Use Multi-Instance Capability on the Firepower 4100/9300

First Published: 2019-07-30

Last Modified: 2026-03-10

Use Multi-Instance Capability on the Firepower 4100/9300

Multi-instance capability lets you run container instances that use a subset of resources of the security module/engine. Multi-instance capability is only supported for the Firepower Threat Defense; it is not supported for the ASA.

Note For multi-instance clustering, see Deploying a Cluster for Firepower Threat Defense for Scalability and High Availability or the FXOS configuration guide.

Note This document covers the latest FXOS version features; see History for Multi-Instance Capability, on page 51 for details about feature changes. If you have an old version installed, see the procedures in the FXOS configuration guide for your version.

About Multi-Instance Capability

The Firepower chassis includes a supervisor and up to three security modules on which you can install logical devices. A logical device lets you run one application instance (Secure Firewall Threat Defense or ASA). 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. The application type determines whether you can run a single instance (native) or multiple instances (container).

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 Firewall Threat Defense using Firewall Management Center; it is not supported for the ASA or the Firewall Threat Defense using Firewall Device Manager.

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 Firewall Threat Defense feature support. Multiple context mode, due to shared resources, supports more contexts on a given platform. Multiple context mode is not available on the Firewall Threat Defense.

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 6 and Add a VLAN Subinterface for Container Instances, on page 37.

Note This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 4 for more information.

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 (Firewall Threat Defense-using-Firewall Management Center 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 Firewall Threat Defense-using-Firewall Management Center devices. To use this interface, you must configure its IP address and other parameters at the Firewall Threat Defense 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 Firewall Device Manager and Security Cloud Control does not support clustering.

Note This chapter discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 4 for more information.

See the following table for interface type support for the Firewall Threat Defense and ASA applications in standalone and cluster deployments.

Table 1: Interface Type Support

Application		Data	Data: Subinterface	Data-Sharing	Data-Sharing: Subinterface	Mgmt	Eventing	Cluster (EtherChannel only)	Cluster: Subinterface
Firewall Threat Defense	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.

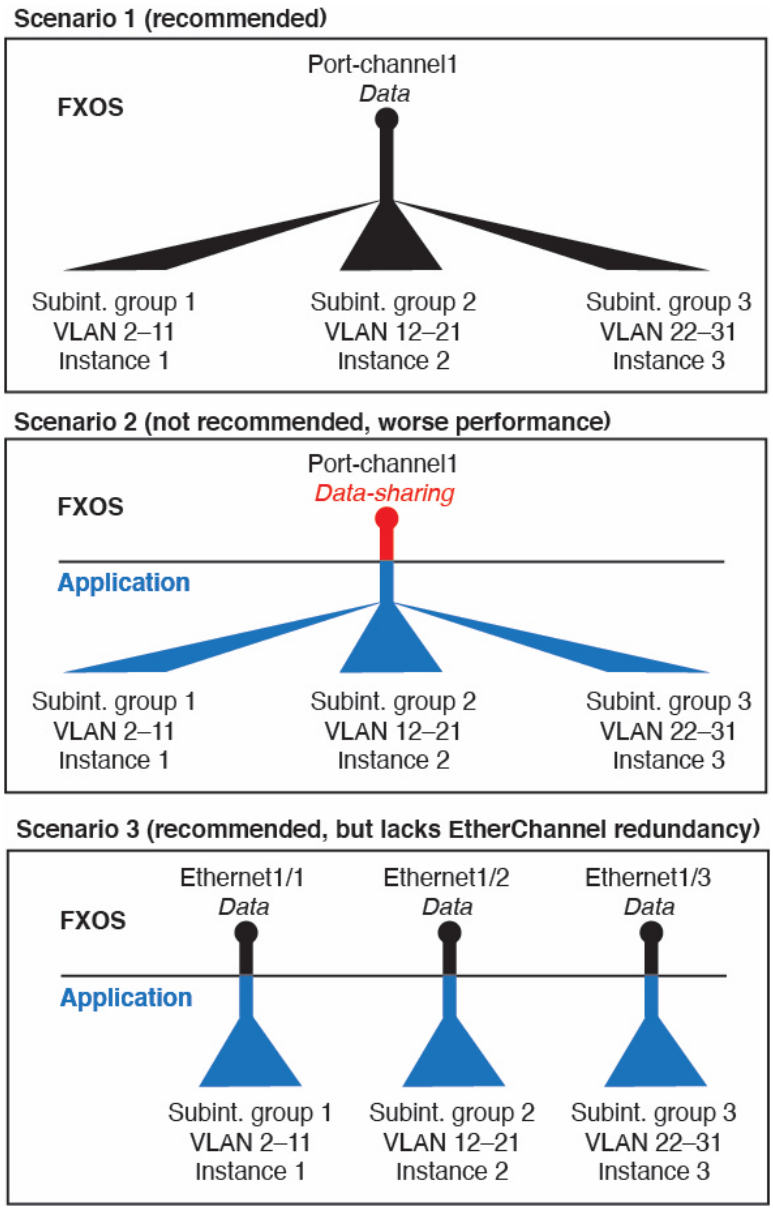
VLAN Subinterfaces

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.

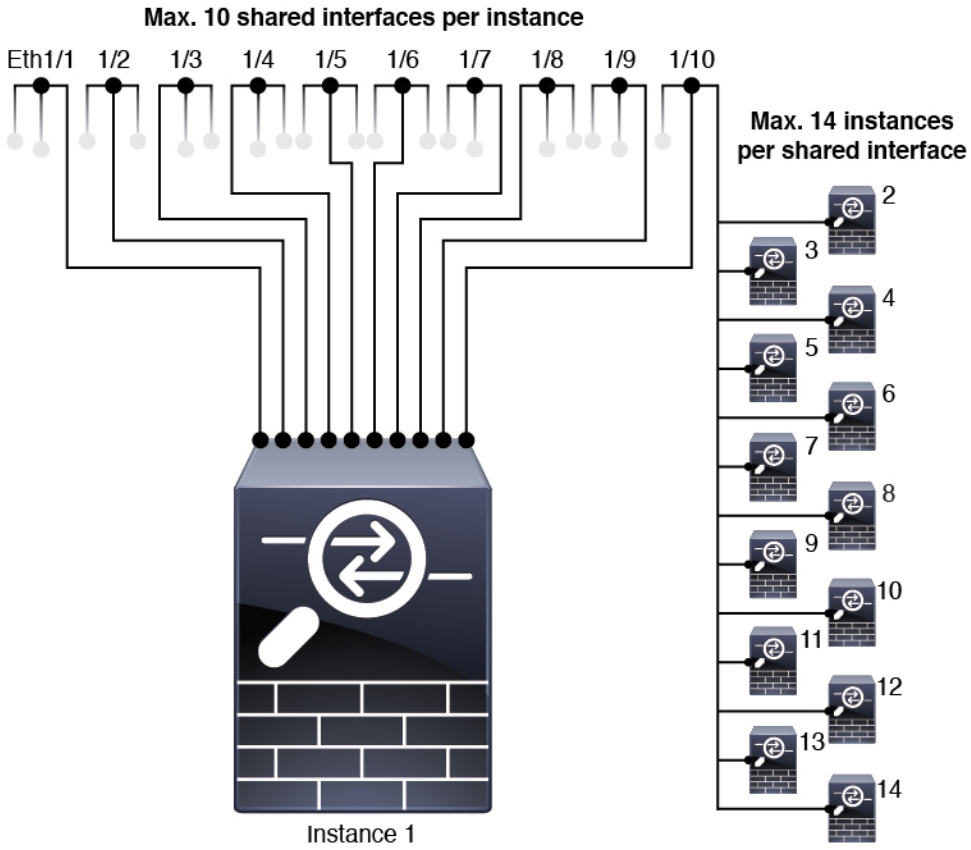
The default state of an interface within the application depends on the type of interface. For example, the physical interface or EtherChannel is disabled by default within the application, but a subinterface is enabled by default.

