMC, perform the following steps *on each chassis in the cluster*.

When you upgraded from the FMC, the startup version in the FXOS configuration was not updated, and the standalone package was not installed on the chassis. Both of these items need to be set manually so the new node can join the cluster using the correct image version.

> **Note** If you only applied a patch release, you can skip this step. Cisco does not provide standalone packages for patches.

 a) Install the running FTD image on the chassis using the **System** > **Updates** page.

 b) Click **Logical Devices** and click the Set Version icon (). For a Firepower 9300 with multiple modules, set the version for each module.

  The **Startup Version** shows the original package you deployed with. The **Current Version** shows the version you upgraded to.

 c) In the **New Version** drop-down menu, choose the version that you uploaded. This version should match the **Current Version** displayed, and will set the startup version to match the new version.

 d) On the new chassis, make sure the new image package is installed.

Step 2   On an existing cluster chassis Firepower Chassis Manager, click **Logical Devices**.

Step 3   Click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

Step 4   Connect to the Firepower Chassis Manager on the new chassis, and click **Add** > **Cluster**.

Step 5   For the **Device Name**, provide a name for the logical device.

Step 6   Click **OK**.

Step 7   In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

Step 8   Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

  • **Chassis ID**—Enter a unique chassis ID.

  • **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. This feature is only configurable using the FMC FlexConfig feature.

  • **Cluster Key**—(Not prefilled) Enter the same cluster key.

  • **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

  Click **OK**.

Step 9   Click **Save**.

  The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.
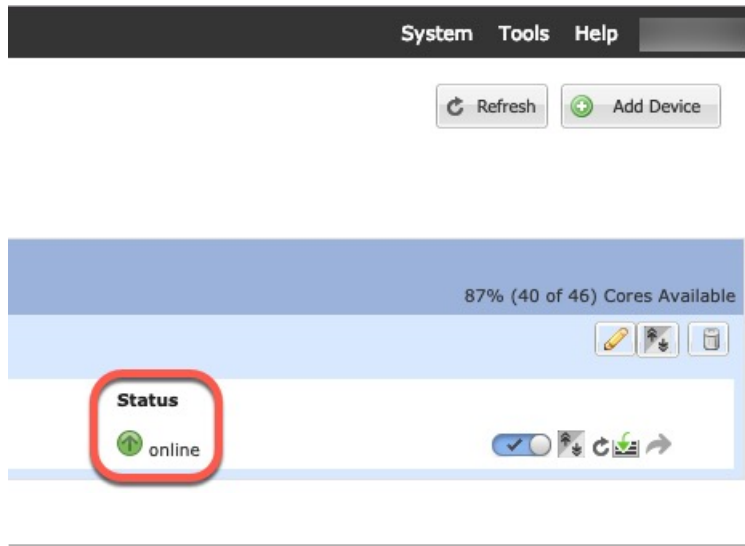
Configure Radware DefensePro

The Cisco Firepower 4100/9300 chassis can support multiple services (for example, a firewall and a third-party DDoS application) on a single blade. These applications and services can be linked together to form a Service Chain.

About Radware DefensePro

In the current supported Service Chaining configuration, the third-party Radware DefensePro virtual platform can be installed to run in front of the ASA firewall, or in front of FTD. Radware DefensePro is a KVM-based virtual platform that provides distributed denial-of-service (DDoS) detection and mitigation capabilities on the Firepower 4100/9300 chassis. When Service Chaining is enabled on your Firepower 4100/9300 chassis, traffic from the network must first pass through the DefensePro virtual platform before reaching the main ASA or FTD firewall.

✎

Note
- The Radware DefensePro virtual platform may be referred to as *Radware vDP* (virtual DefensePro), or simply *vDP*.

- The Radware DefensePro virtual platform may occasionally be referred to as a Link Decorator.

Prerequisites for Radware DefensePro

Prior to deploying Radware DefensePro on your Firepower 4100/9300 chassis, you must configure the Firepower 4100/9300 chassis to use an NTP Server with the **etc/UTC** Time Zone. For more information about setting the date and time in your Firepower 4100/9300 chassis, see Setting the Date and Time, on page 101.

Guidelines for Service Chaining

Models

- ASA—The Radware DefensePro (vDP) platform is supported with ASA on the following models:

  - Firepower 9300

  - Firepower 4115

  - Firepower 4120

  - Firepower 4125

  - Firepower 4140

  - Firepower 4145

  - Firepower 4150

  Note The Radware DefensePro platform is not currently supported with ASA on Firepower 4110 devices.

- FTD—The Radware DefensePro platform is supported with FTD on the following models:

  - Firepower 9300

  - Firepower 4110—Note you must deploy the decorator at the same time as the logical device. You cannot install the decorator after the logical device is already configured on the device.

  - Firepower 4112

  - Firepower 4115

  - Firepower 4120—Note you must deploy the decorator at the same time as the logical device. You cannot install the decorator after the logical device is already configured on the device.

  - Firepower 4125

  - Firepower 4140

  - Firepower 4145

  - Firepower 4150

  Note You must use the CLI to deploy Radware DefensePro for all FTD platforms; the Firepower Chassis Manager does not yet support this functionality.

Additional Guidelines

- Service Chaining is not supported in an inter-chassis cluster configuration. However, the Radware DefensePro (vDP) application can be deployed in a standalone configuration in an inter-chassis cluster scenario.

Configure Radware DefensePro on a Standalone Logical Device

The following procedure shows how to install Radware DefensePro in a single Service Chain in front of a standalone ASA or FTD logical device.

> **Note**
>
> Once you set the vDP application and commit the change at the end of this procedure, the logical device (ASA or FTD) will reboot.
>
> If you are installing Radware vDP in front of ASA on a Firepower 4120 or 4140 security appliance, you must use the FXOS CLI to deploy the decorator. For full CLI instructions on how to install and configure Radware DefensePro in a service chain in front of ASA on Firepower 4100 devices, refer to the FXOS CLI configuration guide.

Before you begin

- Download the vDP image from Cisco.com (see Downloading Images from Cisco.com, on page 56) and then upload that image to the Firepower 4100/9300 chassis (see Uploading an Image to the Security Appliance, on page 56).

- You can deploy the Radware DefensePro application in a standalone configuration on an intra-chassis cluster; for intra-chassis clustering, see Configure Radware DefensePro on an Intra-Chassis Cluster, on page 243.

Procedure

Step 1 If you want to use a separate management interface for vDP, enable the interface and set it to be the mgmt type according to Configure a Physical Interface, on page 169. Otherwise, you can share the application management interface.

Step 2 Choose **Logical Devices** to open the Logical Devices page.

The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown.

Step 3 Create a standalone ASA or FTD Logical Device (see Add a Standalone ASA, on page 204 or Add a Standalone FTD for the FMC, on page 207).

Step 4 In the **Decorators** area, select vDP. The Radware: Virtual DefensePro - Configuration window appears. Configure the following fields under the **General Information** tab.

Step 5 If you have more than one vDP version uploaded to the Firepower 4100/9300 chassis, select the version you want to use in the **Version** drop-down.

Step 6	If you have a resource configurable Radware DefensePro application, a list of supported resource profiles appears under the **Resource Profile** drop-down. Select the resource profile you want to assign to the device. If you do not select a resource profile, the default setting is used.
Step 7	Under the **Management Interface** drop-down, choose the management interface you created in step 1 of this procedure.
Step 8	Select the default **Address Type**, IPv4 only, IPv6 only, or IPv4 and IPv6.
Step 9	Configure the following fields, based on your **Address Type** selection from the previous step.

a) In the **Management IP** field, configure a local IP address.

b) IPv4 only: Enter a **Network Mask**.

IPv6 only: Enter a **Prefix Length**.

c) Enter a **Network Gateway** address.

Step 10	Click the checkbox next to each data port that you want to assign to the device.
Step 11	Click **OK**.
Step 12	Click **Save**.

The FXOS deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the specified security module.

What to do next

Set a password for the DefensePro application. Note that the application does not come online until you set a password. For more information, see the Radware DefensePro DDoS Mitigation User Guide on cisco.com.

Configure Radware DefensePro on an Intra-Chassis Cluster

The following procedure shows how to install the Radware DefensePro image, and configure it in a Service Chain in front of an ASA or FTD intra-chassis cluster.

Note Service Chaining is not supported in an inter-chassis cluster configuration. However, the Radware DefensePro application can be deployed in a standalone configuration in an inter-chassis cluster scenario.

Before you begin

- Download the vDP image from Cisco.com (see Downloading Images from Cisco.com, on page 56) and then upload that image to the Firepower 4100/9300 chassis (see Uploading an Image to the Security Appliance, on page 56).

Procedure

Step 1	If you want to use a separate management interface for vDP, enable the interface and set it to be the mgmt type according to Configure a Physical Interface, on page 169. Otherwise, you can share the application management interface.

Step 2   Configure an ASA or FTD intra-chassis cluster (see Create an ASA Cluster, on page 222 or Create a FTD Cluster, on page 229).

Note that before you click **Save** at the end of the procedure to configure the intra-chassis cluster, you must first follow the following steps to add a vDP decorator to the cluster.

Step 3   In the **Decorators** area, select vDP. The **Radware: Virtual DefensePro - Configuration** dialog box appears. Configure the following fields under the **General Information** tab.

Step 4   If you have more than one vDP version uploaded to the Firepower 4100/9300 chassis, select the vDP version you want to use in the **Version** drop-down.

Step 5   If you have a resource configurable Radware DefensePro application, a list of supported resource profiles appears under the Resource Profile drop-down. Select the resource profile you want to assign to the device. If you do not select a resource profile, the default setting is used.

Step 6   Under the **Management Interface** drop-down, choose a management interface.

Step 7   Click the checkbox next to each data port that you want to assign to the vDP decorator.

Step 8   Click the **Interface Information** tab.

Step 9   Select the **Address Type** to be used, IPv4 only, IPv6 only, or IPv4 and IPv6.

Step 10   Configure the following fields for each Security Module. Note that the fields that display depend on your **Address Type** selection from the previous step.

a) In the **Management IP** field, configure a local IP address.
b) IPv4 only: Enter a **Network Mask**.

    IPv6 only: Enter a **Prefix Length**.

c) Enter a **Network Gateway** address.

Step 11   Click **OK**.

Step 12   Click **Save**.

The FXOS deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the specified security module.

Step 13   Choose **Logical Devices** to open the Logical Devices page.

Step 14   Scroll through the list of configured logical devices to the entries for vDP. Verify their Attributes listed in the **Management IP** column.

- If the **CLUSTER-ROLE** element displays as *unknown* for the DefensePro instances, you must enter the DefensePro application and configure the Control unit IP address to complete the creation of the vDP cluster.

- If the **CLUSTER-ROLE** element displays as *primary* or *secondary* for the DefensePro instances, the applications are online and formed in a cluster.

What to do next

Set a password for the DefensePro application. Note that the application does not come online until you set a password. For more information, see the Radware DefensePro DDoS Mitigation User Guide on cisco.com.

Open UDP/TCP Ports and Enable vDP Web Services

The Radware APSolute Vision Manager interfaces communicate with the Radware vDP appliation using various UDP/TCP ports. In orer for the vDP application to communicate with the APSolute Vision Manager, you must ensure that these ports are accessible and not blocked by your firewall. For more information on which specific ports to open, see the following tables in the APSolute Vision User Guide:

- **Ports for APSolute Vision Server-WBM Communication and Operating System**

- **Communication Ports for APSolute Vision Server with Radware Devices**

In order for Radware APSolute Vision to manage the Virtual DefensePro application deployed on the FXOS chassis, you must enable the vDP web service using the FXOS CLI.

Procedure

Step 1   From the FXOS CLI, connect to the vDP application instance.

connect module *slot* **console**

connect vdp

Step 2   Enable vDP web services.

manage secure-web status set enable

Step 3   Exit the vDP application console and return to the FXOS module CLI.

Ctrl]

Configure TLS Crypto Acceleration

The following topics discuss TLS crypto acceleration, how to enable it, and how to view its status using the FMCr.

The following table maps the FTD and the FXOS version with the required TSL Crypto:

Note   When FXOS 2.6.1 is upgraded to FXOS 2.7.x and above, FTD 6.4 does not automatically enable crypto as 6.4 is not compatible with TLS crypto.

FTD	FXOS	Crypto
6.4	2.6	Support for only one container instance (Phase 1)
6.4	2.7 and above	NA
6.5 and above	2.7 and above	Support for upto 16 container instances (Phase 2)

About TLS Crypto Acceleration

The Firepower 4100/9300 support Transport Layer Security cryptographic acceleration, which performs Transport Layer Security/Secure Sockets Layer (TLS/SSL) encryption and decryption in hardware, which greatly accelerates the following:

- TLS/SSL encryption and decryption

- VPN, including TLS/SSL and IPsec

TLS cryptographic acceleration is automatically enabled on native instances and cannot be disabled. You can enable TLS crypto acceleration on up to 16 FTD container instances per security engine/module as well.

Guidelines and Limitations for TLS Crypto Acceleration

Keep the following in mind if your FTD has TLS crypto acceleration enabled.

Inspection engine failure

If the inspection engine is configured to preserve connections and the inspection engine fails unexpectedly, TLS/SSL traffic is dropped until the engine restarts.

This behavior is controlled by the FTD command **configure snort preserve-connection {enable | disable}** command.

HTTP-only performance

Using TLS crypto acceleration on an FTD container instance that is not decrypting traffic can affect performance. We recommend you enable TLS crypto acceleration *only* on FTD container instances that decrypt TLS/SSL traffic.

Federal Information Processing Standards (FIPS)

If TLS crypto acceleration and Federal Information Processing Standards (FIPS) are both enabled, connections with the following options fail:

- RSA keys less than 2048 bytes in size

- Rivest cipher 4 (RC4)

- Single Data Encryption Standard (single DES)

- Merkle–Damgard 5 (MD5)

- SSL v3

FIPS is enabled when you configure the FMC and FTDs to operate in a security certifications compliance mode. To allow connections when operating in those modes, you can either disable TLS crypto acceleration on the FTD container instance or you can configure web browsers to accept more secure options.

For more information:

- Common Criteria.

High Availability (HA) and clustering

If you have high availability (HA) or clustered FTDs, you must enable TLS crypto acceleration on each FTD individually. One device's TLS crypto acceleration configuration is not shared with the other devices in the HA pair or cluster.

TLS heartbeat

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by RFC6520. TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When an FTD managed by FMC with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the FTD takes the action specified by the FMC setting for **Decryption Errors** in the SSL policy's **Undecryptable Actions**:

- Block

- Block with reset

To determine whether applications are using TLS heartbeat, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

If TLS crypto acceleration is disabled on an FTD container instance, you can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) in the FMC to determine how to handle TLS heartbeats.

For more information about TLS heartbeat, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

TLS/SSL oversubscription

TLS/SSL oversubscription is a state where an FTD is overloaded with TLS/SSL traffic. Any FTD can experience TLS/SSL oversubscription but only the FTDs that support TLS crypto acceleration provide a configurable way to handle it.

When an FTD managed by an FMC with TLS crypto acceleration enabled is oversubscribed, any packet received by the FTD is acted on according to the setting for **Handshake Errors** in the SSL policy's **Undecryptable Actions**:

- Inherit default action

- Do not decrypt

- Block

- Block with reset

If the setting for **Handshake Errors** in the SSL policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

If a significant amount of oversubscription is occurring, you have the following options:

- Upgrade to an FTD with more TLS/SSL processing capacity.

- Change your SSL policies to add **Do Not Decrypt** rules for traffic that is not a high priority to decrypt.

For more information about TLS oversubscription, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

Passive and inline tap sets not supported

TLS/SSL traffic cannot be decrypted on passive or inline tap set interfaces when TLS crypto acceleration is enabled.

Enable TLS Crypto Acceleration for Container Instances

TLS crypto acceleration is automatically enabled when you deploy a logical instance as discussed in Add a Standalone FTD for the FMC, on page 207.

TLS crypto acceleration is enabled on all native instances and cannot be disabled.

View the Status of TLS Crypto Acceleration

This topic discusses how you can determine if TLS crypto acceleration is enabled.

Perform the following task in the FMC.

Procedure

Step 1	Log in to the FMC.
Step 2	Click **Devices** > **Device Management**.
Step 3	Click **Edit** ( ) to edit a managed device.
Step 4	Click **Device** page. TLS crypto acceleration status is displayed in the General section.

