



Logical Devices

- [About Logical Devices, on page 1](#)
- [Requirements and Prerequisites for Logical Devices, on page 9](#)
- [Guidelines and Limitations for Logical Devices, on page 17](#)
- [Add a Standalone Logical Device, on page 22](#)
- [Add a High Availability Pair, on page 35](#)
- [Add a Cluster, on page 36](#)
- [Configure Radware DefensePro, on page 58](#)
- [Configure TLS Crypto Acceleration, on page 63](#)
- [Enable FTD Link State Synchronization, on page 66](#)
- [Manage Logical Devices, on page 67](#)
- [Logical Devices Page, on page 77](#)
- [Examples for Inter-Site Clustering, on page 80](#)
- [History for Logical Devices, on page 84](#)

About Logical Devices

A logical device lets you run one application instance (either ASA or Firepower Threat Defense) and also one optional decorator application (Radware DefensePro) to form a service chain.

When you add a logical device, you also define the application instance type and version, assign interfaces, and configure bootstrap settings that are pushed to the application configuration.



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

Standalone and Clustered Logical Devices

You can add the following logical device types:

- **Standalone**—A standalone logical device operates as a standalone unit or as a unit in a High Availability pair.

- Cluster—A clustered logical device lets you group multiple units together, providing all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. Multiple module devices, like the Firepower 9300, support intra-chassis clustering. For the Firepower 9300, all three modules must participate in the cluster, for both native and container instances. The FDM does not support clustering.

Logical Device Application Instances: Container and Native

Application instances run in the following deployment types:

- Native instance—A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance.
- Container instance—A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances. Multi-instance capability is only supported for the Firepower Threat Defense using FMC; it is not supported for the ASA or the Firepower Threat Defense using FDM.



Note Multi-instance capability is similar to ASA multiple context mode, although the implementation is different. Multiple context mode partitions a single application instance, while multi-instance capability allows independent container instances. Container instances allow hard resource separation, separate configuration management, separate reloads, separate software updates, and full 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.

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](#) and [Add a VLAN Subinterface for Container Instances](#).



Note This document discusses *FXOS* VLAN subinterfaces only. You can separately create subinterfaces within the FTD application. See [FXOS Interfaces vs. Application Interfaces](#) for more information.

How the Chassis Classifies Packets

Each packet that enters the chassis must be classified, so that the chassis can determine to which instance to send a packet.

- **Unique Interfaces**—If only one instance is associated with the ingress interface, the chassis classifies the packet into that instance. For bridge group member interfaces (in transparent mode or routed mode), inline sets, or passive interfaces, this method is used to classify packets at all times.
- **Unique MAC Addresses**—The chassis automatically generates unique MAC addresses for all interfaces, including shared interfaces. If multiple instances share an interface, then the classifier uses unique MAC addresses assigned to the interface in each instance. An upstream router cannot route directly to an instance without unique MAC addresses. You can also set the MAC addresses manually when you configure each interface within the application.




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

Classification Examples

Packet Classification with a Shared Interface Using MAC Addresses

The following figure shows multiple instances sharing an outside interface. The classifier assigns the packet to Instance C because Instance C includes the MAC address to which the router sends the packet.


Figure 1: 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 2: 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 3: 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 4: 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.

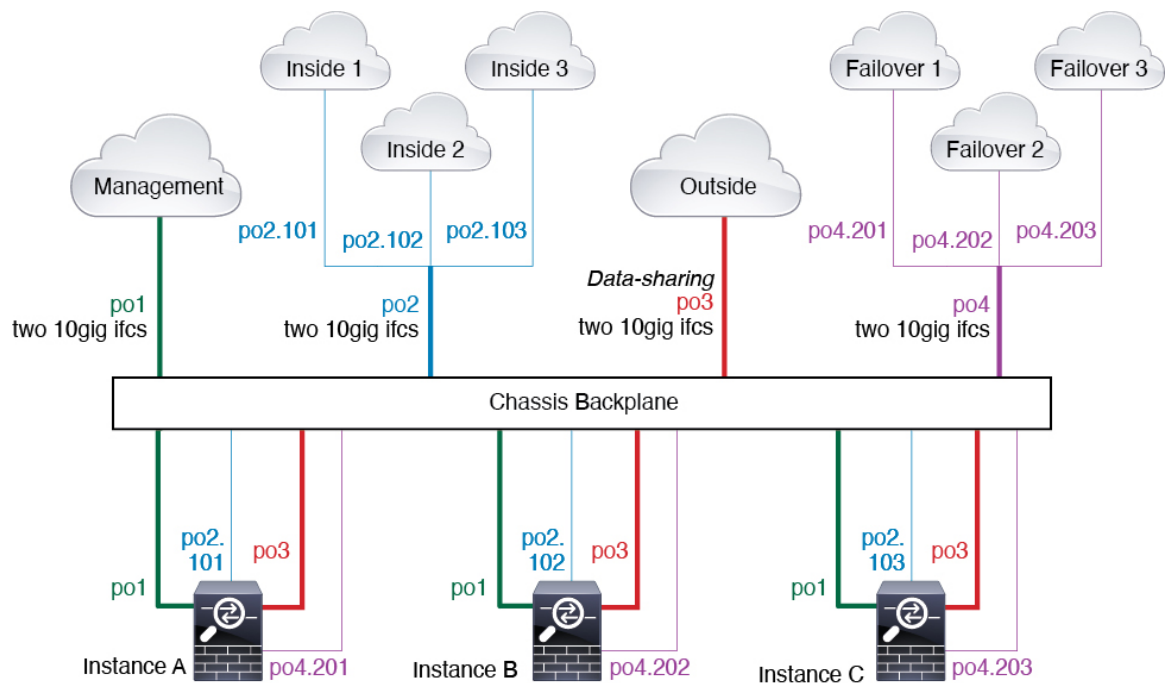
Figure 5: Cascading instances



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 Gigabit 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 Gigabit 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 Gigabit 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 Gigabit 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:

`A2xx.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 (`xyyy`) to match the chassis native form:

`A24D.00zz.zzzz`

For a prefix of 1009 (`03F1`), the MAC address is:

`A2F1.03zz.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 16](#). To add a resource profile, see [Add a Resource Profile for Container Instances](#).

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 15](#) and [Add a High Availability Pair, on page 35](#).

Container Instances and Clustering

You can create a cluster of container instances using one container instance per security module/engine. See [Requirements and Prerequisites for Clustering, on page 11](#) for detailed requirements.

Requirements and Prerequisites for Logical Devices

See the following sections for requirements and prerequisites.

Requirements and Prerequisites for Hardware and Software Combinations

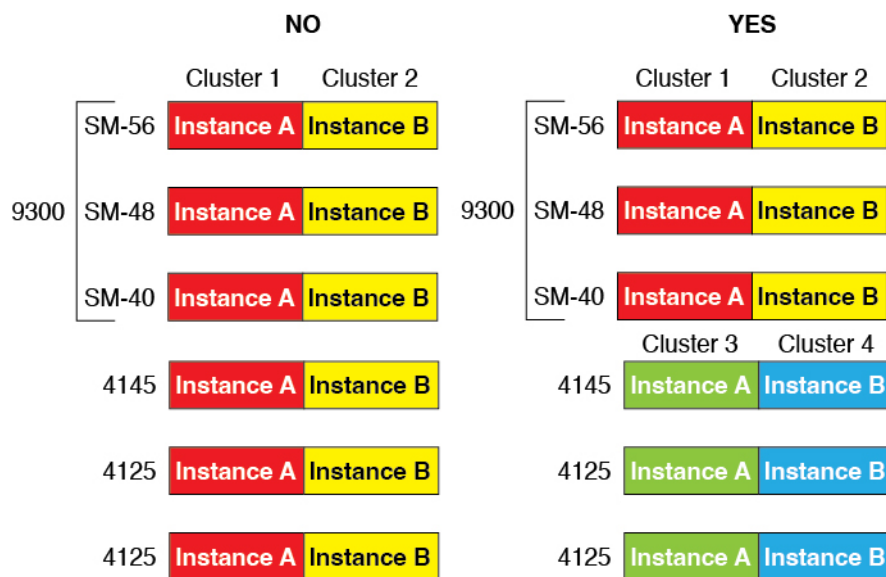
The Firepower 4100/9300 supports multiple models, security modules, application types, and high availability and scalability features. See the following requirements for allowed combinations.

Firepower 9300 Requirements

The Firepower 9300 includes 3 security module slots and multiple types of security modules. See the following requirements:

- Security Module Types—You can install modules of different types in the Firepower 9300. For example, you can install the SM-48 as module 1, SM-40 as module 2, and SM-56 as module 3.

- Native and Container instances—When you install a container instance on a security module, that module can only support other container instances. A native instance uses all of the resources for a module, so you can only install a single native instance on a module. You can use native instances on some modules, and container instances on the other module. For example, you can install a native instance on module 1 and module 2, but container instances on module 3.
- Native instance Clustering—All security modules in the cluster, whether it is intra-chassis or inter-chassis, must be the same type. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots. For example, you can install 2 SM-40s in chassis 1, and 3 SM-40s in chassis 2. You cannot use clustering if you install 1 SM-48 and 2 SM-40s in the same chassis.
- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.



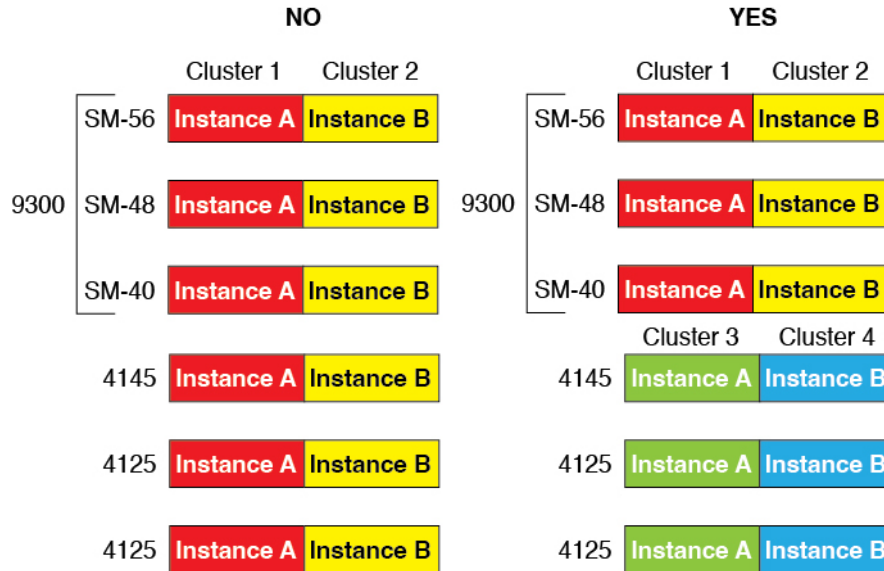
- High Availability—High Availability is only supported between same-type modules on the Firepower 9300. However, the two chassis can include mixed modules. For example, each chassis has an SM-40, SM-48, and SM-56. You can create High Availability pairs between the SM-40 modules, between the SM-48 modules, and between the SM-56 modules.
- ASA and FTD application types—You can install different application types on separate modules in the chassis. For example, you can install ASA on module 1 and module 2, and FTD on module 3.
- ASA or FTD versions—You can run different versions of an application instance type on separate modules, or as separate container instances on the same module. For example, you can install the FTD 6.3 on module 1, FTD 6.4 on module 2, and FTD 6.5 on module 3.

Firepower 4100 Requirements

The Firepower 4100 comes in multiple models. See the following requirements:

- Native and Container instances—When you install a container instance on a Firepower 4100, that device can only support other container instances. A native instance uses all of the resources for a device, so you can only install a single native instance on the device.

- Native instance Clustering—All chassis in the cluster must be the same model.
- Container instance Clustering—You can create a cluster using instances on different model types. For example, you can create a cluster using an instance on a Firepower 4145 and a 4125. You *cannot* mix the Firepower 9300 and the Firepower 4100 in the same cluster, however.



- High Availability—High Availability is only supported between same-type models.
- ASA and FTD application types—The Firepower 4100 can only run a single application type.
- The FTD container instance versions—You can run different versions of Firepower Threat Defense as separate container instances on the same module.

Requirements and Prerequisites for Clustering

Cluster Model Support


- ASA on the Firepower 9300—Maximum 16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. Note that all modules in a chassis must belong to the cluster. Supported for intra-chassis, inter-chassis, and inter-site clustering.
- ASA on the Firepower 4100 series—Maximum 16 chassis. Supported for inter-chassis and inter-site clustering.
- FTD on the Firepower 9300 using FMC—Maximum 6 modules. For example, you can use 2 modules in 3 chassis, or 3 modules in 2 chassis, or any combination that provides a maximum of 6 modules. Note that all modules in a chassis must belong to the cluster. Supported for intra-chassis and inter-chassis clustering.
- FTD on the Firepower 4100 series using FMC—Maximum 6 chassis. Supported for inter-chassis clustering.
- Radware DefensePro—Supported for intra-chassis clustering with the ASA.

- Radware DefensePro—Supported for intra-chassis clustering with the Firepower Threat Defense. Not supported for multi-instance clustering.

Clustering Hardware and Software Requirements

All chassis in a cluster:

- Native instance clustering—For the Firepower 4100: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.
- Container instance clustering—We recommend that you use the same security module or chassis model for each cluster instance. However, you can mix and match container instances on different Firepower 9300 security module types or Firepower 4100 models in the same cluster if required. You cannot mix Firepower 9300 and 4100 instances in the same cluster. For example, you can create a cluster using an instance on a Firepower 9300 SM-56, SM-48, and SM-40. Or you can create a cluster on a Firepower 4145 and a 4125.




- Must run the identical FXOS 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 inter-chassis clustering. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data nodes, and ending with the control node.
- Must use the same NTP server. For Firepower Threat Defense, the FMC must also use the same NTP server. Do not set the time manually.


- ASA: Each FXOS chassis must be registered with the License Authority or satellite server. There is no extra cost for data nodes. For permanent license reservation, you must purchase separate licenses for each chassis. For Firepower Threat Defense, all licensing is handled by the FMC.

Multi-Instance Clustering Requirements

- No intra-security-module/engine clustering—For a given cluster, you can only use a single container instance per security module/engine. You cannot add 2 container instances to the same cluster if they are running on the same module.




- Mix and match clusters and standalone instances—Not all container instances on a security module/engine need to belong to a cluster. You can use some instances as standalone or High Availability nodes. You can also create multiple clusters using separate instances on the same security module/engine.