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

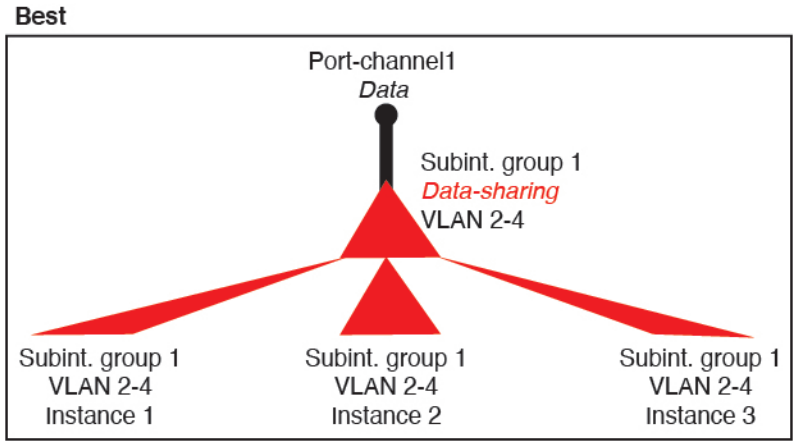
For optimal scalability of the forwarding table, share as few interfaces as possible. Instead, you can create up to 500 VLAN subinterfaces on one or more physical interfaces and then divide the VLANs among the container instances.

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.

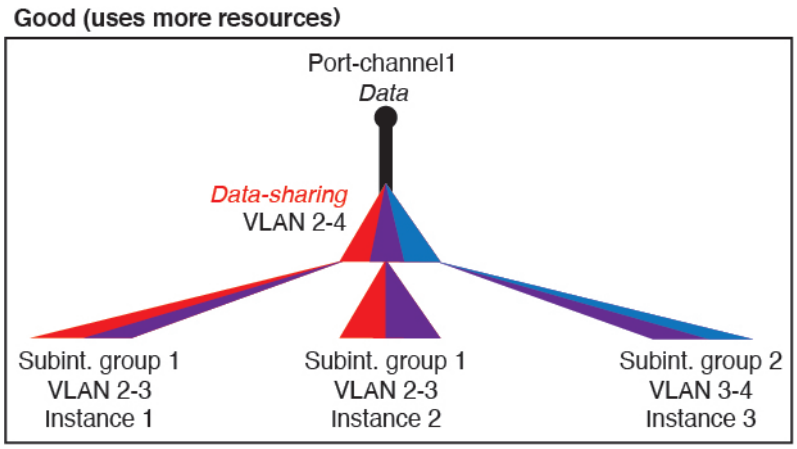
   For example, create a large EtherChannel to bundle all of your like-kind interfaces together, and then share subinterfaces of that EtherChannel: Port-Channel1.2, 3, and 4 instead of Port-Channel2, Port-Channel3, and Port-Channel4. When you share subinterfaces from a single parent, the VLAN group table 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).

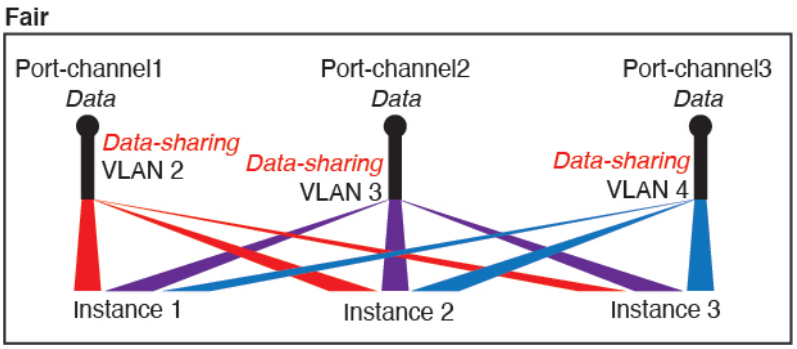
Figure 3: Good: Sharing Multiple Subinterface Groups on One Parent



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.

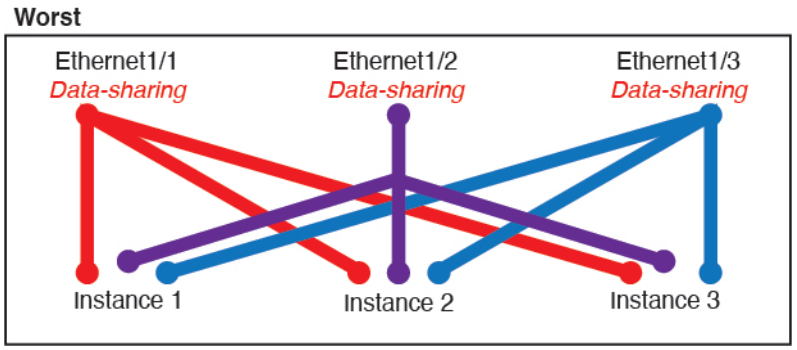
   Figure 4: Fair: Shared Subinterfaces on Separate Parents



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.

- Table 2: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with Three SM-44s, on page 9
- Table 3: Subinterfaces on One Parent and Instances on a Firepower 9300 with Three SM-44s, on page 10
- Table 4: Physical/EtherChannel Interfaces and Instances on a Firepower 9300 with One SM-44, on page 12
- Table 5: Subinterfaces on One Parent and Instances on a Firepower 9300 with One SM-44, on page 13

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 2: 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%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
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%
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 3: 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%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	27%
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%
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 4: 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%

Dedicated Interfaces	Shared Interfaces	Number of Instances	% Forwarding Table Used
32: • 8 • 8 • 8 • 8	**2**	**4:** • Instance 1 • Instance 2 • Instance 3 • Instance 4	25%
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.