Enable FTD Link State Synchronization

The chassis can now synchronize the FTD operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The FTD application interface admin state is not considered. Without synchronization from FTD, data interfaces can be in an Up state physically before the FTD application has completely come online, for example, or can stay Up for a period of time after you initiate an FTD shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the FTD before the FTD can handle it.

This feature is disabled by default, and can be enabled per logical device in FXOS. This feature does not affect non-data interfaces such as Management or Cluster.

When you enable FTD link state synchronization, the **Service State** of an interface in FXOS will be synced with the administrative state of this interface in FTD. For example, if you shut down an interface in FTD, the Service State will show as Disabled. If you shut down the FTD application, all interfaces will show as Disabled. For Hardware Bypass interfaces, administratively shutting down the interface in FTD will set the Service State to Disabled; but shutting down the FTD application or other chassis-level shutdowns, including powering off, keeps the interface pair Enabled.

If you disable FTD link state synchronization, the Service State will always show as Enabled.

Note This feature is not supported for clustering, container instances, or an FTD with a Radware vDP decorator. It is also not supported for the ASA.

Procedure

Step 1 Choose **Logical Devices**, and then for the FTD logical device, choose **Enable Link State** from the drop-down list.
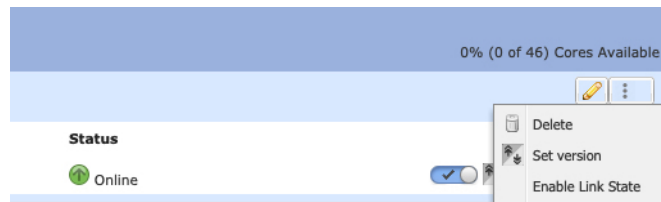
To disable this feature, choose **Disable Link State**.

Step 2 View the current interface state, as well as the last down reason.

show interface expand detail

Example:

```
Firepower # scope eth-uplink
Firepower /eth-uplink # scope fabric a
Firepower /eth-uplink/fabric # show interface expand detail
Interface:
    Port Name: Ethernet1/2
    User Label:
    Port Type: Data
    Admin State: Enabled
    Oper State: Up
    State Reason:
    flow control policy: default
    Auto negotiation: Yes
    Admin Speed: 1 Gbps
    Oper Speed: 1 Gbps
    Admin Duplex: Full Duplex
    Oper Duplex: Full Duplex
    Ethernet Link Profile name: default
    Oper Ethernet Link Profile name: fabric/lan/eth-link-prof-default
    Udld Oper State: Admin Disabled
    Inline Pair Admin State: Enabled
    Inline Pair Peer Port Name:
    Service State: Enabled
    Last Service State Down Reason: None
    Allowed Vlan: All
    Network Control Policy: default
    Current Task:
<...>
```

Manage Logical Devices

You can delete a logical device, convert an ASA to transparent mode, change the interface configuration, and perform other tasks on existing logical devices.

Connect to the Console of the Application

Use the following procedure to connect to the console of the application.

Procedure

Step 1   Connect to the module CLI using a console connection or a Telnet connection.

connect module *slot_number* {**console** | **telnet**}

To connect to the security engine of a device that does not support multiple security modules, always use **1** as the *slot_number*.

The benefits of using a Telnet connection is that you can have multiple sessions to the module at the same time, and the connection speed is faster.

Example:

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit

Firepower-module1>
```

Step 2   Connect to the application console. Enter the appropriate command for your device.

connect asa *name*

connect ftd *name*

connect vdp *name*

To view the instance names, enter the command without a name.

Example:

```
Firepower-module1> connect asa asa1
Connecting to asa(asa1) console... hit Ctrl + A + D  to return to bootCLI
[...]
asa>
```

Example:

```
Firepower-module1> connect ftd ftd1
```

```
Connecting to ftd(ftd-native) console... enter exit to return to bootCLI
[...]
>
```

Step 3   Exit the application console to the FXOS module CLI.

- ASA—Enter **Ctrl-a, d**

- FTD—Enter **exit**

- vDP—Enter **Ctrl-], .**

Step 4   Return to the supervisor level of the FXOS CLI.

Exit the console:

a) Enter **~**

 You exit to the Telnet application.

b) To exit the Telnet application, enter:

 telnet>**quit**

Exit the Telnet session:

a) Enter **Ctrl-], .**

Example

The following example connects to an ASA on security module 1 and then exits back to the supervisor
level of the FXOS CLI.

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit

Firepower-module1>connect asa asa1
asa> ~
telnet> quit
Connection closed.
Firepower#
```

Delete a Logical Device

Procedure

Step 1   Choose **Logical Devices** to open the Logical Devices page.

The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown instead.

Step 2   Click **Delete** for the logical device that you want to delete.

Step 3   Click **Yes** to confirm that you want to delete the logical device.

Step 4   Click **Yes** to confirm that you want to delete the application configuration.

Remove a Cluster Unit

The following sections describe how to remove units temporarily or permanently from the cluster.

Temporary Removal

A cluster unit will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status on the Firepower Chassis Manager **Logical Devices** page:

For FTD using the FMC, you should leave the device in the FMC device list so that it can resume full functionality after you reenable clustering.

- Disable clustering in the application—You can disable clustering using the application CLI. Enter the **cluster remove unit** *name* command to remove any unit other than the one you are logged into. The bootstrap configuration remains intact, as well as the last configuration synced from the control unit, so you can later re-add the unit without losing your configuration. If you enter this command on a data unit to remove the control unit, a new control unit is elected.

  When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster , the Management interface is disabled.

  To reenable clustering, on the ASA enter **cluster group** *name* and then **enable**. To reenable clustering, on the FTD enter **cluster enable**.

- Disable the application instance—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Slider enabled** (). You can later reenable it using the **Slider disabled** ().

  • Shut down the security module/engine—In the Firepower Chassis Manager on the **Security Module/Engine** page, click the **Power Off icon**.

  • Shut down the chassis—In the Firepower Chassis Manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

You can permanently remove a cluster member using the following methods.

For FTD using the FMC, be sure to remove the unit from the FMC device list after you disable clustering on the chassis.

  • Delete the logical device—In the Firepower Chassis Manager on the **Logical Devices** page, click the

  Delete (). You can then deploy a standalone logical device, a new cluster, or even add a new logical device to the same cluster.

  • Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new member of the cluster.

Delete an Application Instance that is not Associated with a Logical Device

When you delete a logical device, you are prompted as to whether you want to also delete the application configuration for the logical device. If you do not delete the application configuration, you will not be able to create a logical device using a different application until that application instance is deleted. You can use the following procedure to delete an application instance from a security module/engine when it is no longer associated with a logical device.

Procedure

Step 1   Choose **Logical Devices** to open the Logical Devices page.

The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown instead. Below the list of logical devices, you can see a list of application instances that are not associated with a logical device.

Step 2   Click **Delete** for the application instance that you want to delete.

Step 3   Click **Yes** to confirm that you want to delete the application instance.

Change an Interface on a FTD Logical Device

You can allocate or unallocate an interface, or replace a management interface on the FTD logical device. You can then sync the interface configuration in the FMC or the FDM.

Adding a new interface, or deleting an unused interface has minimal impact on the FTD configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the FTD configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Policies that refer to security zones are not affected. You can also edit

the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the FMC or the FDM.

For the FMC: Deleting an interface will delete any configuration associated with that interface.

For the FDM: You can migrate the configuration from one interface to another interface before you delete the old interface.

Before you begin

- Configure your interfaces, and add any EtherChannels according to Configure a Physical Interface, on page 169 and Add an EtherChannel (Port Channel), on page 170.

- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.

- If you want to replace the management or eventing interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the FTD device reboots (management interface changes cause a reboot), and you sync the configuration in the FMC or the FDM, you can add the (now unallocated) management interface to the EtherChannel as well.

- For clustering or High Availability, make sure you add or remove the interface on all units before you sync the configuration in the FMC or the FDM. We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. Note that new interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

Step 1 In the Firepower Chassis Manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Allocate a new data interface by selecting the interface in the **Data Ports** area.

Do not delete any interfaces yet.
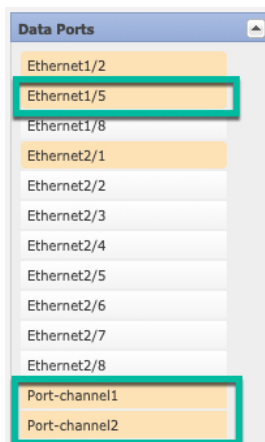
Step 4 Replace the management or eventing interface:

For these types of interfaces, the device reboots after you save your changes.

a) Click the device icon in the center of the page.

b) On the **General** or **Cluster Information** tab, choose the new **Management Interface** from the drop-down list.

c) On the **Settings** tab, choose the new **Eventing Interface** from the drop-down list.

d) Click **OK**.

If you change the IP address of the Management interface, then you must also change the IP address for the device in the FMC: go to **Devices** > **Device Management** > **Device/Cluster**. In the **Management** area, set the IP address to match the bootstrap configuration address.

Step 5   Click **Save**.

Step 6   Sync the interfaces in the FMC.

a) Log into the FMC.

b) Select **Devices** > **Device Management** and click **Edit** (✎) for your FTD device. The **Interfaces** page is selected by default.

c) Click the **Sync Device** button on the top left of the **Interfaces** page.

d) After the changes are detected, you will see a red banner on the **Interfaces** page indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.

e) If you plan to delete an interface, manually transfer any interface configuration from the old interface to the new interface.

Because you have not yet deleted any interfaces, you can refer to the existing configuration. You will have additional opportunity to fix the configuration after you delete the old interface and re-run the validation. The validation will show you all locations in which the old interface is still used.

f) Click **Validate Changes** to make sure your policy will still work with the interface changes.

If there are any errors, you need to change your policy and rerun the validation.

g) Click **Save**.

h) Select the devices and click **Deploy** to deploy the policy to the assigned devices. The changes are not active until you deploy them.

Step 7   Sync and migrate the interfaces in the FDM.

a) Log into the FDM.

b) Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.
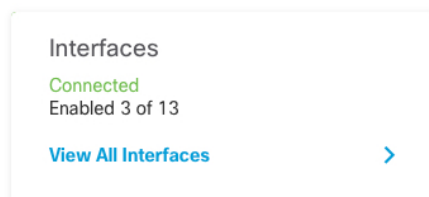
c) Click the **Scan Interfaces icon**.

d) Wait for the interfaces to scan, and then click **OK**.
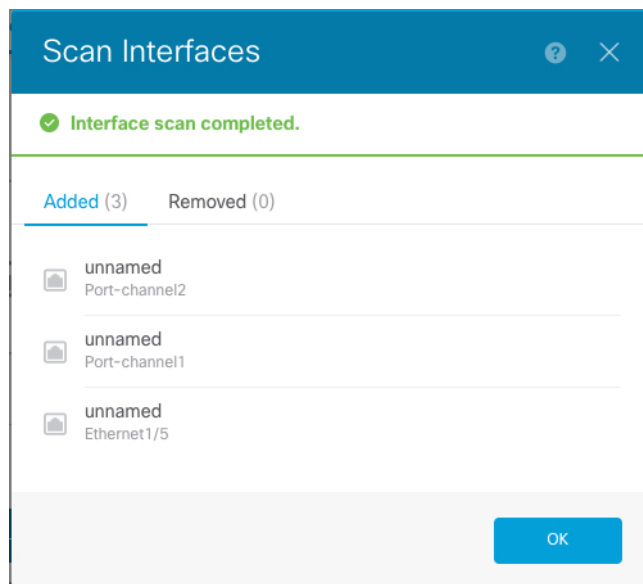
e) Configure the new interfaces with names, IP addresses, and so on.

If you want to use the existing IP address and name of an interface that you want to delete, then you need to reconfigure the old interface with a dummy name and IP address so that you can use those settings on the new interface.

f) To replace an old interface with a new interface, click the Replace icon for the old interface.

Replace icon

This process replaces the old interface with the new interface in all configuration settings that refer to the interface.

g) Choose the new interface from the **Replacement Interface** drop-down list.
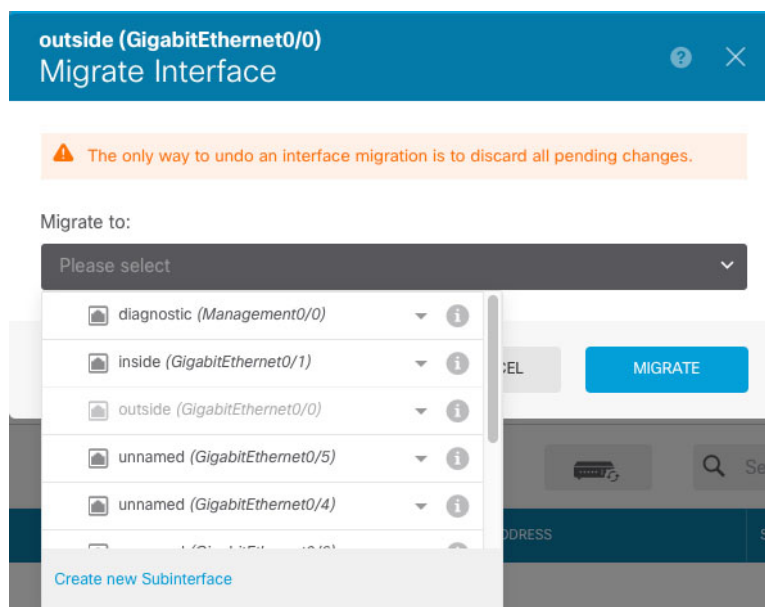
h) A message appears on the **Interfaces** page. Click the link in the message.

A task to migrate an interface was created.
Open the task list.
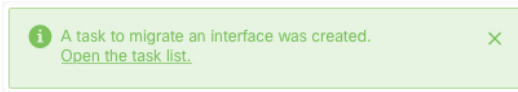
i)   Check the **Task List** to ensure that the migration was successful.

Task List

8 total	0 running	7 completed	1 failures

Delete all finished tasks

Name	Start Time	End Time	Status	Actions
Config migration from source interface outside to destination interface outside_2	06 Jun 2019 12:37 PM	06 Jun 2019 12:37 PM	✅ Migration is successful	

Step 8   In the Firepower Chassis Manager, unallocate a data interface by de-selecting the interface in the **Data Ports** area.

Data Ports

Ethernet1/2
Ethernet1/5
Ethernet1/8
Ethernet2/1
Ethernet2/2
Ethernet2/3
Ethernet2/4
Ethernet2/5
Ethernet2/6
Ethernet2/7
Ethernet2/8
Port-channel1
Port-channel2
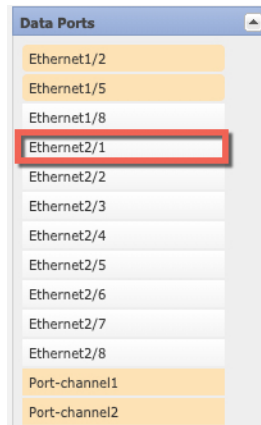
Step 9   Click **Save**.

Step 10   Sync the interfaces again in the FMC or the FDM.

Figure 16: FDM Scan Interfaces
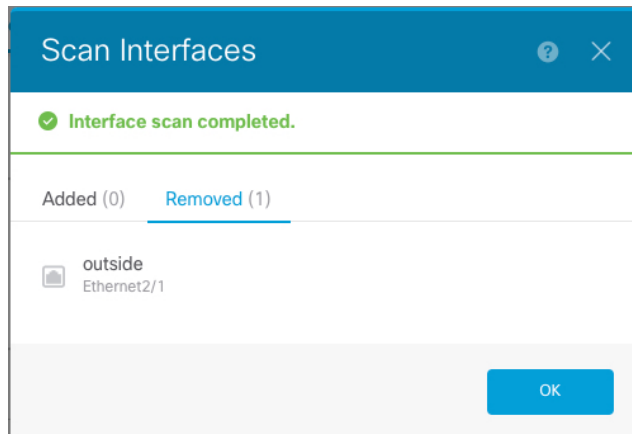
Change an Interface on an ASA Logical Device

You can allocate, unallocate, or replace a management interface on an ASA logical device. ASDM discovers the new interfaces automatically.

Adding a new interface, or deleting an unused interface has minimal impact on the ASA configuration. However, if you remove an allocated interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an allocated interface to an EtherChannel), and the interface is used in your security policy, removal will impact the ASA configuration. In this case, the ASA configuration retains the original commands so that you can make any necessary adjustments. You can manually remove the old interface configuration in the ASA OS.

> **Note** You can edit the membership of an allocated EtherChannel without impacting the logical device.