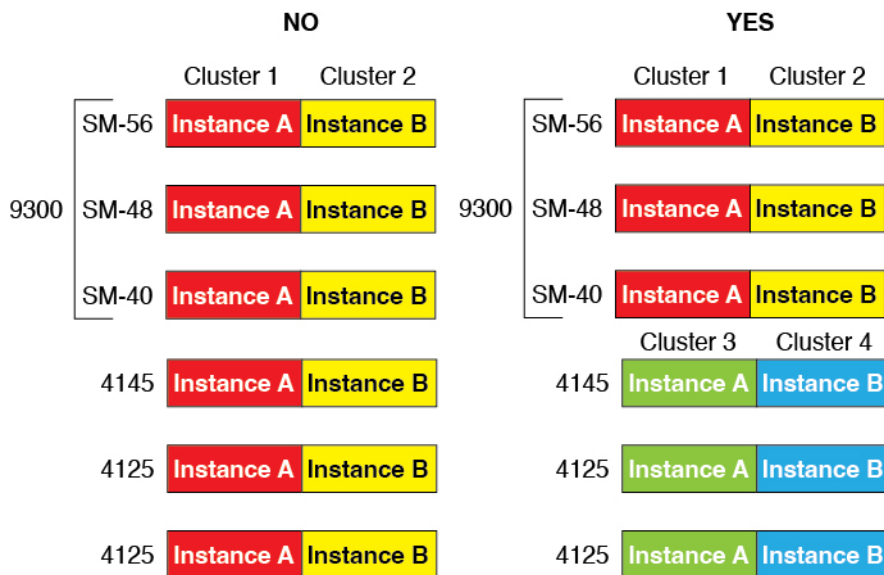
- All 3 modules in a Firepower 9300 must belong to the cluster—For the Firepower 9300, a cluster requires a single container instance on all 3 modules. You cannot create a cluster using instances on module 1 and 2, and then use a native instance on module 3, or example.



- Match resource profiles—We recommend that each node in the cluster use the same resource profile attributes; however, mismatched resources are allowed when changing cluster nodes to a different resource profile, or when using different models.
- Dedicated cluster control link—For inter-chassis clustering, each cluster needs a dedicated cluster control link. For example, each cluster can use a separate subinterface on the same cluster-type EtherChannel, or use separate EtherChannels.



- No shared interfaces—Shared-type interfaces are not supported with clustering. However, the same Management and Eventing interfaces can be 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 for Inter-Chassis Clustering

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.

- For supported switch characteristics, see [Cisco FXOS Compatibility](#).

Sizing the Data Center Interconnect for Inter-Site Clustering

You should reserve bandwidth on the data center interconnect (DCI) for cluster control link traffic equivalent to the following calculation:

$$\frac{\text{\# of cluster members per site}}{2} \times \text{cluster control link size per member}$$

If the number of members differs at each site, use the larger number for your calculation. The minimum bandwidth for the DCI should not be less than the size of the cluster control link for one member.

For example:

- For 4 members at 2 sites:
 - 4 cluster members total
 - 2 members at each site
 - 5 Gbps cluster control link per member

Reserved DCI bandwidth = 5 Gbps ($2/2 \times 5$ Gbps).

- For 6 members at 3 sites, the size increases:
 - 6 cluster members total
 - 3 members at site 1, 2 members at site 2, and 1 member at site 3
 - 10 Gbps cluster control link per member

Reserved DCI bandwidth = 15 Gbps ($3/2 \times 10$ Gbps).

- For 2 members at 2 sites:
 - 2 cluster members total
 - 1 member at each site
 - 10 Gbps cluster control link per member

Reserved DCI bandwidth = 10 Gbps ($1/2 \times 10$ Gbps = 5 Gbps; but the minimum bandwidth should not be less than the size of the cluster control link (10 Gbps)).

Requirements and Prerequisites for High Availability

- The two units in a High Availability Failover configuration must:
 - Be on a separate chassis; intra-chassis High Availability for the Firepower 9300 is not supported.
 - Be the same model.
 - Have the same interfaces assigned to the High Availability logical devices.

- Have the same number and types of interfaces. All interfaces must be preconfigured in FXOS identically before you enable High Availability.
- High Availability is only supported between same-type modules on the Firepower 9300; but the two chassis can include mixed modules. For example, each chassis has an SM-56, SM-48, and SM-40. You can create High Availability pairs between the SM-56 modules, between the SM-48 modules, and between the SM-40 modules.
- For container instances, each unit must use the same resource profile attributes.
- For other High Availability system requirements, see the application configuration guide chapter for High Availability.

Requirements and Prerequisites for Container Instances

Supported Application Types

- The FTD using FMC

Maximum Container Instances and Resources per Model

For each container instance, you can specify the number of CPU cores to assign to the instance. RAM is dynamically allocated according to the number of cores, and disk space is set to 40 GB per instance.

Table 1: Maximum Container Instances and Resources per Model

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 4110	3	22	53 GB	125.6 GB
Firepower 4112	3	22	78 GB	308 GB
Firepower 4115	7	46	162 GB	308 GB
Firepower 4120	3	46	101 GB	125.6 GB
Firepower 4125	10	62	162 GB	644 GB
Firepower 4140	7	70	222 GB	311.8 GB
Firepower 4145	14	86	344 GB	608 GB
Firepower 4150	7	86	222 GB	311.8 GB
Firepower 9300 SM-24 security module	7	46	226 GB	656.4 GB
Firepower 9300 SM-36 security module	11	70	222 GB	640.4 GB
Firepower 9300 SM-40 security module	13	78	334 GB	1359 GB

Model	Max. Container Instances	Available CPU Cores	Available RAM	Available Disk Space
Firepower 9300 SM-44 security module	14	86	218 GB	628.4 GB
Firepower 9300 SM-48 security module	15	94	334 GB	1341 GB
Firepower 9300 SM-56 security module	18	110	334 GB	1314 GB

FMC Requirements

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

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the Firepower Threat Defense and ASA.

High Availability

- Configure high availability within the application configuration.
- You can use any data interfaces as the failover and state links. Data-sharing interfaces are not supported.

Multi-Instance and Context Mode

- Multiple context mode is only supported on the ASA.
- Enable multiple context mode in the ASA after you deploy.
- Multi-instance capability with container instances is only available for the Firepower Threat Defense using FMC.
- For Firepower Threat Defense container instances, a single FMC must manage all instances on a security module/engine.
- You can enable TLS crypto acceleration on up to 16 container instances.
- For Firepower Threat Defense container instances, the following features are not supported:
 - Radware DefensePro link decorator

- FMC UCAPL/CC mode
- Flow offload to hardware

Clustering Guidelines and Limitations

Switches for Inter-Chassis Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- Some switches do not support dynamic port priority with LACP (active and standby links). You can disable dynamic port priority to provide better compatibility with Spanned EtherChannels.
- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

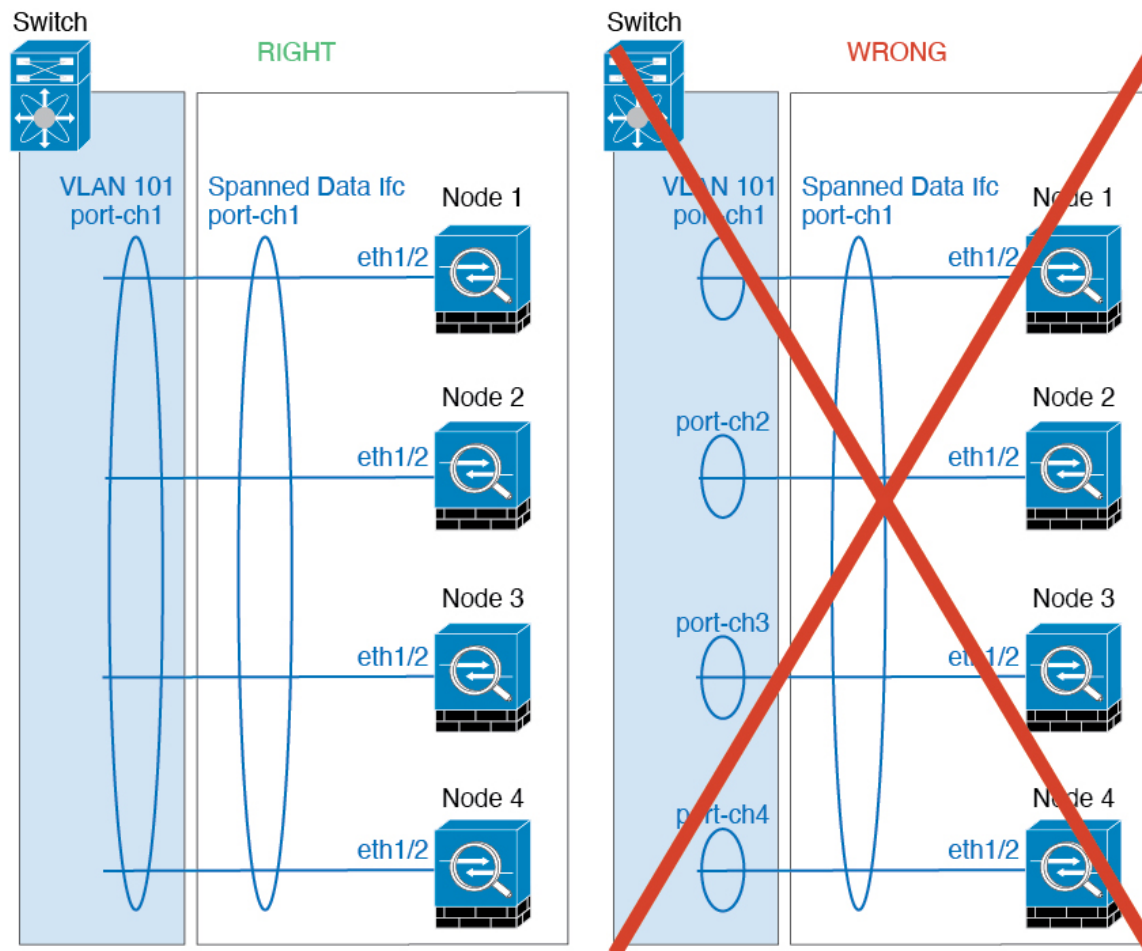

```
router(config)# port-channel id hash-distribution fixed
```

Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.
- Firepower 4100/9300 clusters support LACP graceful convergence. So you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default.

Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

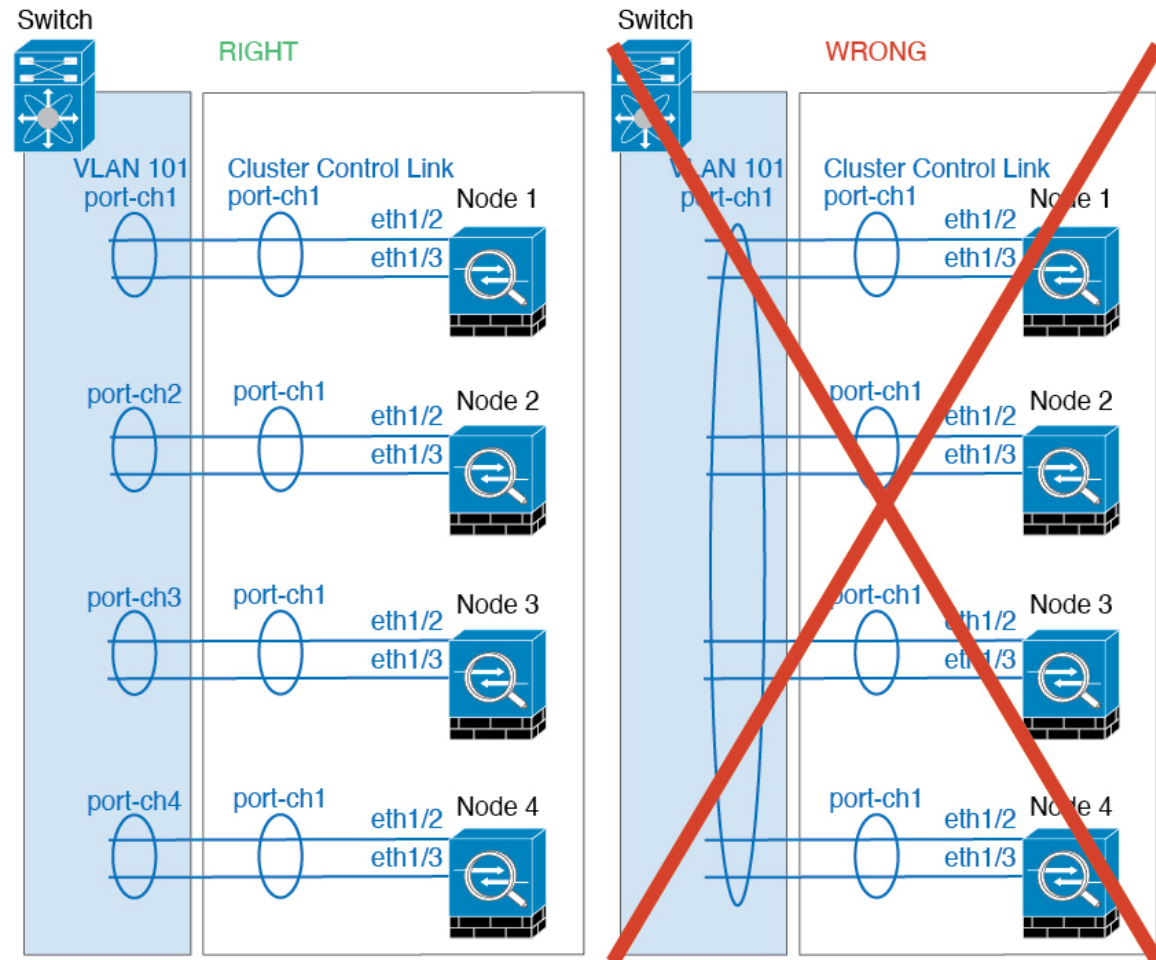
EtherChannels for Inter-Chassis Clustering

- In Catalyst 3750-X Cisco IOS software versions earlier than 15.1(1)S2, the cluster unit did not support connecting an EtherChannel to a switch stack. With default switch settings, if the cluster unit EtherChannel is connected cross stack, and if the control unit switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to a more stable switch software version, such as 15.1(1)S2.
- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.
 - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.



- Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels

on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.



Inter-Site Clustering

See the following guidelines for inter-site clustering:

- The cluster control link latency must be less than 20 ms round-trip time (RTT).
- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.
- The does not encrypt forwarded data traffic on the cluster control link because it is a dedicated link, even when used on a Data Center Interconnect (DCI). If you use Overlay Transport Virtualization (OTV), or are otherwise extending the cluster control link outside of the local administrative domain, you can configure encryption on your border routers such as 802.1AE MacSec over OTV.
- The cluster implementation does not differentiate between members at multiple sites for incoming connections; therefore, connection roles for a given connection may span across sites. This is expected

behavior. However, if you enable director localization, the local director role is always chosen from the same site as the connection owner (according to site ID). Also, the local director chooses a new owner at the same site if the original owner fails (Note: if the traffic is asymmetric across sites, and there is continuous traffic from the remote site after the original owner fails, then a node from the remote site might become the new owner if it receives a data packet within the re-hosting window.).

- For director localization, the following traffic types do not support localization: NAT or PAT traffic; SCTP-inspected traffic; Fragmentation owner query.
- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.
- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.
- For transparent mode, if the cluster is connected to an HSRP router, you must add the router HSRP MAC address as a static MAC address table entry on the . When adjacent routers use HSRP, traffic destined to the HSRP IP address will be sent to the HSRP MAC Address, but return traffic will be sourced from the MAC address of a particular router's interface in the HSRP pair. Therefore, the MAC address table is typically only updated when the ARP table entry for the HSRP IP address expires, and the sends an ARP request and receives a reply. Because the 's ARP table entries expire after 14400 seconds by default, but the MAC address table entry expires after 300 seconds by default, a static MAC address entry is required to avoid MAC address table expiration traffic drops.
- For routed mode using Spanned EtherChannel, configure site-specific MAC addresses. Extend the data VLANs across the sites using OTV, or something similar. You need to create filters to prevent traffic that is destined to the global MAC address from being sent over the DCI to the other site. If the cluster becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's cluster nodes. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

Additional Guidelines

- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- We recommend connecting EtherChannels to a VSS, vPC, StackWise, or StackWise Virtual for redundancy.

- Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.
- For decrypted TLS/SSL connections, the decryption states are not synchronized, and if the connection owner fails, then decrypted connections will be reset. New connections will need to be established to a new unit. Connections that are not decrypted (they match a do-not-decrypt rule) are not affected and are replicated correctly.

Defaults

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Add a Standalone Logical Device

Standalone logical devices can be used alone or as high availability units. For more information about high availability usage, see [Add a High Availability Pair, on page 35](#).

Add a Standalone ASA

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

You can deploy a routed or transparent firewall mode ASA from the Firepower 4100/9300 chassis.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.



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

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).
- Gather the following information:

- Interface IDs for this device
- Management interface IP address and network mask
- Gateway IP address

Procedure

Step 1

Choose **Logical Devices**.

Step 2

Click **Add > Standalone**, and set the following parameters:

- a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

- b) For the **Template**, choose **Cisco: Adaptive Security Appliance**.
 c) Choose the **Image Version**.
 d) Click **OK**.

You see the Provisioning - *device name* window.

Step 3

Expand the **Data Ports** area, and click each port that you want to assign to the device.

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces on the ASA, including setting the IP addresses.

Step 4

Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5

On the **General Information** page, complete the following:

- a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.
 b) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

- c) Choose the management interface **Address Type: IPv4 only, IPv6 only, or IPv4 and IPv6**.
 d) Configure the **Management IP** address.

Set a unique IP address for this interface.

- e) Enter a **Network Mask** or **Prefix Length**.
- f) Enter a **Network Gateway** address.

Step 6 Click the **Settings** tab.

Cisco: Adaptive Security Appliance - Bootstrap Configuration

General Information **Settings**

Firewall Mode:

Password:

Confirm Password:

Step 7 Choose the **Firewall Mode: Routed** or **Transparent**.

In routed mode, the ASA is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.

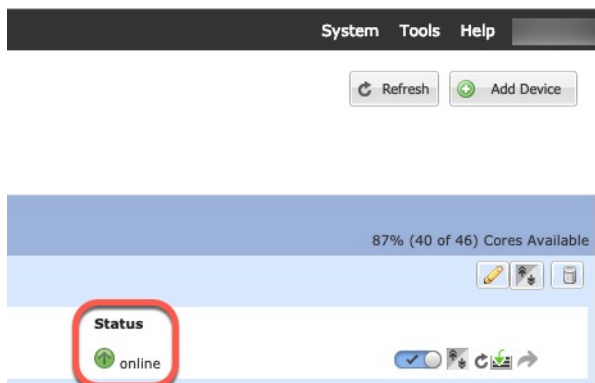
Step 8 Enter and confirm a **Password** for the admin user and for the enable password.

The pre-configured ASA admin user/password and enable password is useful for password recovery; if you have FXOS access, you can reset the admin user password/enable password if you forget it.

Step 9 Click **OK** to close the configuration dialog box.

Step 10 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.



Step 11 See the ASA configuration guide to start configuring your security policy.

Add a Standalone FTD for the FMC

Standalone logical devices work either alone or in a High Availability pair. On the Firepower 9300 with multiple security modules, you can deploy either a cluster or standalone devices. The cluster must use all modules, so you cannot mix and match a 2-module cluster plus a single standalone device, for example.

You can use native instances on some modules, and container instances on the other module(s).

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.



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

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).
- You can later enable management from a data interface; but you must assign a Management interface to the logical device even if you don't intend to use it after you enable data management. See the **configure network management-data-interface** command in the [FTD command reference](#) for more information.
- You must also configure at least one Data type interface. Optionally, you can also create a firepower-eventing interface to carry all event traffic (such as web events). See [Interface Types](#) 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](#).
- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the **Reinitialize icon**. An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance. See [Reinitializing a Security Module/Engine](#) for more information.
- Gather the following information:
 - Interface IDs for this device
 - Management interface IP address and network mask
 - Gateway IP address
 - FMC IP address and/or NAT ID of your choosing
 - DNS server IP address
 - Firepower Threat Defense hostname and domain name

Procedure

Step 1

Choose **Logical Devices**.

Step 2

Click **Add > Standalone**, and set the following parameters:

The image shows two side-by-side screenshots of configuration dialog boxes. The left dialog is titled "Add Standalone" and the right is titled "Add Device". Both dialogs have the following fields and values:

- Device Name: FTD_Instance2
- Template: Cisco Firepower Threat Defense
- Image Version: 6.5.0.1159 (left) / 6.4.0.42 (right)
- Instance Type: Container
- Usage: Standalone (selected), Cluster

Both dialogs include a warning icon and text: "Before you add the first container instance, you must reinitialize the security module/engine so that the disk has the correct formatting. You only need to perform this action once." Each dialog has "OK" and "Cancel" buttons at the bottom.

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.

c) Choose the **Image Version**.

d) Choose the **Instance Type**: **Container** or **Native**.


A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.

e) Click **OK**.


You see the Provisioning - *device name* window.


Step 3

Expand the **Data Ports** area, and click each interface that you want to assign to the device.



You can only assign data and data-sharing interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in FMC, including setting the IP addresses.

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

Hardware Bypass-capable ports are shown with the following icon: . For certain interface modules, you can enable the Hardware Bypass feature for Inline Set interfaces only (see the FMC configuration guide). Hardware Bypass ensures that traffic continues to flow between an inline interface pair during a power outage. This feature can be used to maintain network connectivity in the case of software or hardware failures. If you do not assign both interfaces in a Hardware Bypass pair, you see a warning message to make sure your assignment is intentional. You do not need to use the Hardware Bypass feature, so you can assign single interfaces if you prefer.

Step 4 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5 On the **General Information** page, complete the following:

- a) (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.
- b) For a container instance, specify the **Resource Profile**.

If you later assign a different resource profile, then the instance will reload, which can take approximately 5 minutes. Note that for established High Availability pairs, if you assign a different-sized resource profile, be sure to make all members the same size as soon as possible.

- c) Choose the **Management Interface**.
This interface is used to manage the logical device. This interface is separate from the chassis management port.
- d) Choose the management interface **Address Type: IPv4 only, IPv6 only, or IPv4 and IPv6**.
- e) Configure the **Management IP** address.
Set a unique IP address for this interface.
- f) Enter a **Network Mask** or **Prefix Length**.
- g) Enter a **Network Gateway** address.

Step 6 On the **Settings** tab, complete the following:

Cisco Firepower Threat Defense - Bootstrap Configuration

General Information **Settings** Agreement

Management type of application instance:	FMC
Firepower Management Center IP:	10.89.5.35
Search domains:	cisco.com
Firewall Mode:	Routed
DNS Servers:	10.89.5.67
Firepower Management Center NAT ID:	test
Fully Qualified Hostname:	ftd2.cisco.com
Registration Key:
Confirm Registration Key:
Password:
Confirm Password:
Eventing Interface:	

Cisco Firepower Threat Defense - Bootstrap Configuration

General Information **Settings** Agreement

Registration Key:
Confirm Registration Key:
Password:
Confirm Password:
Firepower Management Center IP:	10.89.5.35
Permit Expert mode for FTD SSH sessions:	yes
Search domains:	cisco.com
Firewall Mode:	Routed
DNS Servers:	10.89.5.67
Firepower Management Center NAT ID:	test
Fully Qualified Hostname:	ftd2.cisco.com
Eventing Interface:	

- For a native instance, in the **Management type of application instance** drop-down list, choose **FMC**.
Native instances also support FDM as a manager. After you deploy the logical device, you cannot change the manager type.
- Enter the **Firepower Management Center IP** of the managing FMC. If you do not know the FMC IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.
- For a container instance, **Permit Expert mode from FTD SSH sessions: Yes** or **No**. Expert Mode provides Firepower 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 session 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 Firepower 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 Firepower 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 Firepower Threat Defense uses DNS if you specify a hostname for the FMC, for example.

- g) Enter the **Fully Qualified Hostname** for the Firepower Threat Defense.
- h) Enter a **Registration Key** to be shared between the FMC and the device during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the Firepower Threat Defense.

- i) Enter a **Password** for the Firepower 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 FMC configuration guide to add the Firepower Threat Defense as a managed device and start configuring your security policy.

Add a Standalone FTD for the FDM

You can use the FDM with a native instance. Container instances are not supported. Standalone logical devices work either alone or in a High Availability pair.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.



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

- Configure a management interface to use with the logical device. The management interface is required. Note that this management interface is not the same as the chassis management port that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).
- You must also configure at least one Data type interface.
- Gather the following information:
 - Interface IDs for this device
 - Management interface IP address and network mask
 - Gateway IP address
 - DNS server IP address
 - FTD hostname and domain name

Procedure

Step 1

Choose **Logical Devices**.

Step 2

Click **Add > Standalone**, and set the following parameters:

a) Provide a **Device Name**.

This name is used by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

b) For the **Template**, choose **Cisco Firepower Threat Defense**.

c) Choose the **Image Version**.

d) Choose the **Instance Type: Native**.

Container instances are not supported with the FDM.

e) Click **OK**.

You see the Provisioning - *device name* window.

Step 3

Expand the **Data Ports** area, and click each interface that you want to assign to the device.

You can only assign data interfaces that you previously enabled on the **Interfaces** page. You will later enable and configure these interfaces in the FDM, including setting the IP addresses.

Step 4 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 5 On the **General Information** page, complete the following:

- (For the Firepower 9300) Under **Security Module Selection** click the security module that you want to use for this logical device.
- Choose the **Management Interface**.
This interface is used to manage the logical device. This interface is separate from the chassis management port.
- Choose the management interface **Address Type**: **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.
- Configure the **Management IP** address.
Set a unique IP address for this interface.
- Enter a **Network Mask** or **Prefix Length**.
- Enter a **Network Gateway** address.

Step 6 On the **Settings** tab, complete the following:

The screenshot shows the 'Cisco Firepower Threat Defense - Bootstrap Configuration' dialog box with the 'Settings' tab selected. The 'General Information' tab is also visible. The 'Settings' tab contains the following fields:

- Management type of application instance: **LOCALLY_MANAGED** (dropdown)
- Firepower Management Center IP: (empty text box)
- Search domains: **cisco.com** (text box)
- Firewall Mode: **Routed** (dropdown)
- DNS Servers: **10.8.9.6** (text box)
- Firepower Management Center NAT ID: (empty text box)
- Fully Qualified Hostname: **ftd.example.cisco.com** (text box)
- Registration Key: (empty text box)
- Confirm Registration Key: (empty text box)
- Password: (masked text box)
- Confirm Password: (masked text box)
- Eventing Interface: (empty dropdown)

At the bottom of the dialog box are 'OK' and 'Cancel' buttons.


- In the **Management type of application instance** drop-down list, choose **LOCALLY_MANAGED**.
Native instances also support the Firepower Management Center as a manager. If you change the manager after you deploy the logical device, then your configuration is erased and the device is reinitialized.
- Enter the **Search Domains** as a comma-separated list.
- The **Firewall Mode** only supports **Routed** mode.
- Enter the **DNS Servers** as a comma-separated list.
- Enter the **Fully Qualified Hostname** for the Firepower Threat Defense.
- Enter a **Password** for the Firepower Threat Defense admin user for CLI access.

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

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can start configuring the security policy in the application.



Step 10 See the FDM configuration guide to start configuring your security policy.

Add a High Availability Pair

FTD or ASA High Availability (also known as failover) is configured within the application, not in FXOS. However, to prepare your chassis for high availability, see the following steps.

Before you begin

See [Requirements and Prerequisites for High Availability, on page 15](#).

Procedure

Step 1 Allocate the same interfaces to each logical device.

Step 2 Allocate 1 or 2 data interfaces for the failover and state link(s).

These interfaces exchange high availability traffic between the 2 chassis. We recommend that you use a 10 GB data interface for a combined failover and state link. If you have available interfaces, you can use separate failover and state links; the state link requires the most bandwidth. You cannot use the management-type interface for the failover or state link. We recommend that you use a switch between the chassis, with no other device on the same network segment as the failover interfaces.

For container instances, data-sharing interfaces are not supported for the failover link. We recommend that you create subinterfaces on a parent interface or EtherChannel, and assign a subinterface for each instance to use as a failover link. Note that you must use all subinterfaces on the same parent as failover links. You cannot use one subinterface as a failover link and then use other subinterfaces (or the parent interface) as regular data interfaces.

Step 3 Enable High Availability on the logical devices.

Step 4 If you need to make interface changes after you enable High Availability, perform the changes on the standby unit first, and then perform the changes on the active unit.

Note For the ASA, if you remove an interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an interface to an EtherChannel), then the ASA configuration retains the original commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration in the ASA OS.

Add a Cluster

Clustering lets you group multiple devices together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices. The Firepower 9300, which includes multiple modules, supports intra-chassis clustering where you group all modules within a single chassis into a cluster. You can also use inter-chassis clustering, where multiple chassis are grouped together; inter-chassis clustering is the only option for single module devices like the Firepower 4100 series.

About Clustering on the Firepower 4100/9300 Chassis

When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

- For native instance clustering: Creates a *cluster-control link* (by default, port-channel 48) for unit-to-unit communication.

For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

For intra-chassis clustering (Firepower 9300 only), this link utilizes the Firepower 9300 backplane for cluster communications.

For inter-chassis clustering, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings. Some parts of the bootstrap configuration may be user-configurable within the application if you want to customize your clustering environment.

- Assigns data interfaces to the cluster as *Spanned* interfaces.

For intra-chassis clustering, spanned interfaces are not limited to EtherChannels, like it is for inter-chassis clustering. The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For inter-chassis clustering, you must use Spanned EtherChannels for all data interfaces.



Note Individual interfaces are not supported, with the exception of a management interface.

- Assigns a management interface to all units in the cluster.

Primary and Secondary Unit Roles

One member of the cluster is the primary unit. The primary unit is determined automatically. All other members are secondary units.

You must perform all configuration on the primary unit only; the configuration is then replicated to the secondary units.

Cluster Control Link

For native instance clustering: The cluster control link is automatically created using the Port-channel 48 interface.

For multi-instance clustering: You should pre-configure subinterfaces on one or more cluster-type EtherChannels; each instance needs its own cluster control link.

For intra-chassis clustering, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications for intra-chassis clustering. For inter-chassis clustering, you must add one or more interfaces to the EtherChannel.

For a 2-member inter-chassis cluster, do not directly connect the cluster control link from one chassis to the other chassis. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Cluster control link traffic includes both control and data traffic.

Size the Cluster Control Link for Inter-Chassis Clustering

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster control link can handle the worst-case scenarios.

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features. For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.




Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster Control Link Redundancy for Inter-Chassis Clustering

The following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the

EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.



Cluster Control Link Reliability for Inter-Chassis Clustering

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

Cluster Control Link Network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: `127.2.chassis_id.slot_id`. For multi-instance clusters, which typically use different VLAN subinterfaces of the same EtherChannel, the same IP address can be used for different clusters because of VLAN separation. You can customize this IP address when you deploy the cluster. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed. For inter-site traffic, Cisco recommends using Overlay Transport Virtualization (OTV).

Management Network

We recommend connecting all units to a single management network. This network is separate from the cluster control link.

Management Interface

You must assign a Management type interface to the cluster. This interface is a special *individual* interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit.