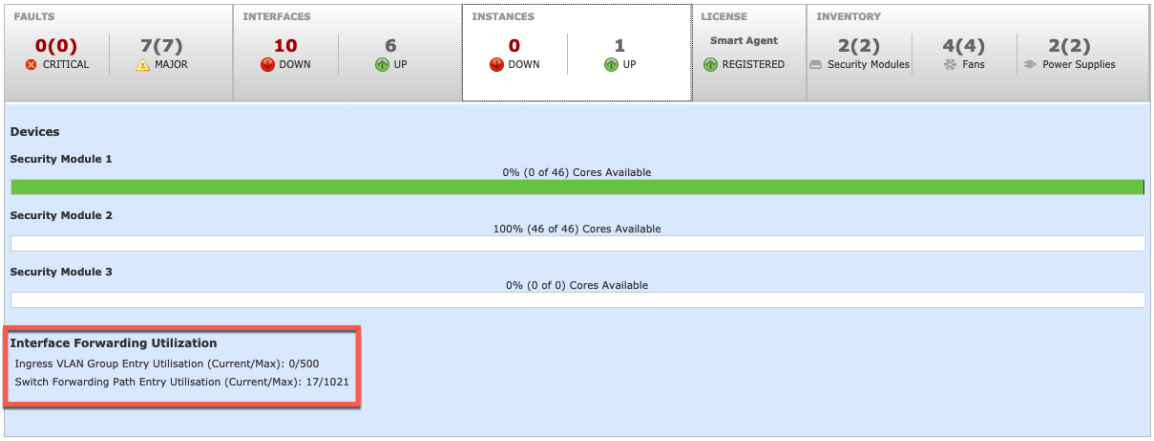
Table 5: Subinterfaces on One Parent and Instances on a Firepower 9300 with One SM-44

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
112: • 112 (8 ea.)	**0**	**14:** • Instance 1-Instance 14	17%

Dedicated Subinterfaces	Shared Subinterfaces	Number of Instances	% Forwarding Table Used
224: • 224 (16 ea.)	**0**	**14:** • Instance 1-Instance 14	17%
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%
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:



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.

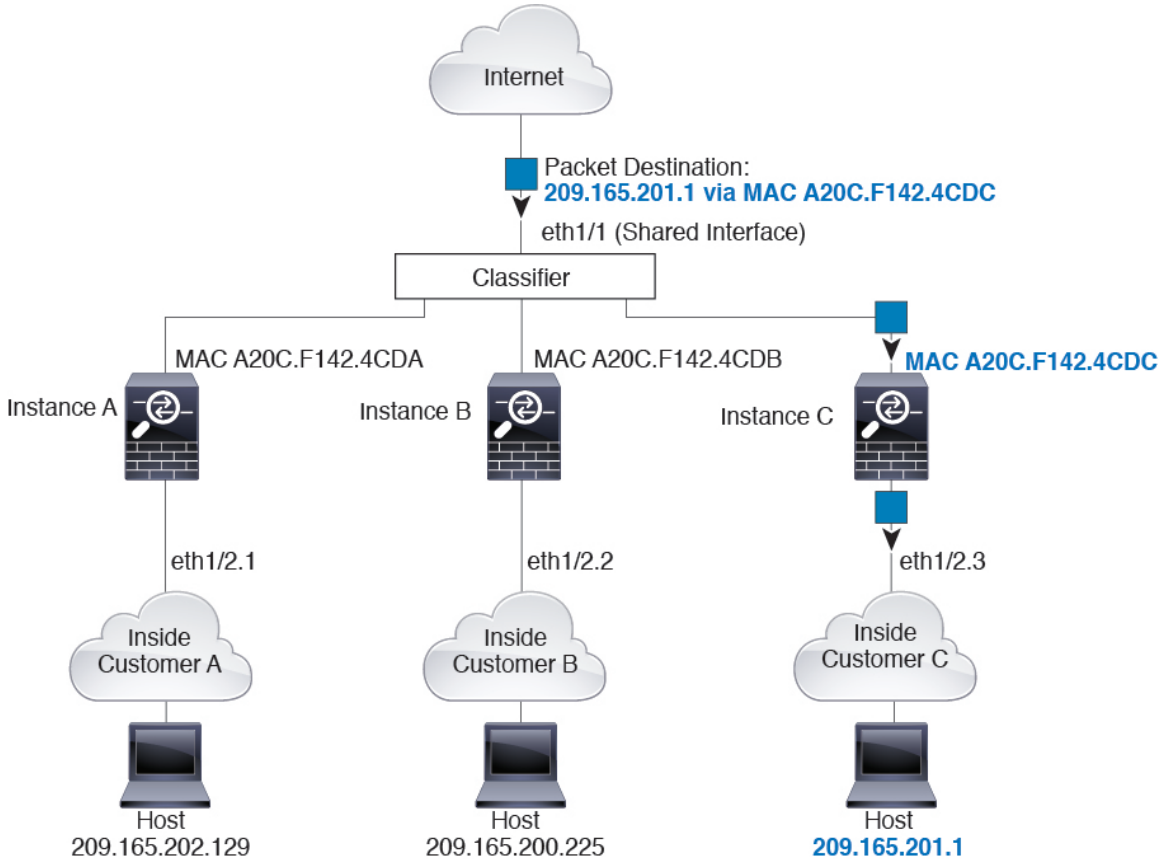
Note If the destination MAC address is a multicast or broadcast MAC address, the packet is duplicated and delivered to each instance.

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.

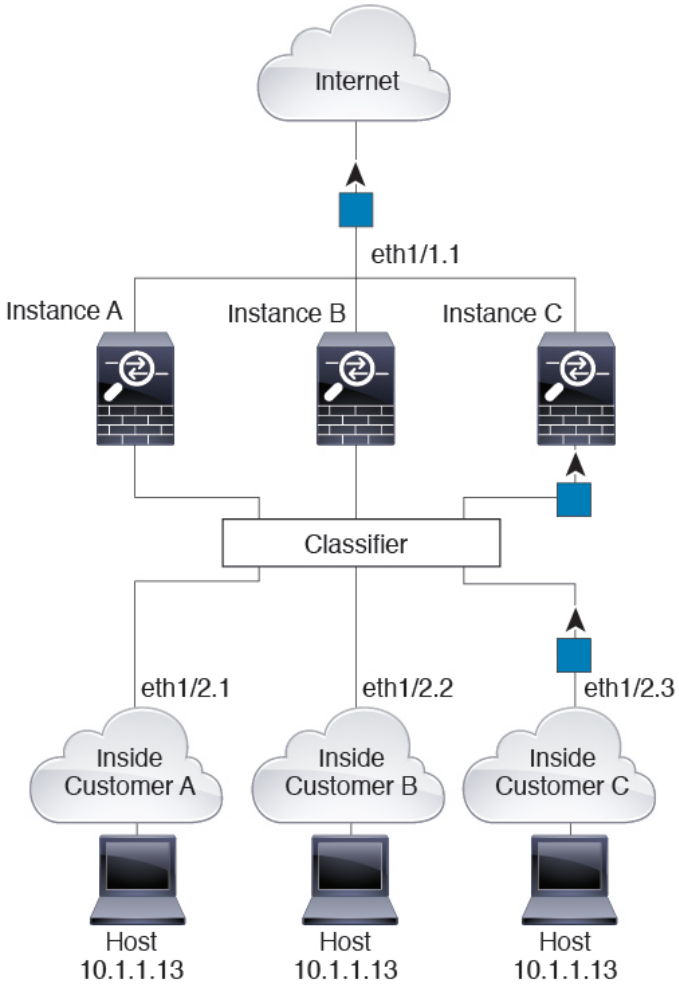
Figure 6: Packet Classification with a Shared Interface Using MAC Addresses



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.