Before you begin

- Configure your interfaces and add any EtherChannels according to Configure a Physical Interface, on page 169 and Add an EtherChannel (Port Channel), on page 170.

- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.

- If you want to replace the management interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the ASA reloads (management interface changes cause a reload), you can add the (now unallocated) management interface to the EtherChannel as well.

- For clustering or failover, make sure you add or remove the interface on all units. We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. New interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

Step 1	In the Firepower Chassis Manager, choose **Logical Devices**.
Step 2	Click the **Edit** icon at the top right to edit the logical device.
Step 3	Unallocate a data interface by de-selecting the interface in the **Data Ports** area.
Step 4	Allocate a new data interface by selecting the interface in the **Data Ports** area.
Step 5	Replace the management interface:

For this type of interface, the device reloads after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General/Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) Click **OK**.

Step 6	Click **Save**.

Modify or Recover Bootstrap Settings for a Logical Device

You can modify bootstrap settings for a logical device. You can then immediately restart the application instance using those new settings or save the changes and restart the application instance using those new settings at a later time.

Procedure

Step 1	In the Firepower Chassis Manager, choose **Logical Devices**.
Step 2	Click the **Edit** icon at the top right to edit the logical device.
Step 3	Click the device icon in the center of the page.
Step 4	Modify the logical device settings as required.
Step 5	Click **OK**.
Step 6	Click **Restart Now** to save the changes and restart the application instance now. Click **Restart Later** to save the changes without restarting the application instance.

Note	If you selected **Restart Later**, you can restart the application instance when you are ready by clicking **Restart Instance** from the Logical Devices page.

Logical Devices Page

Use the **Logical Devices** page of the Firepower Chassis Manager to create, edit, and delete logical devices. The **Logical Devices** page includes an informational area for the logical device(s) installed on each Firepower 4100/9300 chassis security module/engine.

The header for each logical device area provides the following information:

• The unique name of the logical device.

• The logical device mode, either Standalone or Clustered.

• **Status**—Shows the state of the logical device:

  • ok—The logical device configuration is complete.

  • incomplete-configuration—The logical device configuration is incomplete.

Each logical device area provides the following information:

• **Application**—Shows the application running on the security module.

• **Version**—Shows the software version number of the application running on the security module.

> **Note** Updates to FTD logical devices are done using FMC and are not reflected on the **Logical Devices** > **Edit** and **System** > **Updates** pages in Firepower Chassis Manager. On these pages, the version shown indicates the software version (CSP image) that was used to create the FTD logical device.

• **Resource Profile**—Shows the resource profile assigned to the logical device/application instance.

• **Management IP**—Shows the local IP address assigned as the logical device Management IP.

• **Gateway**—Shows the network gateway address assigned to the application instance.

• **Management Port**—Shows the management port assigned to the application instance.

• **Status**—Shows the state of the application instance:

  • Online—The application is running and operating.

  • Offline—The application is stopped and inoperable.

  • Installing—The application installation is in progress.

  • Not Installed—The application is not installed.

  • Install Failed—The application installation failed.

  • Starting—The application is starting up.

  • Start Failed—The application failed to start up.

  • Started—The application started successfully, and is waiting for app agent heartbeat.

  • Stopping—The application is in the process of stopping.

  • Stop Failed—The application was unable to be brought offline.

  • Not Responding—The application is unresponsive.

  • Updating—The application software update is in progress.

  • Update Failed—The application software update failed.

  • Update Succeeded—The application software update succeeded.

> • Unsupported—The installed application is not supported.

If a security module is not present or is in a fault state, that information is shown in the status field. You can hover over the information icon to see additional information for a fault. For more information on security module faults, see About FXOS Security Modules/Security Engine, on page 271.

- **Expanded Information Area**—Shows additional attributes for the application instance that is currently running.

✎ **Note** If you modify the bootstrap settings for an application without immediately restarting the application instance, the Attributes fields show information for the application that is currently running and will not reflect the changes that were made until the application is restarted.

- **Ports**—Shows the names and types of interfaces assigned to the application instance.

- **Cluster Operation Status**—Shows the management URL assigned to the application instance.

- **Management IP/Firepower Management IP**—Shows the management IP address assigned to the application instance.

- **Cluster Role**—Shows the cluster role for the application instance, control or data.

- **Cluster IP**—Shows the IP address assigned to the application instance.

- **HA Role**—Shows the high-availability role for the application instance, active or standby.

- **Management URL**—Shows the URL of the management application assigned to the application instance.

- **UUID**—Shows the universally unique identifier for the application instance.

From the **Logical Devices** page of the Firepower Chassis Manager, you can perform the following functions on a logical device:

- **Refresh**—Refreshes the information on the Logical Devices page.

- **Add Device**—Allows you to create a logical device.

- **Edit**—Allows you to edit an existing logical device.

- **Set Version**—Allows you to upgrade or downgrade the software on a logical device.

- **Delete**—Deletes a logical device.

- **Show Configuration**—Opens a dialog box showing the configuration information in JSON format for a logical device or cluster. You can copy the configuration information and use it when creating additional devices that are part of a cluster.

- **Enable/Disable**—Enables or disables an application instance.

- **Upgrade/Downgrade**—Allows you to upgrade or downgrade an application instance.

- **Restart Instance**—Allows you to restart the application instance. If you have modified the device bootstrap information but have not yet restarted the application instance, you can click Restart Instance

to clear the existing management bootstrap information and restart the application instance using the new bootstrap information.

- **Reinstall Instance**—Allows you to reinstall the application instance.

- **Go To Device Manager**—Provides a link to the FMC or ASDM defined for the application instance.

- **Enable/Disable Link State**—Enable or disable FTD link state synchronization. For more information, see .

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

Spanned EtherChannel Routed Mode Example with Site-Specific MAC Addresses

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and an inside network at each site (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the inside and outside networks. Each EtherChannel is spanned across all chassis in the cluster.

The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters blocking the global MAC address to prevent traffic from traversing the DCI to the other site when the traffic is destined for the cluster. If the cluster nodes at one site become unreachable, you must remove the filters so traffic can be sent to the other site's cluster nodes. You should use VACLs to filter the global MAC address.Be sure to disable ARP inspection.

The cluster acts as the gateway for the inside networks. The global virtual MAC, which is shared across all cluster nodes, is used only to receive packets. Outgoing packets use a site-specific MAC address from each DC cluster. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address.

In this scenario:

- All egress packets sent from the cluster use the site MAC address and are localized at the data center.

- All ingress packets to the cluster are sent using the global MAC address, so they can be received by any of the nodes at both sites; filters at the OTV localize the traffic within the data center.
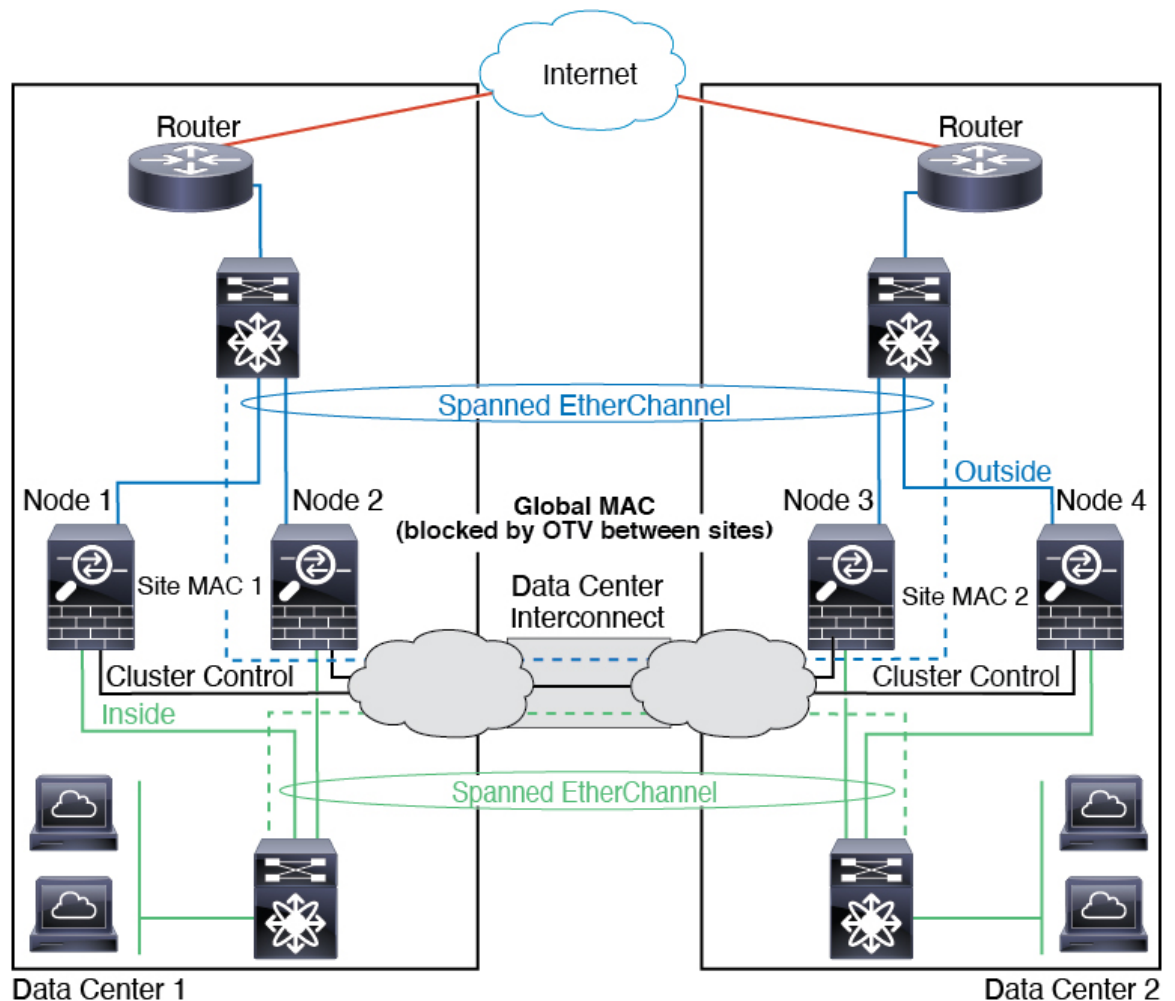
Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. By assigning a higher cost route across the DCI, traffic stays within each data center unless all cluster members at a given site go down. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the cluster members at the other site.

The implementation of the switches at each site can include:

- Inter-site VSS, vPC, StackWise, or StackWise Virtual—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster nodes at each Data Center to only connect to the local switch, while the redundant switch traffic goes across the DCI. In this case, connections

are for the most part kept local to each datacenter. You can optionally connect each node to both switches across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.

- Local VSS, vPC, StackWise, or StackWise Virtual at each site—For better switch redundancy, you can install 2 separate redundant switch pairs at each site. In this case, although the cluster nodes still have a spanned EtherChannel with Data Center 1 chassis connected only to both local switches, and Data Center 2 chassis connected to those local switches, the spanned EtherChannel is essentially "split." Each local redundant switch system sees the spanned EtherChannel as a site-local EtherChannel.

Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. A good practice to avoid unintended MAC address flapping is to statically add the gateway routers real MAC addresses to the ASA MAC address table. Without these entries, if the gateway at site 1 communicates with the gateway at site 2, that traffic might pass through the ASA and attempt to reach

site 2 from the inside interface and cause problems. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.

History for Logical Devices

Feature Name	Platform Releases	Feature Information
Synchronization between the FTD operational link state and the physical link state	2.9.1	The chassis can now synchronize the FTD operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The FTD application interface admin state is not considered. Without synchronization from FTD, data interfaces can be in an Up state physically before the FTD application has completely come online, for example, or can stay Up for a period of time after you initiate an FTD shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the FTD before the FTD can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS. **Note** This feature is not supported for clustering, container instances, or an FTD with a Radware vDP decorator. It is also not supported for the ASA. New/Modified Firepower Chassis Manager screens: **Logical Devices > Enable Link State** New/Modified FXOS commands: **set link-state-sync enabled**, **show interface expand detail**

Feature Name	Platform Releases	Feature Information
FTD configuration backup and restore using FMC for container instances	2.9.1	You can now use the FMC backup/restore tool on an FTD container instance. New/Modified FMC screens: **System > Tools > Backup/Restore > Managed Device Backup** New/Modified FTD CLI commands: **restore** Supported platforms: Firepower 4100/9300 **Note** Requires Firepower 6.7.
Multi-instance clustering	2.8.1	You can now create a cluster using container instances. On the Firepower 9300, you must include one container instance on each module in the cluster. You cannot add more than one container instance to the cluster per security engine/module. We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. New/modified screens: • **Logical Devices** > **Add Cluster** • **Interfaces** > **All Interfaces** > **Add New** drop-down menu **> Subinterface > Type** field **Note** Requires Firepower 6.6 or later.
Support for FTD with FDM	2.7.1	You can now deploy a native FTD instance and specify FDM management. Container instances are not supported. New/modified Firepower Chassis Manager screens: **Logical Devices** > **Add Device** > **Settings** > **Management type of application instance** **Note** Requires FTD 6.5 or later.
TLS crypto acceleration for multiple container instances	2.7.1	TLS crypto acceleration is now supported on multiple container instances (up to 16) on a Firepower 4100/9300 chassis. Previously, you could enable TLS crypto acceleration for only *one* container instance per module/security engine. New instances have this feature enabled by default. However, the upgrade does *not* enable acceleration on existing instances. Instead, use the **enter hw-crypto** and then the **set admin-state enabled** FXOS commands. New/Modified Firepower Chassis Manager screens: **Logical Devices** > **Add Device** > **Settings** > **Hardware Crypto** drop-down menu **Note** Requires FTD 6.5 or later.

Feature Name	Platform Releases	Feature Information
Firepower 4115, 4125, and 4145	2.6.1	We introduced the Firepower 4115, 4125, and 4145. **Note** Requires ASA 9.12(1). Firepower 6.4.0 requires FXOS 2.6.1.157. No modified screens.
Firepower 9300 SM-40, SM-48, and SM-56 support	2.6.1	We introduced the following three security modules: SM-40, SM-48, and SM-56. **Note** The SM-40 and SM-48 require ASA 9.12(1). The SM-56 requires ASA 9.12(2) and FXOS 2.6.1.157. All modules require FTD 6.4 and FXOS 2.6.1.157. No modified screens.
Support for ASA and FTD on separate modules of the same Firepower 9300	2.6.1	You can now deploy ASA and FTD logical devices on the same Firepower 9300. **Note** Requires ASA 9.12(1). Firepower 6.4.0 requires FXOS 2.6.1.157. No modified screens.
For the FTD bootstrap configuration, you can now set the NAT ID for the FMC in the Firepower Chassis Manager	2.6.1	You can now set the FMC NAT ID in the Firepower Chassis Manager. Previously, you could only set the NAT ID within the FXOS CLI or FTD CLI. Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration. New/Modified screens: **Logical Devices** > **Add Device** > **Settings** > **Firepower Management Center NAT ID** field
Support for SSL hardware acceleration on one FTD container instance on a module/security engine	2.6.1	You can now enable SSL hardware acceleration for one container instance on a module/security engine. SSL hardware acceleration is disabled for other container instances, but enabled for native instances. See the FMC configuration guide for more information. New/Modified commands: **config hwCrypto enable**, **show hwCrypto** No modified screens.

