Logical Devices

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 must install the same application instance type (ASA or Firepower Threat Defense) on all modules in the chassis; different types are not supported at this time. Note that modules can run different versions of an application instance type.

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 module application instances belong to a single logical device.

Note For the Firepower 9300, all modules must belong to the cluster. You cannot create a standalone logical device on one security module and then create a cluster using the remaining 2 security modules.

Requirements and Prerequisites for Logical Devices

See the following sections for requirements and prerequisites.

Requirements and Prerequisites for Clustering

Inter-Chassis Clustering Hardware and Software Requirements

All chassis in a cluster:

- All security modules 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.
- Must run the identical FXOS software except at the time of an image upgrade.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in inter-chassis clustering. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the slave units, and ending with the master.
- Must 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 slave units. For Firepower Threat Defense, all licensing is handled by the Firepower Management Center.

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 9300 chassis.
- For a list of supported switches, see Cisco FXOS Compatibility.

Sizing the Data Center Interconnect for Inter-Site Clustering

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

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

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

For example:

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

  Reserved DCI bandwidth = 5 Gbps (2/2 x 5 Gbps).
- For 6 members at 3 sites, the size increases:
  - 6 cluster members total
  - 3 members at site 1, 2 members at site 2, and 1 member at site 3
  - 10 Gbps cluster control link per member

  Reserved DCI bandwidth = 15 Gbps (3/2 x 10 Gbps).
- For 2 members at 2 sites:
  - 2 cluster members total
  - 1 member at each site
  - 10 Gbps cluster control link per member

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

Guidelines and Limitations for Logical Devices

See the following sections for guidelines and limitations.

General Guidelines and Limitations

Firewall Mode

You can set the firewall mode to routed or transparent in the bootstrap configuration for the FTD.

High Availability

- Configure high availability within the application configuration.
- You can use any data interfaces as the failover and state links.

- For more information, see

Context Mode

- Multiple context mode is only supported on the ASA.

Clustering Guidelines and Limitations

Switches for Inter-Chassis Clustering

- For the ASR 9006, if you want to set a non-default MTU, set the ASR 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 ASR *IPv4* MTU.
- 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.
- 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. 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.
- 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 **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.
- 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.

EtherChannels for Inter-Chassis Clustering

- For connecting switches, set the EtherChannel mode to Active; On mode is not supported on the Firepower 9300 chassis, even for the cluster control link.

- FXOS EtherChannels have the LACP rate set to fast by default. 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.

- 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 master switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.

- Spanned vs. Device-Local EtherChannel Configuration—Be sure to configure the switch appropriately for Spanned EtherChannels vs. Device-local EtherChannels.

  - Spanned EtherChannels—For cluster unit *Spanned* EtherChannels, which span across all members of the cluster, the interfaces are combined into a single EtherChannel on the switch. Make sure each interface is in the same channel group on the switch.

