

Clustering for the Firepower Threat Defense

Clustering lets you group multiple Firepower Threat Defense units together as a single logical device. Clustering is only supported for the Firepower Threat Defense device on the Firepower 9300. 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.

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

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 23.

- About Clustering on the Firepower 4100/9300 Chassis, on page 1
- Licenses for Clustering, on page 5
- Requirements and Prerequisites for Clustering, on page 6
- Clustering Guidelines and Limitations, on page 6
- Configure Clustering, on page 7
- Remove a Cluster Node, on page 13
- FMC: Manage Cluster Members, on page 15
- FMC: Monitoring the Cluster, on page 17
- Examples for Clustering, on page 18
- Reference for Clustering, on page 22
- History for Clustering, on page 32

About Clustering on the Firepower 4100/9300 Chassis

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

- Creates a *cluster-control link* (by default, port-channel 48) for node-to-node communication.

  For a cluster isolated to security modules within one Firepower 9300 chassis, this link utilizes the Firepower 9300 backplane for cluster communications.

- 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.

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

  For a cluster isolated to security modules within one Firepower 9300 chassis, spanned interfaces are not limited to EtherChannels. 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.



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

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

See the following sections for more information about clustering.

Bootstrap Configuration

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.

Cluster Members

Cluster members work together to accomplish the sharing of the security policy and traffic flows.

One member of the cluster is the **control** unit. The control unit is determined automatically. All other members are **data** units.

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

Some features do not scale in a cluster, and the control unit handles all traffic for those features. .

Cluster Control Link

The cluster control link is automatically created using the Port-channel 48 interface.

For a cluster isolated to security modules within one Firepower 9300 chassis, this interface has no member interfaces. This Cluster type EtherChannel utilizes the Firepower 9300 backplane for cluster communications. For clustering with multiple chassis, you must add one or more interfaces to the EtherChannel.

For a cluster with two chassis, 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

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:

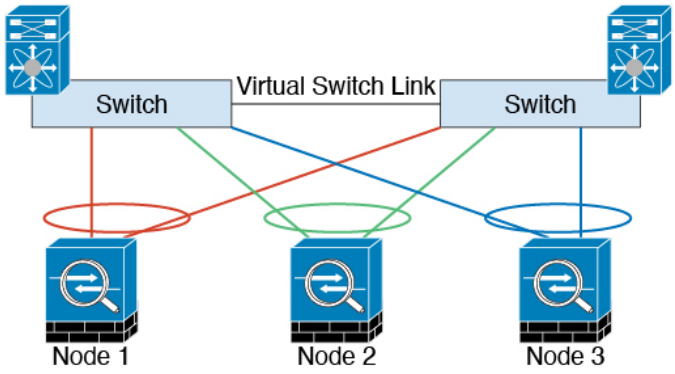
- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

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

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*. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed.

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. This Management logical interface is separate from the other interfaces on the device. It is used to set up and register the device to the Firepower Management Center. It runs a separate SSH server and uses its own local authentication, IP address, and static routing. Each cluster member uses a separate IP address on the management network that you set as part of the bootstrap configuration.

The management interface is shared between the Management logical interface and the *Diagnostic* logical interface. The Diagnostic logical interface is optional and is not configured as part of the bootstrap configuration. The Diagnostic interface can be configured along with the rest of the data interfaces. If you choose to configure the Diagnostic interface, configure a Main cluster IP address as a fixed address for the cluster that always belongs to the current control unit. You also configure a range of addresses so that each unit, including the current control unit, can use a Local address from the range. The Main cluster IP address provides consistent diagnostic access to an address; when a control unit changes, the Main cluster IP address moves to the new control unit, so access to the cluster continues seamlessly. For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current control 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 control unit, uses the Local IP address to connect to the server.