Feature Name	Platform Releases	Feature Information
Multi-instance capability for FTD	2.4.1	You can now deploy multiple logical devices, each with a FTD container instance, on a single security engine/module. Formerly, you could only deploy a single native application instance. Native instances are still also supported. For the Firepower 9300, you can use a native instance on some modules, and container instances on the other module(s).
		To provide flexible physical interface use, you can create VLAN subinterfaces in FXOS and also share interfaces between multiple instances. When you deploy a container instance, you must specify the number of CPU cores assigned; RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance. This resource management lets you customize performance capabilities for each instance.
		You can use High Availability using a container instance on 2 separate chassis; for example, if you have 2 chassis, each with 10 instances, you can create 10 High Availability pairs. Clustering is not supported.
		Note Multi-instance capability is similar to ASA multiple context mode, although the implementation is different. Multiple context mode partitions a single application instance, while multi-instance capability allows independent container instances. Container instances allow hard resource separation, separate configuration management, separate reloads, separate software updates, and full FTD feature support. Multiple context mode, due to shared resources, supports more contexts on a given platform. Multiple context mode is not available on the FTD.
		Note Requires FTD Version 6.3 or later.
		New/Modified Firepower Chassis Manager screens:
		Overview > **Devices**
		Interfaces > **All Interfaces** > **Add New** drop-down menu **> Subinterface**
		Interfaces > **All Interfaces** > **Type**
		Logical Devices > **Add Device**
		Platform Settings > **Mac Pool**
		Platform Settings > **Resource Profiles**
		New/Modified FMC screens:
		Devices > **Device Management** > **Edit** icon **> Interfaces** tab
Support for transparent mode deployment for an ASA logical device	2.4.1	You can now specify transparent or routed mode when you deploy the ASA.
		New/modified Firepower Chassis Manager screens:
		Logical Devices > **Add Device** > **Settings**
		New/Modified options: **Firewall Mode** drop-down list

Feature Name	Platform Releases	Feature Information
Cluster control link customizable IP Address	2.4.1	By default, the cluster control link uses the 127.2.0.0/16 network. You can now set the network when you deploy the cluster in FXOS. The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: 127.2.*chassis_id.slot_id*. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. Therefore, you can now set a custom /16 subnet for the cluster control link in FXOS except for loopback (127.0.0.0/8) and multicast (224.0.0.0/4) addresses. New/modified screens: **Logical Devices** > **Add Device** > **Cluster Information** > **CCL Subnet IP** field
For the FTD bootstrap configuration, you can now set the NAT ID for the FMC at the FXOS CLI	2.4.1	You can now set the FMC NAT ID at the FXOS CLI. Previously, you could only set the NAT ID within the FTD CLI. Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration. New/Modified commands: **enter bootstrap-key NAT_ID**
Inter-site clustering improvement for the ASA	2.1.1	You can now configure the site ID for each Firepower 4100/9300 chassis when you deploy the ASA cluster. Previously, you had to configure the site ID within the ASA application; this new feature eases initial deployment. Note that you can no longer set the site ID within the ASA configuration. Also, for best compatibility with inter-site clustering, we recommend that you upgrade to ASA 9.7(1) and FXOS 2.1.1, which includes several improvements to stability and performance. We modified the following screen: **Logical Devices** > **Configuration**
Inter-chassis clustering for 6 FTD modules on the Firepower 9300	2.1.1	You can now enable inter-chassis clustering for the FTD on the Firepower 9300. You can include up to 6 modules. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules. We modified the following screen: **Logical Devices** > **Configuration**
Support for FTD clustering on the Firepower 4100	2.1.1	You can cluster up to 6 chassis in an FTD cluster.
Support for 16 Firepower 4100 chassis in an ASA cluster	2.0.1	You can cluster up to 16 chassis in an ASA cluster.
Support for ASA clustering on the Firepower 4100	1.1.4	You can cluster up to 6 chassis in an ASA cluster.

Feature Name	Platform Releases	Feature Information
Support for intra-chassis clustering on the FTD on the Firepower 9300	1.1.4	The Firepower 9300 supports intra-chassis clustering with the FTD application. We modified the following screen: **Logical Devices** > **Configuration**
Inter-chassis clustering for 16 ASA modules on the Firepower 9300	1.1.3	You can now enable inter-chassis clustering for the ASA. You can include up to 16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. We modified the following screen: **Logical Devices** > **Configuration**
Intra-chassis Clustering for the ASA on the Firepower 9300	1.1.1	You can cluster all ASA security modules within the Firepower 9300 chassis. We introduced the following screen: **Logical Devices** > **Configuration**

Security Module/Engine Management

About FXOS Security Modules/Security Engine

From the Security Modules/Security Engine page of the Firepower Chassis Manager, you can view the status of a security module/engine and can perform various functions on the security module/engine:

The Security Modules/Security Engine page provides the following information:

- Hardware State—Shows the state of the security module/engine hardware.

    - Up—The security module/engine has powered up successfully and is not showing any hardware faults, even if the security module/engine does not have a logical device associated with it.

    - Booting Up—The security module/engine is in the process of powering up.

    - Restart—The security module/engine is in the process of being restarted.

    - Down—The security module/engine is not powered on or a hardware fault is preventing the security module/engine from starting successfully.

    - Mismatch—The security module has been decommissioned or a new security module was installed into the slot. Use the Acknowledge function to return the security module to a functioning state.

    - Empty—A security module is not installed in that slot.

- Service State—Shows the state of the software on the security module/engine:

    - Not-available—The security module has been removed from the chassis slot. Reinstall the security module to return it to its normal operational state.

    - Online—The security module/engine is installed and is in normal operation mode.

- Not Responding—The security module/engine is unresponsive.

- Token Mismatch—Indicates that a security module other than the one previously configured has been installed into the chassis slot. This could also be caused by a software installation error. Use the Reinitialize function to return the security module to a functioning state.

- Fault—The security module/engine is in a fault state. Review the system fault listing for more information about what might be causing the fault state. You can also hover over the information icon for a fault to see additional information.

Security Module Faults

- Failsafe Mode—the security module is in failsafe mode. Applications are blocked from starting in this mode. Connect to the security module for troubleshooting or to disable failsafe mode. The app-instance can also be deleted.

- HDD Error—the security module disk drive has errors. Please verify that the disk drive is present and replace the faulty disk drive if the fault does not clear.

- Filesystem Error—disk partitions on the security module are not compatible. Reboot the security module for possible recovery. If the fault persists, please reinitialize the slot after backing up your data on an external device.

- Format Failure—automatic format failed on the security module disk drive. Reinitialize the security module to reformat.

- Power—Shows the power status of the security module/engine:

  - On—Use the Power off/on function to toggle the power status for the security module/engine.

  - Off—Use the Power off/on function to toggle the power status for the security module/engine.

- Application—Shows the logical device type that is installed on the security module/engine.

From the Security Modules/Security Engine page of the Firepower Chassis Manager, you can perform the following functions on a security module/engine:

- Decommission (security modules only)—Decommissioning a security module places the security module into maintenance mode. You can also decommission and then acknowledge a security module in order to correct certain fault states. See Decommissioning a Security Module, on page 273.

- Acknowledge—Brings newly installed security modules online. See Acknowledge a Security Module/Engine, on page 273.

- Power Cycle—Restarts the security module/engine. See Power-Cycling a Security Module/Engine, on page 274.

- Reinitialize—Reformats the security module/engine hard disk, removing all deployed applications and configurations from the security module/engine, and then restarts the system. After reinitialization is complete, if a logical device is configured for the security module/engine, the FXOS will reinstall the application software, redeploy the logical device, and auto start the application. See Reinitializing a Security Module/Engine, on page 274.

> **Warning**   All application data on the security module/engine is deleted during reinitialization. Please back up all application data before reinitializing a security module/engine.

- Power off/on—Toggles the power state for the security module/engine. See Power-Cycling a Security Module/Engine, on page 274.

Decommissioning a Security Module

When you decommission a security module, the security module object is deleted from the configuration and the security module becomes unmanaged. Any logical devices or software running on the security module will become inactive.

You can decommission a security module if you want to temporarily discontinue use of the security module.

Procedure

Step 1   Choose **Security Modules** to open the Security Modules page.

Step 2   To decommission a security module, click **Decommission** for that security module.

Step 3   Click **Yes** to verify that you want to decommission the specified security module.

Acknowledge a Security Module/Engine

When a new security module is installed into the chassis, or when an existing module is replaced with one with a different product ID (PID), you must acknowledge the security module before you can begin using it.

If the security module is showing a status of "mismatch" or "token mismatch," this is an indication that the security module installed in the slot has data on it that does not match what was previously installed in the slot. If the security module has existing data on it and you are sure you want to use it in the new slot (in other words, the security module wasn't inadvertently installed into the wrong slot), you must reinitialize the security module before you can deploy a logical device to it.

Procedure

Step 1   Choose **Security Modules/Security Engine** to open the Security Modules/Security Engine page.

Step 2   Click **Acknowledge** for the security module/engine that you want to acknowledge.

Step 3   Click **Yes** to verify that you want to acknowledge the specified security module/engine.

Power-Cycling a Security Module/Engine

Follow these steps to power-cycle a security module/engine.

Procedure

Step 1   Choose **Security Modules/Security Engine** to open the Security Modules/Security Engine page.

Step 2   Click **Power Cycle** for the security module/engine that you want to reboot.

Step 3   Do one of the following:

- Click **Safe Power Cycle** to have the system wait for up to five minutes for the application running on the security module/engine to shut down before the system power-cycles the specified security module/engine.

- Click **Power Cycle Immediately** to have the system power-cycle the specified security module/engine immediately.

Reinitializing a Security Module/Engine

When a security module/engine is reinitialized, the security module/engine hard disk is formatted and all installed application instances, configurations, and data are removed. After reinitialization has completed, if a logical device is configured for the security module/engine, FXOS will reinstall the application software, redeploy the logical device, and auto start the application.

⚠

Caution   All application data on the security module/engine is deleted during reinitialization. Back up all application data before reinitializing a security module/engine.

Procedure

Step 1   Choose **Security Modules/Security Engine** to open the Security Modules/Security Engine page.

Step 2   Click **Reinitialize** for the security module/engine that you want to reinitialize.

Step 3   Click **Yes** to verify that you want to reinitialize the specified security module/engine.

The security module/engine is restarted and all data on the security module is deleted. This process can take several minutes.

Acknowledge a Network Module

When a new network module is installed into the chassis, or when an existing module is replaced with one with a different product ID (PID), you must acknowledge the network module before you can begin using it.

Procedure

Step 1   Enter `scope fabric-interconnect` mode:

`scope fabric-interconnect`

Step 2   Enter the **`acknowledge`** command after installing a new module or replacing a network module with another that is not the same type (that is, with a different PID):

`acknowledge`

Example:
```
FPR1 /fabric-interconnect # acknowledge
  fault  Fault
  slot   Card Config Slot Id  <======
```

Step 3   Enter the **`acknowledge slot`** to acknowledge the inserted slot.

`acknowledge slot`

Example:
```
FPR1 /fabric-interconnect # acknowledg slot 2
  0-4294967295  Slot Id
```

Step 4   Commit the configuration:

`commit-buffer`

Taking a Network Module Offline or Online

Follow these steps to use CLI commands to take a network module offline, or to bring it back online; used for example, when performing module online insertion and removal (OIR).

Note

- If removing and replacing a network module, follow the instructions in the "Maintenance and Upgrades" chapter of the appropriate Install Guide for your device. See https://www.cisco.com/c/en/us/support/security/firepower-ngfw/products-installation-guides-list.html.

- If performing a network module online insertion and removal (OIR) on a 8 port 1G Copper FTW Network Module (FPR-NM-8X1G-F FTW), note that the network module LED stays off until you bring the card online using this procedure. The LED first flashes amber, then changes to green once the network module is discovered and the application comes online.

Note

If you remove a FTW network module and acknowledge the slot, the network module ports are deleted from the FTD logical device. In this case, you must delete the hardware bypass inline set configurations using FMC before reinserting the network module. After reinserting the network module, you must:

- Configure the network module ports as adminstrative online state using Firepower Chassis Manager or FXOS Command Line Interface (CLI).

- Add the network module ports to the FTD logical device and reconfigure the ports using FMC.

If you remove the network module without acknowledging the slot, the inline set configuration is retained and ports display as down in FMC. Once you reinsert the network module, the previous configuration is restored.

For more information about hardware bypass for inline sets, see Hardware Bypass Pairs, on page 156.

Procedure

Step 1　Use the following commands to enter `/fabric-interconnect` mode and then enter `/card` mode for the module to be taken offline:

```
scope fabric-interconnect a
scope card ID
```

Step 2　You can use the `show detail` command to view information about this card, including its current status.

Step 3　To take the module offline, enter:

```
set adminstate offline
```

Step 4　Enter the `commit-buffer` command to save the configuration change.

You can use the `show detail` command again to confirm that the module is offline.

Step 5　To bring the network module back online, enter:

```
set adminstate online
commit-buffer
```

Example

```
FP9300-A# scope fabric-interconnect a
FP9300-A /fabric-interconnect # scope card 2
FP9300-A /fabric-interconnect/card # show detail

Fabric Card:
    Id: 2
    Description: Firepower 4x40G QSFP NM
    Number of Ports: 16
    State: Online
    Vendor: Cisco Systems, Inc.
    Model: FPR-NM-4X40G
    HW Revision: 0
    Serial (SN): JAD191601DE
    Perf: N/A
    Admin State: Online
    Power State: Online
    Presence: Equipped
    Thermal Status: N/A
    Voltage Status: N/A
FP9300-A /fabric-interconnect/card # set adminstate offline
FP9300-A /fabric-interconnect/card* # commit-buffer
FP9300-A /fabric-interconnect/card # show detail

Fabric Card:
    Id: 2
    Description: Firepower 4x40G QSFP NM
    Number of Ports: 16
    State: Offline
    Vendor: Cisco Systems, Inc.
    Model: FPR-NM-4X40G
    HW Revision: 0
    Serial (SN): JAD191601DE
    Perf: N/A
    Admin State: Offline
    Power State: Off
    Presence: Equipped
    Thermal Status: N/A
    Voltage Status: N/A
FP9300-A /fabric-interconnect/card #
```

Blade Health Monitoring

Failsafe is engaged on a security module or engine when a specified number of unexpected application restarts are detected on a blade to prevent an endless boot loop condition, which can cause further side effects in a redundant HA or Cluster deployment.

Blade platform performs health checks periodically and reports it to the MIO. If the blade is in failed state, you will be notified with faults and error messages.

Faults and Error Messages

You can view the faults and error messages in the Overview page of the platform if there are any issues with the blade.

- Overview page—Security Module shows the fault symbol with the operational state as Fault.

- Security Module page—Service State in the blade will show as Fault. The 'i' icon displays the error message when you hover over.

- Logical Devices page—If logical devices are available and the security module goes faulty, the "i" icon displays the error message when you hover over.

Note   You can configure and manage failsafe settings from FXOS CLI.

CHAPTER 12

Configuration Import/Export

About Configuration Import/Export

You can use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis to a remote server or your local computer. You can later import that configuration file to quickly apply the configuration settings to your Firepower 4100/9300 chassis to return to a known good configuration or to recover from a system failure.

Guidelines and Restrictions

- Beginning with FXOS 2.6.1, the encryption key is now configurable. You must set the encryption key before you can export a configuration. The same encryption key must be set on the system when importing that configuration. If you modified the encryption key so that it no longer matches what was used during export, the import operation will fail. Make sure you keep track of the encryption key for each exported configuration.

- Do not modify the contents of the configuration file. If a configuration file is modified, configuration import using that file might fail.

- Application-specific configuration settings are not contained in the configuration file. You must use the configuration backup tools provided by the application to manage application-specific settings and configurations.

- When you import a configuration to the Firepower 4100/9300 chassis, all existing configuration on the Firepower 4100/9300 chassis (including any logical devices) are deleted and completely replaced by the configuration contained in the import file.

- Except in an RMA scenario, we recommend you only import a configuration file to the same Firepower 4100/9300 chassis where the configuration was exported.

- The platform software version of the Firepower 4100/9300 chassis where you are importing should be the same version as when the export was taken. If not, the import operation is not guaranteed to be

successful. We recommend you export a backup configuration whenever the Firepower 4100/9300 chassis is upgraded or downgraded.

• The Firepower 4100/9300 chassis where you are importing must have the same Network Modules installed in the same slots as when the export was taken.

• The Firepower 4100/9300 chassis where you are importing must have the correct software application images installed for any logical devices defined in the export file that you are importing.

• If the configuration file being imported contains a logical device whose application has an End-User License Agreement (EULA), you must accept the EULA for that application on the Firepower 4100/9300 chassis before you import the configuration or the operation will fail.

• To avoid overwriting existing backup files, change the file name in the backup operation or copy the existing file to another location.

Note   You must backup the logicl APP separately as the FXOS import/export will backup only the FXOS configuration. The FXOS configuration import will cause logical device reboot and it rebuilds the device with the factory default configuration.

Setting an Encryption Key for Configuration Import/Export

When exporting configurations, FXOS encrypts sensitive data such as passwords and keys.

Beginning with FXOS 2.6.1, the encryption key is now configurable. You must set the encryption key before you can export a configuration. The same encryption key must be set on the system when importing that configuration. If you have modified the encryption key so that it no longer matches what was used during export, the import operation will fail. Make sure that you keep track of the encryption key that is used for each exported configuration.

You can set the encryption key on either the Export page or the Import page. However, once set, the same key is used for both exporting and importing.

If you are importing a configuration into FXOS 2.6.1 or later that was exported from an FXOS release prior to 2.6.1, the system will not check the encryption key and will allow the import.