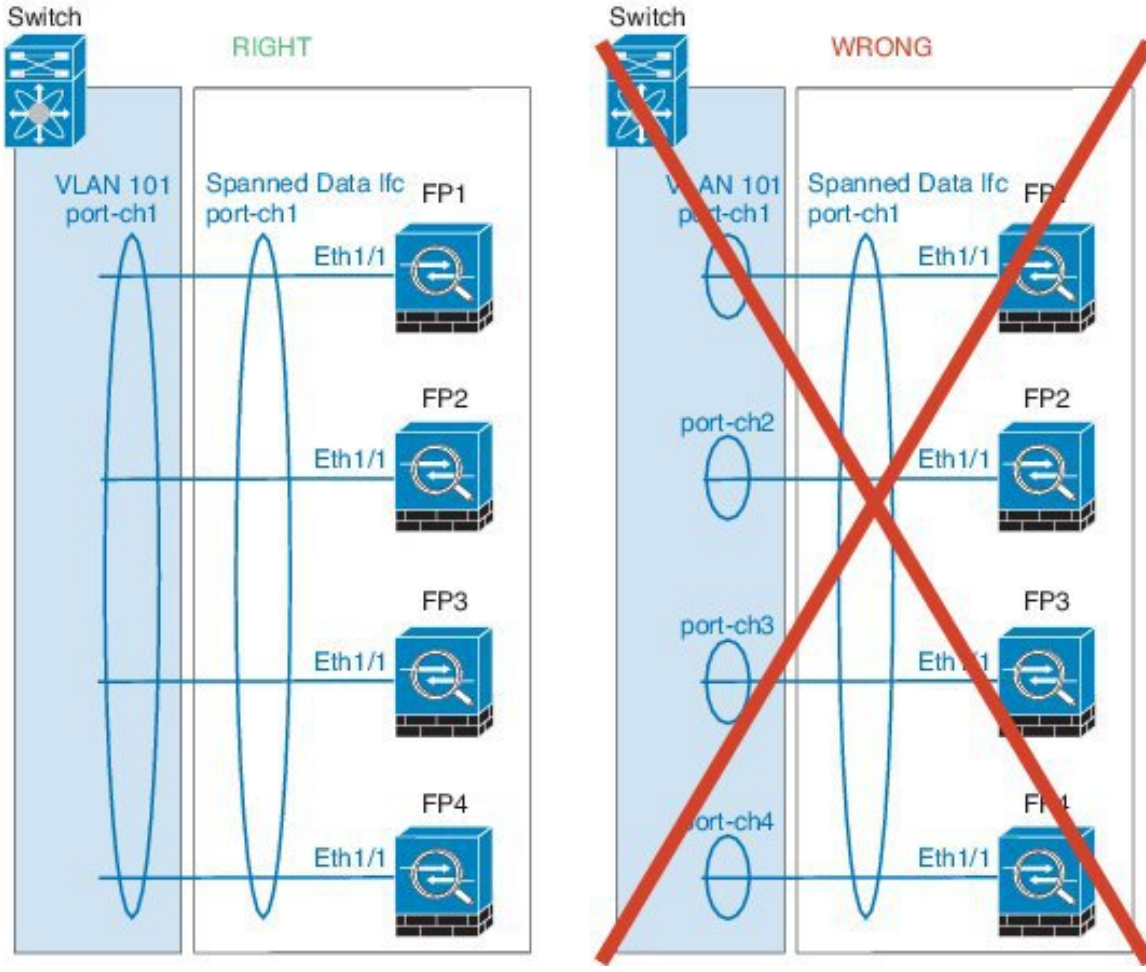

  - Device-local EtherChannels—For cluster unit *Device-local* EtherChannels including any EtherChannels configured for the cluster control link, be sure to configure discrete EtherChannels on the switch; do not combine multiple cluster unit EtherChannels into one EtherChannel on the switch.