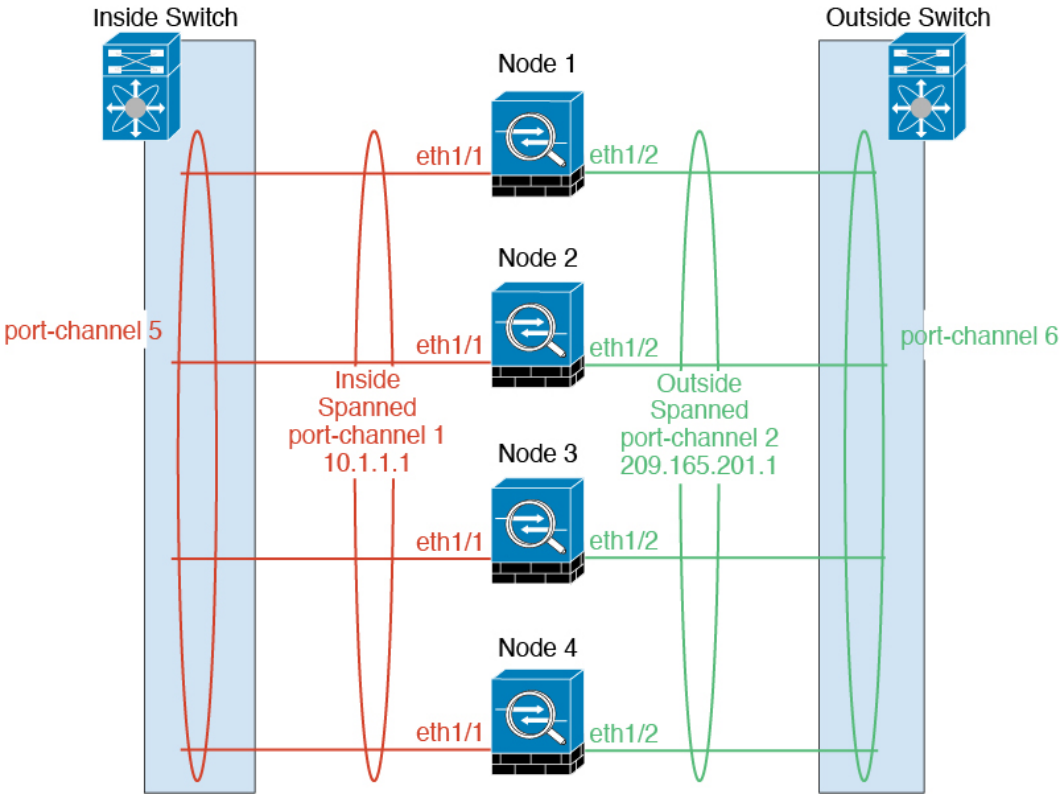
Cluster Interfaces

For a cluster isolated to security modules within one Firepower 9300 chassis, you can assign both physical interfaces or EtherChannels (also known as port channels) to the cluster. Interfaces assigned to the cluster are Spanned interfaces that load-balance traffic across all members of the cluster.

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

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.



Connecting to a Redundant Switch System

We recommend connecting EtherChannels to a redundant switch system such as a VSS, vPC, StackWise, or StackWise Virtual system to provide redundancy for your interfaces.

Configuration Replication

All nodes in the cluster share a single configuration. You can only make configuration changes on the control node (with the exception of the bootstrap configuration), and changes are automatically synced to all other nodes in the cluster.

Licenses for Clustering

The Firepower Threat Defense uses Smart Licensing. You assign licenses to the cluster as a whole, not to individual units. However, each unit of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add a cluster member to the FMC, you can specify the feature licenses you want to use for the cluster. If you choose different feature licenses for each device before you form them into a cluster in the FMC, then the licenses you chose for the control unit are used for the cluster. You can modify licenses for the cluster in the **Devices** > **Device Management** > **Cluster** > **License** area.

Note If you add the cluster before the FMC is licensed (and running in Evaluation mode), then when you license the FMC, you can experience traffic disruption when you deploy policy changes to the cluster. Changing to licensed mode causes all data units to leave the cluster and then rejoin.