For the ASA, the Main cluster IP address is a fixed address for the cluster that always belongs to the current primary unit. You must configure a range of addresses so that each unit, including the current primary unit, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a primary unit changes, the Main cluster IP address moves to the new primary unit, so management of the cluster continues seamlessly. The Local IP address is used for routing, and is also useful for troubleshooting. For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current primary unit. To manage an individual member, you can connect to

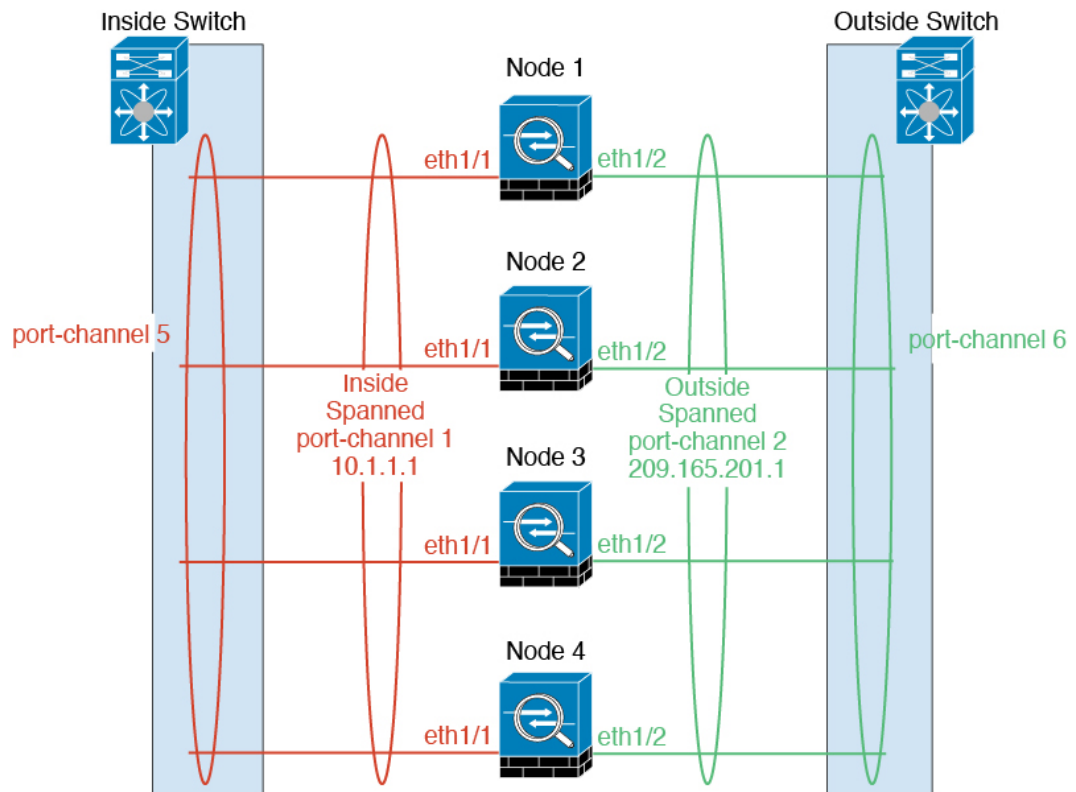
the Local IP address. For outbound management traffic such as TFTP or syslog, each unit, including the primary unit, uses the Local IP address to connect to the server.

For the Firepower Threat Defense, assign a management IP address to each unit on the same network. Use these IP addresses when you add each unit to the FMC.

Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel. A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface. The EtherChannel inherently provides load balancing as part of basic operation.

For multi-instance clusters, each cluster requires dedicated data EtherChannels; you cannot use shared interfaces or VLAN subinterfaces.



Inter-Site Clustering

For inter-site installations, you can take advantage of clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID.

Site IDs work with site-specific MAC addresses and IP addresses. Packets egressing the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. This feature prevents the switches from learning the same global MAC address from both

sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address. Site-specific MAC addresses and IP address are supported for routed mode using Spanned EtherChannels only.

Site IDs are also used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

See the following sections for more information about inter-site clustering:

- Sizing the Data Center Interconnect—[Requirements and Prerequisites for Clustering, on page 11](#)
- Inter-Site Guidelines—[Clustering Guidelines and Limitations, on page 18](#)
- Inter-Site Examples—[Examples for Inter-Site Clustering, on page 80](#)

Add an ASA Cluster

You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering. For inter-chassis clustering, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment

Create an ASA Cluster

Set the scope to the image version.

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For inter-chassis clustering, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- Gather the following information:
 - Management interface ID, IP address, and network mask
 - Gateway IP address

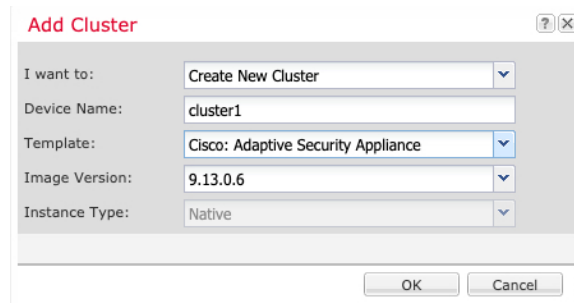
Procedure

Step 1

Configure interfaces.

Step 2 Choose **Logical Devices**.

Step 3 Click **Add > Cluster**, and set the following parameters:



Add Cluster

I want to: Create New Cluster

Device Name: cluster1

Template: Cisco: Adaptive Security Appliance

Image Version: 9.13.0.6

Instance Type: Native

OK Cancel

a) Choose **I want to:** > **Create New Cluster**

b) Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

c) For the **Template**, choose **Cisco Adaptive Security Appliance**.

d) Choose the **Image Version**.

e) For the **Instance Type**, only the **Native** type is supported.

f) Click **OK**.

You see the Provisioning - *device name* window.

Step 4 Choose the interfaces you want to assign to this cluster.

All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

Step 5 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Cisco: Adaptive Security Appliance - Bootstrap Configuration ? X

Cluster Information Settings

Security Module

Security Module-1,Security Module-2,Security Module-3

Interface Information

Chassis ID:

Site ID:

Cluster Key:

Confirm Cluster Key:

Cluster Group Name:

Management Interface:

CCL Subnet IP:

DEFAULT

Address Type:

IPv4

Management IP Pool: -

Virtual IPv4 Address:

Network Mask:

Network Gateway:

OK Cancel

- a) For inter-chassis clustering, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.
This field only appears if you added a member interface to cluster control link Port-Channel 48.
- b) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8.
- c) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.
The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.
- d) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.
The name must be an ASCII string from 1 to 38 characters.
- e) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

- f) (Optional) Set the **CCL Subnet IP** as *a.b.0.0*.

By default, the cluster control link uses the 127.2.0.0/16 network. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. In this case, specify any /16 network address on a unique network for the cluster, except for loopback (127.0.0.0/8), multicast (224.0.0.0/4), and internal (169.254.0.0/16) addresses. If you set the value to 0.0.0.0, then the default network is used.

The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: *a.b.chassis_id.slot_id*.

- g) Choose the **Address Type** for the management interface.

This information is used to configure a management interface in the ASA configuration. Set the following information:

- **Management IP Pool**—Configure a pool of Local IP addresses, one of which will be assigned to each cluster unit for the interface, by entering the starting and ending addresses separated by a hyphen.

Include at least as many addresses as there are units in the cluster. Note that for the Firepower 9300, you must include 3 addresses per chassis, even if you do not have all module slots filled. If you plan to expand the cluster, include additional addresses. The Virtual IP address (known as the Main cluster IP address) that belongs to the current control unit is *not* a part of this pool; be sure to reserve an IP address on the same network for the Main cluster IP address. You can use IPv4 and/or IPv6 addresses.

- **Network Mask or Prefix Length**

- **Network Gateway**

- **Virtual IP address**—Set the management IP address of the current control unit. This IP address must be on the same network as the cluster pool addresses, but not be part of the pool.

Step 7

On the **Settings** page, complete the following.

Cisco: Adaptive Security Appliance - Bootstrap Configuration

General Information **Settings**

Firewall Mode:

Password:

Confirm Password:

- a) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the FTD is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices.

The firewall mode is only set at initial deployment. If you re-apply the bootstrap settings, this setting is not used.


- b) Enter and confirm a **Password** for the admin user and for the enable password.

The pre-configured ASA admin user is useful for password recovery; if you have FXOS access, you can reset the admin user password if you forget it.

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for intra-chassis clustering start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.




Step 10 For inter-chassis clustering, add the next chassis to the cluster:

- On the first chassis of the Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.
- Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.
- Choose **I want to: > Join an Existing Cluster**.
- Click **OK**.
- In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.
- Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:
 - Chassis ID**—Enter a unique chassis ID.
 - Site ID**—Enter the correct site ID.
 - Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

- Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Step 11 Connect to the control unit ASA to customize your clustering configuration.

Add More Cluster Members

Add or replace the ASA cluster member.




Note This procedure only applies to adding or replacing a *chassis*; if you are adding or replacing a module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- Make sure your existing cluster has enough IP addresses in the management IP address pool for this new member. If not, you need to edit the existing cluster bootstrap configuration on each chassis before you add this new member. This change causes a restart of the logical device.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.
- For multiple context mode, enable multiple context mode in the ASA application on the first cluster member; additional cluster members will inherit the multiple context mode configuration automatically.

Procedure

- Step 1** On an existing cluster the Firepower Chassis Manager, choose **Logical Devices** to open the **Logical Devices** page.
- Step 2** Click the Show Configuration icon () at the top right; copy the displayed cluster configuration.
- Step 3** Connect to the Firepower Chassis Manager on the new chassis, and click **Add > Cluster**.

Add Cluster

I want to: Join Existing Cluster

Device Name: cluster1

OK Cancel

Step 4 Choose **I want to:** > **Join Existing Cluster**

Step 5 For the **Device Name**, provide a name for the logical device.

Step 6 Click **OK**.

Step 7 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.


Step 8 Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—Enter the correct site ID.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Add a FTD Cluster

In native mode: You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering.

In multi-instance mode: You can add one or more clusters on a single Firepower 9300 chassis as intra-chassis clusters (you must include an instance on each module), or add one or more clusters on multiple chassis for inter-chassis clustering.

For inter-chassis clustering, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

Create a FTD Cluster

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For inter-chassis clustering, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, or for container instances, a container instance in each slot, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- For container instances, if you do not want to use the default profile, add a resource profile according to [Add a Resource Profile for Container Instances](#).
- For container instances, before you can install a container instance for the first time, you must reinitialize the security module/engine so that the disk has the correct formatting. Choose **Security Modules** or **Security Engine**, and click the Reinitialize icon (🔄). An existing logical device will be deleted and then reinstalled as a new device, losing any local application configuration. If you are replacing a native instance with container instances, you will need to delete the native instance in any case. You cannot automatically migrate a native instance to a container instance. See [Reinitializing a Security Module/Engine](#) for more information.
- Gather the following information:
 - Management interface ID, IP addresses, and network mask
 - Gateway IP address
 - FMC IP address and/or NAT ID of your choosing
 - DNS server IP address
 - FTD hostname and domain name

Procedure

- Step 1** Configure interfaces.
- Step 2** Choose **Logical Devices**.
- Step 3** Click **Add > Cluster**, and set the following parameters:

Figure 6: Native Cluster

Add Cluster [?] [X]

I want to:

Device Name:

Template:

Image Version:

Instance Type:

OK Cancel

Figure 7: Multi-Instance Cluster

Add Cluster [?] [X]

I want to:

Device Name:

Template:

Image Version:

Instance Type:

Resource Profile:

SM 1 - 46 Cores Available
SM 2 - 46 Cores Available
SM 3 - Module offline. No information available

i Before you add the first container instance, you must reinitialize the security module/engine so that the disk has the correct formatting. You only need to perform this action once.

OK Cancel

Add Device [?] [X]

Device Name:

Template:

Image Version:

Instance Type:

Usage: Standalone Cluster

Do you want to: Create New Cluster Join Existing Cluster

OK Cancel

- Choose **I want to:** > **Create New Cluster**
- Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

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

A native instance uses all of the resources (CPU, RAM, and disk space) of the security module/engine, so you can only install one native instance. A container instance uses a subset of resources of the security module/engine, so you can install multiple container instances.


- f) (Container Instance only) For the **Resource Type**, choose one of the resource profiles from the drop-down list.

For the Firepower 9300, this profile will be applied to each instance on each security module. You can set different profiles per security module later in this procedure; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model. We recommend choosing the correct profile before you create the cluster. If you need to create a new profile, cancel out of the cluster creation, and add one using [Add a Resource Profile for Container Instances](#).

- g) Click **OK**.

You see the Provisioning - *device name* window.

Step 4 Choose the interfaces you want to assign to this cluster.



For native mode clustering: All valid interfaces are assigned by default. If you defined multiple Cluster type interfaces, deselect all but one.

For multi-instance clustering: Choose each data interface you want to assign to the cluster, and also choose the Cluster type port-channel or port-channel subinterface.

Step 5 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Figure 8: Native Cluster

The screenshot shows the 'Cisco Firepower Threat Defense - Bootstrap Configuration' dialog box with the 'Cluster Information' tab selected. The 'Security Module' field contains 'Security Module-1, Security Module-2, Security Module-3'. The 'Interface Information' section includes the following fields:

Chassis ID:	1
Site ID:	1
Cluster Key:	••••
Confirm Cluster Key:	••••
Cluster Group Name:	cluster1
Management Interface:	Ethernet1/4
CCL Subnet IP:	Eg:x.x.0.0

At the bottom of the dialog are 'OK' and 'Cancel' buttons.

Figure 9: Multi-Instance Cluster

The screenshot shows the 'Cisco Firepower Threat Defense - Bootstrap Configuration' dialog box. It has a title bar with a question mark and a close button. Below the title bar are tabs for 'Cluster Information', 'Settings', 'Interface Information', and 'Agreement'. The 'Cluster Information' tab is active. The dialog is divided into two main sections: 'Security Module(SM) and Resource Profile Selection' and 'Interface Information'. In the first section, there are three rows for Security Module 1, 2, and 3. Each row has a dropdown menu for 'Resource Profile Selection' set to 'Default-Small' and a note indicating 'SM 1 - 46 Cores Available' for each. In the second section, there are input fields for 'Chassis ID' (1), 'Site ID' (1), 'Cluster Key' (masked with dots), 'Confirm Cluster Key' (masked with dots), 'Cluster Group Name' (mi-cluster-1), 'Management Interface' (Ethernet1/4), and 'CCL Subnet IP' (Eg:x.x.0.0). At the bottom are 'OK' and 'Cancel' buttons.

- a) (Container Instance for the Firepower 9300 only) In the **Security Module (SM) and Resource Profile Selection** area, you can set a different resource profile per module; for example, if you are using different security module types, and you want to use more CPUs on a lower-end model.
- b) For inter-chassis clustering, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

This field only appears if you added a member interface to cluster control link Port-Channel 48.

- c) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8. FlexConfig feature. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.
- d) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.
The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.
- e) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.
The name must be an ASCII string from 1 to 38 characters.
- f) Choose the **Management Interface**.

This interface is used to manage the logical device. This interface is separate from the chassis management port.

If you assign a Hardware Bypass-capable interface as the Management interface, you see a warning message to make sure your assignment is intentional.

- g) (Optional) Set the **CCL Subnet IP** as *a.b.0.0*.