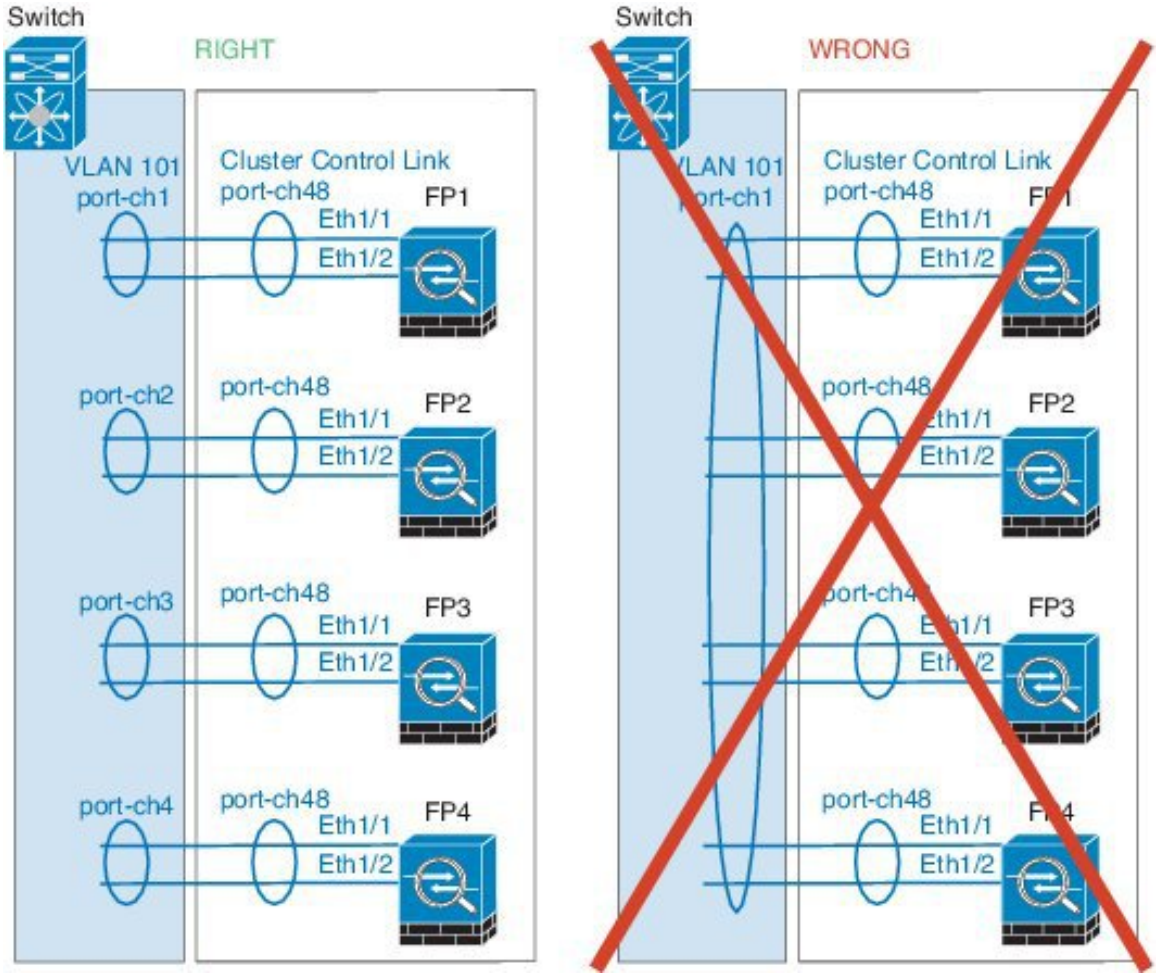

Inter-Site Clustering

See the following guidelines for inter-site clustering:

- The cluster control link latency must be less than 20 ms round-trip time (RTT).
- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.
- The 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.
- 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 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 units. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

Additional Guidelines

- We recommend connecting EtherChannels to a VSS or vPC 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.

Defaults

The cluster control link uses Port-channel 48.

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 9.

Add a Standalone ASA

Standalone logical devices work either alone or in a High Availability pair. On multiple module devices, like the Firepower 9300, 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 firewall mode ASA from the Firepower 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 9300 chassis.
- 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 interface that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).

Procedure

Step 1 Choose **Logical Devices**.

The **Logical Devices** page shows a list of logical devices on the chassis.

Step 2 Click **Add Device**.

The **Add Device** dialog box appears.

Step 3 For the **Device Name**, provide a name for the logical device.

This name is used by the Firepower 9300 chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the security module/engine configuration.

Step 4 For the **Template**, choose **Cisco Adaptive Security Appliance**.

Step 5 Choose the **Image Version**.

Step 6 For the **Device Mode**, click the **Standalone** radio button.

Step 7 Click **OK**.

You see the Provisioning - *device name* window.

Step 8 Expand the **Data Ports** area, and click each port that you want to assign to the device.

Step 9 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 change most values in the application CLI configuration.

Step 10 On the **General Information** tab, complete the following:

a) (On multiple module devices, like 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**.
c) Choose the management interface **Address Type**, **IPv4 only**, **IPv6 only**, or **IPv4 and IPv6**.
d) Configure the **Management IP** address.
e) Enter a **Network Mask** or **Prefix Length**.
f) Enter a **Network Gateway** address.

Step 11 Click the **Settings** tab.

Step 12 Enter and confirm a **Password** for the admin user.

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 13 Click **OK** to close the configuration dialog box.

Step 14 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 specified security module/engine.

Add a High Availability Pair

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

- For High Availability system requirements, see.

Procedure

Step 1 Each logical device should be on a separate chassis; intra-chassis High Availability for the Firepower 9300 is not recommended and may not be supported.

Step 2 Allocate the same interfaces to each logical device.

Step 3 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.

Step 4 Enable High Availability on the logical devices.

Step 5 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.

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.

Note The FTD does not support a cluster across multiple chassis (inter-chassis); only intra-chassis clustering is supported.

About Clustering on the Firepower 9300 Chassis

The cluster consists of multiple devices acting as a single logical unit. When you deploy a cluster on the Firepower 9300 chassis, it does the following:

- Creates a *cluster-control link* (by default, port-channel 48) for unit-to-unit communication. For intra-chassis clustering, 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 Firepower 9300 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.

The following sections provide more detail about clustering concepts and implementation.

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

The cluster control link is automatically created using the Port-channel 48 interface. For intra-chassis clustering, this interface has no member interfaces. For inter-chassis clustering, you must add one or more interfaces to the EtherChannel. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications for intra-chassis clustering.

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) or Virtual Port Channel (vPC) environment. All links in the EtherChannel are active. When the switch is part of a VSS or vPC, then you can connect Firepower 9300 chassis interfaces within the same EtherChannel to separate switches in the VSS or vPC. 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.