Requirements and Prerequisites for Clustering

Cluster Model Support

The FTD supports clustering on the following models:

- Firepower 9300—You must include all modules installed in the chassis for a maximum 3-unit, intra-chassis cluster.

User Roles

- Admin
- Access Admin
- Network Admin

Clustering Guidelines and Limitations

- 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.

Configure Clustering

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit. You can then add the units to the FMC and group them into a cluster.

FXOS: Add a FTD Cluster

You can add a cluster to a single Firepower 9300 chassis that is isolated to security modules within the chassis.

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.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, 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.
- 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** > **ClusterAdd Device**, and set the following parameters:

a) 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.

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

c) Choose the **Image Version**.
d) For the **Instance Type**, .
e) Click **OK**.

You see the Provisioning - *device name* window.

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

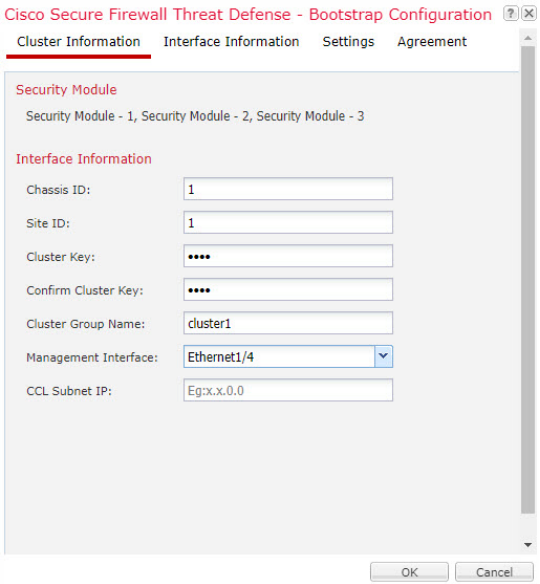
All valid interfaces are assigned by default.

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 1:



a) 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.

b) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.

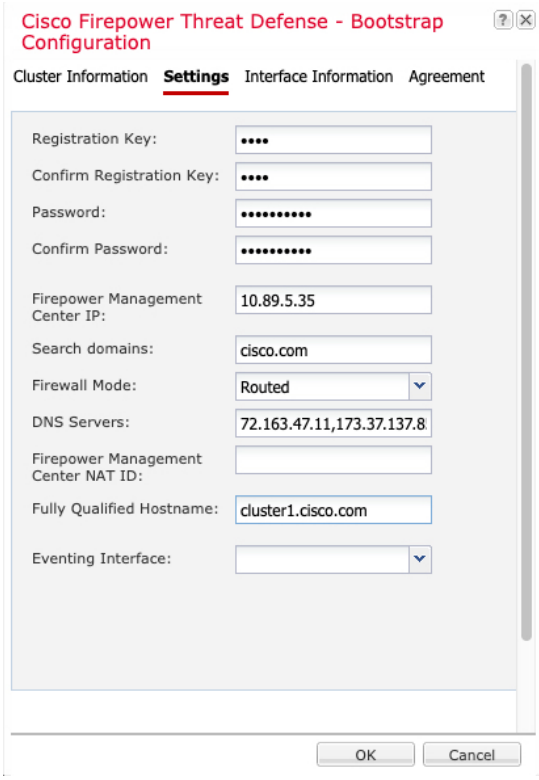
The name must be an ASCII string from 1 to 38 characters.

Important From 2.4.1, spaces in cluster group name will be considered as special characters and may result in error while deploying the logical devices. To avoid this issue, you must rename the cluster group name without a space.

c) Choose the **Management Interface**.

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

Step 7 On the **Settings** page, complete the following.



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

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

d) (Optional) In the **Search Domains** field, enter a comma-separated list of search domains for the management network.

e) (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.

f) (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.