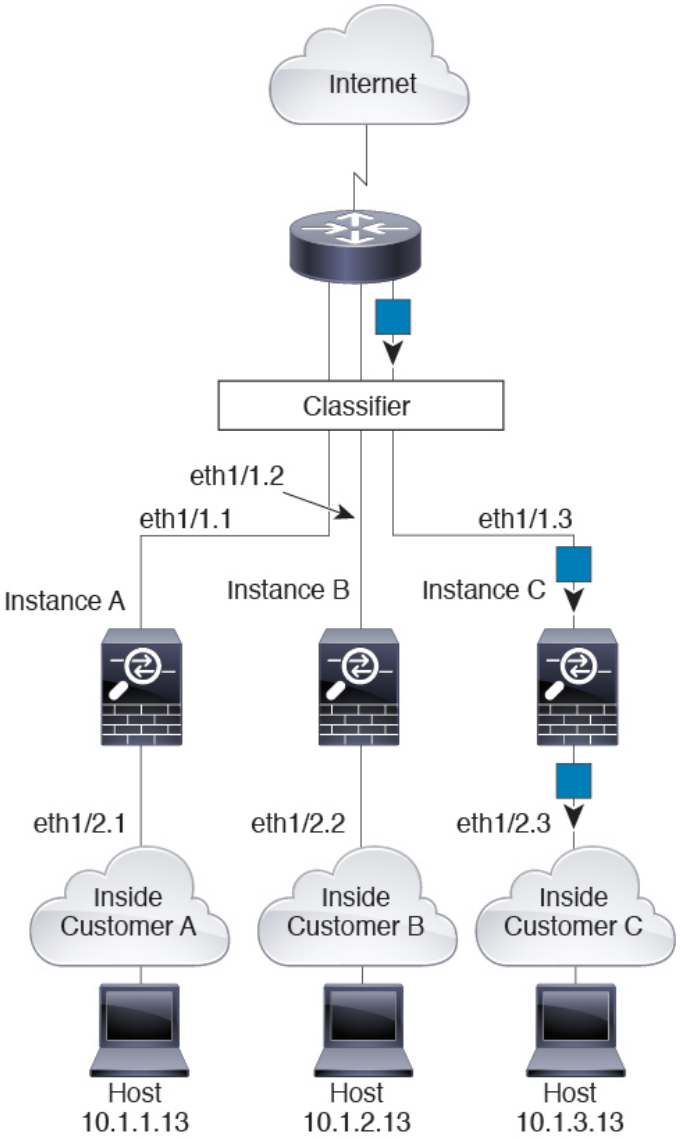
Figure 7: Incoming Traffic from Inside Networks



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.

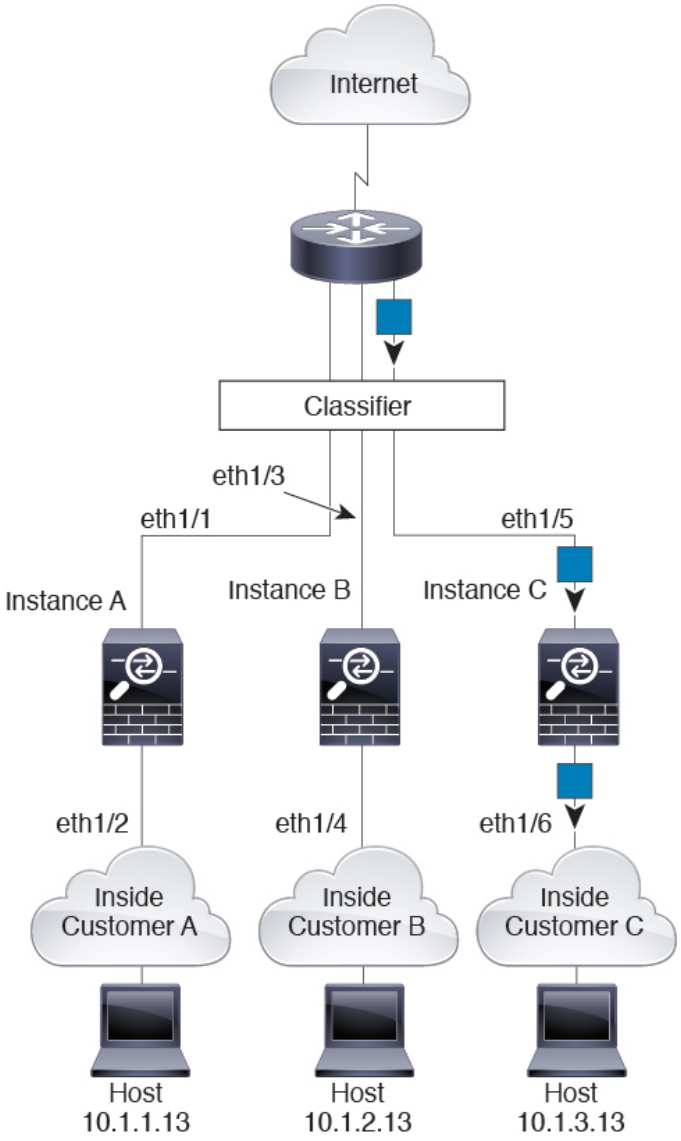
Figure 8: Transparent Firewall Instances