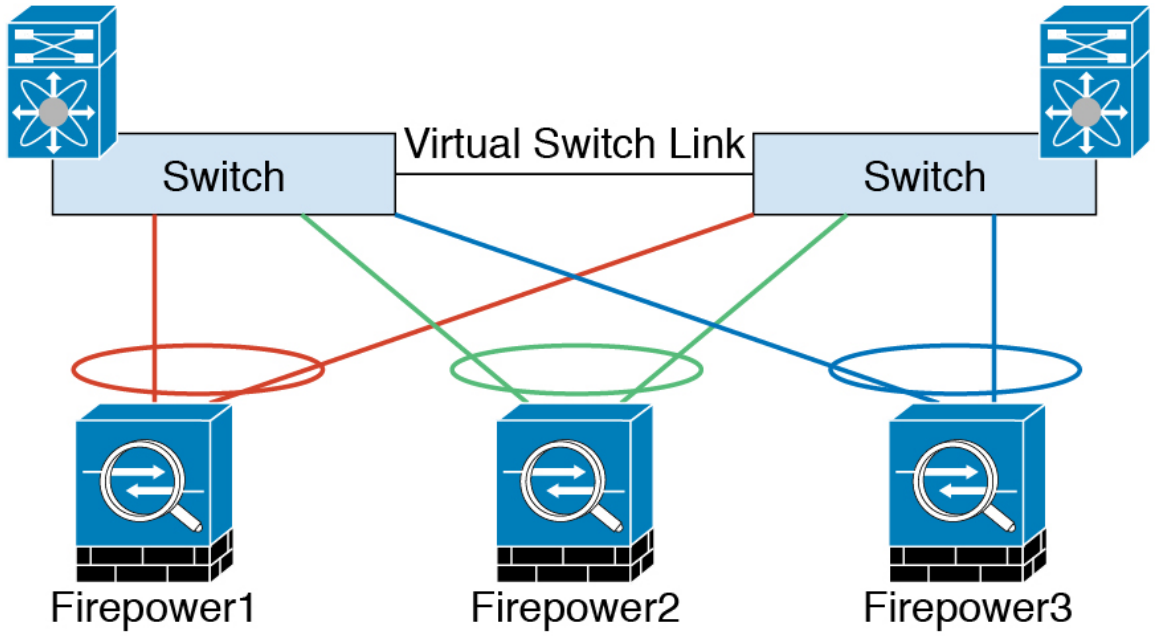

Cluster Control Link Reliability for Inter-Chassis Clustering

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 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*. You cannot set this IP address manually, either in FXOS or within the application. 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.

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.