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

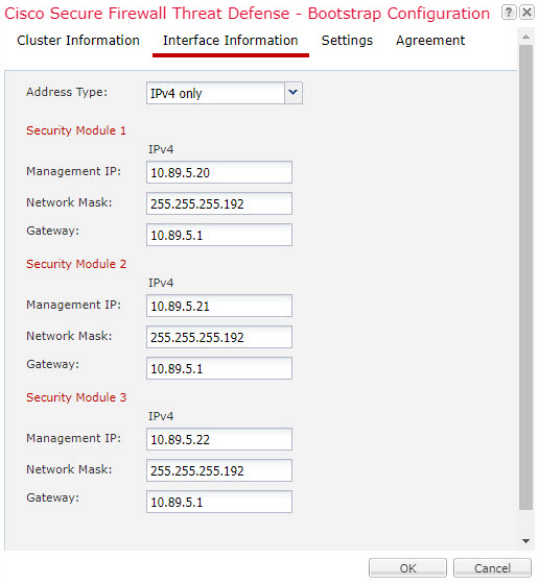
Valid characters are the letters from a to z, the digits from 0 to 9, the dot (.), and the hyphen (-); maximum number of characters is 253.

h) (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.

Step 8 On the **Interface Information** page, configure a management IP address for each security module in the cluster. Select the type of address from the **Address Type** drop-down list and then complete the following for each security module.

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



a) In the **Management IP** field, configure an IP address.

Specify a unique IP address on the same network for each module.

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

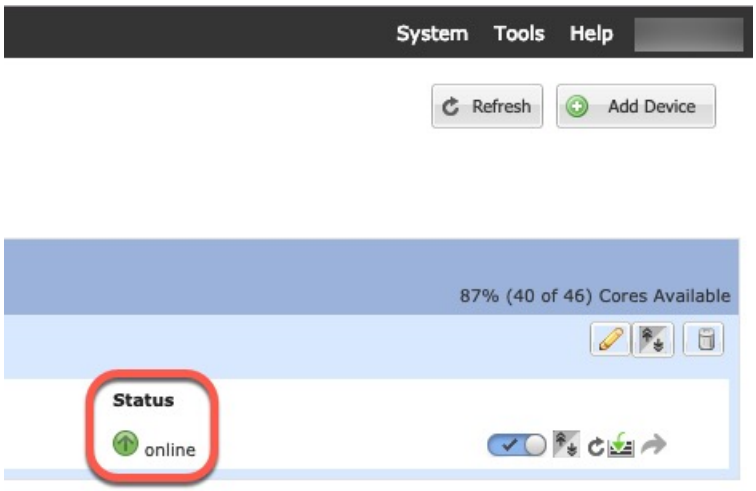
c) Enter a **Network Gateway** address.

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

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

Step 11 Click **Save**.

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



FMC: Add a Cluster

Add the logical devices to the FMC, and then group them into a cluster.

Before you begin

- Refer to the Firepower Chassis Manager **Logical Devices** screen to see which unit is the control unit.
- All cluster units must be in a successfully formed cluster on FXOS prior to adding them to the FMC.

Procedure

Step 1 In the FMC, choose **Devices** > **Device Management**, and choose **Add** > **Add Device** to add each unit as a separate managed device using the management IP addresses you assigned when you deployed the cluster.

Step 2 Choose **Add** > **Add Cluster** to group the units into a cluster.

a) Choose the control device from the drop-down list.

All other eligible members are added to the data devices box.

b) Specify a **Name** for the cluster.

c) Click **OK**.

The cluster object is added to the **Devices** screen, with the member units underneath. The current control unit is indicated by its role after the unit name.

Note If you add more units to the cluster later on the FXOS chassis, then you must add each unit to the FMC, and then add them as data nodes of the cluster as soon as possible.

Step 3 To configure device-specific settings, click **Edit** ( ) for the cluster; you can only configure the cluster as a whole, and not member units in the cluster.

Step 4 On the **Devices** > **Device Management** > **Cluster**, you can see **General**, **License**, and **Health** settings. This is useful for setting license entitlements.

Step 5 On the **Devices** > **Device Management** > **Devices**, you can choose each member in the cluster from the top right drop-down menu.

If you change the management IP address in the device configuration, you must match the new address in the FMC so that it can reach the device on the network; edit the **Host** address in the **Management** area.