By default, the cluster control link uses the 127.2.0.0/16 network. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. In this case, specify any /16 network address on a unique network for the cluster, except for loopback (127.0.0.0/8), multicast (224.0.0.0/4), and internal (169.254.0.0/16) addresses. If you set the value to 0.0.0.0, then the default network is used.

The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: *a.b.chassis_id.slot_id*.

Step 7

On the **Settings** page, complete the following.

The image shows two screenshots of the Cisco Firepower Threat Defense - Bootstrap Configuration dialog box, specifically the Settings tab. The left screenshot shows the following fields filled out: Registration Key (****), Confirm Registration Key (****), Password (*****), Confirm Password (*****), Firepower Management Center IP (10.89.5.35), Permit Expert mode for FTD SSH sessions (yes), Search domains (cisco.com), Firewall Mode (Routed), DNS Servers (10.89.4.5), Firepower Management Center NAT ID (test), Fully Qualified Hostname (ftd1.cisco.com), and Eventing Interface (blank). The right screenshot shows the following fields filled out: Registration Key (****), Confirm Registration Key (****), Password (*****), Confirm Password (*****), Firepower Management Center IP (10.89.5.35), Search domains (cisco.com), Firewall Mode (Routed), DNS Servers (72.163.47.11,173.37.137.8), Firepower Management Center NAT ID (blank), Fully Qualified Hostname (cluster1.cisco.com), and Eventing Interface (blank).

- a) In the **Registration Key** field, enter the key to be shared between the FMC and the cluster members during registration.

You can choose any text string for this key between 1 and 37 characters; you will enter the same key on the FMC when you add the Firepower Threat Defense.

- b) Enter a **Password** for the Firepower Threat Defense admin user for CLI access.
- c) In the **Firepower Management Center IP** field, enter the IP address of the managing FMC. If you do not know the FMC IP address, leave this field blank and enter a passphrase in the **Firepower Management Center NAT ID** field.
- d) (Optional) For a container instance, **Permit Expert mode from FTD SSH sessions**: **Yes** or **No**. Expert Mode provides Firepower 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 session 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 Firepower Threat Defense CLI.

- e) (Optional) In the **Search Domains** field, enter a comma-separated list of search domains for the management network.
- f) (Optional) From the **Firewall Mode** drop-down list, choose **Transparent** or **Routed**.

In routed mode, the Firepower 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.

- g) (Optional) In the **DNS Servers** field, enter a comma-separated list of DNS servers.

The Firepower Threat Defense uses DNS if you specify a hostname for the FMC, for example.

- h) (Optional) In the **Firepower Management Center NAT ID** field, enter a passphrase that you will also enter on the FMC when you add the cluster as a new device.

Normally, you need both IP addresses (along with a registration key) for both routing purposes and for authentication: the FMC specifies the device IP address, and the device specifies the FMC IP address. However, if you only know one of the IP addresses, which is the minimum requirement for routing purposes, then you must also specify a unique NAT ID on both sides of the connection to establish trust for the initial communication and to look up the correct registration key. You can specify any text string as the NAT ID, from 1 to 37 characters. The FMC and device use the registration key and NAT ID (instead of IP addresses) to authenticate and authorize for initial registration.

- i) (Optional) In the **Fully Qualified Hostname** field, enter a fully qualified name for the Firepower Threat Defense device.

Valid characters are the letters from a to z, the digits from 0 to 9, the dot (.), and the hyphen (-); maximum number of characters is 253.

- j) (Optional) From the **Eventing Interface** drop-down list, choose the interface on which events should be sent. If not specified, the management interface will be used.

To specify a separate interface to use for events, you must configure an interface as a *firepower-eventing* interface. If you assign a Hardware Bypass-capable interface as the Eventing interface, you see a warning message to make sure your assignment is intentional.

Step 8

On the **Interface Information** page, configure a management IP address for each security module in the cluster. Select the type of address from the **Address Type** drop-down list and then complete the following for each security module.

Note You must set the IP address for all 3 module slots in a chassis, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

Cisco Firepower Threat Defense - Bootstrap Configuration

Cluster Information Settings **Interface Information** Agreement

Address Type: IPv4 only

Security Module 1
IPv4

Management IP: 10.89.5.20

Network Mask: 255.255.255.192

Gateway: 10.89.5.1

Security Module 2
IPv4

Management IP: 10.89.5.21

Network Mask: 255.255.255.192

Gateway: 10.89.5.1

Security Module 3
IPv4

Management IP: 10.89.5.22

Network Mask: 255.255.255.192

Gateway: 10.89.5.1

OK Cancel


- a) In the **Management IP** field, configure an IP address.
Specify a unique IP address on the same network for each module.
- b) Enter a **Network Mask** or **Prefix Length**.
- c) Enter a **Network Gateway** address.

Step 9 On the **Agreement** tab, read and accept the end user license agreement (EULA).

Step 10 Click **OK** to close the configuration dialog box.

Step 11 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for intra-chassis clustering start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.

**Step 12**


For inter-chassis clustering, add the next chassis to the cluster:

- a) On the first chassis of the Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.
- b) Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.
- c) Choose **I want to: > Join an Existing Cluster**.
- d) Click **OK**.
- e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.
- f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:
 - **Chassis ID**—Enter a unique chassis ID.
 - **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FMC FlexConfig feature.
 - **Cluster Key**—(Not pre-filled) Enter the same cluster key.
 - **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

Click **OK**.

- g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status as online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



- Step 13** Add the control unit to the FMC using the management IP address.
All cluster units must be in a successfully-formed cluster on FXOS prior to adding them to FMC.
The FMC then automatically detects the data units.

Add More Cluster Nodes

Add or replace the Firepower Threat Defense cluster node in an existing cluster. When you add a new cluster node in FXOS, the FMC adds the node automatically.



Note The FXOS steps in this procedure only apply to adding a new *chassis*; if you are adding a new module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- In the case of a replacement, you must delete the old cluster node from the FMC. When you replace it with a new node, it is considered to be a new device on the FMC.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.

Procedure

- Step 1** If you previously upgraded the Firepower Threat Defense image using the FMC, perform the following steps *on each chassis in the cluster*.
- When you upgraded from the FMC, the startup version in the FXOS configuration was not updated, and the standalone package was not installed on the chassis. Both of these items need to be set manually so the new node can join the cluster using the correct image version.

Note If you only applied a patch release, you can skip this step. Cisco does not provide standalone packages for patches.

- a) Install the running Firepower Threat Defense image on the chassis using the **System > Updates** page.
- b) Click **Logical Devices** and click the Set Version icon (🔧). For a Firepower 9300 with multiple modules, set the version for each module.

The **Startup Version** shows the original package you deployed with. The **Current Version** shows the version you upgraded to.

- c) In the **New Version** drop-down menu, choose the version that you uploaded. This version should match the **Current Version** displayed, and will set the startup version to match the new version.
- d) On the new chassis, make sure the new image package is installed.

Step 2 On an existing cluster chassis Firepower Chassis Manager, click **Logical Devices**.

Step 3 Click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.

Step 4 Connect to the Firepower Chassis Manager on the new chassis, and click **Add > Cluster**.

Step 5 For the **Device Name**, provide a name for the logical device.

Step 6 Click **OK**.

Step 7 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

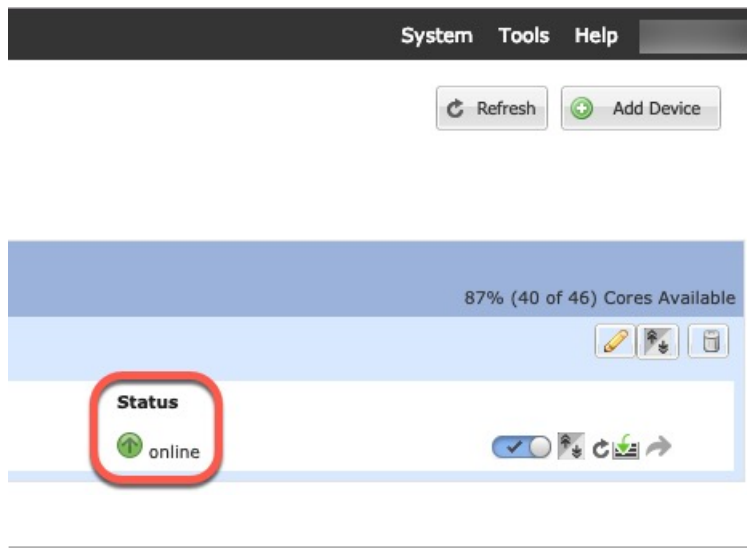
Step 8 Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:

- **Chassis ID**—Enter a unique chassis ID.
- **Site ID**—For inter-site clustering, enter the site ID for this chassis between 1 and 8. This feature is only configurable using the FMC FlexConfig feature.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.
- **Management IP**—Change the management address for each module to be a unique IP address on the same network as the other cluster members.

Click **OK**.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Configure Radware DefensePro

The Cisco Firepower 4100/9300 chassis can support multiple services (for example, a firewall and a third-party DDoS application) on a single blade. These applications and services can be linked together to form a Service Chain.

About Radware DefensePro

In the current supported Service Chaining configuration, the third-party Radware DefensePro virtual platform can be installed to run in front of the ASA firewall, or in front of Firepower Threat Defense. Radware DefensePro is a KVM-based virtual platform that provides distributed denial-of-service (DDoS) detection and mitigation capabilities on the Firepower 4100/9300 chassis. When Service Chaining is enabled on your Firepower 4100/9300 chassis, traffic from the network must first pass through the DefensePro virtual platform before reaching the main ASA or Firepower Threat Defense firewall.



Note

- The Radware DefensePro virtual platform may be referred to as *Radware vDP* (virtual DefensePro), or simply *vDP*.
- The Radware DefensePro virtual platform may occasionally be referred to as a Link Decorator.

Prerequisites for Radware DefensePro

Prior to deploying Radware DefensePro on your Firepower 4100/9300 chassis, you must configure the Firepower 4100/9300 chassis to use an NTP Server with the **etc/UTC** Time Zone. For more information about setting the date and time in your Firepower 4100/9300 chassis, see [Setting the Date and Time](#).

Guidelines for Service Chaining

Models

- ASA—The Radware DefensePro (vDP) platform is supported with ASA on the following models:
 - Firepower 9300
 - Firepower 4110
 - Firepower 4115
 - Firepower 4120
 - Firepower 4125
 - Firepower 4140
 - Firepower 4145
 - Firepower 4150
- FTD—The Radware DefensePro platform is supported with Firepower Threat Defense on the following models:
 - Firepower 9300
 - Firepower 4110—Note you must deploy the decorator at the same time as the logical device. You cannot install the decorator after the logical device is already configured on the device.
 - Firepower 4112
 - Firepower 4115
 - Firepower 4120—Note you must deploy the decorator at the same time as the logical device. You cannot install the decorator after the logical device is already configured on the device.
 - Firepower 4125
 - Firepower 4140
 - Firepower 4145
 - Firepower 4150

Additional Guidelines

- Service Chaining is not supported in an inter-chassis cluster configuration. However, the Radware DefensePro (vDP) application can be deployed in a standalone configuration in an inter-chassis cluster scenario.

Configure Radware DefensePro on a Standalone Logical Device

The following procedure shows how to install Radware DefensePro in a single Service Chain in front of a standalone ASA or Firepower Threat Defense logical device.



Note Once you set the vDP application and commit the change at the end of this procedure, the logical device (ASA or Firepower Threat Defense) will reboot.

If you are installing Radware vDP in front of ASA on a Firepower 4120 or 4140 security appliance, you must use the FXOS CLI to deploy the decorator. For full CLI instructions on how to install and configure Radware DefensePro in a service chain in front of ASA on Firepower 4100 devices, refer to the FXOS CLI configuration guide.

Before you begin

- Download the vDP image from Cisco.com (see [Downloading Images from Cisco.com](#)) and then upload that image to the Firepower 4100/9300 chassis (see [Uploading an Image to the Security Appliance](#)).
- You can deploy the Radware DefensePro application in a standalone configuration on an intra-chassis cluster; for intra-chassis clustering, see [Configure Radware DefensePro on an Intra-Chassis Cluster, on page 61](#).

Procedure

-
- Step 1** If you want to use a separate management interface for vDP, enable the interface and set it to be the mgmt type according to [Configure a Physical Interface](#). Otherwise, you can share the application management interface.
- Step 2** Choose **Logical Devices** to open the Logical Devices page.
- The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown.
- Step 3** Create a standalone ASA or Firepower Threat Defense Logical Device (see [Add a Standalone ASA, on page 22](#) or [Add a Standalone FTD for the FMC, on page 25](#)).
- Step 4** In the **Decorators** area, select vDP. The Radware: Virtual DefensePro - Configuration window appears. Configure the following fields under the **General Information** tab.
- Step 5** If you have more than one vDP version uploaded to the Firepower 4100/9300 chassis, select the version you want to use in the **Version** drop-down.
- Step 6** If you have a resource configurable Radware DefensePro application, a list of supported resource profiles appears under the **Resource Profile** drop-down. Select the resource profile you want to assign to the device. If you do not select a resource profile, the default setting is used.
- Step 7** Under the **Management Interface** drop-down, choose the management interface you created in step 1 of this procedure.
- Step 8** Select the default **Address Type**, IPv4 only, IPv6 only, or IPv4 and IPv6.
- Step 9** Configure the following fields, based on your **Address Type** selection from the previous step.
- In the **Management IP** field, configure a local IP address.
 - IPv4 only: Enter a **Network Mask**.
 - IPv6 only: Enter a **Prefix Length**.
 - Enter a **Network Gateway** address.
- Step 10** Click the checkbox next to each data port that you want to assign to the device.

Step 11 Click **OK**.

Step 12 Click **Save**.

The FXOS deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the specified security module.

What to do next

Set a password for the DefensePro application. Note that the application does not come online until you set a password. For more information, see the Radware DefensePro DDoS Mitigation User Guide on cisco.com.

Configure Radware DefensePro on an Intra-Chassis Cluster

The following procedure shows how to install the Radware DefensePro image, and configure it in a Service Chain in front of an ASA or Firepower Threat Defense intra-chassis cluster.



Note Service Chaining is not supported in an inter-chassis cluster configuration. However, the Radware DefensePro application can be deployed in a standalone configuration in an inter-chassis cluster scenario.

Before you begin

- Download the vDP image from Cisco.com (see [Downloading Images from Cisco.com](#)) and then upload that image to the Firepower 4100/9300 chassis (see [Uploading an Image to the Security Appliance](#)).

Procedure

- Step 1** If you want to use a separate management interface for vDP, enable the interface and set it to be the mgmt type according to [Configure a Physical Interface](#). Otherwise, you can share the application management interface.
- Step 2** Configure an ASA or Firepower Threat Defense intra-chassis cluster (see [Create an ASA Cluster, on page 40](#) or [Create a FTD Cluster, on page 47](#)).
- Note that before you click **Save** at the end of the procedure to configure the intra-chassis cluster, you must first follow the following steps to add a vDP decorator to the cluster.
- Step 3** In the **Decorators** area, select vDP. The **Radware: Virtual DefensePro - Configuration** dialog box appears. Configure the following fields under the **General Information** tab.
- Step 4** If you have more than one vDP version uploaded to the Firepower 4100/9300 chassis, select the vDP version you want to use in the **Version** drop-down.
- Step 5** If you have a resource configurable Radware DefensePro application, a list of supported resource profiles appears under the Resource Profile drop-down. Select the resource profile you want to assign to the device. If you do not select a resource profile, the default setting is used.
- Step 6** Under the **Management Interface** drop-down, choose a management interface.
- Step 7** Click the checkbox next to each data port that you want to assign to the vDP decorator.