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 9: Inline Sets

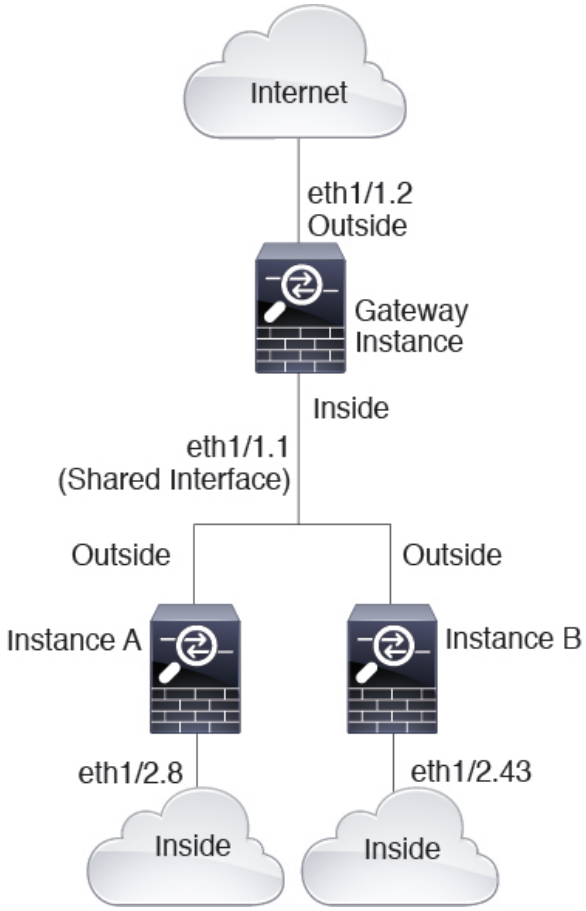


Cascading Container Instances

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 10: Cascading Instances



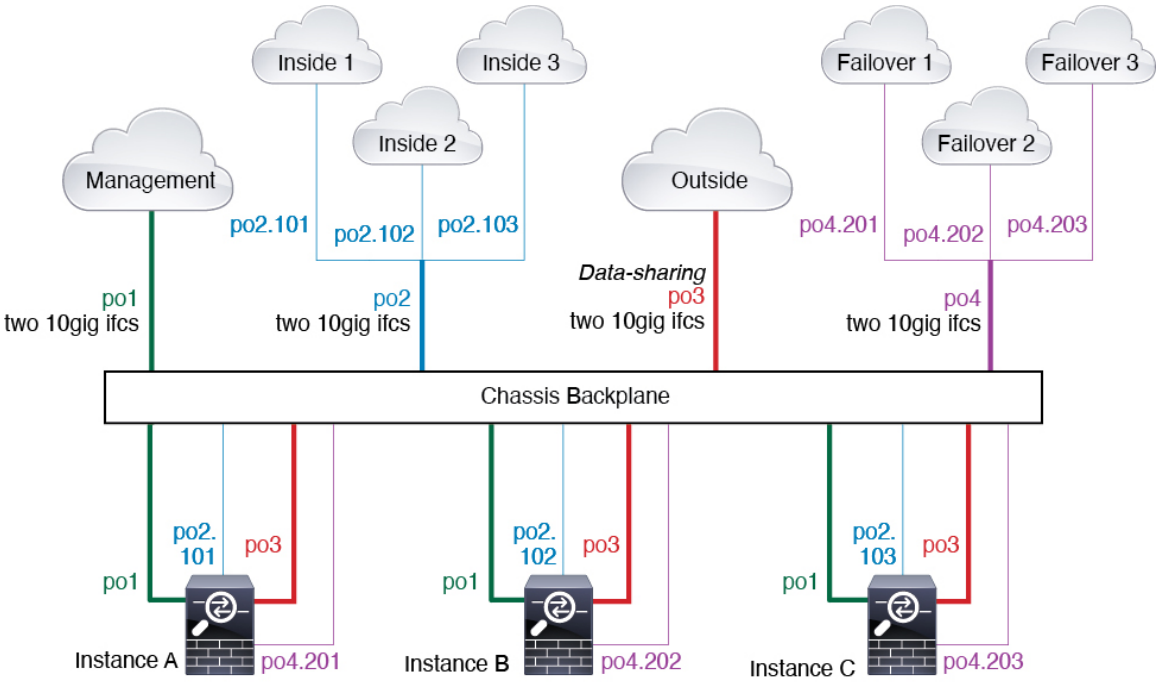
Note Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.

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 24. To add a resource profile, see Add a Resource Profile for Container Instances, on page 31.

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 25 and Add a High Availability Pair, on page 44.

End-to-End Procedure

This procedure configures all necessary elements for a multi-instance environment.

Note For multi-instance clustering, see Deploying a Cluster for Firepower Threat Defense for Scalability and High Availability or the FXOS configuration guide.

Procedure

Step 1 Add a Resource Profile for Container Instances, on page 31.

The chassis includes a default resource profile called "Default-Small," which includes the minimum number of cores. If you do not want to use only this profile, you need to add profiles before you add the container instances.

Step 2 (Optional) Add a MAC Pool Prefix and View MAC Addresses for Container Instance Interfaces, on page 33.

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. You can optionally define the prefix used in generation.

Step 3 Configure Interfaces, on page 34.

You can use interfaces with container instances the same way as you would with native instances. However, VLAN subinterfaces and data-shared interfaces are only available for container instances and provide scalability and flexibility for your deployment. Be sure to read about limitations on shared interfaces in this guide.

Step 4 Add a Standalone Firewall Threat Defense for the Firewall Management Center, on page 38.

Container instances are only supported for Firewall Threat Defense as a standalone device or a failover pair; clustering is not supported.

Step 5 Add a High Availability Pair, on page 44.

If you want to deploy a High Availability pair, refer to the requirements in this section.

Licenses for Container Instances

All licenses are consumed per security engine/chassis (for the Firepower 4100) or per security module (for the Firepower 9300), and not per container instance. See the following details:

- Base licenses are automatically assigned: one per security module/engine.
- Feature licenses are manually assigned to each instance; but you only consume one license per feature per security module/engine. For example, for the Firepower 9300 with 3 security modules, you only need one URL Filtering license per module for a total of 3 licenses, regardless of the number of instances in use.

For example:

Table 6: Sample License Usage for Container Instances on a Firepower 9300

Firepower 9300	Instance	Licenses
Security Module 1	Instance 1	Base, URL Filtering, Malware
	Instance 2	Base, URL Filtering
	Instance 3	Base, URL Filtering
Security Module 2	Instance 4	Base, Threat
	Instance 5	Base, URL Filtering, Malware, Threat
Security Module 3	Instance 6	Base, Malware, Threat
	Instance 7	Base, Threat

Table 7: Total Number of Licenses

Base	URL Filtering	Malware	Threat
3	2	3	2

Requirements and Prerequisites for Container Instances

For information about high-availability or clustering requirements with multi-instance, see Requirements and Prerequisites for High Availability, on page 25 and see Requirements and Prerequisites for Clustering, on page 26.

Supported Application Types

- The Firewall Threat Defense using Firewall Management Center

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 8: 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

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
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
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

Firewall Management Center Requirements

For all instances on a Firepower 4100 chassis or Firepower 9300 module, you must use the same Firewall Management Center due to the licensing implementation.

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 container instances: Do not use cascading instances (using a shared interface) with High Availability. After a failover occurs and the standby unit rejoins, MAC addresses can overlap temporarily and cause an outage. You should instead use unique interfaces for the gateway instance and inside instance using an external switch to pass traffic between the instances.
- For other High Availability system requirements, see the application configuration guide chapter for High Availability.

Requirements and Prerequisites for Clustering

Cluster Model Support

The Firewall Threat Defense supports clustering on the following models:

- Firepower 9300—You must include all modules installed in the chassis for a maximum 3-unit cluster. You can include up to 6 nodes in the cluster. 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. Supports clustering with multiple chassis and clustering isolated to security modules within one chassis.
- Firepower 4100—Supported for up to 6 nodes using clustering with multiple chassis.

User roles

- Admin
- Access Admin
- Network Admin

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.