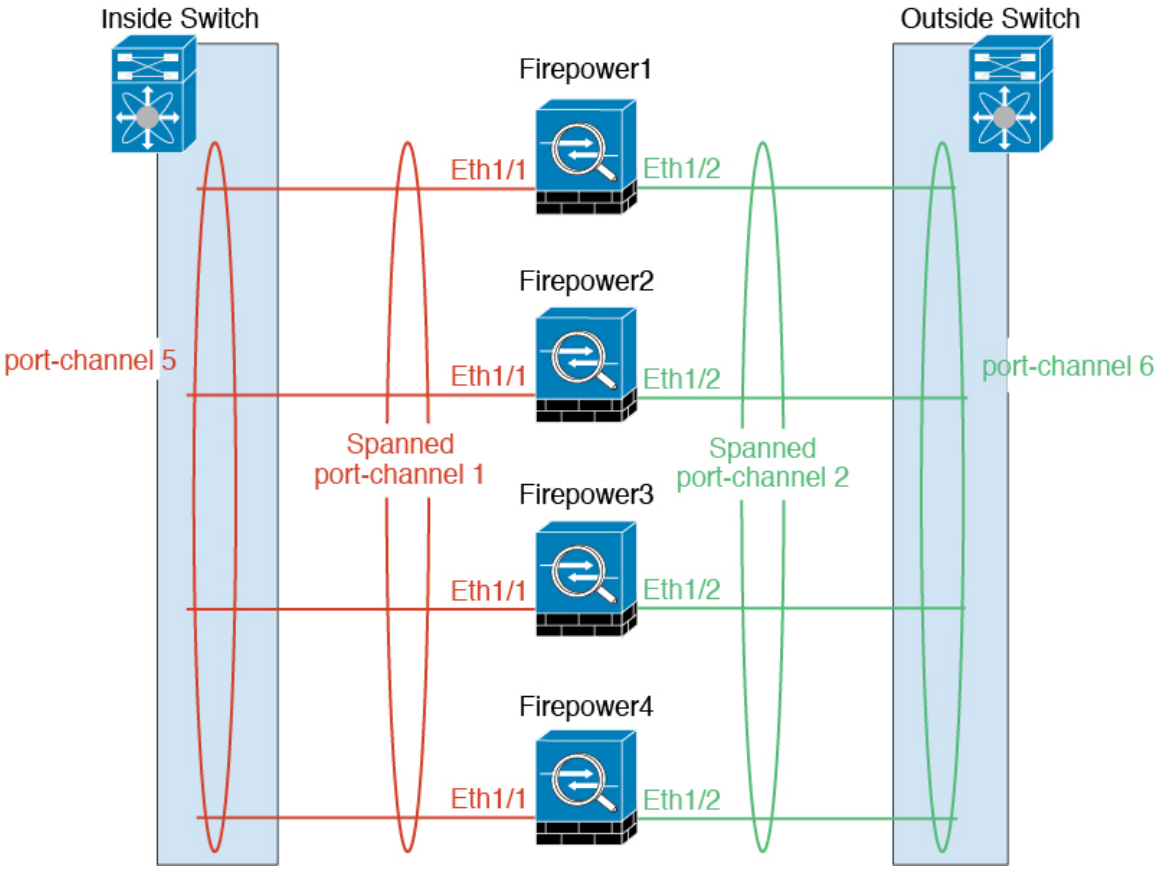

Inter-Site Clustering

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. Packets sourced from the cluster use a site-specific MAC address, while packets received by the cluster use a global MAC 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 are supported for routed mode using Spanned EtherChannels only.

Site IDs are also used to enable flow mobility using LISP inspection.

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

- Sizing the Data Center Interconnect—Requirements and Prerequisites for Clustering, on page 2
- Inter-Site Guidelines—Clustering Guidelines and Limitations, on page 4
- Inter-Site Examples—Examples for Inter-Site Clustering, on page 24

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

Deploy the cluster on the Firepower 9300 chassis.

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

You can deploy a routed firewall mode ASA from the Firepower 9300 chassis.

Before you begin

- You must enable clustering for all 3 module slots in a Firepower 9300 chassis, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.
- On the **Interfaces** tab, the port-channel 48 cluster type interface shows the **Operation State** as **failed** if it does not include any member interfaces. For intra-chassis clustering, this EtherChannel does not require any member interfaces, and you can ignore this Operational State.

Procedure

Step 1 Add at least one Data type interface or EtherChannel (also known as a port-channel) before you deploy the cluster. See Add an EtherChannel (Port Channel) or Configure a Physical Interface.

You can also add data interfaces to the cluster after you deploy it.

For inter-chassis clustering, all data interfaces must be EtherChannels with at least one member interface. Add the same EtherChannels on each chassis.

Step 2 Add a Management type interface or EtherChannel. See Add an EtherChannel (Port Channel) or Configure a Physical Interface.

For inter-chassis clustering, add the same management interface on each chassis. The management interface is required. Note that this management interface is not the same as the chassis management interface that is used only for chassis management (and that appears at the top of the **Interfaces** tab as **MGMT**).

Step 3 For inter-chassis clustering, add a member interface to port-channel 48, which is used as the cluster control link.

If you do not include a member interface, then when you deploy the logical device, the Firepower Chassis Manager assumes that this cluster is an intra-chassis cluster and does not show the **Chassis ID** field. Add the same member interfaces on each chassis.

Step 4 Choose **Logical Devices**.

The **Logical Devices** page shows a list of logical devices on the chassis.

Step 5 Click **Add Device**.

The **Add Device** dialog box appears.

Step 6 For the **Device Name**, provide a name for the logical device.

This name is used by the Firepower 9300 chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the security module/engine configuration.

Step 7 For the **Template**, choose **Cisco Adaptive Security Appliance**.

Step 8 Choose the ASA **Image Version**.

Step 9 For the **Device Mode**, click the **Cluster** radio button.

Step 10 Click the **Create New Cluster** radio button.

Step 11 Click **OK**.

If you have any standalone devices configured, you are prompted to replace them with a new cluster. You see the Provisioning - *device name* window.

All interfaces are assigned to the cluster by default.

Step 12 Click the device icon in the center of the screen.

The **ASA Configuration** dialog box appears with the **Cluster Information** tab selected.

Step 13 In the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

Step 14 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.

Step 15 Set the **Cluster Group Name**, which is the cluster group name in the security module configuration.

The name must be an ASCII string from 1 to 38 characters.

Step 16 Click **Management Interface** and choose the management interface you created earlier.

Step 17 Choose the **Address Type** for the management interface.

This information is used to configure a management interface in the security module configuration.

a) In the **Management IP Pool** field, 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 master 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.

b) Enter a **Network Mask** or **Prefix Length**.

c) Enter a **Network Gateway**.

d) Enter a **Virtual IP address**.

   This IP address must be on the same network as the cluster pool addresses, but not be part of the pool.