- Step 8** Click the **Interface Information** tab.
- Step 9** Select the **Address Type** to be used, IPv4 only, IPv6 only, or IPv4 and IPv6.
- Step 10** Configure the following fields for each Security Module. Note that the fields that display depend on your **Address Type** selection from the previous step.
- In the **Management IP** field, configure a local IP address.
 - IPv4 only: Enter a **Network Mask**.
IPv6 only: Enter a **Prefix Length**.
 - Enter a **Network Gateway** address.

Step 11 Click **OK**.

Step 12 Click **Save**.

The FXOS deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the specified security module.

Step 13 Choose **Logical Devices** to open the Logical Devices page.

Step 14 Scroll through the list of configured logical devices to the entries for vDP. Verify their Attributes listed in the **Management IP** column.

- If the **CLUSTER-ROLE** element displays as *unknown* for the DefensePro instances, you must enter the DefensePro application and configure the Control unit IP address to complete the creation of the vDP cluster.
- If the **CLUSTER-ROLE** element displays as *primary* or *secondary* for the DefensePro instances, the applications are online and formed in a cluster.

What to do next

Set a password for the DefensePro application. Note that the application does not come online until you set a password. For more information, see the Radware DefensePro DDoS Mitigation User Guide on cisco.com.

Open UDP/TCP Ports and Enable vDP Web Services

The Radware APSolute Vision Manager interfaces communicate with the Radware vDP application using various UDP/TCP ports. In order for the vDP application to communicate with the APSolute Vision Manager, you must ensure that these ports are accessible and not blocked by your firewall. For more information on which specific ports to open, see the following tables in the APSolute Vision User Guide:

- **Ports for APSolute Vision Server-WBM Communication and Operating System**
- **Communication Ports for APSolute Vision Server with Radware Devices**

In order for Radware APSolute Vision to manage the Virtual DefensePro application deployed on the FXOS chassis, you must enable the vDP web service using the FXOS CLI.

Procedure

- Step 1** From the FXOS CLI, connect to the vDP application instance.

```
connect module slot console
```

```
connect vdp
```

Step 2 Enable vDP web services.

```
manage secure-web status set enable
```

Step 3 Exit the vDP application console and return to the FXOS module CLI.

```
Ctrl ]
```

Configure TLS Crypto Acceleration

The following topics discuss TLS crypto acceleration, how to enable it, and how to view its status using the FMCr.

The following table maps the Firepower Threat Defense and the FXOS version with the required TSL Crypto:



Note When FXOS 2.6.1 is upgraded to FXOS 2.7.x and above, FTD 6.4 does not automatically enable crypto as 6.4 is not compatible with TLS crypto.

FTD	FXOS	Crypto
6.4	2.6	Support for only one container instance (Phase 1)
6.4	2.7 and above	NA
6.5 and above	2.7 and above	Support for upto 16 container instances (Phase 2)

About TLS Crypto Acceleration

The Firepower 4100/9300 support Transport Layer Security cryptographic acceleration, which performs Transport Layer Security/Secure Sockets Layer (TLS/SSL) encryption and decryption in hardware, which greatly accelerates the following:

- TLS/SSL encryption and decryption
- VPN, including TLS/SSL and IPsec

TLS cryptographic acceleration is automatically enabled on native instances and cannot be disabled. You can enable TLS crypto acceleration on up to 16 FTD container instances per security engine/module as well.

Guidelines and Limitations for TLS Crypto Acceleration

Keep the following in mind if your Firepower Threat Defense has TLS crypto acceleration enabled.

Inspection engine failure

If the inspection engine is configured to preserve connections and the inspection engine fails unexpectedly, TLS/SSL traffic is dropped until the engine restarts.

This behavior is controlled by the Firepower Threat Defense command **configure snort preserve-connection {enable | disable}** command.

HTTP-only performance

Using TLS crypto acceleration on an FTD container instance that is not decrypting traffic can affect performance. We recommend you enable TLS crypto acceleration *only* on FTD container instances that decrypt TLS/SSL traffic.

Federal Information Processing Standards (FIPS)

If TLS crypto acceleration and Federal Information Processing Standards (FIPS) are both enabled, connections with the following options fail:

- RSA keys less than 2048 bytes in size
- Rivest cipher 4 (RC4)
- Single Data Encryption Standard (single DES)
- Merkle–Damgard 5 (MD5)
- SSL v3

FIPS is enabled when you configure the FMC and Firepower Threat Defenses to operate in a security certifications compliance mode. To allow connections when operating in those modes, you can either disable TLS crypto acceleration on the FTD container instance or you can configure web browsers to accept more secure options.

For more information:

- [Common Criteria](#).

High Availability (HA) and clustering

If you have high availability (HA) or clustered Firepower Threat Defenses, you must enable TLS crypto acceleration on each Firepower Threat Defense individually. One device's TLS crypto acceleration configuration is not shared with the other devices in the HA pair or cluster.

TLS heartbeat

Some applications use the *TLS heartbeat* extension to the Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) protocols defined by [RFC6520](#). TLS heartbeat provides a way to confirm the connection is still alive—either the client or server sends a specified number of bytes of data and requests the other party echo the response. If this is successful, encrypted data is sent.

When an Firepower Threat Defense managed by FMC with TLS crypto acceleration enabled encounters a packet that uses the TLS heartbeat extension, the Firepower Threat Defense takes the action specified by the FMC setting for **Decryption Errors** in the SSL policy's **Undecryptable Actions**:

- Block

- Block with reset

To determine whether applications are using TLS heartbeat, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

If TLS crypto acceleration is disabled on an FTD container instance, you can configure a **Max Heartbeat Length** in a Network Analysis Policy (NAP) in the FMC to determine how to handle TLS heartbeats.

For more information about TLS heartbeat, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

TLS/SSL oversubscription

TLS/SSL oversubscription is a state where an Firepower Threat Defense is overloaded with TLS/SSL traffic. Any Firepower Threat Defense can experience TLS/SSL oversubscription but only the Firepower Threat Defenses that support TLS crypto acceleration provide a configurable way to handle it.

When an Firepower Threat Defense managed by an FMC with TLS crypto acceleration enabled is oversubscribed, any packet received by the Firepower Threat Defense is acted on according to the setting for **Handshake Errors** in the SSL policy's **Undecryptable Actions**:

- Inherit default action
- Do not decrypt
- Block
- Block with reset

If the setting for **Handshake Errors** in the SSL policy's **Undecryptable Actions** is **Do Not decrypt** and the associated access control policy is configured to inspect the traffic, inspection occurs; decryption does *not* occur.

If a significant amount of oversubscription is occurring, you have the following options:

- Upgrade to an Firepower Threat Defense with more TLS/SSL processing capacity.
- Change your SSL policies to add **Do Not Decrypt** rules for traffic that is not a high priority to decrypt.

For more information about TLS oversubscription, see the chapter on troubleshooting TLS/SSL rules in the *Firepower Management Center Configuration Guide*.

Passive and inline tap sets not supported

TLS/SSL traffic cannot be decrypted on passive or inline tap set interfaces when TLS crypto acceleration is enabled.

Enable TLS Crypto Acceleration for Container Instances

TLS crypto acceleration is automatically enabled when you deploy a logical instance as discussed in [Add a Standalone FTD for the FMC, on page 25](#).


TLS crypto acceleration is enabled on all native instances and cannot be disabled.

View the Status of TLS Crypto Acceleration

This topic discusses how you can determine if TLS crypto acceleration is enabled.

Perform the following task in the FMC.

Procedure

- Step 1** Log in to the FMC.
 - Step 2** Click **Devices > Device Management**.
 - Step 3** Click **Edit** () to edit a managed device.
 - Step 4** Click **Device** page. TLS crypto acceleration status is displayed in the General section.
-

Enable FTD Link State Synchronization

The chassis can now synchronize the Firepower Threat Defense operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The Firepower Threat Defense application interface admin state is not considered. Without synchronization from Firepower Threat Defense, data interfaces can be in an Up state physically before the Firepower Threat Defense application has completely come online, for example, or can stay Up for a period of time after you initiate an Firepower Threat Defense shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the Firepower Threat Defense before the Firepower Threat Defense can handle it.

This feature is disabled by default, and can be enabled per logical device in FXOS. This feature does not affect non-data interfaces such as Management or Cluster.

When you enable Firepower Threat Defense link state synchronization, the **Service State** of an interface in FXOS will be synced with the administrative state of this interface in Firepower Threat Defense. For example, if you shut down an interface in Firepower Threat Defense, the Service State will show as Disabled. If you shut down the Firepower Threat Defense application, all interfaces will show as Disabled. For Hardware Bypass interfaces, administratively shutting down the interface in Firepower Threat Defense will set the Service State to Disabled; but shutting down the Firepower Threat Defense application or other chassis-level shutdowns, including powering off, keeps the interface pair Enabled.


If you disable Firepower Threat Defense link state synchronization, the Service State will always show as Enabled.



Note This feature is not supported for clustering, container instances, or an Firepower Threat Defense with a Radware vDP decorator. It is also not supported for the ASA.

Procedure

- Step 1** Choose **Logical Devices**, and then for the Firepower Threat Defense logical device, choose **Enable Link State** from the drop-down list.



To disable this feature, choose **Disable Link State**.

- Step 2** View the current interface state, as well as the last down reason.

show interface expand detail

Example:

```
Firepower # scope eth-uplink
Firepower /eth-uplink # scope fabric a
Firepower /eth-uplink/fabric # show interface expand detail
Interface:
  Port Name: Ethernet1/2
  User Label:
  Port Type: Data
  Admin State: Enabled
  Oper State: Up
  State Reason:
  flow control policy: default
  Auto negotiation: Yes
  Admin Speed: 1 Gbps
  Oper Speed: 1 Gbps
  Admin Duplex: Full Duplex
  Oper Duplex: Full Duplex
  Ethernet Link Profile name: default
  Oper Ethernet Link Profile name: fabric/lan/eth-link-prof-default
  Uddl Oper State: Admin Disabled
  Inline Pair Admin State: Enabled
  Inline Pair Peer Port Name:
  Service State: Enabled
  Last Service State Down Reason: None
  Allowed Vlan: All
  Network Control Policy: default
  Current Task:
  <...>
```

Manage Logical Devices

You can delete a logical device, convert an ASA to transparent mode, change the interface configuration, and perform other tasks on existing logical devices.

Connect to the Console of the Application

Use the following procedure to connect to the console of the application.

Procedure

Step 1 Connect to the module CLI using a console connection or a Telnet connection.

connect module *slot_number* { **console** | **telnet** }

To connect to the security engine of a device that does not support multiple security modules, always use **1** as the *slot_number*.

The benefits of using a Telnet connection is that you can have multiple sessions to the module at the same time, and the connection speed is faster.

Example:

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.
```

```
CISCO Serial Over LAN:
Close Network Connection to Exit
```

```
Firepower-module1>
```

Step 2 Connect to the application console. Enter the appropriate command for your device.

connect asa *name*

connect ftd *name*

connect vdp *name*

To view the instance names, enter the command without a name.

Example:

```
Firepower-module1> connect asa asal
Connecting to asa(asal) console... hit Ctrl + A + D to return to bootCLI
[...]
asa>
```

Example:

```
Firepower-module1> connect ftd ftd1
Connecting to ftd(ftd-native) console... enter exit to return to bootCLI
[...]
>
```

Step 3 Exit the application console to the FXOS module CLI.

- ASA—Enter **Ctrl-a, d**

- FTD—Enter **exit**
- vDP—Enter **Ctrl-], .**

Step 4 Return to the supervisor level of the FXOS CLI.

Exit the console:

- Enter **~**
You exit to the Telnet application.
- To exit the Telnet application, enter:
telnet>**quit**

Exit the Telnet session:

- Enter **Ctrl-], .**

Example

The following example connects to an ASA on security module 1 and then exits back to the supervisor level of the FXOS CLI.

```
Firepower# connect module 1 console
Telnet escape character is '~'.
Trying 127.5.1.1...
Connected to 127.5.1.1.
Escape character is '~'.