Note   If the platform software version to which you are importing is not the same version as when the export was taken, the import operation is not guaranteed to be successful. We recommend that you export a backup configuration whenever the Firepower 4100/9300 chassis is upgraded or downgraded.

Use the 'Set Version' option and export a backup configuration whenever the FTD logical appliance is upgraded to a new software so that the new startup version matches the software release of the upgraded version.

Procedure

Step 1   Choose **System** > **Configuration** > **Export**.

Step 2 Under **Encryption**, enter a key to use for encrypting/decrypting sensitive data in the **Key** field. The encryption key must be 4-40 characters in length.

Step 3 Click **Save Key**.

The encryption key is set and will be used to encrypt/decrypt sensitive data when exporting and importing configurations. The system displays *Set:Yes* next to the **Key** field to indicate that the encryption key has been set.

Exporting an FXOS Configuration File

Use the configuration export feature to export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis to a remote server or your local computer.

Before you begin

Review the About Configuration Import/Export.

Procedure

Step 1 Choose **System** > **Configuration** > **Export**.

Step 2 To export a configuration file to your local computer, click **Export Locally**.
The configuration file is created and, depending on your browser, the file might be automatically downloaded to your default download location or you might be prompted to save the file.

Step 3 To export the configuration file to a previously configured remote server, click **Export** for the Remote Configuration you want to use.
The configuration file is created and exported to the specified location.

Step 4 To export the configuration file to a new remote server:

a) Under On-Demand Export, click **Add On-Demand Configuration**.

b) Choose the protocol to use when communicating with the remote server. This can be one of the following: FTP, TFTP, SCP, or SFTP.

c) Enter the hostname or IP address of the location where the backup file should be stored. This can be a server, storage array, local drive, or any read/write media that the Firepower 4100/9300 chassis can access through the network.

   If you use a hostname rather than an IP address, you must configure a DNS server.

d) If you are using a non-default port, enter the port number in the **Port** field.

e) Enter the username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.

f) Enter the password for the remote server username. This field does not apply if the protocol is TFTP.

g) In the **Location** field, enter the full path to where you want the configuration file exported including the filename.

h) Click **OK**.
The Remote Configuration is added to the On-Demand Export table.

i) Click **Export** for the Remote Configuration you want to use.

The configuration file is created and exported to the specified location.

Scheduling Automatic Configuration Export

Use the scheduled export feature to automatically export an XML file containing logical device and platform configuration settings for your Firepower 4100/9300 chassis to a remote server or your local computer. You can schedule the exports to be run daily, weekly, or every two weeks. The configuration export will be executed according to the schedule based on the when the scheduled export feature is enabled. So, for example, if you enable weekly scheduled export on a Wednesday at 10:00pm, the system will trigger a new export every Wednesday at 10:00pm.

Please review the About Configuration Import/Export for important information about using the configuration export feature.

Procedure

Step 1	Choose **System** > **Configuration** > **Export**.
Step 2	Click **Schedule Export**. You see the **Configure Scheduled Export** dialog box.
Step 3	Choose the protocol to use when communicating with the remote server. This can be one of the following: FTP, TFTP, SCP, or SFTP.
Step 4	To enable the scheduled export, check the **Enable** check box.

> **Note** You can enable or disable the schedule export at a later time using this check box; however, you will need to specify the password again when enabling or disabling the scheduled export.

Step 5	Enter the hostname or IP address of the location where the backup file should be stored. This can be a server, storage array, local drive, or any read/write media that the Firepower 4100/9300 chassis can access through the network. If you use a hostname rather than an IP address, you must configure a DNS server.
Step 6	If you are using a non-default port, enter the port number in the **Port** field.
Step 7	Enter the username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.
Step 8	Enter the password for the remote server username. This field does not apply if the protocol is TFTP.
Step 9	In the **Location** field, enter the full path to where you want the configuration file exported including the filename. If you omit the filename, the export procedure assigns a name to the file.
Step 10	Choose the schedule on which you would like to have the configuration automatically exported. This can be one of the following: Daily, Weekly, or BiWeekly.
Step 11	Click **OK**. The scheduled export is created. If you enabled the scheduled export, the system will automatically export a configuration file to the specified location according to the schedule that you selected.

Setting a Configuration Export Reminder

Use the Export Reminder feature to have the system generate a fault when a configuration export hasn't been executed in a certain number of days.

By default, the export reminder is enabled with a frequency of 30 days.

Note　If the reminder frequency is smaller than the number of days in the scheduled export policy (daily, weekly, or bi-weekly), you will receive an export-reminder fault message ("`Config backup may be outdated`"). For example, if your export schedule is weekly, and the reminder frequency is five days, this fault message will be issued every five days if no configuration has been exported in that time.

Procedure

Step 1	Choose **System** > **Configuration** > **Export**.
Step 2	To enable the configuration export reminder, check the check box under **Reminder to trigger an export**.
Step 3	Enter the number of days, between 1 and 365, that the system should wait between configuration exports before generating a reminder fault.
Step 4	Click **Save Reminder**.

Importing a Configuration File

You can use the configuration import feature to apply configuration settings that were previously exported from your Firepower 4100/9300 chassis. This feature allows you to return to a known good configuration or to recover from a system failure.

Before you begin

Review the About Configuration Import/Export.

Procedure

Step 1	Choose **System** > **Tools** > **Import/Export**.
Step 2	To import from a local configuration file:
	a) Click **Choose File** to navigate to and select the configuration file that you want to import.
	b) Click **Import**. A confirmation dialog box opens asking you to confirm that you want to proceed and warn you that the chassis might need to restart.
	c) Click **Yes** to confirm that you want to import the specified configuration file.

The existing configuration is deleted and the configuration specified in the import file is applied to the Firepower 4100/9300 chassis. If there is a breakout port configuration change during the import, the Firepower 4100/9300 chassis will need to restart.

Step 3   To import the configuration file from a previously configured remote server:

a)   In the Remote Import table, click **Import** for the Remote Configuration you want to use.
A confirmation dialog box opens asking you to confirm that you want to proceed and warning you that the chassis might need to restart.

b)   Click **Yes** to confirm that you want to import the specified configuration file.
The existing configuration is deleted and the configuration specified in the import file is applied to the Firepower 4100/9300 chassis. If there is a breakout port configuration change during the import, the Firepower 4100/9300 chassis will need to restart.

Step 4   To import from a configuration file on a new remote server:

a)   Under Remote Import, click **Add Remote Configuration**.

b)   Choose the protocol to use when communicating with the remote server. This can be one of the following: FTP, TFTP, SCP, or SFTP.

c)   If you are using a non-default port, enter the port number in the **Port** field.

d)   Enter the hostname or IP address of the location where the backup file is stored. This can be a server, storage array, local drive, or any read/write media that the Firepower 4100/9300 chassis can access through the network.

If you use a hostname rather than an IP address, you must configure a DNS server.

e)   Enter the username the system should use to log in to the remote server. This field does not apply if the protocol is TFTP.

f)   Enter the password for the remote server username. This field does not apply if the protocol is TFTP.

g)   In the **File Path** field, enter the full path to the configuration file including the file name.

h)   Click **Save**.
The Remote Configuration is added to the Remote Import table.

i)   Click **Import** for the Remote Configuration you want to use.
A confirmation dialog box opens asking you to confirm that you want to proceed and warning you that the chassis might need to restart.

j)   Click **Yes** to confirm that you want to import the specified configuration file.
The existing configuration is deleted and the configuration specified in the import file is applied to the Firepower 4100/9300 chassis. If there is a breakout port configuration change during the import, the Firepower 4100/9300 chassis will need to restart.

Troubleshooting

Packet Capture

The Packet Capture tool is a valuable asset for use in debugging connectivity and configuration issues and for understanding traffic flows through your Firepower 4100/9300 chassis. You can use the Packet Capture tool to log traffic that is going through specific interfaces on your Firepower 4100/9300 chassis.

You can create multiple packet capture sessions, and each session can capture traffic on multiple interfaces. For each interface included in a packet capture session, a separate packet capture (PCAP) file will be created.

Backplane Port Mappings

The Firepower 4100/9300 chassis uses the following mappings for internal backplane ports:

Security Module	Port Mapping	Description
Security Module 1/Security Engine	Ethernet1/9	Internal-Data0/0
Security Module 1/Security Engine	Ethernet1/10	Internal-Data0/1
Security Module 2	Ethernet1/11	Internal-Data0/0
Security Module 2	Ethernet1/12	Internal-Data0/1

Security Module	Port Mapping	Description
Security Module 3	Ethernet1/13	Internal-Data0/0
Security Module 3	Ethernet1/14	Internal-Data0/1

Guidelines and Limitations for Packet Capture

The Packet Capture tool has the following limitations:

- Can capture only up to 100 Mbps.

- Packet capture sessions can be created even when there is not enough storage space available to run the packet capture session. You should verify that you have enough storage space available before you start a packet capture session.

- For packet capture sessions on a single-wide 4x100Gbps or 2x100Gbps network module (part numbers FPR-NM-4X100G and FPR-NM-2X100G respectively), if the module `adminstate` is set to `off`, the capture session is automatically disabled with an "Oper State Reason: Unknown Error." You will have to restart the capture session after the module `adminstate` is set to `on` again.

  With all other network modules, packet capture sessions continue across module `adminstate` changes.

- Does not support multiple active packet capturing sessions.

- Captures only at the ingress stage of the internal switch.

- Filters are not effective on packets that cannot be understood by the internal switch (for example Security Group Tag and Network Service Header packets).

- You can only capture packets for one subinterface per session, even if you have multiple subinterfaces on one or more parents.

- You cannot capture packets for an EtherChannel as a whole or for subinterfaces of an EtherChannel. However, for an EtherChannel allocated to a logical device, you can capture packets on each member interface of the EtherChannel. If you allocate a subinterface, but not the parent interface, then you cannot capture packets on member interfaces.

- You cannot copy or export a PCAP file while the capture session is still active.

- When you delete a packet capture session, all packet capture files associated with that session are also deleted.

Creating or Editing a Packet Capture Session

Procedure

Step 1 Choose **Tools** > **Packet Capture**.

The **Capture Session** tab displays a list of currently configured packet capture sessions. If no packet capture sessions are currently configured, a message stating so is displayed instead.

Step 2 Do one of the following:

- To create a packet capture session, click the **Capture Session** button.

- To edit an existing packet capture session, click the **Edit** button for that session.

The left side of the window lets you select a specific application instance and then shows a representation of that instance. This representation is used to select the interfaces on which you would like to capture packets. The right side of the window contains fields for defining the packet capture session.

Step 3   **Select an Instance** from the drop-down menu.

Step 4   Click the interfaces on which you want to capture traffic. Selected interfaces show a check mark.

Step 5   For subinterfaces, click the icon to the left of the parent interface to view subinterfaces in the **Subinterface selection** column. Click one subinterface in the column; you can only capture packets for one subinterface per capture session, even if you have multiple subinterfaces on one or more parents.

In the case of multiple subinterfaces, the icon will be labled **Subinterfaces(n)**; for a single subinterface, it will be labeled with the subinterface ID. If the parent interface is also allocated to the instance, you can either choose the parent interface or a subinterface; you cannot choose both. If the parent is not allocated, it will be grayed out. Subinterfaces for EtherChannels are not supported.

Step 6   To capture traffic from the logical device going out over the backplane ports:

a)   Click the box representing the application instance.

The **Capture On**, **Application Port**, and **Application Capture Direction** fields are made available on the right side of the **Configure Packet Capture Session** window.

b)   Select the backplane port you wish to capture traffic on or select **All Backplane Ports** from the **Capture On** drop-down list.

Step 7   Enter a name for the packet capture session in the **Session Name** field.

Step 8   Specify the buffer size to use for this packet capture session by selecting one of the pre-defined values from the **Buffer Size** list, or by selecting **Custom in MB** and then entering the desired buffer size. The specified buffer size must be between 1 and 2048 MB.

Step 9   Specify the length of the packet that you want to capture in the **Snap Length** field. Valid values are from 64 to 9006 bytes. The default snap length is 1518 bytes.

Step 10   Specify whether you want to overwrite existing PCAP files or append data to the PCAP files when this packet capture session is executed.

Step 11   To capture traffic between the application instance and a specific interface:

a)   Click the box representing the logical device.

b)   From the **Capture On** drop-down list, choose the application type (for example, **asa**).

c)   Select the **Application Port** that you would like to capture traffic coming from or going to.

d)   To capture only the traffic going from the logical device toward the specified interface, click the **Egress Packets** option next to **Application Capture Direction**.

Note   If you choose **Egress Packets**, traffic will be captured only on the selected backplane ports—traffic will not be captured on physical ports even if you have selected them.

e)   To capture traffic coming from or going to the specified interface, click the **All Packets** option next to **Application Capture Direction**.

Step 12   To filter the traffic being captured:

a)   Click the **Apply Filter** option for the **Capture Filter** field.

You are given a set of fields for configuring the filter.

b) If you need to create the filter, click **Create Filter**.

You see the **Create Packet Filter** dialog box. For more information, see Configuring Filters for Packet Capture, on page 288.

c) Select the filter you want to use from the **Apply** drop-down list.

d) Select the interface to which you want to apply the filter from the **To** drop-down list.

e) To apply additional filters, click **Apply Another Filter** and then repeat the steps above to apply the additional filter.

Step 13 Do one of the following:

- To save this packet capture session and run it now, click the **Save and Run** button. This option is only available if no other packet capture sessions are currently running.

- To save this packet capture session so that it can be ran at a later time, click the **Save** button.

You see the **Capture Session** tab with your session listed along with any other sessions that have been created. If you selected **Save and Run**, your packet capture session will be capturing packets. You will need to stop capturing before you can download the PCAP files from your session.

Configuring Filters for Packet Capture

You can create filters to limit the traffic that is included in a packet capture session. You can select which interfaces should use a specific filter while creating a packet capture session.

Note   If you modify or delete a filter that is applied to a packet capture session that is currently running, the changes will not take affect until you disable that session and then reenable it.

Procedure

Step 1 Choose **Tools** > **Packet Capture**.

The **Capture Session** tab displays a list of currently configured packet capture sessions. If no packet capture sessions are currently configured, a message stating so is displayed instead.

Step 2 Do one of the following:

- To create a filter, click the **Add Filter** button.

- To edit an existing filter, click the **Edit** button for that fitler.

You see the **Create or Edit Packet Filter** dialog box.

Step 3 Enter a name for the packet capture filter in the **Filter Name** field.

Step 4 To filter on a specific protocol, select it from the **Protocol** list, or select**Custom** and then enter the desired protocol. The custom protocol must be an IANA defined protocol in decimal format (0-255).

Step 5   To filter on a specific EtherType, select it from the **EtherType** list, or select**Custom** and then enter the desired EtherType. The custom EhterType must be an IANA defined EtherType in decimal format (for example, IPv4 = 2048, IPv6 = 34525, ARP = 2054, and SGT = 35081).

Step 6   To filter traffic based on an Inner VLAN (VLAN ID while ingressing the port) or Outer VLAN (VLAN ID added by the Firepower 4100/9300 chassis), enter the VLAN ID in the specified field.

Step 7   To filter traffic from a specific source or destination, enter the IP address and port or enter the MAC address in the specified source or destination fields.

Note	You can filter using IPv4 or IPv6 addresses, but you cannot filter on both in the same packet capture session.

Step 8   Click **Save** to save the filter,

You see the **Filter List** tab with your filter listed along with any other filters that have been created.

Starting and Stopping a Packet Capture Session

Procedure

Step 1   Choose **Tools** > **Packet Capture**.

The **Capture Session** tab displays a list of currently configured packet capture sessions. If no packet capture sessions are currently configured, a message stating so is displayed instead.

Step 2   To start a packet capture session, click the **Enable Session** button for that session and then click **Yes** to confirm.

Note	You cannot start a packet capture session while another session is running.

The PCAP files for the interfaces included in the session will start collecting traffic. If the session is configured to overwrite session data, the existing PCAP data will be erased. If not, data will be appended to the existing file (if any).

While the packet capture session is running, the file size for the individual PCAP files will increase as traffic is captured. Once the Buffer Size limit is reached, the system will start dropping packets and you will see the Drop Count field increase.

Step 3   To stop a packet capture session, click the **Disable Session** button for that session and then click **Yes** to confirm.

After the session has been disabled, you can then download the PCAP files (see Downloading a Packet Capture File, on page 289).

Downloading a Packet Capture File

You can download the Packet Capture (PCAP) files from a session to your local computer so that they can be analyzed using a network packet analyzer.