Step 18 Click the **Settings** tab.

Step 19 Enter and confirm a **Password** for the admin user.

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 20 Click **OK** to close the ASA Configuration dialog box.

Step 21 Click **Save**.

The Firepower 9300 chassis supervisor deploys the cluster by downloading the specified software version and pushing the cluster bootstrap configuration and management interface settings to each security module.

Step 22 For inter-chassis clustering, add the next chassis to the cluster:

a) On the first chassis Firepower Chassis Manager, click the **Show Cluster Details** 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 **Join an Existing Cluster**.
d) Click the **Copy config** check box, and click **OK**. If you uncheck this check box, you must manually enter the settings to match the first chassis configuration.
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.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

g) Click **Save**.

Step 23 Connect to the master unit ASA to customize your clustering configuration.

Add More Cluster Members

Add or replace an 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.
- 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 chassis 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 Device**.

Step 4 For the **Device Name**, provide a name for the logical device.

Step 5 For the **Template**, choose **Cisco Adaptive Security Appliance**.

Step 6 For the **Image Version**, choose the ASA software version.

Step 7 For the **Device Mode**, click the **Cluster** radio button.

Step 8 Choose **Join an Existing Cluster**.

Step 9 Click the **Copy config** check box, and click **OK**. If you uncheck this check box, you must manually enter the settings to match the first chassis configuration.

Step 10 In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.

Step 11 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.
- **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

Step 12 Click **Save**.

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.

connect module *slot_number* **console**

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

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.

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

You might want to use the FXOS module CLI for troubleshooting purposes.

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

a) Enter ~

You exit to the Telnet application.

b) To exit the Telnet application, enter:

telnet>**quit**

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
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.

Change the ASA to Transparent Firewall Mode

You can only deploy a routed firewall mode ASA from the Firepower 9300 chassis. To change the ASA to transparent firewall mode, complete the initial deployment, and then change the firewall mode within the ASA CLI. For standalone ASAs, because changing the firewall mode erases the configuration, you must then redeploy the configuration from the Firepower 9300 chassis to regain the bootstrap configuration. The ASA then remains in transparent mode with a working bootstrap configuration. For clustered ASAs, the configuration is not erased, so you do not need to redeploy the bootstrap configuration from FXOS.

Procedure

Step 1 Connect to the ASA console according to Connect to the Console of the Application, on page 17. For a cluster, connect to the primary unit. For a failover pair, connect to the active unit.

Step 2 Enter configuration mode:

enable

configure terminal

By default, the enable password is blank.

Step 3 Set the firewall mode to transparent:

firewall transparent

Step 4 Save the configuration:

write memory

For a cluster or failover pair, this configuration is replicated to secondary units:

```
asa(config)# firewall transparent
asa(config)# write memory
Building configuration...
Cryptochecksum: 9f831dfb 60dffa8c 1d939884 74735b69

3791 bytes copied in 0.160 secs
[OK]
asa(config)#
Beginning configuration replication to Slave unit-1-2
 End Configuration Replication to slave.

asa(config)#
```

Step 5 On the Firepower Chassis Manager **Logical Devices** page, click the **Edit** icon to edit the ASA.

The **Provisioning** page appears.

Step 6 Click the device icon to edit the bootstrap configuration. Change any value in your configuration, and click **OK**.

You must change the value of at least one field, for example, the **Password** field.

You see a warning about changing the bootstrap configuration; click **Yes**.

Step 7 For an inter-chassis cluster or for a failover pair, repeat steps 5 through 7 to redeploy the bootstrap configuration on each chassis.

Wait several minutes for the chassis/security modules to reload, and for the ASA to become operational again. The ASA now has an operational bootstrap configuration, but remains in transparent mode.

Change an Interface on a Firepower Threat Defense Logical Device

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

Before you begin

- Configure your interfaces, and add any EtherChannels according to Configure a Physical Interface and Add an EtherChannel (Port Channel).
- You can edit the membership of an allocated EtherChannel without affecting the logical device or requiring a sync on the Firepower Management Center.
- 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 firepower 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 Firepower Threat Defense device reboots (management interface changes cause a reboot), and you sync the configuration in the Firepower Management Center, 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 Firepower Management Center. We recommend that you make the interface changes on the slave/standby unit(s) first, and then on the master/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 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 or eventing interface:

For these types of interfaces, the device reboots after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General/Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) On the **Settings** tab, choose the new **Eventing Interface** from the drop-down list.
d) Click **OK**.

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

Step 6 Click **Save**.

Step 7 Log into the Firepower Management Center.