CISCO Serial Over LAN:
Close Network Connection to Exit

Firepower-module1>connect asa asa1
asa> ~
telnet> quit
Connection closed.
Firepower#
```

Delete a Logical Device

Procedure

- Step 1** Choose **Logical Devices** to open the Logical Devices page.
The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown instead.
- Step 2** Click **Delete** for the logical device that you want to delete.
- Step 3** Click **Yes** to confirm that you want to delete the logical device.

Step 4 Click **Yes** to confirm that you want to delete the application configuration.

Remove a Cluster Unit

The following sections describe how to remove units temporarily or permanently from the cluster.

Temporary Removal

A cluster unit will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status on the Firepower Chassis Manager **Logical Devices** page:

The screenshot shows the Firepower Chassis Manager interface. At the top right, there are icons for edit, refresh, print, and back. Below this is a table with two columns: 'Management Port' and 'Status'. The first row shows 'Ethernet1/4' under 'Management Port' and 'online' under 'Status'. To the right of the 'Status' cell, there is a blue slider control that is currently turned on (checked), and several other icons including a refresh, a power off, and a right arrow. Below the table, there is a section titled 'Attributes' with the following details:



Attributes	
Cluster Operational Status	: not-in-cluster
FIREPOWER-MGMT-IP	: 10.89.5.20
CLUSTER-ROLE	: none
CLUSTER-IP	: 127.2.1.1
MGMT-URL	: https://10.89.5.35/
UUID	: 8e459170-451d-11e9-8475-f22f06c32630

For FTD using the FMC, you should leave the device in the FMC device list so that it can resume full functionality after you reenables clustering.

- Disable clustering in the application—You can disable clustering using the application CLI. Enter the **cluster remove unit** *name* command to remove any unit other than the one you are logged into. The bootstrap configuration remains intact, as well as the last configuration synced from the control unit, so you can later re-add the unit without losing your configuration. If you enter this command on a data unit to remove the control unit, a new control unit is elected.

When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster, the Management interface is disabled.

To reenables clustering, on the ASA enter **cluster group** *name* and then **enable**. To reenables clustering, on the FTD enter **cluster enable**.

- Disable the application instance—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Slider enabled** (). You can later reenables it using the **Slider disabled** (.
- Shut down the security module/engine—In the Firepower Chassis Manager on the **Security Module/Engine** page, click the **Power Off icon**.
- Shut down the chassis—In the Firepower Chassis Manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

You can permanently remove a cluster member using the following methods.

For FTD using the FMC, be sure to remove the unit from the FMC device list after you disable clustering on the chassis.

- Delete the logical device—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Delete** (🗑️). You can then deploy a standalone logical device, a new cluster, or even add a new logical device to the same cluster.
- Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new member of the cluster.

Delete an Application Instance that is not Associated with a Logical Device

When you delete a logical device, you are prompted as to whether you want to also delete the application configuration for the logical device. If you do not delete the application configuration, you will not be able to create a logical device using a different application until that application instance is deleted. You can use the following procedure to delete an application instance from a security module/engine when it is no longer associated with a logical device.

Procedure

-
- Step 1** Choose **Logical Devices** to open the Logical Devices page.
- The Logical Devices page shows a list of configured logical devices on the chassis. If no logical devices have been configured, a message stating so is shown instead. Below the list of logical devices, you can see a list of application instances that are not associated with a logical device.
- Step 2** Click **Delete** for the application instance that you want to delete.
- Step 3** Click **Yes** to confirm that you want to delete the application instance.
-

Change an Interface on a FTD Logical Device

You can allocate or unallocate an interface, or replace a management interface on the Firepower Threat Defense logical device. You can then sync the interface configuration in the FMC or the FDM.

Adding a new interface, or deleting an unused interface has minimal impact on the Firepower Threat Defense configuration. However, deleting an interface that is used in your security policy will impact the configuration. Interfaces can be referenced directly in many places in the Firepower Threat Defense configuration, including access rules, NAT, SSL, identity rules, VPN, DHCP server, and so on. Policies that refer to security zones are not affected. You can also edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the FMC or the FDM.

For the FMC: Deleting an interface will delete any configuration associated with that interface.

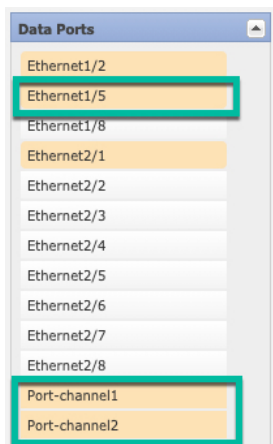
For the FDM: You can migrate the configuration from one interface to another interface before you delete the old interface.

Before you begin

- Configure your interfaces, and add any EtherChannels according to [Configure a Physical Interface](#) and [Add an EtherChannel \(Port Channel\)](#).
- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.
- If you want to replace the management or eventing interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the FTD device reboots (management interface changes cause a reboot), and you sync the configuration in the FMC or the FDM, you can add the (now unallocated) management interface to the EtherChannel as well.
- For clustering or High Availability, make sure you add or remove the interface on all units before you sync the configuration in the FMC or the FDM. We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. Note that new interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

- Step 1** In the Firepower Chassis Manager, choose **Logical Devices**.
- Step 2** Click the **Edit** icon at the top right to edit the logical device.
- Step 3** Allocate a new data interface by selecting the interface in the **Data Ports** area.
- Do not delete any interfaces yet.



- Step 4** Replace the management or eventing interface:
- For these types of interfaces, the device reboots after you save your changes.
- Click the device icon in the center of the page.
 - On the **General** or **Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
 - On the **Settings** tab, choose the new **Eventing Interface** from the drop-down list.
 - Click **OK**.

If you change the IP address of the Management interface, then you must also change the IP address for the device in the FMC: go to **Devices > Device Management > Device/Cluster**. In the **Management** area, set the IP address to match the bootstrap configuration address.

Step 5 Click **Save**.

Step 6 Sync the interfaces in the FMC.


- a) Log into the FMC.
- b) Select **Devices > Device Management** and click **Edit** (🔧) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.
- c) Click the **Sync Device** button on the top left of the **Interfaces** page.
- d) After the changes are detected, you will see a red banner on the **Interfaces** page indicating that the interface configuration has changed. Click the **Click to know more** link to view the interface changes.
- e) If you plan to delete an interface, manually transfer any interface configuration from the old interface to the new interface.

Because you have not yet deleted any interfaces, you can refer to the existing configuration. You will have additional opportunity to fix the configuration after you delete the old interface and re-run the validation. The validation will show you all locations in which the old interface is still used.

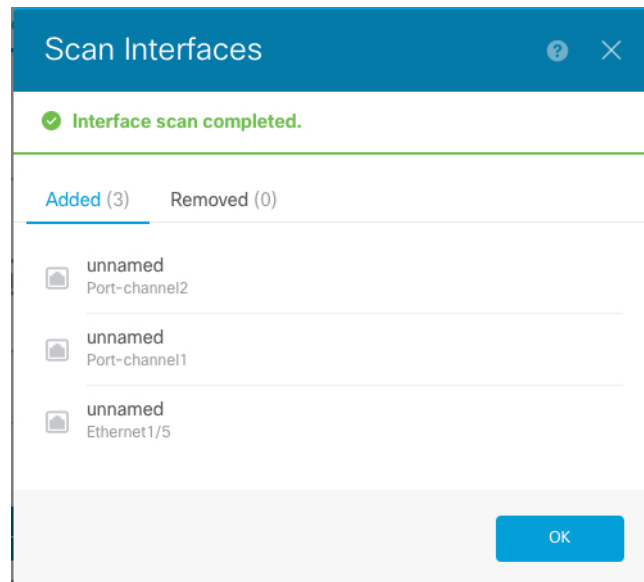
- f) Click **Validate Changes** to make sure your policy will still work with the interface changes.
If there are any errors, you need to change your policy and rerun the validation.
- g) Click **Save**.
- h) Select the devices and click **Deploy** to deploy the policy to the assigned devices. The changes are not active until you deploy them.

Step 7 Sync and migrate the interfaces in the FDM.

- a) Log into the FDM.
- b) Click **Device**, then click the **View All Interfaces** link in the **Interfaces** summary.



- c) Click the **Scan Interfaces icon**.
- d) Wait for the interfaces to scan, and then click **OK**.



- e) Configure the new interfaces with names, IP addresses, and so on.

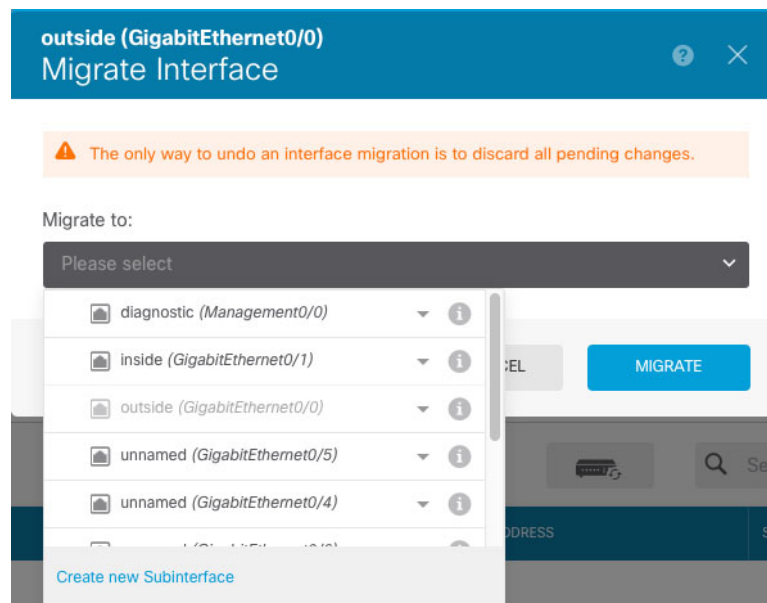
If you want to use the existing IP address and name of an interface that you want to delete, then you need to reconfigure the old interface with a dummy name and IP address so that you can use those settings on the new interface.

- f) To replace an old interface with a new interface, click the Replace icon for the old interface.

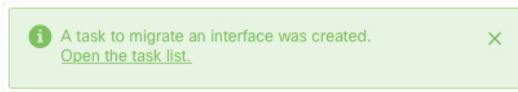
Replace icon

This process replaces the old interface with the new interface in all configuration settings that refer to the interface.

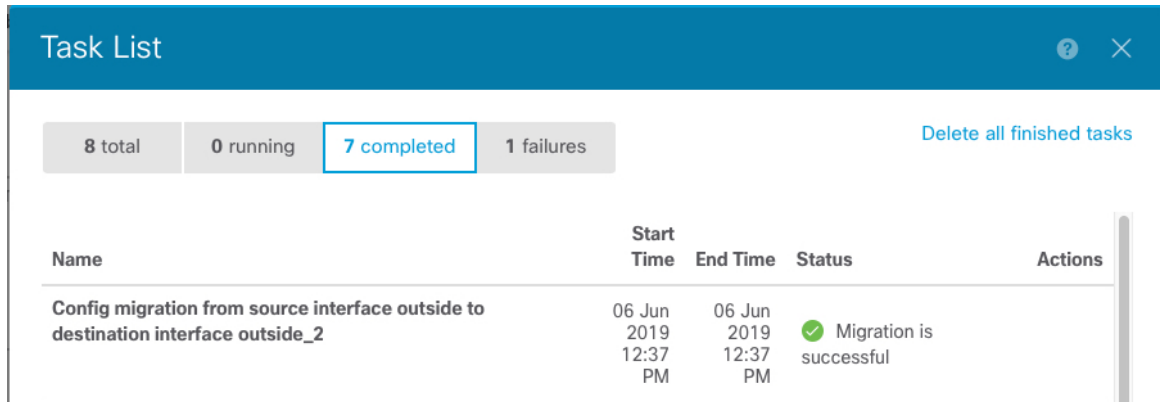
- g) Choose the new interface from the **Replacement Interface** drop-down list.



- h) A message appears on the **Interfaces** page. Click the link in the message.




- i) Check the **Task List** to ensure that the migration was successful.

A screenshot of the 'Task List' window. The title bar is blue with 'Task List' and a close button. Below the title bar, there are four buttons: '8 total', '0 running', '7 completed' (highlighted with a blue border), and '1 failures'. To the right of these buttons is a link 'Delete all finished tasks'. Below this is a table with columns: Name, Start Time, End Time, Status, and Actions. The table contains one row with the following data:


Name	Start Time	End Time	Status	Actions
Config migration from source interface outside to destination interface outside_2	06 Jun 2019 12:37 PM	06 Jun 2019 12:37 PM	✓ Migration is successful	

- Step 8** In the Firepower Chassis Manager, unallocate a data interface by de-selecting the interface in the **Data Ports** area.



- Step 9** Click **Save**.
- Step 10** Sync the interfaces again in the FMC or the FDM.

Figure 10: FDM Scan Interfaces



Change an Interface on an ASA Logical Device

You can allocate, unallocate, or replace a management interface on an ASA logical device. ASDM discovers the new interfaces automatically.

Adding a new interface, or deleting an unused interface has minimal impact on the ASA configuration. However, if you remove an allocated interface in FXOS (for example, if you remove a network module, remove an EtherChannel, or reassign an allocated interface to an EtherChannel), and the interface is used in your security policy, removal will impact the ASA configuration. In this case, the ASA configuration retains the original commands so that you can make any necessary adjustments. You can manually remove the old interface configuration in the ASA OS.



Note You can edit the membership of an allocated EtherChannel without impacting the logical device.

Before you begin

- Configure your interfaces and add any EtherChannels according to [Configure a Physical Interface](#) and [Add an EtherChannel \(Port Channel\)](#).
- If you want to add an already-allocated interface to an EtherChannel (for example, all interfaces are allocated by default to a cluster), you need to unallocate the interface from the logical device first, then add the interface to the EtherChannel. For a new EtherChannel, you can then allocate the EtherChannel to the device.
- If you want to replace the management interface with a management EtherChannel, then you need to create the EtherChannel with at least 1 unallocated data member interface, and then replace the current management interface with the EtherChannel. After the ASA reloads (management interface changes cause a reload), you can add the (now unallocated) management interface to the EtherChannel as well.
- For clustering or failover, make sure you add or remove the interface on all units. We recommend that you make the interface changes on the data/standby unit(s) first, and then on the control/active unit. New interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

- Step 1** In the Firepower Chassis Manager, choose **Logical Devices**.
- Step 2** Click the **Edit** icon at the top right to edit the logical device.
- Step 3** Unallocate a data interface by de-selecting the interface in the **Data Ports** area.
- Step 4** Allocate a new data interface by selecting the interface in the **Data Ports** area.
- Step 5** Replace the management interface:
- For this type of interface, the device reloads after you save your changes.
- Click the device icon in the center of the page.
 - On the **General/Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
 - Click **OK**.
- Step 6** Click **Save**.
-

Modify or Recover Bootstrap Settings for a Logical Device

You can modify bootstrap settings for a logical device. You can then immediately restart the application instance using those new settings or save the changes and restart the application instance using those new settings at a later time.

Procedure

- Step 1** In the Firepower Chassis Manager, choose **Logical Devices**.
- Step 2** Click the **Edit** icon at the top right to edit the logical device.
- Step 3** Click the device icon in the center of the page.
- Step 4** Modify the logical device settings as required.
- Step 5** Click **OK**.
- Step 6** Click **Restart Now** to save the changes and restart the application instance now. Click **Restart Later** to save the changes without restarting the application instance.

Note If you selected **Restart Later**, you can restart the application instance when you are ready by clicking **Restart Instance** from the Logical Devices page.

Logical Devices Page

Use the **Logical Devices** page of the Firepower Chassis Manager to create, edit, and delete logical devices. The **Logical Devices** page includes an informational area for the logical device(s) installed on each Firepower 4100/9300 chassis security module/engine.

The header for each logical device area provides the following information:

- The unique name of the logical device.
- The logical device mode, either Standalone or Clustered.
- **Status**—Shows the state of the logical device:
 - ok—The logical device configuration is complete.
 - incomplete-configuration—The logical device configuration is incomplete.

Each logical device area provides the following information:

- **Application**—Shows the application running on the security module.
- **Version**—Shows the software version number of the application running on the security module.



Note Updates to Firepower Threat Defense logical devices are done using FMC and are not reflected on the **Logical Devices > Edit** and **System > Updates** pages in Firepower Chassis Manager. On these pages, the version shown indicates the software version (CSP image) that was used to create the Firepower Threat Defense logical device.

- **Resource Profile**—Shows the resource profile assigned to the logical device/application instance.
- **Management IP**—Shows the local IP address assigned as the logical device Management IP.
- **Gateway**—Shows the network gateway address assigned to the application instance.
- **Management Port**—Shows the management port assigned to the application instance.
- **Status**—Shows the state of the application instance:
 - Online—The application is running and operating.
 - Offline—The application is stopped and inoperable.
 - Installing—The application installation is in progress.
 - Not Installed—The application is not installed.
 - Install Failed—The application installation failed.
 - Starting—The application is starting up.