FMC: Configure Cluster, Data, and Diagnostic Interfaces

This procedure configures basic parameters for each data interface that you assigned to the cluster when you deployed it in FXOS. For inter-chassis clustering, data interfaces are always Spanned EtherChannel interfaces. For the cluster control link interface for inter-chassis clustering, you must increase the MTU from the default. You can also configure the Diagnostic interface, which is the only interface that can run as an individual interface.

Note When using Spanned EtherChannels for inter-chassis clustering, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

Step 1 Choose **Devices** > **Device Management**, and click **Edit** ( ) next to the cluster.

Step 2 Click **Interfaces**.

Step 3 Configure the cluster control link.

For inter-chassis clustering, set the cluster control link MTU to be at least 100 bytes higher than the highest MTU of the data interfaces. 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. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 bytes. For example, because the maximum MTU is 9184, then the highest data interface MTU can be 9084, while the cluster control link can be set to 9184.

The cluster control link interface is Port-Channel48 by default.

a) Click **Edit** ( ) for the cluster control link interface.
b) On the **General** page, in the **MTU** field, enter a value between 1400 and 9184. We suggest using the maximum, 9184.
c) Click **OK**.

Step 4 Configure data interfaces.

a) (Optional) Configure VLAN subinterfaces on the data interface. The rest of this procedure applies to the subinterfaces.

b) Click **Edit** ( ) for the data interface.

c) Configure the name, IP address, and other parameters.

Note If the cluster control link interface MTU is not at least 100 bytes higher than the data interface MTU, you will see an error that you must reduce the MTU of the data interface. See step Step 3, on page 12 to increase the cluster control link MTU, after which you can continue configuring the data interfaces.

d) For inter-chassis clusters, set a manual global MAC address for the EtherChannel. Click **Advanced**, and in the **Active Mac Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Do not set the **Standby Mac Address**; it is ignored.

You must configure a MAC address for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

e) Click **OK**. Repeat the above steps for other data interfaces.

Step 5 (Optional) Configure the Diagnostic interface.

The Diagnostic interface is the only interface that can run in Individual interface mode. You can use this interface for syslog messages or SNMP, for example.

a) On **Devices** > **Device Management** > **Interfaces**, click **Edit** ( ) for the Diagnostic interface.
b) On the **IPv4**, enter the **IP Address** and mask. This IP address is a fixed address for the cluster, and always belongs to the current control unit.
c) For the **Start Address** and **End Address**, enter the start and end of an IP address pool, one of which will be assigned to each cluster unit for the interface.

Include at least as many addresses as there are units in the cluster. The Virtual IP address is not a part of this pool, but needs to be on the same network. You cannot determine the exact Local address assigned to each unit in advance.

d) For the **Mask**, enter the subnet mask for the cluster IP pool.
e) On the **IPv6**, under **Basic > Cluster IPv6 Pool** specify the *ipv6-address/prefix-length* for the **Address** and the number of addresses in the pool for the **Count**.
f) Configure other interface settings as normal.

Step 6 Click **Save**.

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

Remove a Cluster Node

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

Temporary Removal

A cluster node 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:



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

- Disable the application instance—In the Firepower Chassis Manager on the **Logical Devices** page, click the **Slider enabled** (). You can later reenable it using the **Slider disabled** ().

- Shut down the security module/engine—In the Firepower Chassis Manager on the **Security Module/Engine** page, click the **Power Off icon**.

- Shut down the chassis—In the Firepower Chassis Manager on the **Overview** page, click the **Shut Down icon**.

Permanent Removal

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

- 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 node of the cluster.

FMC: Manage Cluster Members

After you deploy the cluster, you can change the configuration and manage cluster members.

Add or Replace a Cluster Member

You can add a new cluster member to an existing cluster, for example, when you add an additional module to the Firepower 9300 device, or replace a unit. This procedure also applies to a unit that was reinitialized; in this case, although the hardware remains the same, it appears to be a new member.