Procedure

Step 1   Choose **Tools** > **Packet Capture**.

The **Capture Session** tab displays a list of currently configured packet capture sessions. If no packet capture sessions are currently configured, a message stating so is displayed instead.

Step 2   To download the PCAP file for a specific interface from a packet capture session, click the **Download** button that corresponds to that interface.

> **Note**   You cannot download a PCAP file while a packet capture session is running.

Depending on your browser, the specified PCAP file is either automatically downloaded to your default download location or you are prompted to save the file.

Deleting Packet Capture Sessions

You can delete an individual packet capture session if it is not currently running or you can delete all inactive packet capture sessions.

Procedure

Step 1   Choose **Tools** > **Packet Capture**.

The **Capture Session** tab displays a list of currently configured packet capture sessions. If no packet capture sessions are currently configured, a message stating so is displayed instead.

Step 2   To delete a specific packet capture session, click the **Delete** button that corresponds to that session.

Step 3   To delete all inactive packet capture sessions, click the **Delete All Sessions** button above the list of packet capture sessions.

Testing Network Connectivity

Before you begin

To test basic network connectivity by pinging another device on the network with its host name or IPv4 address, use the **ping** command. To ping another device on the network with its host name or IPv6 address, use the **ping6** command.

To trace the route to another device on the network with its host name or IPv4 address, use the **traceroute** command. To trace the route to another device on the network with its host name or IPv6 address, use the **traceroute6** command.

- The **ping** and **ping6** commands are available in `local-mgmt` mode.

- The **ping** command is also available in `module` mode.

• The **traceroute** and **traceroute6** commands are available in `local-mgmt` mode.

• The **traceroute** command is also available in `module` mode.

Procedure

Step 1   Connect to `local-mgmt` or `module` mode by entering one of the following commands:

• **connect local-mgmt**

• **connect module** *module-ID* {**console** | **telnet**}

Example:

```
FP9300-A# connect local-mgmt
FP9300-A(local-mgmt)#
```

Step 2   To test basic network connectivity by pinging another device on the network with its host name or IPv4 address:

ping {*hostname*|*IPv4_address*} [**count** *number_packets*] | [**deadline** *seconds*] | [**interval** *seconds*] | [**packet-size** *bytes*]

Example:

This example shows how to connect to ping another device on the network twelve times:

```
FP9300-A(local-mgmt)# ping 198.51.100.10 count 12
PING 198.51.100.10 (198.51.100.10) from 203.0.113.5 eth0: 56(84) bytes of data.
64 bytes from 198.51.100.10: icmp_seq=1 ttl=61 time=0.264 ms
64 bytes from 198.51.100.10: icmp_seq=2 ttl=61 time=0.219 ms
64 bytes from 198.51.100.10: icmp_seq=3 ttl=61 time=0.234 ms
64 bytes from 198.51.100.10: icmp_seq=4 ttl=61 time=0.205 ms
64 bytes from 198.51.100.10: icmp_seq=5 ttl=61 time=0.216 ms
64 bytes from 198.51.100.10: icmp_seq=6 ttl=61 time=0.251 ms
64 bytes from 198.51.100.10: icmp_seq=7 ttl=61 time=0.223 ms
64 bytes from 198.51.100.10: icmp_seq=8 ttl=61 time=0.221 ms
64 bytes from 198.51.100.10: icmp_seq=9 ttl=61 time=0.227 ms
64 bytes from 198.51.100.10: icmp_seq=10 ttl=61 time=0.224 ms
64 bytes from 198.51.100.10: icmp_seq=11 ttl=61 time=0.261 ms
64 bytes from 198.51.100.10: icmp_seq=12 ttl=61 time=0.261 ms

--- 198.51.100.10 ping statistics ---
12 packets transmitted, 12 received, 0% packet loss, time 11104ms
rtt min/avg/max/mdev = 51.005/51.062/51.164/0.064 ms

FP9300-A(local-mgmt)#
```

Step 3   To trace the route to another device on the network using its host name or IPv4 address:

traceroute {*hostname* | *IPv4_address*}

Example:

```
FP9300-A(local-mgmt)# traceroute 198.51.100.10
traceroute to 198.51.100.10 (198.51.100.10), 30 hops max, 40 byte packets
 1  198.51.100.57 (198.51.100.57)  0.640 ms  0.737 ms  0.686 ms
 2  net1-gw1-13.cisco.com (198.51.100.101)  2.050 ms  2.038 ms  2.028 ms
```

```
    3  net1-sec-gw2.cisco.com (198.51.100.201)  0.540 ms  0.591 ms  0.577 ms
    4  net1-fp9300-19.cisco.com (198.51.100.108)  0.336 ms  0.267 ms  0.289 ms

FP9300-A(local-mgmt)#
```

Step 4   (Optional) Enter **exit** to exit `local-mgmt` mode and return to the top-level mode.

Troubleshooting Management Interface Status

During initialization and configuration, if you suspect the management interface has not come up for some reason (for example, you cannot access the Chassis Manager), use the **show mgmt-port** command in the `local-mgmt` shell to determine the status of the management interface.

Note   Do not use the **show interface brief** command in the `fxos` shell as it currently displays incorrect information.

Procedure

Step 1   Connect to `local-mgmt` mode by entering the following command:

- **connect local-mgmt**

Example:

```
firepower# connect local-mgmt
firepower(local-mgmt)#
```

Step 2   Use the **show mgmt-port** command to determine the status of the management interface.

Example:

```
firepower(local-mgmt)# show mgmt-port
eth0      Link encap:Ethernet  HWaddr b0:aa:77:2f:f0:a9
          inet addr:10.89.5.14  Bcast:10.89.5.63  Mask:255.255.255.192
          inet6 addr: fe80::b2aa:77ff:fe2f:f0a9/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:3210912 errors:0 dropped:0 overruns:0 frame:0
          TX packets:705434 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1648941394 (1.5 GiB)  TX bytes:138386379 (131.9 MiB)

firepower(local-mgmt)#
```

You also can use the **show mgmt-ip-debug** command; however, it produces an extensive listing of interface-configuration information.

Determine Port Channel Status

You can follow these steps to determine the status of currently defined port channels.

Procedure

Step 1   Enter `/eth-uplink/fabric` mode by entering the following commands:

- **scope eth-uplink**

- **scope fabric** {**a** | **b**}

Example:

```
FP9300-A# scope eth-uplink
FP9300-A /eth-uplink # scope fabric a
FP9300-A /eth-uplink/fabric #
```

Step 2   Enter the **show port-channel** command to display a list current port channels with the administrative state and operational state for each.

Example:

```
FP9300-A /eth-uplink/fabric # show port-channel

Port Channel:
    Port Channel Id Name            Port Type          Admin
 State Oper State      State Reason
    --------------- --------------- ------------------ -----
------ --------------- -----------
    10              Port-channel10  Data               Enabl
ed    Failed          No operational members
    11              Port-channel11  Data               Enabl
ed    Failed          No operational members
    12              Port-channel12  Data               Disab
led    Admin Down      Administratively down
    48              Port-channel48  Cluster            Enabl
ed    Up

FP9300-A /eth-uplink/fabric #
```

Step 3   Enter `/port-channel` mode to display individual port-channel and port information by entering the following command:

- **scope port-channel** *ID*

Example:

```
FP9300-A /eth-uplink/fabric/port-channel # top
FP9300-A# connect fxos
Cisco Firepower Extensible Operating System (FX-OS) Software
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2017, Cisco Systems, Inc. All rights reserved.

The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license.
```

```
                    <--- remaining lines removed for brevity --->

          FP9300-A(fxos)#
```

Step 4   Enter the **show** command to display status information for the specified port channel.

Example:

```
FP9300-A /eth-uplink/fabric/port-channel # show

Port Channel:
    Port Channel Id Name            Port Type         Admin
 State Oper State      State Reason
    --------------- --------------- ----------------- -----
------ --------------- ------------
    10              Port-channel10  Data              Enabl
ed    Failed          No operational members

FP9300-A /eth-uplink/fabric/port-channel #
```

Step 5   Enter the **show member-port** command to display status information for the port channel's member port(s).

Example:

```
FP9300-A /eth-uplink/fabric/port-channel # show member-port

Member Port:
    Port Name       Membership        Oper State        State Reas
on
    --------------- ----------------- ----------------- ----------
--
    Ethernet2/3     Suspended         Failed            Suspended
    Ethernet2/4     Suspended         Failed            Suspended

FP9300-A /eth-uplink/fabric/port-channel #
```

A port channel does not come up until you assign it to a logical device. If the port channel is removed from the logical device, or the logical device is deleted, the port channel reverts to a Suspended state.

Step 6   To view additional port channel and LACP information, exit `/eth-uplink/fabric/port-channel` mode and enter `fxos` mode by entering the following commands:

   • **top**

   • **connect fxos**

Example:

Step 7   Enter the **show port-channel summary** command to display summary information for the current port channels.

Example:

```
FP9300-A(fxos)# show port-channel summary
Flags:  D - Down        P - Up in port-channel (members)
        I - Individual  H - Hot-standby (LACP only)
        s - Suspended   r - Module-removed
        S - Switched    R - Routed
        U - Up (port-channel)
        M - Not in use. Min-links not met
------------------------------------------------------------
```

```
                  ------------------
Group Port-          Type     Protocol  Member Ports
      Channel
--------------------------------------------------------------
                  ------------------
10    Po10(SD)       Eth      LACP      Eth2/3(s)    Eth2/4(s)

11    Po11(SD)       Eth      LACP      Eth2/1(s)    Eth2/2(s)

12    Po12(SD)       Eth      LACP      Eth1/4(D)    Eth1/5(D)

48    Po48(SU)       Eth      LACP      Eth1/1(P)    Eth1/2(P)
```

Additional **show port-channel** and **show lacp** commands are available in `fxos` mode. You can use these commands to display a variety of port channel and LACP information such as capacity, traffic, counters, and usage.

What to do next

See for information about creating port channels.

Recovering from a Software Failure

Before you begin

In the event of software failure that prevents the system from booting successfully, you can use the following procedure to boot a new version of software. To complete this process you need to TFTP boot a kickstart image, download new system and manager images, and then boot using the new images.

The recovery images for a specific FXOS version can be obtained from Cisco.com at one of the following locations:

- Firepower 9300—https://software.cisco.com/portal/pub/download/portal/
  select.html?&mdfid=286287252&flowid=77282&softwareid=286287263

- Firepower 4100 Series—https://software.cisco.com/portal/pub/download/portal/
  select.html?&mdfid=286305187&flowid=79423&softwareid=286287263

The recovery images include three separate files. For example, below are the current recovery images for FXOS 2.1.1.64.

```
Recovery image (kickstart) for FX-OS 2.1.1.64.
fxos-k9-kickstart.5.0.3.N2.4.11.63.SPA

Recovery image (manager) for FX-OS 2.1.1.64.
fxos-k9-manager.4.1.1.63.SPA

Recovery image (system) for FX-OS 2.1.1.64.
fxos-k9-system.5.0.3.N2.4.11.63.SPA
```

Procedure

Step 1 Access ROMMON:

a) Connect to the console port.

b) Reboot the system.

The system will start loading and during the process display a countdown timer.

c) Press the **Escape** key during the countdown to enter ROMMON mode.

Example:

```
Cisco System ROMMON, version 1.0.09, RELEASE SOFTWARE
Copright (c) 1994-2015 by Cisco Systems, Inc.
Compiled Sun 01/01/1999 23:59:59:59.99 by user

Current image running: Boot ROM0
Last reset cause: LocalSoft
DIMM Slot 0 : Present
DIMM Slot 1 : Present
No USB drive !!

Platform FPR9K-SUP with 16384 Mbytes of main memory
MAC Address aa:aa:aa:aa:aa:aa

find the string ! boot bootflash:/installables/switch/fxos-k9-kickstart.5.0.3.N2.0.00.00.SPA

  bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.0.00.00.SPA

Use BREAK, ESC or CTRL+L to interrupt boot.
use SPACE to begin boot immediately.
Boot interrupted.

rommon 1  >
```

Step 2 TFTP boot a kickstart image:

a) Verify that the management IP address, management netmask, and gateway IP address are set correctly. You can see their values using the **set** command. You can test the connectivity to the TFTP server using the **ping** command.

```
rommon 1 > set
    ADDRESS=
    NETMASK=
    GATEWAY=
    SERVER=
    IMAGE=
    PS1="ROMMON ! > "
rommon > address <ip-address>
rommon > netmask <network-mask>
rommon > gateway <default-gateway>
```

b) Copy the kickstart image to a TFTP directory that is accessible from your Firepower 4100/9300 chassis.

Note The kickstart image version number will not match the bundle version number. Information showing the mapping between your FXOS version and the kickstart image can be found on the Cisco.com software download page.

c) Boot the image from ROMMON using the boot command:

```
boot tftp://<IP address>/<path to image>
```

Note   You can also boot the kickstart from ROMMON using a FAT32 formatted USB media device inserted into the USB slot on the front panel of the Firepower 4100/9300 chassis. If the USB device is inserted while the system is running, you will need to reboot the system before it will recognize the USB device.

The system will display a series of #'s indicating that the image is being received and will then load the kickstart image.

Example:

```
rommon 1 > set
    ADDRESS=
    NETMASK=
    GATEWAY=
    SERVER=
    IMAGE=
    PS1="ROMMON ! > "

rommon 2 > address 10.0.0.2
rommon 3 > netmask 255.255.255.0
rommon 4 > gateway 10.0.0.1
rommon 5 > ping 10.0.0.2
..!!!!!!!!!!
Success rate is 100 percent (10/10)
rommon 6 > ping 192.168.1.2
..!!!!!!!!!!
Success rate is 100 percent (10/10)

rommon 7 > boot tftp://192.168.1.2/fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA
            ADDRESS: 10.0.0.2
            NETMASK: 255.255.255.0
            GATEWAY: 10.0.0.1
             SERVER: 192.168.1.2
              IMAGE: fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA

      TFTP_MACADDR: aa:aa:aa:aa:aa:aa
...........................................................................

Receiving fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA from 192.168.1.2

###########################################################################
###########################################################################
#########################################################################

File reception completed.
```

Step 3   Download the recovery system and manager images that match the kickstart image you just loaded to the Firepower 4100/9300 chassis:

a)   To download the recovery system and manager images you will need to set the management IP address and gateway. You cannot download these images via USB.

```
switch(boot)# config terminal
switch(boot)(config)# interface mgmt 0
switch(boot)(config-if)# ip address <ip address> <netmask>
switch(boot)(config-if)# no shutdown
switch(boot)(config-if)# exit
switch(boot)(config)# ip default-gateway <gateway>
```

```
switch(boot)(config)# exit
```

b) Copy the recovery system and manager images from the remote server to the bootflash:

switch(boot)# **copy** *URL* **bootflash:**

Specify the URL for the file being imported using one of the following syntax:

- **ftp**://*username@hostname*/*path*/*image_name*

- **scp**://*username@hostname*/*path*/*image_name*

- **sftp**://*username@hostname*/*path*/*image_name*

- **tftp**://*hostname*/*path*/*image_name*

Example:

```
switch(boot)# copy
  scp://<username>@192.168.1.2/recovery_images/fxos-k9-system.5.0.3.N2.4.11.69.SPA
  bootflash:

switch(boot)# copy
  scp://<username>@192.168.1.2/recovery_images/fxos-k9-manager.4.1.1.69.SPA
  bootflash:
```

c) After the images have been successfully copied to the Firepower 4100/9300 chassis, make a symlink to the manager image from nuova-sim-mgmt-nsg.0.1.0.001.bin. This link tells the load mechanism which manager image to load. The symlink name should always be nuova-sim-mgmt-nsg.0.1.0.001.bin regardless of what image you are trying to load.

```
switch(boot)# copy bootflash:<manager-image>
  bootflash:nuova-sim-mgmt-nsg.0.1.0.001.bin
```

Example:

```
switch(boot)# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.

switch(boot)(config)# interface mgmt 0
switch(boot)(config-if)# ip address 10.0.0.2 255.255.255.0
switch(boot)(config-if)# no shutdown
switch(boot)(config-if)# exit
switch(boot)(config)# ip default-gateway 10.0.0.1
switch(boot)(config)# exit
switch(boot)# copy
  tftp://192.168.1.2/recovery_images/fxos-k9-system.5.0.3.N2.4.11.69.SPA
  bootflash:
Trying to connect to tftp server......
Connection to server Established. Copying Started.....
/
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...

switch(boot)# copy
  tftp://192.168.1.2/recovery_images/fxos-k9-manager.4.1.1.69.SPA
  bootflash:
Trying to connect to tftp server......
Connection to server Established. Copying Started.....
/
```

```
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...

switch(boot)# copy bootflash:fxos-k9-manager.4.1.1.69.SPA
  bootflash:nuova-sim-mgmt-nsg.0.1.0.001.bin

Copy complete, now saving to disk (please wait)...

switch(boot)#
```

Step 4   Load the system image that you just downloaded:

```
switch(boot)# load bootflash:<system-image>
```

Example:

```
switch(boot)# load bootflash:fxos-k9-system.5.0.3.N2.4.11.69.SPA
Uncompressing system image: bootflash:/fxos-k9-system.5.0.3.N2.4.11.69.SPA


Manager image digital signature verification successful
...
System is coming up ... Please wait ...


Cisco FPR Series Security Appliance
FP9300-A login:
```

Step 5   After the recovery images have loaded, enter the following commands to prevent the system from trying to load the prior images:

Note   This step should be performed immediately after loading the recovery images.

```
FP9300-A# scope org
FP9300-A /org # scope fw-platform-pack default
FP9300-A /org/fw-platform-pack # set platform-bundle-version ""
Warning: Set platform version to empty will result software/firmware incompatibility issue.
FP9300-A /org/fw-platform-pack* # commit-buffer
```

Step 6   Download and install the Platform Bundle image that you want to use on your Firepower 4100/9300 chassis. For more information, see Image Management, on page 55.