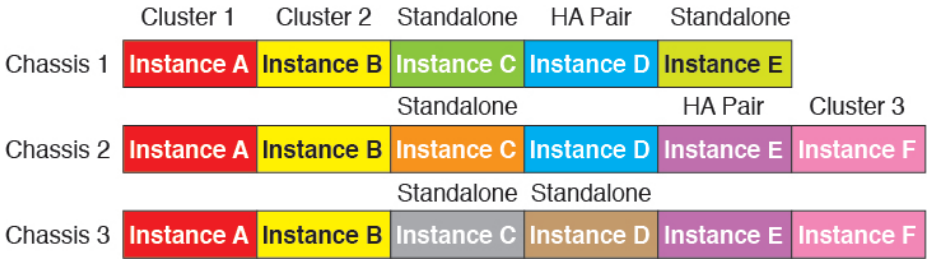
- Must run the identical FXOS and application software except at the time of an image upgrade. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- Must include the same interface configuration for interfaces you assign to the cluster, 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 clusters with multiple chassis. 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 Firewall Threat Defense, the Firewall Management Center 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 Firewall Threat Defense, all licensing is handled by the Firewall Management Center.

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.

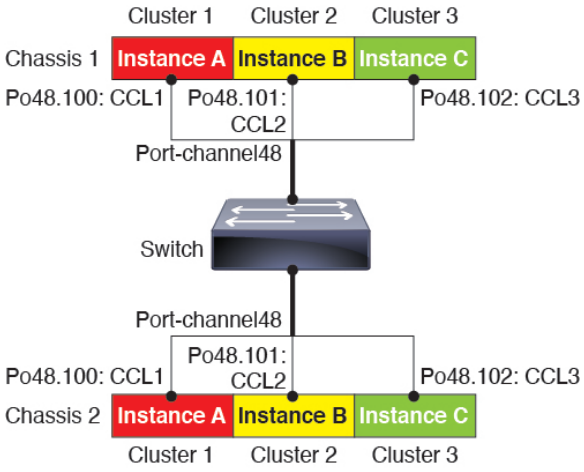


- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.



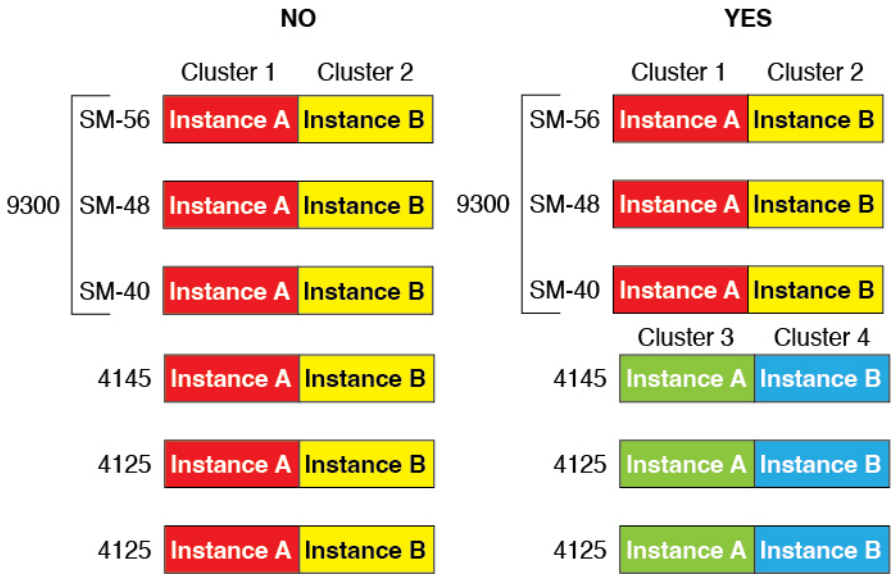
- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.

- Dedicated cluster control link—For clusters with multiple chassis, 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.



- 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.



- Maximum 6 nodes—You can use up to six container instances in a cluster.

Switch Requirements

- 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.

Guidelines and Limitations

General Guidelines

- Multi-instance capability with container instances is only available for the Firewall Threat Defense using Firewall Management Center.
- For Firewall Threat Defense container instances, a single Firewall Management Center must manage all instances on a security module/engine.
- You can enable TLS crypto acceleration on up to 16 container instances.
- For Firewall Threat Defense container instances, the following features are not supported:
  - Radware DefensePro link decorator

- Firewall Management Center UCAPL/CC mode
- Flow offload to hardware

VLAN Subinterfaces

- This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the Firewall Threat Defense application. See FXOS Interfaces vs. Application Interfaces, on page 4 for more information.
- Subinterfaces (and the parent interfaces) can only be assigned to container instances.



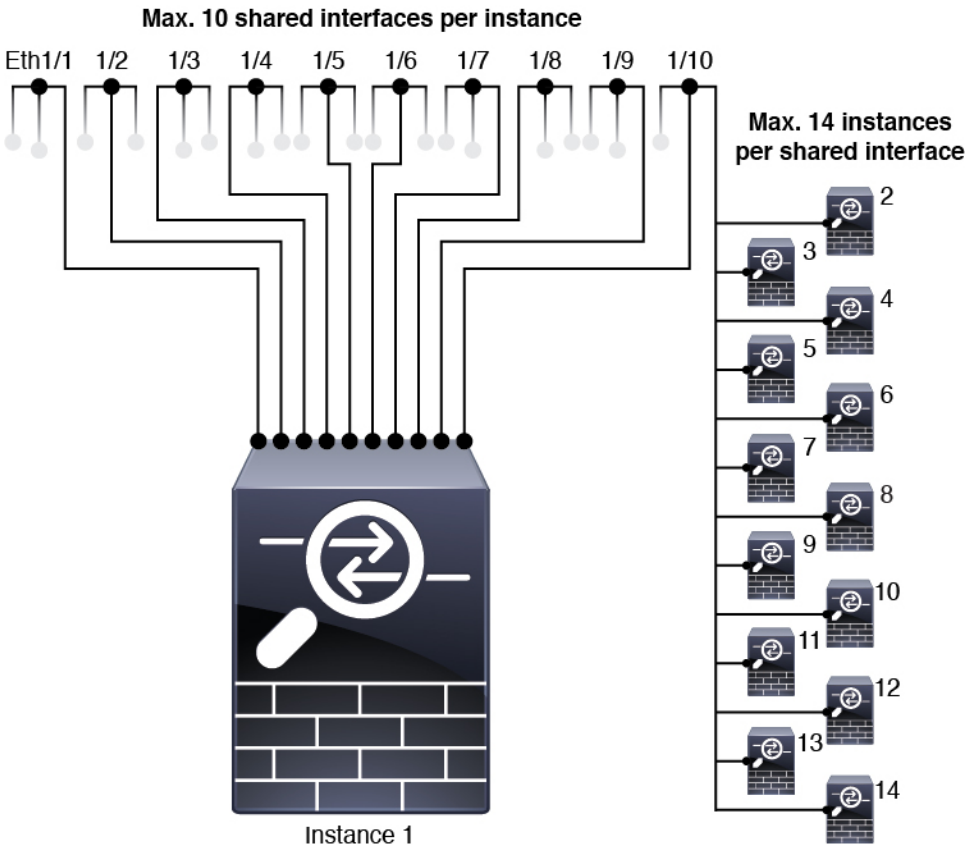
Note If you assign a parent interface to a container instance, it only passes untagged (non-VLAN) traffic. Do not assign the parent interface unless you intend to pass untagged traffic. For Cluster type interfaces, the parent interface cannot be used.