Before you begin

- Add the units to the cluster on the FXOS chassis, and make sure they are in the FXOS cluster before you add them to the FMC. Make sure the interface configuration is the same as other chassis.

Procedure

Step 1 In the case of a replacement or a reinitialized unit, you must delete the old cluster member from the Firepower Management Center.

When you replace it with a new or reinitialized unit, it is considered to be a new device on the Firepower Management Center

a) In the FMC, choose **Devices** > **Device Management**, and click **Delete** () next to the data unit.
b) Confirm that you want to delete the unit.

The unit is removed from the cluster and from the FMC devices list.

Step 2 Choose **Devices** > **Device Management**, and choose **Add** > **Add Device** to add the new logical device.

Step 3 Choose **Add** > **Add Cluster**.

Step 4 Choose the current control device from the drop-down list.

When you choose a control device that is already in a cluster, then the existing cluster name is auto-filled, and all eligible data devices are added to the data devices box, including the new unit you just added to the FMC.

Step 5 Click **Add**, and then **Deploy**.

The cluster is updated to include the new member(s).

Deactivate a Member

You may want to deactivate a member in preparation for deleting the unit, or temporarily for maintenance. This procedure is meant to temporarily deactivate a member; the unit will still appear in the FMC device list.

To deactivate a member other than the unit you are logged into, perform the following steps at the Firepower Threat Defense CLI.

Note When a unit becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, reenable 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. You must use the console for any further configuration.

Procedure

Step 1 Access the Firepower Threat Defense CLI.

Step 2 Remove the unit from the cluster:

cluster remove unit *unit_name*

The bootstrap configuration remains intact, as well as the last configuration synched from the control unit, so that 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.

To view member names, enter **cluster remove unit ?**, or enter the **show cluster info** command.

Example:

```
> cluster remove unit ?

Current active units in the cluster:
ftd1
ftd2
ftd3

> cluster remove unit ftd2
WARNING: Clustering will be disabled on unit ftd2. To bring it back
to the cluster please logon to that unit and re-enable clustering
```

Step 3 To reenable clustering, see Rejoin the Cluster, on page 16.

Rejoin the Cluster

If a unit was removed from the cluster, for example for a failed interface, you must manually rejoin the cluster by accessing the unit CLI. Make sure the failure is resolved before you try to rejoin the cluster. See Rejoining the Cluster, on page 28 for more information about why a unit can be removed from a cluster.

Procedure

Step 1 Access the CLI of the unit that needs to rejoin the cluster, either from the console port or using SSH to the Management interface. Log in with the username **admin** and the password you set during initial setup.

Step 2 Enable clustering:

cluster enable

Delete a Data Unit

If you need to permanently remove a cluster member (for example, if you remove a module on the Firepower 9300), then you should delete it from the FMC.

Do not delete the member if it is still a healthy part of the cluster, or if you only want to disable the member temporarily. To delete it permanently from the cluster in FXOS, see Remove a Cluster Node, on page 13. If you remove it from the FMC, and it is still part of the cluster, it will continue to pass traffic, and could even become the control unit—a control unit that the FMC can no longer manage.

Before you begin

To manually deactivate the unit, see Deactivate a Member, on page 15. Before you delete a unit, the unit must be inactive, either manually or because of a health failure.

Procedure

Step 1 In the FMC for the data unit you want to delete, choose **Devices** > **Device Management**, and click **Delete** () .

Step 2 Confirm that you want to delete the unit.

The unit is removed from the cluster and from the FMC devices list.

FMC: Monitoring the Cluster

You can monitor the cluster in Firepower Management Center and at the Firepower Threat Defense CLI.

- **Devices** > **Device Management** > *cluster_name*.

  When you expand the cluster on the devices listing page, you can see all member units, including the control unit shown with its role next to the IP address.

- **show cluster** {**access-list** [*acl_name*] | **conn** [**count**] | **cpu** [**usage**] | **history** | **interface-mode** | **memory** | **resource usage** | **service-policy** | **traffic** | **xlate count**}

  To view aggregated data for the entire cluster or other information, use the **show cluster** command.