Step 8 Select **Devices** > **Device Management** and click the edit icon ( ) for your FTD device. The **Interfaces** tab is selected by default.

Step 9 Click the **Sync Interfaces from device** button on the top left of the **Interfaces** tab.

Step 10 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

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.

Before you begin

- Configure your interfaces and add any EtherChannels according to Configure a Physical Interface and Add an EtherChannel (Port Channel).
- You can edit the membership of an allocated EtherChannel without impacting the logical device.
- 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 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), 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.
- 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 slave/standby unit(s) first, and then on the master/active unit. New interfaces are added in an administratively down state, so they do not affect interface monitoring.

Procedure

Step 1 In the Firepower Chassis Manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Unallocate a data interface by de-selecting the interface in the **Data Ports** area.

Step 4 Allocate a new data interface by selecting the interface in the **Data Ports** area.

Step 5 Replace the management interface:

For this type of interface, the device reloads after you save your changes.

a) Click the device icon in the center of the page.
b) On the **General/Cluster Information** tab, choose the new **Management Interface** from the drop-down list.
c) Click **OK**.

Step 6 Click **Save**.

Modify or Recover Bootstrap Settings for a Logical Device

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

Procedure

Step 1 In the Firepower Chassis Manager, choose **Logical Devices**.

Step 2 Click the **Edit** icon at the top right to edit the logical device.

Step 3 Click the device icon in the center of the page.

Step 4 Modify the logical device settings as required.

Step 5 Click **OK**.

Step 6 Click **Save** to save the changes and restart the application instance.

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 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:

- **Security Module**—Shows the security module.
- **Ports**—Shows the ports assigned to the application instance.
- **Application**—Shows the application running on the security module.
- **Version**—Shows the software version number of the application running on the security module.
- **Management IP**—Shows the local IP address assigned as the logical device Management IP.
- **Management URL**—Shows the management URL assigned to the application instance.
- **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.

- **Attributes**—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.

  - **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, master or slave.
  - **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:

- **Add Device**—Allows you to create a logical device.
- **Edit**—Allows you to edit an existing logical device.
- **Update 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.
- **Go To Device Manager**—Provides a link to the Firepower Management Center or ASDM defined for the application instance.

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

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—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster units at each Data Center to only connect to the local switch, while the VSS/vPC traffic goes across the DCI. In this case, connections are for the most part kept local to each datacenter. You can optionally connect each unit 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 at each site—For better switch redundancy, you can install 2 separate VSS/vPC pairs at each site. In this case, although the cluster units 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 VSS/vPC sees the spanned EtherChannel as a site-local EtherChannel.