Example:

```
FP9300-A# scope firmware
FP9300-A /firmware # show download-task

Download task:
    File Name Protocol Server          Port      Userid          State
    --------- -------- --------------- ---------- --------------- -----
    fxos-k9.2.1.1.73.SPA
            Tftp     192.168.1.2               0                Downloaded
FP9300-A /firmware # show package fxos-k9.2.1.1.73.SPA detail
Firmware Package fxos-k9.2.1.1.73.SPA:
    Version: 2.1(1.73)
    Type: Platform Bundle
    State: Active
Time Stamp: 2012-01-01T07:40:28.000
Build Date: 2017-02-28 13:51:08 UTC
```

```
FP9300-A /firmware #
```

Recovering from a Corrupted File System

Before you begin

If the Supervisor's onboard flash becomes corrupted and the system is no longer able to start successfully, you can use the following procedure to recover the system. To complete this process you need to TFTP boot a kickstart image, reformat the flash, download new system and manager images, and then boot using the new images.

Note   This procedure includes reformatting the system flash. As a result, you will need to completely reconfigure your system after it has been recovered.

The recovery images for a specific FXOS version can be obtained from Cisco.com at one of the following locations:

- Firepower 9300—https://software.cisco.com/portal/pub/download/portal/select.html?&mdfid=286287252&flowid=77282&softwareid=286287263

- Firepower 4100 Series—https://software.cisco.com/portal/pub/download/portal/select.html?&mdfid=286305187&flowid=79423&softwareid=286287263

The recovery images include three separate files. For example, below are the recovery images for FXOS 2.1.1.64.

```
Recovery image (kickstart) for FX-OS 2.1.1.64.
fxos-k9-kickstart.5.0.3.N2.4.11.63.SPA

Recovery image (manager) for FX-OS 2.1.1.64.
fxos-k9-manager.4.1.1.63.SPA

Recovery image (system) for FX-OS 2.1.1.64.
fxos-k9-system.5.0.3.N2.4.11.63.SPA
```

Procedure

Step 1   Access ROMMON:

a) Connect to the console port.
b) Reboot the system.

   The system will start loading and during the process display a countdown timer.

c) Press the **Escape** key during the countdown to enter ROMMON mode.

Example:

```
Cisco System ROMMON, version 1.0.09, RELEASE SOFTWARE
Copright (c) 1994-2015 by Cisco Systems, Inc.
```

```
Compiled Sun 01/01/1999 23:59:59:59.99 by user

Current image running: Boot ROM0
Last reset cause: LocalSoft
DIMM Slot 0 : Present
DIMM Slot 1 : Present
No USB drive !!

Platform FPR9K-SUP with 16384 Mbytes of main memory
MAC Address aa:aa:aa:aa:aa:aa

find the string ! boot bootflash:/installables/switch/fxos-k9-kickstart.5.0.3.N2.0.00.00.SPA

  bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.0.00.00.SPA

Use BREAK, ESC or CTRL+L to interrupt boot.
use SPACE to begin boot immediately.
Boot interrupted.

rommon 1  >
```

Step 2   TFTP boot a kickstart image:

a) Verify that the management IP address, management netmask, and gateway IP address are set correctly. You can see their values using the **set** command. You can test the connectivity to the TFTP server using the **ping** command.

```
rommon 1 > set
     ADDRESS=
     NETMASK=
     GATEWAY=
     SERVER=
     IMAGE=
     PS1="ROMMON ! > "
rommon > address <ip-address>
rommon > netmask <network-mask>
rommon > gateway <default-gateway>
```

b) Copy the kickstart image to a TFTP directory that is accessible from your Firepower 4100/9300 chassis.

Note   The kickstart image version number will not match the bundle version number. Information showing the mapping between your FXOS version and the kickstart image can be found on the Cisco.com software download page.

c) Boot the image from ROMMON using the boot command:

boot tftp://<IP address>**/**<path to image>

Note   You can also boot the kickstart from ROMMON using a USB media device inserted into the USB slot on the front panel of the Firepower 4100/9300 chassis. If the USB device is inserted while the system is running, you will need to reboot the system before it will recognize the USB device.

The system will display a series of #'s indicating that the image is being received and will then load the kickstart image.

Example:

```
rommon 1 > set
     ADDRESS=
```

```
            NETMASK=
            GATEWAY=
            SERVER=
            IMAGE=
            PS1="ROMMON ! > "

rommon 2 > address 10.0.0.2
rommon 3 > netmask 255.255.255.0
rommon 4 > gateway 10.0.0.1
rommon 5 > ping 10.0.0.2
..!!!!!!!!!!!
Success rate is 100 percent (10/10)
rommon 6 > ping 192.168.1.2
..!!!!!!!!!!!
Success rate is 100 percent (10/10)

rommon 7 > boot tftp://192.168.1.2/fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA
              ADDRESS: 10.0.0.2
              NETMASK: 255.255.255.0
              GATEWAY: 10.0.0.1
               SERVER: 192.168.1.2
                IMAGE: fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA

        TFTP_MACADDR: aa:aa:aa:aa:aa:aa
..........................................................................

Receiving fxos-k9-kickstart.5.0.3.N2.1.11.1.SPA from 192.168.1.2

##############################################################################
##############################################################################
#############################################################################

File reception completed.
```

Step 3 After the kickstart image has loaded, reformat the flash using the **init system** command.

The **init system** command erases the contents of the flash including all software images downloaded to the system and all configurations on the system. The command takes approximately 20-30 minutes to complete.

Example:

```
switch(boot)# init system

This command is going to erase your startup-config, licenses as well as the contents of
your bootflash:.

Do you want to continue? (y/n)  [n] y

Detected 32GB flash...
Initializing the system
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):          done
Initializing startup-config and licenses
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):          done
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):          done
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):          done
Formatting bootflash:
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):          done
Formatting SAM partition:
```

```
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):        done
Formatting Workspace partition:
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):        done
Formatting Sysdebug  partition:
mke2fs 1.35 (28-Feb-2004)
Checking for bad blocks (read-only test):        done
```

Step 4   Download the recovery images to the Firepower 4100/9300 chassis:

a)   To download the recovery images you will need to set the management IP address and gateway. You cannot download these images via USB.

```
switch(boot)# config terminal
switch(boot)(config)# interface mgmt 0
switch(boot)(config-if)# ip address <ip address> <netmask>
switch(boot)(config-if)# no shutdown
switch(boot)(config-if)# exit
switch(boot)(config)# ip default-gateway <gateway>
switch(boot)(config)# exit
```

b)   Copy all three recovery images from the remote server to the bootflash:

switch(boot)# **copy** *URL* **bootflash:**

Specify the URL for the file being imported using one of the following syntax:

- **ftp**://*username@hostname*/*path*/*image_name*

- **scp**://*username@hostname*/*path*/*image_name*

- **sftp**://*username@hostname*/*path*/*image_name*

- **tftp**://*hostname*/*path*/*image_name*

Example:

```
switch(boot)# copy
  scp://<username>@192.168.1.2/recovery_images/fxos-k9-kickstart.5.0.3.N2.4.11.69.SPA
  bootflash:
```

```
switch(boot)# copy
  scp://<username>@192.168.1.2/recovery_images/fxos-k9-system.5.0.3.N2.4.11.69.SPA
  bootflash:
```

```
switch(boot)# copy
  scp://<username>@192.168.1.2/recovery_images/fxos-k9-manager.4.1.1.69.SPA
  bootflash:
```

c)   After the images have been successfully copied to the Firepower 4100/9300 chassis, make a symlink to the manager image from nuova-sim-mgmt-nsg.0.1.0.001.bin. This link tells the load mechanism which manager image to load. The symlink name should always be nuova-sim-mgmt-nsg.0.1.0.001.bin regardless of what image you are trying to load.

```
switch(boot)# copy bootflash:<manager-image>
  bootflash:nuova-sim-mgmt-nsg.0.1.0.001.bin
```

Example:

```
switch(boot)# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.

switch(boot)(config)# interface mgmt 0
switch(boot)(config-if)# ip address 10.0.0.2 255.255.255.0
switch(boot)(config-if)# no shutdown
switch(boot)(config-if)# exit
switch(boot)(config)# ip default-gateway 10.0.0.1
switch(boot)(config)# exit
switch(boot)# copy
  tftp://192.168.1.2/recovery_images/fxos-k9-kickstart.5.0.3.N2.4.11.69.SPA
  bootflash:
Trying to connect to tftp server......
Connection to server Established. Copying Started.....
/
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...

switch(boot)# copy
  tftp://192.168.1.2/recovery_images/fxos-k9-system.5.0.3.N2.4.11.69.SPA
  bootflash:
Trying to connect to tftp server......
Connection to server Established. Copying Started.....
/
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...

switch(boot)# copy
  tftp://192.168.1.2/recovery_images/fxos-k9-manager.4.1.1.69.SPA
  bootflash:
Trying to connect to tftp server......
Connection to server Established. Copying Started.....
/
TFTP get operation was successful
Copy complete, now saving to disk (please wait)...

switch(boot)# copy bootflash:fxos-k9-manager.4.1.1.69.SPA
  bootflash:nuova-sim-mgmt-nsg.0.1.0.001.bin

Copy complete, now saving to disk (please wait)...

switch(boot)#
```

Step 5   Reload the switch:

```
switch(boot)# reload
```

Example:

```
switch(boot)# reload
This command will reboot this supervisor module. (y/n) ? y
[ 1866.310313] Restarting system.

!!  Rommon image verified successfully  !!


Cisco System ROMMON, Version 1.0.11, RELEASE SOFTWARE
Copyright (c) 1994-2016  by Cisco Systems, Inc.
Compiled Wed 11/23/2016 11:23:23.47 by builder
Current image running: Boot ROM1
Last reset cause: ResetRequest
DIMM Slot 0 : Present
```

```
DIMM Slot 1 : Present
No USB drive !!
BIOS has been locked !!

Platform FPR9K-SUP with 16384 Mbytes of main memory
MAC Address: bb:aa:77:aa:aa:bb

autoboot: Can not find autoboot file 'menu.lst.local'
          Or can not find correct boot string !!
rommon 1 >
```

Step 6   Boot from the kickstart and system images:

```
rommon 1 > boot <kickstart-image> <system-image>
```

Note   You will likely see license manager failure messages while the system image is loading. These messages can be safely ignored.

Example:

```
rommon 1 > dir
Directory of: bootflash:\

  01/01/12  12:33a <DIR>           4,096  .
  01/01/12  12:33a <DIR>           4,096  ..
  01/01/12  12:16a <DIR>          16,384  lost+found
  01/01/12  12:27a          34,333,696  fxos-k9-kickstart.5.0.3.N2.4.11.69.SPA
  01/01/12  12:29a         330,646,465  fxos-k9-manager.4.1.1.69.SPA
  01/01/12  12:31a         250,643,172  fxos-k9-system.5.0.3.N2.4.11.69.SPA
  01/01/12  12:34a         330,646,465  nuova-sim-mgmt-nsg.0.1.0.001.bin
          4 File(s) 946,269,798 bytes
          3 Dir(s)

rommon 2 > boot fxos-k9-kickstart.5.0.3.N2.4.11.69.SPA fxos-k9-system.5.0.3.N2.4.11.69.SPA
  !!   Kickstart Image verified successfully   !!

Linux version: 2.6.27.47 (security@cisco.com) #1 SMP Thu Nov 17 18:22:00 PST 2016
[    0.000000] Fastboot Memory at 0c100000 of size 201326592
Usage: init 0123456SsQqAaBbCcUu

INIT: version 2.86 booting

POST INIT Starts at Sun Jan  1 00:27:32 UTC 2012
S10mount-ramfs.supnuovaca Mounting /isan 3000m
Mounted /isan
Creating /callhome..
Mounting /callhome..
Creating /callhome done.
Callhome spool file system init done.
Platform is BS or QP MIO: 30
FPGA Version 0x00010500 FPGA Min Version 0x00000600
Checking all filesystems..r.r..r done.
Warning: switch is starting up with default configuration
Checking NVRAM block device ... done
.
FIPS power-on self-test passed
Unpack CMC Application software
Loading system software
Uncompressing system image: bootflash:/fxos-k9-system.5.0.3.N2.4.11.69.SPA


Manager image digital signature verification successful
```

```
...

System is coming up ... Please wait ...
nohup: appending output to `nohup.out'

          ---- Basic System Configuration Dialog ----

  This setup utility will guide you through the basic configuration of
  the system. Only minimal configuration including IP connectivity to
  the Fabric interconnect and its clustering mode is performed through these steps.

  Type Ctrl-C at any time to abort configuration and reboot system.
  To back track or make modifications to already entered values,
  complete input till end of section and answer no when prompted
  to apply configuration.


  You have chosen to setup a new Security Appliance. Continue? (y/n):
```

Step 7   After the images have loaded, the system will prompt you to enter initial configuration settings. For more information, see Initial Configuration Using Console Port, on page 8.

Step 8   Download the Platform Bundle image that you want to use on your Firepower 4100/9300 chassis. For more information, see Image Management, on page 55.

Example:

```
FP9300-A# scope firmware
FP9300-A /firmware # show download-task

Download task:
    File Name Protocol Server          Port       Userid          State
    --------- -------- --------------- ---------- --------------- -----
    fxos-k9.2.1.1.73.SPA
            Tftp     192.168.1.2               0                 Downloaded
FP9300-A /firmware # show package fxos-k9.2.1.1.73.SPA detail
Firmware Package fxos-k9.2.1.1.73.SPA:
    Version: 2.1(1.73)
    Type: Platform Bundle
    State: Active
Time Stamp: 2012-01-01T07:40:28.000
Build Date: 2017-02-28 13:51:08 UTC
FP9300-A /firmware #
```

Step 9   Install the Platform Bundle image you downloaded in the previous step:

Note   Installation process typically takes between 15 and 20 minutes.

a) Enter auto-install mode:

Firepower-chassis /firmware # **scope auto-install**

b) Install the FXOS platform bundle:

Firepower-chassis /firmware/auto-install # **install platform platform-vers** *version_number*

version_number is the version number of the FXOS platform bundle you are installing--for example, 2.1(1.73).

c) The system will first verify the software package that you want to install. It will inform you of any incompatibility between currently installed applications and the specified FXOS platform software package.

It will also warn you that any existing sessions will be terminated and that the system will need to be rebooted as part of the upgrade.

Enter **yes** to confirm that you want to proceed with verification.

d) Enter **yes** to confirm that you want to proceed with installation, or enter **no** to cancel the installation.

The FXOS unpacks the bundle and upgrades/reloads the components.

e) To monitor the upgrade process:

- Enter **scope firmware**.

- Enter **scope auto-install**.

- Enter **show fsm status expand**.