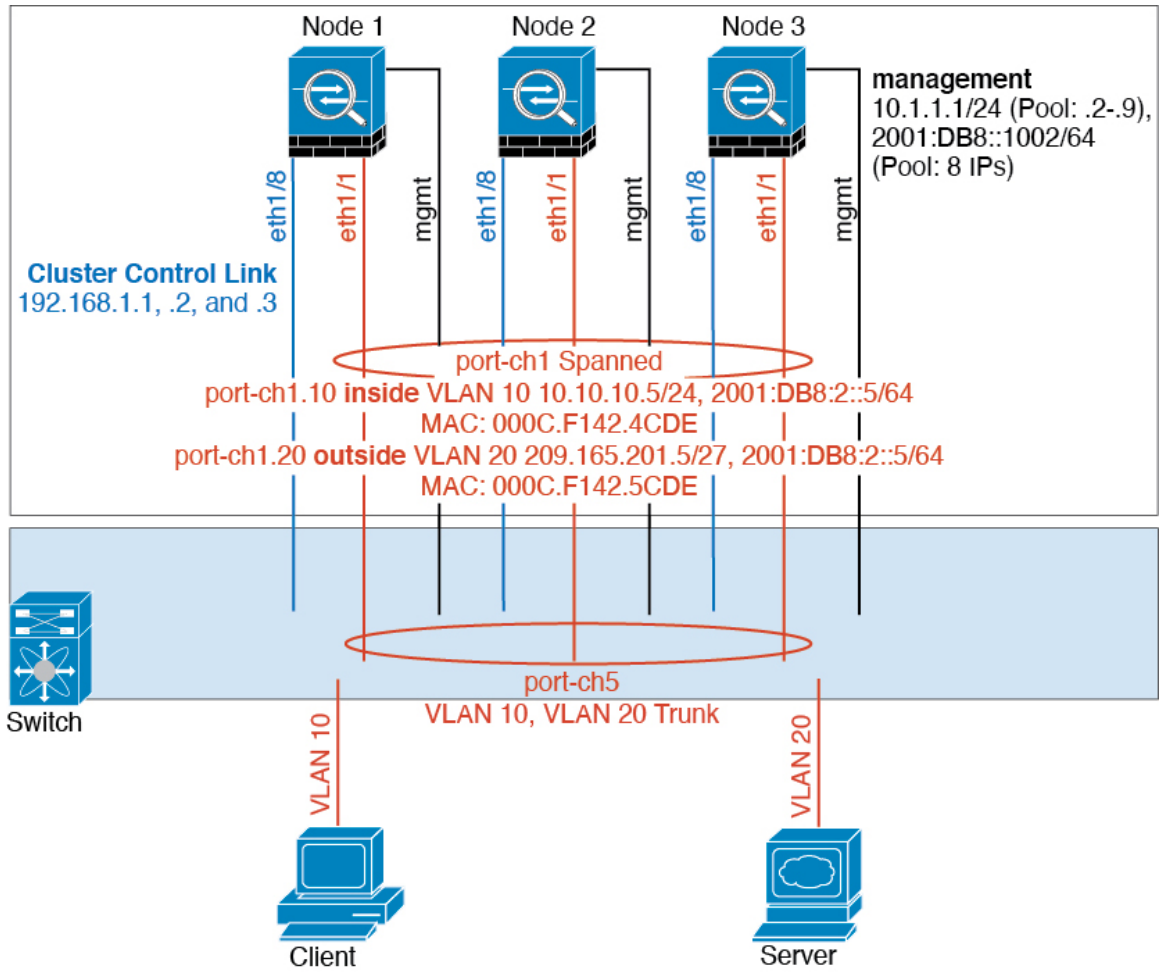
- **show cluster info** [**auto-join** | **clients** | **conn-distribution** | **flow-mobility counters** | **goid** [*options*] | **health** | **incompatible-config** | **loadbalance** | **old-members** | **packet-distribution** | **trace** [*options*] | **transport** { **asp** | **cp**}]

  To view cluster information, use the **show cluster info** command.