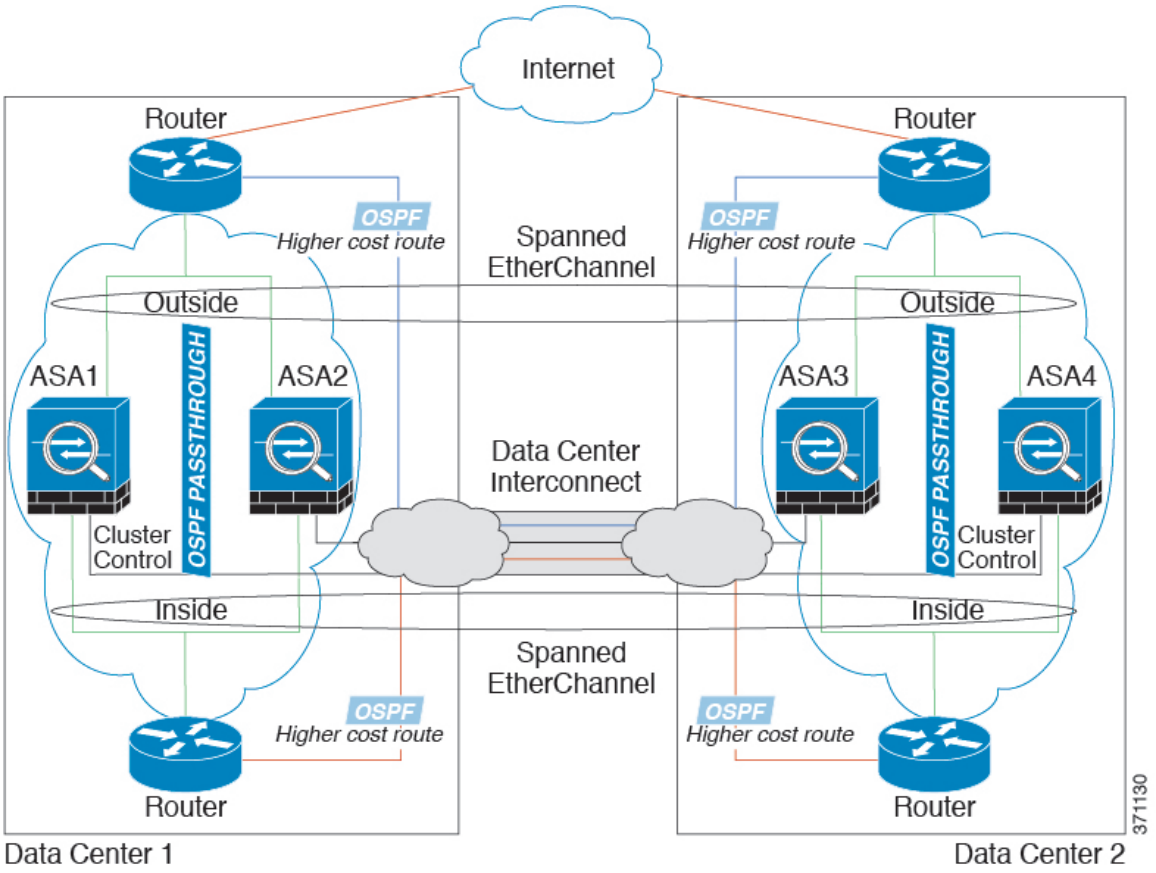

Spanned EtherChannel Transparent Mode East-West Inter-Site Example

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.

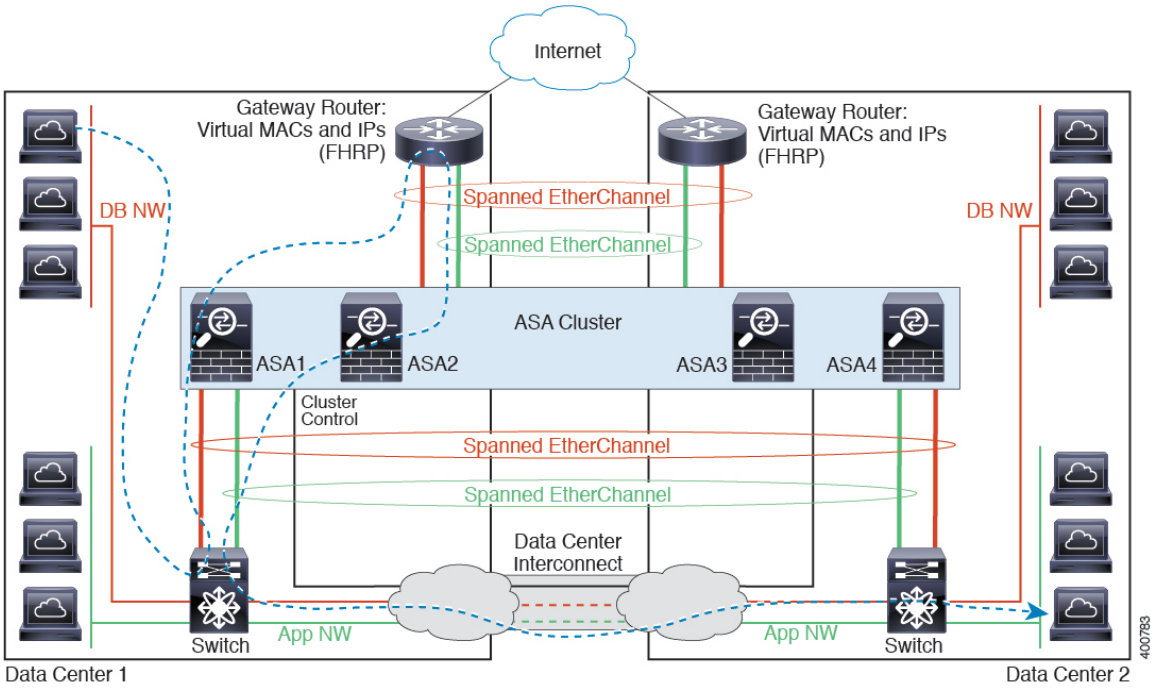

See Spanned EtherChannel Transparent Mode North-South Inter-Site Example, on page 25 for information about vPC/VSS options.

History for Logical Devices

Feature Name	Platform Releases	Feature Information
Inter-chassis clustering for 6 ASA modules	1.1.3	You can now enable inter-chassis clustering for the ASA. You can include up to 6 modules in up to 6 chassis. We modified the following screen: **Logical Devices** > **Configuration**
Intra-chassis Clustering for the Cisco ASA	1.1.1	You can cluster all ASA security modules within the Firepower 9300 chassis. We introduced the following screen: **Logical Devices** > **Configuration**