 - Start Failed—The application failed to start up.
 - Started—The application started successfully, and is waiting for app agent heartbeat.
 - Stopping—The application is in the process of stopping.
 - Stop Failed—The application was unable to be brought offline.
 - Not Responding—The application is unresponsive.
 - Updating—The application software update is in progress.
 - Update Failed—The application software update failed.

- **Update Succeeded**—The application software update succeeded.
- **Unsupported**—The installed application is not supported.

If a security module is not present or is in a fault state, that information is shown in the status field. You can hover over the information icon to see additional information for a fault. For more information on security module faults, see [About FXOS Security Modules/Security Engine](#).

- **Expanded Information Area**—Shows additional attributes for the application instance that is currently running.



Note If you modify the bootstrap settings for an application without immediately restarting the application instance, the Attributes fields show information for the application that is currently running and will not reflect the changes that were made until the application is restarted.

- **Ports**—Shows the names and types of interfaces assigned to the application instance.
- **Cluster Operation Status**—Shows the management URL assigned to the application instance.
- **Management IP/Firepower Management IP**—Shows the management IP address assigned to the application instance.
- **Cluster Role**—Shows the cluster role for the application instance, control or data.
- **Cluster IP**—Shows the IP address assigned to the application instance.
- **HA Role**—Shows the high-availability role for the application instance, active or standby.
- **Management URL**—Shows the URL of the management application assigned to the application instance.
- **UUID**—Shows the universally unique identifier for the application instance.

From the **Logical Devices** page of the Firepower Chassis Manager, you can perform the following functions on a logical device:

- **Refresh**—Refreshes the information on the Logical Devices page.
- **Add Device**—Allows you to create a logical device.
- **Edit**—Allows you to edit an existing logical device.
- **Set Version**—Allows you to upgrade or downgrade the software on a logical device.
- **Delete**—Deletes a logical device.
- **Show Configuration**—Opens a dialog box showing the configuration information in JSON format for a logical device or cluster. You can copy the configuration information and use it when creating additional devices that are part of a cluster.
- **Enable/Disable**—Enables or disables an application instance.
- **Upgrade/Downgrade**—Allows you to upgrade or downgrade an application instance.

- **Restart Instance**—Allows you to restart the application instance. If you have modified the device bootstrap information but have not yet restarted the application instance, you can click Restart Instance to clear the existing management bootstrap information and restart the application instance using the new bootstrap information.
- **Reinstall Instance**—Allows you to reinstall the application instance.
- **Go To Device Manager**—Provides a link to the FMC or ASDM defined for the application instance.
- **Enable/Disable Link State**—Enable or disable Firepower Threat Defense link state synchronization. For more information, see [Enable FTD Link State Synchronization, on page 66](#).

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

Spanned EtherChannel Routed Mode Example with Site-Specific MAC Addresses


The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and an inside network at each site (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the inside and outside networks. Each EtherChannel is spanned across all chassis in the cluster.

The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters blocking the global MAC address to prevent traffic from traversing the DCI to the other site when the traffic is destined for the cluster. If the cluster nodes at one site become unreachable, you must remove the filters so traffic can be sent to the other site's cluster nodes. You should use VACLs to filter the global MAC address. Be sure to disable ARP inspection.

The cluster acts as the gateway for the inside networks. The global virtual MAC, which is shared across all cluster nodes, is used only to receive packets. Outgoing packets use a site-specific MAC address from each DC cluster. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address.

In this scenario:

- All egress packets sent from the cluster use the site MAC address and are localized at the data center.
- All ingress packets to the cluster are sent using the global MAC address, so they can be received by any of the nodes at both sites; filters at the OTV localize the traffic within the data center.



Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

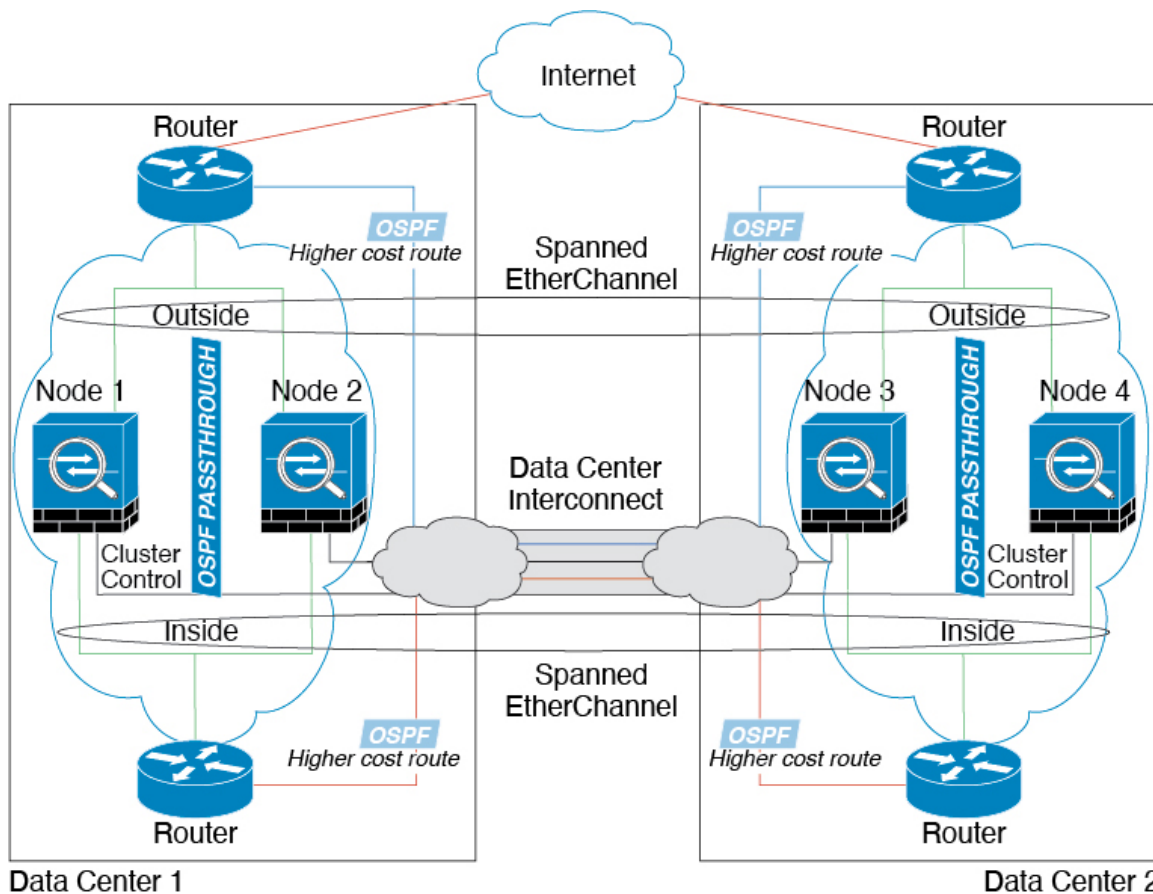
The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. By assigning a higher cost route across the DCI, traffic stays within each data center unless all cluster members at a given site go down. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the cluster members at the other site.

The implementation of the switches at each site can include:

- Inter-site VSS, vPC, StackWise, or StackWise Virtual—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster nodes at each Data Center to only connect to the local switch, while the redundant switch traffic goes across the DCI. In this case, connections

are for the most part kept local to each datacenter. You can optionally connect each node to both switches across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.

- Local VSS, vPC, StackWise, or StackWise Virtual at each site—For better switch redundancy, you can install 2 separate redundant switch pairs at each site. In this case, although the cluster nodes still have a spanned EtherChannel with Data Center 1 chassis connected only to both local switches, and Data Center 2 chassis connected to those local switches, the spanned EtherChannel is essentially “split.” Each local redundant switch system sees the spanned EtherChannel as a site-local EtherChannel.




Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. A good practice to avoid unintended MAC address flapping is to statically add the gateway routers real MAC addresses to the ASA MAC address table. Without these entries, if the gateway at site 1 communicates with the gateway at site 2, that traffic might pass through the ASA and attempt to reach

site 2 from the inside interface and cause problems. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.



History for Logical Devices

Feature Name	Platform Releases	Feature Information
Synchronization between the Firepower Threat Defense operational link state and the physical link state	2.9.1	<p>The chassis can now synchronize the Firepower Threat Defense operational link state with the physical link state for data interfaces. Currently, interfaces will be in an Up state as long as the FXOS admin state is up and the physical link state is up. The Firepower Threat Defense application interface admin state is not considered. Without synchronization from Firepower Threat Defense, data interfaces can be in an Up state physically before the Firepower Threat Defense application has completely come online, for example, or can stay Up for a period of time after you initiate an Firepower Threat Defense shutdown. For inline sets, this state mismatch can result in dropped packets because external routers may start sending traffic to the Firepower Threat Defense before the Firepower Threat Defense can handle it. This feature is disabled by default, and can be enabled per logical device in FXOS.</p> <p>Note This feature is not supported for clustering, container instances, or an Firepower Threat Defense with a Radware vDP decorator. It is also not supported for the ASA.</p> <p>New/Modified Firepower Chassis Manager screens: Logical Devices > Enable Link State</p> <p>New/Modified FXOS commands: set link-state-sync enabled, show interface expand detail</p>
Firepower Threat Defense configuration backup and restore using FMC for container instances	2.9.1	<p>You can now use the FMC backup/restore tool on an Firepower Threat Defense container instance.</p> <p>New/Modified FMC screens: System > Tools > Backup/Restore > Managed Device Backup</p> <p>New/Modified Firepower Threat Defense CLI commands: restore</p> <p>Supported platforms: Firepower 4100/9300</p> <p>Note Requires Firepower 6.7.</p>
Multi-instance clustering	2.8.1	<p>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.</p> <p>New/modified screens:</p> <ul style="list-style-type: none"> • Logical Devices > Add Cluster • Interfaces > All Interfaces > Add New drop-down menu > Subinterface > Type field <p>Note Requires Firepower 6.6 or later.</p>

Feature Name	Platform Releases	Feature Information
Support for Firepower Threat Defense with FDM	2.7.1	<p>You can now deploy a native Firepower Threat Defense instance and specify FDM management. Container instances are not supported.</p> <p>New/modified Firepower Chassis Manager screens:</p> <p>Logical Devices > Add Device > Settings > Management type of application instance</p> <p>Note Requires Firepower Threat Defense 6.5 or later.</p>
TLS crypto acceleration for multiple container instances	2.7.1	<p>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 <i>one</i> container instance per module/security engine.</p> <p>New instances have this feature enabled by default. However, the upgrade does <i>not</i> enable acceleration on existing instances. Instead, use the enter hw-crypto and then the set admin-state enabled FXOS commands.</p> <p>New/Modified Firepower Chassis Manager screens:</p> <p>Logical Devices > Add Device > Settings > Hardware Crypto drop-down menu</p> <p>Note Requires Firepower Threat Defense 6.5 or later.</p>
Firepower 4115, 4125, and 4145	2.6.1	<p>We introduced the Firepower 4115, 4125, and 4145.</p> <p>Note Requires ASA 9.12(1). Firepower 6.4.0 requires FXOS 2.6.1.157.</p> <p>No modified screens.</p>
Firepower 9300 SM-40, SM-48, and SM-56 support	2.6.1	<p>We introduced the following three security modules: SM-40, SM-48, and SM-56.</p> <p>Note The SM-40 and SM-48 require ASA 9.12(1). The SM-56 requires ASA 9.12(2) and FXOS 2.6.1.157.</p> <p>All modules require Firepower Threat Defense 6.4 and FXOS 2.6.1.157.</p> <p>No modified screens.</p>
Support for ASA and Firepower Threat Defense on separate modules of the same Firepower 9300	2.6.1	<p>You can now deploy ASA and Firepower Threat Defense logical devices on the same Firepower 9300.</p> <p>Note Requires ASA 9.12(1). Firepower 6.4.0 requires FXOS 2.6.1.157.</p> <p>No modified screens.</p>

Feature Name	Platform Releases	Feature Information
For the Firepower Threat Defense bootstrap configuration, you can now set the NAT ID for the FMC in the Firepower Chassis Manager	2.6.1	<p>You can now set the FMC NAT ID in the Firepower Chassis Manager. Previously, you could only set the NAT ID within the FXOS CLI or Firepower 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.</p> <p>New/Modified screens:</p> <p>Logical Devices > Add Device > Settings > Firepower Management Center NAT ID field</p>
Support for SSL hardware acceleration on one Firepower Threat Defense container instance on a module/security engine	2.6.1	<p>You can now enable SSL hardware acceleration for one container instance on a module/security engine. SSL hardware acceleration is disabled for other container instances, but enabled for native instances. See the FMC configuration guide for more information.</p> <p>New/Modified commands: config hwCrypto enable, show hwCrypto</p> <p>No modified screens.</p>

Feature Name	Platform Releases	Feature Information
<p>Multi-instance capability for Firepower Threat Defense</p>	<p>2.4.1</p>	<p>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).</p> <p>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.</p> <p>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.</p> <p>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.</p> <p>Note Requires Firepower Threat Defense Version 6.3 or later.</p> <p>New/Modified Firepower Chassis Manager screens:</p> <p>Overview > Devices</p> <p>Interfaces > All Interfaces > Add New drop-down menu > Subinterface</p> <p>Interfaces > All Interfaces > Type</p> <p>Logical Devices > Add Device</p> <p>Platform Settings > Mac Pool</p> <p>Platform Settings > Resource Profiles</p> <p>New/Modified FMC screens:</p> <p>Devices > Device Management > Edit icon > Interfaces tab</p>
<p>Support for transparent mode deployment for an ASA logical device</p>	<p>2.4.1</p>	<p>You can now specify transparent or routed mode when you deploy the ASA.</p> <p>New/modified Firepower Chassis Manager screens:</p> <p>Logical Devices > Add Device > Settings</p> <p>New/Modified options: Firewall Mode drop-down list</p>

Feature Name	Platform Releases	Feature Information
Cluster control link customizable IP Address	2.4.1	<p>By default, the cluster control link uses the 127.2.0.0/16 network. You can now set the network when you deploy the cluster in FXOS. The chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: <code>127.2.chassis_id.slot_id</code>. However, some networking deployments do not allow 127.2.0.0/16 traffic to pass. Therefore, you can now set a custom /16 subnet for the cluster control link in FXOS except for loopback (127.0.0.0/8) and multicast (224.0.0.0/4) addresses.</p> <p>New/modified screens:</p> <p>Logical Devices > Add Device > Cluster Information > CCL Subnet IP field</p>
For the Firepower Threat Defense bootstrap configuration, you can now set the NAT ID for the FMC at the FXOS CLI	2.4.1	<p>You can now set the FMC NAT ID at the FXOS CLI. Previously, you could only set the NAT ID within the Firepower 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.</p> <p>New/Modified commands: enter bootstrap-key NAT_ID</p>
Inter-site clustering improvement for the ASA	2.1.1	<p>You can now configure the site ID for each Firepower 4100/9300 chassis when you deploy the ASA cluster. Previously, you had to configure the site ID within the ASA application; this new feature eases initial deployment. Note that you can no longer set the site ID within the ASA configuration. Also, for best compatibility with inter-site clustering, we recommend that you upgrade to ASA 9.7(1) and FXOS 2.1.1, which includes several improvements to stability and performance.</p> <p>We modified the following screen: Logical Devices > Configuration</p>
Inter-chassis clustering for 6 Firepower Threat Defense modules on the Firepower 9300	2.1.1	<p>You can now enable inter-chassis clustering for the Firepower Threat Defense on the Firepower 9300. You can include up to 6 modules. For example, you can use 1 module in 6 chassis, or 2 modules in 3 chassis, or any combination that provides a maximum of 6 modules.</p> <p>We modified the following screen: Logical Devices > Configuration</p>
Support for Firepower Threat Defense clustering on the Firepower 4100	2.1.1	You can cluster up to 6 chassis in an Firepower Threat Defense cluster.
Support for 16 Firepower 4100 chassis in an ASA cluster	2.0.1	You can cluster up to 16 chassis in an ASA cluster.
Support for ASA clustering on the Firepower 4100	1.1.4	You can cluster up to 6 chassis in an ASA cluster.

Feature Name	Platform Releases	Feature Information
Support for intra-chassis clustering on the Firepower Threat Defense on the Firepower 9300	1.1.4	The Firepower 9300 supports intra-chassis clustering with the Firepower Threat Defense application. We modified the following screen: Logical Devices > Configuration
Inter-chassis clustering for 16 ASA modules on the Firepower 9300	1.1.3	You can now enable inter-chassis clustering for the ASA. You can include up to 16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules. We modified the following screen: Logical Devices > Configuration
Intra-chassis Clustering for the ASA on the Firepower 9300	1.1.1	You can cluster all ASA security modules within the Firepower 9300 chassis. We introduced the following screen: Logical Devices > Configuration