- 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 Firewall Threat Defense 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 Firewall Threat Defense inline sets or passive interfaces.
  - You cannot use a data-sharing interface for the failover link.

Default MAC Addresses

- 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 21.

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 24.

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.

Changing the resource profile after you assign it is disruptive. See the following guidelines:

- 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 change the resource profile settings after you add the Firewall Threat Defense instance to the Firewall Management Center, then update the inventory for each unit on the Firewall Management Center.

  Choose **Devices** > **Device Management**, click **Edit** ( ) for the instance, then click **Refresh** ( ) on the **Device** > **Inventory Details** area.
- If you assign a different profile to an instance, it reboots.
- If you assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:
  1. Break high availability.
  2. Assign the new profile to both units.
  3. Re-establish high availability.
- If you assign a different profile to instances in an established cluster, which allows mismatched profiles, then apply the new profile on the data nodes first; after they all come back up, you can apply the new profile to the control node.

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**.

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 21 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.

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

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.
- If you replace an SFP on a port with a different SFP module, the speed, duplex, and auto-negotiation of the interface is not updated automatically. You must manually re-configure the interface.

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 Firewall Threat Defense 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.

If a peer switch connecting to the port over a 50G cable does not support auto-negotiation, ensure to disable auto-negotiation on the switch and the platform interface as well. For example, N9K-C93400LD-H1 does not support auto-negotiation on a 50G cable. Hence, for the port to be connected you must disable the default auto-negotiation on the platform and the switch.

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.

Note
Configuring Debounce Time is not supported on 1G interfaces.

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 Firewall Threat Defense 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 If your interface supports **Auto Negotiation**, click the **Yes** or **No** radio button.

Note
If a port-channel is upgraded from 1G to 10G, ensure that the **Admin Speed** is set to **10gbps** and **Auto Negotiation** is set to **No**. The10G interface members do not support auto negotiation.

Step 10 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 11 To remove an interface from the port channel, click the **Delete** button to the right of the interface in the Member ID list.

Step 12 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 Firewall Threat Defense 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.

Add a Standalone Firewall Threat Defense for the Firewall Management Center

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.

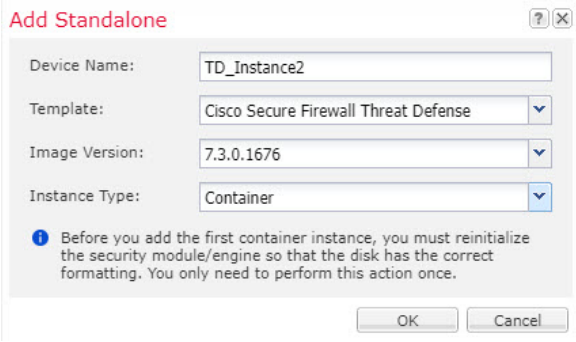
Note For the Firepower 9300, you can install different application types (ASA and Firewall Threat Defense) on separate modules in the chassis. You can also run different versions of an application instance type on separate modules.

- 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 2 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 31.
- 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.
- Gather the following information:
  - Interface IDs for this device
  - Management interface IP address and network mask
  - Gateway IP address
  - Firewall Management Center IP address and/or NAT ID of your choosing
  - DNS server IP address
  - Firewall Threat Defense 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.

Note
You cannot change this name after you add the logical device.

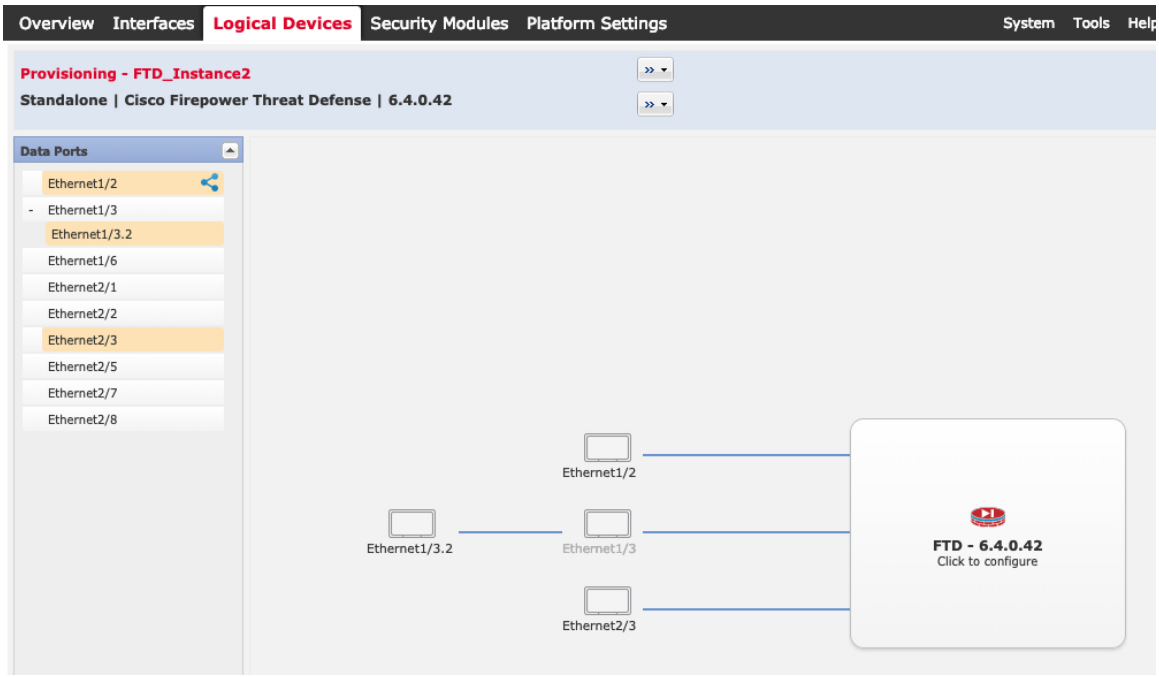
b) For the **Template**, choose **Cisco Firepower Threat Defense**.
c) Choose the **Image Version**.
d) Choose the **Instance Type**: **Container** or **Native**.

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 Firewall Management Center, including setting the IP addresses.