Example:

```
TB10 /firmware/auto-install # show fsm status expand
    FSM Status:
        Affected Object: sys/fw-system/fsm
        Current FSM: Deploy
        Status: In Progress
        Completion Time:
        Progress (%): 98

        FSM Stage:
        Order   Stage Name                                Status        Try
        ------  ----------------------------------------  ------------  ---
        1       DeployWaitForDeploy                       Success       0
        2       DeployResolveDistributableNames           Skip          0
        3       DeployResolveDistributable                Skip          0
        4       DeployResolveImages                       Skip          0
        5       DeployValidatePlatformPack                Success       1
        6       DeployDebundlePort                        Success       0
        7       DeployPollDebundlePort                    Success       1
        8       DeployActivateUCSM                        Success       0
        9       DeployPollActivateOfUCSM                  Success       0
        10      DeployActivateMgmtExt                     Skip          0
        11      DeployPollActivateOfMgmtExt               Skip          0
        12      DeployUpdateIOM                           Skip          0
        13      DeployPollUpdateOfIOM                     Skip          0
        14      DeployActivateIOM                         Skip          0
        15      DeployPollActivateOfIOM                   Skip          0
        16      DeployActivateRemoteFI                    Skip          0
        17      DeployPollActivateOfRemoteFI              Skip          0
        18      DeployWaitForUserAck                      Skip          0
        19      DeployActivateLocalFI                     Success       0
        20      DeployPollActivateOfLocalFI               In Progress   1
```

Note   Do not proceed to the next step until the status of the stages changes from "In Progress" to "Skip" or "Success."

Step 10   If the Platform Bundle image that you installed corresponds with the images you used for recovering your system, you must manually activate the kickstart and system images so that they will be used when loading the system in the future. Automatic activation does not occur when installing a Platform Bundle that has same images as the recovery images that were used.

a) Set the scope for fabric-interconnect a:

```
FP9300-A# scope fabric-interconnect a
```

b) Use the **show version** command to view the running kernel version and the running system version. You will use these strings to activate the images.

```
FP9300-A /fabric-interconnect # show version
```

Note If the Startup-Kern-Vers and Startup-Sys-Vers are already set and match the Running-Kern-Vers and Running-Sys-Vers, you do not need to activate the images and can proceed to Step 11.

c) Enter the following command to activate the images:

```
FP9300-A /fabric-interconnect # activate firmware
  kernel-version <running_kernel_version> system-version <running_system_version>
commit-buffer
```

Note The server status might change to "Disk Failed." You do not need to worry about this message and can continue with this procedure.

d) Use the **show version** command to verify that the startup versions have been set correctly and to monitor the activation status for the images.

Important Do not proceed to the next step until the status changes from "Activating" to "Ready."

```
FP9300-A /fabric-interconnect # show version
```

Example:

```
FP9300-A /firmware # top
FP9300-A# scope fabric-interconnect a
FP9300-A /fabric-interconnect # show version
Fabric Interconnect A:
    Running-Kern-Vers: 5.0(3)N2(4.11.69)
    Running-Sys-Vers: 5.0(3)N2(4.11.69)
    Package-Vers: 2.1(1.73)
    Startup-Kern-Vers:
    Startup-Sys-Vers:
    Act-Kern-Status: Ready
    Act-Sys-Status: Ready
    Bootloader-Vers:

FP9300-A /fabric-interconnect # activate firmware kernel-version
  5.0(3)N2(4.11.69) system-version 5.0(3)N2(4.11.69)
Warning: When committed this command will reset the end-point
FP9300-A /fabric-interconnect* # commit-buffer
FP9300-A /fabric-interconnect # show version
Fabric Interconnect A:
    Running-Kern-Vers: 5.0(3)N2(4.11.69)
    Running-Sys-Vers: 5.0(3)N2(4.11.69)
    Package-Vers: 2.1(1.73)
    Startup-Kern-Vers: 5.0(3)N2(4.11.69)
    Startup-Sys-Vers: 5.0(3)N2(4.11.69)
    Act-Kern-Status: Activating
    Act-Sys-Status: Activating
    Bootloader-Vers:

FP9300-A /fabric-interconnect # show version
Fabric Interconnect A:
```

```
Running-Kern-Vers: 5.0(3)N2(4.11.69)
Running-Sys-Vers: 5.0(3)N2(4.11.69)
Package-Vers: 2.1(1.73)
Startup-Kern-Vers: 5.0(3)N2(4.11.69)
Startup-Sys-Vers: 5.0(3)N2(4.11.69)
Act-Kern-Status: Ready
Act-Sys-Status: Ready
Bootloader-Vers:
```

Step 11   Reboot the system:

Example:

```
FP9300-A /fabric-interconnect # top
FP9300-A# scope chassis 1
FP9300-A /chassis # reboot no-prompt
Starting chassis reboot. Monitor progress with the command "show fsm status"
FP9300-A /chassis #
```

The system will power down each security module/engine before finally powering down and then restarting the Firepower 4100/9300 chassis. This process takes approximately 5-10 minutes.

Step 12   Monitor the system status. The server status should go from "Discovery" to "Config" and then finally to "Ok".

Example:

```
FP9300-A# show server status
Server  Slot Status                      Overall Status       Discovery
------- -------------------------------- -------------------- ---------
1/1     Equipped                         Discovery            In Progress
1/2     Equipped                         Discovery            In Progress
1/3     Empty

FP9300-A# show server status
Server  Slot Status                      Overall Status       Discovery
------- -------------------------------- -------------------- ---------
1/1     Equipped                         Config               Complete
1/2     Equipped                         Config               Complete
1/3     Empty

FP9300-A# show server status
Server  Slot Status                      Overall Status       Discovery
------- -------------------------------- -------------------- ---------
1/1     Equipped                         Ok                   Complete
1/2     Equipped                         Ok                   Complete
1/3     Empty
```

When the Overall Status is "Ok" your system has been recovered. You must still reconfigure your security appliance (including license configuration) and re-create any logical devices. For more information:

- Firepower 9300 Quick Start Guides—http://www.cisco.com/go/firepower9300-quick

- Firepower 9300 Configuration Guides—http://www.cisco.com/go/firepower9300-config

- Firepower 4100 Series Quick Start Guides—http://www.cisco.com/go/firepower4100-quick

- Firepower 4100 Series Configuration Guides—http://www.cisco.com/go/firepower4100-config

Restoring the Factory Default Configuration when the Admin Password is Unknown

This procedure returns your Firepower 4100/9300 chassis system to its default configuration settings, including the admin password. Use this procedure to reset the configurations on your device when the admin password is not known. This procedure erases any installed logical devices as well.

Note This procedure requires console access to the Firepower 4100/9300 chassis.

Procedure

Step 1 Connect your PC to the console port using the provided console cable, and connect to the console using a terminal emulator set for 9600 baud, 8 data bits, no parity, 1 stop bit, and no flow control. For more information on the console cable, see Cisco Firepower 9300 Hardware Installation Guide .

Step 2 Power on the device. When you see the following prompt, press ESC to stop the boot.

Example:

```
!!  Rommon image verified successfully  !!

Cisco System ROMMON, Version 1.0.09, RELEASE SOFTWARE
Copyright (c) 1994-2015  by Cisco Systems, Inc.

Current image running: Boot ROM0
Last reset cause: ResetRequest
DIMM Slot 0 : Present
DIMM Slot 1 : Present
No USB drive !!
BIOS has been locked !!

Platform FPR9K-SUP with 16384 Mbytes of main memory
MAC Address: 00:00:00:00:00:00

find the string ! boot bootflash:/installables/switch/fxos-k9-kickstart.5.0.3.N2.3.14.69.SPA
 bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.3.14.69.SPA

Use BREAK, ESC or CTRL+L to interrupt boot.
Use SPACE to begin boot immediately.
Boot interrupted.
rommon 1 >
```

Step 3 Make a note of the kickstart and system image names:

Example:

```
bootflash:/installables/switch/fxos-k9-kickstart.5.0.3.N2.3.14.69.SPA
bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.3.14.69.SPA
```

Step 4 Load the kickstart image:

rommon 1 > **boot** *kickstart_image*

Example:

```
rommon 1 > boot bootflash:/installables/switch/fxos-k9-kickstart.5.0.3.N2.3.14.69.SPA
 !! Kickstart Image verified successfully !!

Linux version: 2.6.27.47 (security@cisco.com) #1 SMP Tue Nov 24 12:10:28 PST 2015
[ 0.000000] Fastboot Memory at 0c100000 of size 201326592
Usage: init 0123456SsQqAaBbCcUu
INIT: POST INIT Starts at Wed Jun 1 13:46:33 UTC 2016
can't create lock file /var/lock/mtab~302: No such file or directory (use -n flag to override)
S10mount-ramfs.supnuovaca Mounting /isan 3000m
Mounted /isan
TAC support: http://www.cisco.com/tac
Copyright (c) 2002-2015, Cisco Systems, Inc. All rights reserved.
The copyrights to certain works contained in this software are
owned by other third parties and used and distributed under
license. Certain components of this software are licensed under
the GNU General Public License (GPL) version 2.0 or the GNU
Lesser General Public License (LGPL) Version 2.1. A copy of each
such license is available at
http://www.opensource.org/licenses/gpl-2.0.php and
http://www.opensource.org/licenses/lgpl-2.1.php
switch(boot)#
```

Step 5 Enter the config terminal mode:

switch(boot) # **config terminal**

Example:

```
switch(boot)#
switch(boot)# config terminal
Enter configuration commands, one per line.  End with CNTL/Z.
```

Step 6 Reset the password and confirm the change:

switch(boot) (config) # **admin-password erase**

Note This step erases all configurations and returns your system to its default configuration settings.

Example:

```
switch(boot)(config)# admin-password erase
Your password and configuration will be erased!
Do you want to continue? (y/n)  [n] y
```

Step 7 Exit the config terminal mode:

switch(boot) (config) # **exit**

Step 8 Load the system image noted in step 3 of this procedure and configure your system from scratch using the Initial Configuration, on page 7 task flow.

switch(boot) # **load** *system_image*

Example:

```
switch(boot)# load bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.3.14.69.SPA

Uncompressing system image: bootflash:/installables/switch/fxos-k9-system.5.0.3.N2.3.14.69.SPA
```

Generating Troubleshooting Log Files

You can generate log files to help with troubleshooting or to send to Cisco TAC if requested.

Procedure

Step 1   Choose **Tools** > **Troubleshooting Logs**.

Step 2   Choose the type of log file you would like to generate from the drop-down list:

   • Chassis—generates log files to use for troubleshooting chassis hardware issues and software issues including the supervisor and service manager.

   • Module <#>—generates log files to use for troubleshooting security module/engine issues.

Step 3   Click **Generate Log**.

Step 4   Click **Yes** to confirm that you want to generate the log files.

The log files are generated. This process can take some time. While the log files are being generated, a yellow status message is displayed. You can click **Abort Job** in the status message to cancel the log file generation. After the log files have been generated, the status message changes to green and indicates that the job completed successfully.

Step 5   To download a generated log file, navigate to the log file in the **Download Files** list and then click **Download**. The log files are stored under the techsupport folder.

   Note   You might have to click **Refresh** for newly generated files to be shown in the **Download Files** list.

Step 6   To delete a generated log file, navigate to the log file in the **Download Files** list and then click **Delete**.

Enabling Module Core Dumps

Enabling core dumps on a module can help with troubleshooting in the event of a system crash, or to send to Cisco TAC if requested.

Procedure

Step 1   Connect to the desired module; for example:

Firepower# connect module 1 console

Step 2   (Optional) Enter the following command to view current core dump status:

Firepower-module1> show coredump detail

The command output shows current core dump status information, including whether core dump compression is enabled.

Example:

```
Firepower-module1>show coredump detail
Configured status: ENABLED.
ASA Coredump: ENABLED.
Bootup status: ENABLED.
Compress during crash: DISABLED.
```

Note This command is available only when running ASA Logical device on appliance and not when running FTD Logical device on appliance.

Step 3 Use the **config coredump** command to enable or disable core dumps, and to enable or disable core dump compression during a crash.

  • Use **config coredump enable** to enable creation of a core dump during a crash.

  • Use **config coredump disable** to disable core dump creation during a crash.

  • Use **config coredump compress enable** to enable compression of core dumps.

  • Use **config coredump compress disable** to disable core dump compression.

Example:

```
Firepower-module1>config coredump enable
Coredump enabled successfully.
ASA coredump enabled, do 'config coredump disableAsa' to disable
Firepower-module1>config coredump compress enable
WARNING: Enabling compression delays system reboot for several minutes after a system
failure.  Are you sure? (y/n):
y
Firepower-module1>
```

Note Core dump files consume disk space, and if space is running low and compression is not enabled, a core dump file may not be saved even if core dumps are enabled.

Finding the Serial Number of the Firepower 4100/9300 Chassis

You can find details about the Firepower 4100/9300 Chassis and its serial number. Note that serial number of Firepower 4100/9300 Chassis is different than serial numbers of the logical devices.

Procedure

Step 1 Choose **Overview** > **Inventory** > **All**.

The table lists the components installed in the chassis and provides relevant details for those components.

Step 2 Look for the chassis serial number in the **Serial** column.

Rebuild RAID Virtual Drive

RAID (Redundant Array of Independent Disks) is an array, or group, of multiple independent physical drives that provide high performance and fault tolerance. A drive group is a group of physical drives. These drives are managed in partitions known as virtual drives.

RAID drive groups improve data storage reliability and fault tolerance compared to single-drive storage systems. Data loss resulting from a drive failure can be prevented by reconstructing missing data from the remaining drives. RAID improves I/O performance and increases storage subsystem reliability.

If one of your RAID drives has failed or is offline, then the RAID virtual drive is considered to be in a degraded state. Use this procedure to verify whether a RAID virtual drive is in a degraded state, and temporaritly set the local disk configuration protection policy to no to rebuild it if necessary.

Note   When you set the local disk configuration protection policy to no, all data on the disk is destroyed.

Procedure

Step 1   Check the RAID drive status.

  a.   Enter chassis mode:

   scope chassis

  b.   Enter server mode:

   scope server 1

  c.   Enter the raid controller:

   scope raid-controller 1 sas

  d.   View the virtual drive:

   show virtual-drive

   If the RAID virtual drive is degraded, the operability displays as **Degraded**. For example:

```
Virtual Drive:
    ID: 0
    Block Size: 512
    Blocks: 3123046400
    Size (MB): 1524925
    Operability: Degraded
Presence: Equipped
```

Step 2   Set the local disk configuration policy protection to no to rebuild the RAID drive. Note - all data on the disk will be destroyed after you complete this step.

  a.   Enter the organization scope:

   scope org

  b.   Enter the local disk configuration policy scope:

scope local-disk-config-policy ssp-default

c. Set protect to no:

set protect no

d. Commit the configuration:

commit-buffer

Step 3 Wait for the RAID drive to rebuild. Check the RAID rebuild status:

scope chassis 1

show server

When the RAID drive has rebuilt successfully, the slot's overall status displays as **Ok**. For example:

Example:
```
Server:
    Slot    Overall Status        Service Profile
    ------- --------------------- ---------------
        1 Ok                      ssp-sprof-1
```

Step 4 Once the RAID drive has rebuilt successfully, set the local disk configuration policy protection back to yes.

a. Enter the organization scope:

scope org

b. Enter the local disk configuration policy scope:

scope local-disk-config-policy ssp-default

c. Set protect to no:

set protect yes

d. Commit the configuration:

commit-buffer

Identify Issues with the SSD

Use the following procedure to collect information and identify possible issues with the SSD installed on your device. One example symptom of an SSD issue is the Data Management Engine (DME) process failing to start.

Note When you insert a new SSD, only the basic information (Type, Model, SN, etc.) gets populated under inventory after the Blade BIOS detection. Only upon the SSP-OS upgrade completion, the Local Disk data gets populated under inventory. If the SSP-OS upgrade is still under "Updating state", the inventory shows no entry for the Local Disk and no fault messages regarding conection of the SSD.

If the output of the below logging files indicate a problem with the SSD, contact TAC (see https://www.cisco.com/c/en/us/buy/product-returns-replacements-rma.html).

Procedure

Step 1   Connect to the FXOS command shell:

connect fxos

Step 2   Display the nvram logging file:

show logging nvram

Example error output:

```
2020 Oct 22 13:03:26 MDCNGIPSAPL02 %$ VDC-1 %$ Oct 22 13:03:25 %KERN-2-SYSTEM_MSG:
[28175880.598580] EXT3-fs error (device sda4): ext3_get_inode_loc: unable to read inode
block - inode=14, block=6
```

Step 3   Display the logging file:

show logging logfile

Example error output:

```
   2020 Oct 21 21:11:25  (none) kernel: [28118744.718445] EXT3-fs error (device sda4):
ext3_get_inode_loc: unable to read inode block - inode=14, block=6
```

INDEX