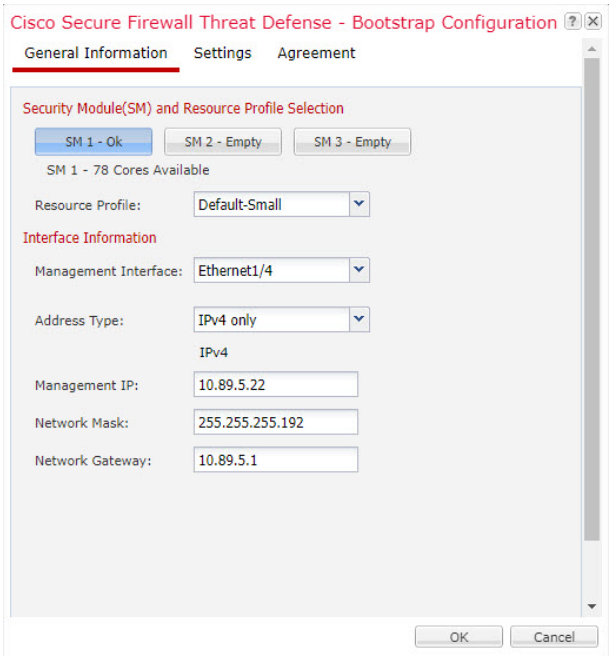
You can only assign up to 10 data-sharing interfaces to a container instance. Also, each data-sharing interface can be assigned to at most 14 container instances. A data-sharing interface is indicated by the sharing icon ( ).

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 Firewall Management Center 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
If you later assign a different profile to instances in an established high-availability pair, which requires the profile to be the same on both units, you must:

1. Break high availability.
2. Assign the new profile to both units.
3. Re-establish high availability.

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**.

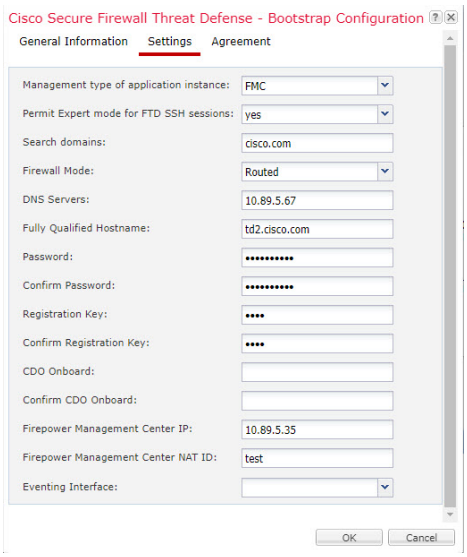
e) Configure the **Management IP** address.

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 Firewall Device Manager 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 Firewall Management Center. If you do not know the Firewall Management Center 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 Firewall Threat Defense 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 Firewall Threat Defense CLI.

d) Enter the **Search Domains** as a comma-separated list.

e) Choose the **Firewall Mode**: **Transparent** or **Routed**.

   In routed mode, the Firewall Threat Defense 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 Firewall Threat Defense uses DNS if you specify a hostname for the Firewall Management Center, for example.

g) Enter the **Fully Qualified Hostname** for the Firewall Threat Defense.

h) Enter a **Registration Key** to be shared between the Firewall Management Center 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 Firewall Management Center when you add the Firewall Threat Defense.

i) Enter a **Password** for the Firewall Threat Defense 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.

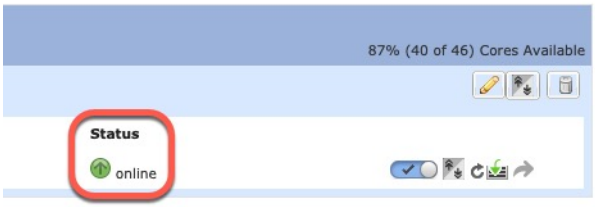
Step 7 On the **Agreement** tab, read and accept the end user license agreement (EULA).

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 Firewall Management Center configuration guide to add the Firewall Threat Defense as a managed device and start configuring your security policy.

Add a High Availability Pair

Firewall Threat Defense 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 25.

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.

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

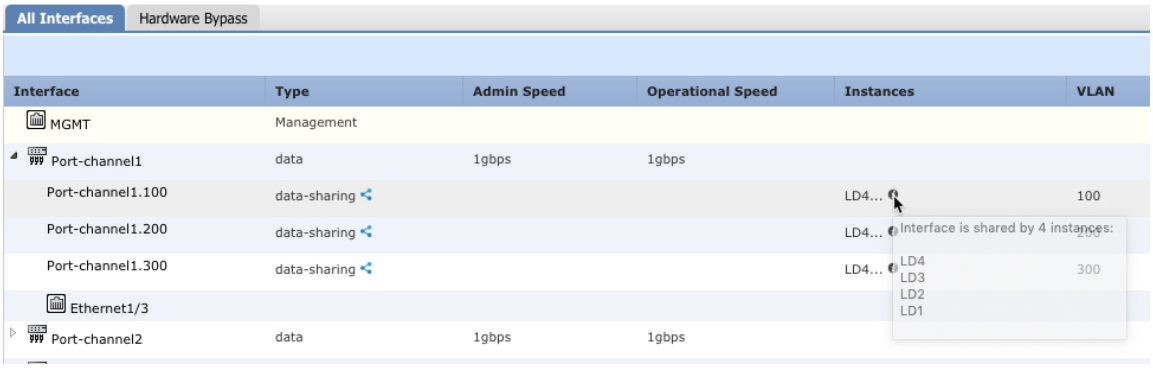
Remove the subinterface from all logical devices in the group at the CLI, and then save the configuration so that the removal is simultaneous. For example:

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 Firewall Chassis Manager, you can hover over each shared subinterface to see which instances it is allocated to.

Figure 11: 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 Multi-Instance Capability

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.

Feature Name	Platform Releases	Feature Information
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.
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.
Firepower 4115, 4125, and 4145	2.6.1	We introduced the Firepower 4115, 4125, and 4145. **Note** Requires Firepower 6.4. No modified screens.
Support for ASA and Firewall Threat Defense on separate modules of the same Firepower 9300	2.6.1	You can now deploy ASA and Firewall Threat Defense logical devices on the same Firepower 9300. **Note** Requires ASA 9.12(1) and Firepower 6.4.0. No modified screens.

Feature Name	Platform Releases	Feature Information
For the Firewall Threat Defense bootstrap configuration, you can now set the NAT ID for the Firewall Management Center in the Firepower Chassis Manager	2.6.1	You can now set the Firewall Management Center NAT ID in the Firepower Chassis Manager. Previously, you could only set the NAT ID within the FXOS CLI or Firewall Threat Defense 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 Secure Firewall Management Center 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 Firepower Threat Defense	2.4.1	You can now deploy multiple logical devices, each with a Firepower Threat Defense 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 Firepower Threat Defense feature support. Multiple context mode, due to shared resources, supports more contexts on a given platform. Multiple context mode is not available on the Firepower Threat Defense. **Note** Requires Firewall Threat Defense 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 Firepower Management Center screens: **Devices** > **Device Management** > **Edit** icon > **Interfaces** tab

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2026 Cisco Systems, Inc. All rights reserved.