Examples for Clustering

These examples include typical deployments.

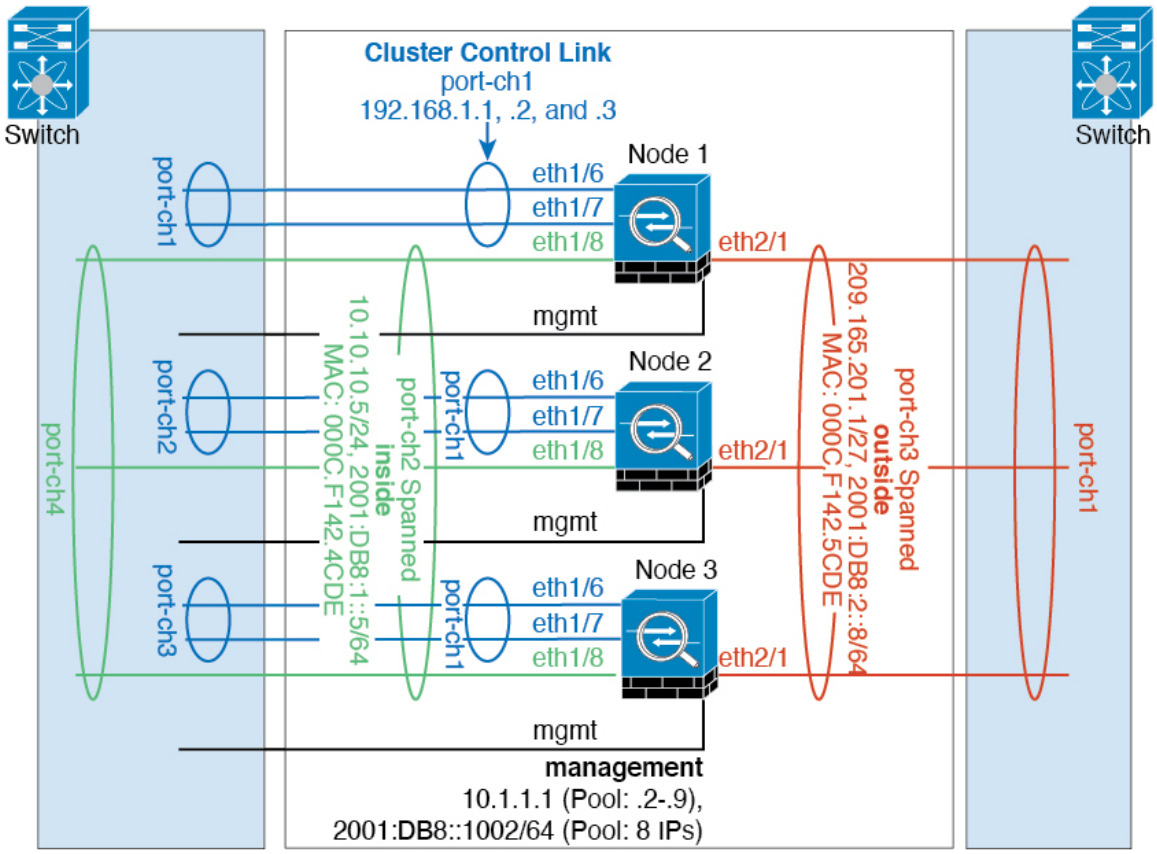
Firewall on a Stick



Data traffic from different security domains are associated with different VLANs, for example, VLAN 10 for the inside network and VLAN 20 for the outside network. Each has a single physical port connected to the external switch or router. Trunking is enabled so that all packets on the physical link are 802.1q encapsulated. The is the firewall between VLAN 10 and VLAN 20.

When using Spanned EtherChannels, all data links are grouped into one EtherChannel on the switch side. If the becomes unavailable, the switch will rebalance traffic between the remaining units.

Traffic Segregation

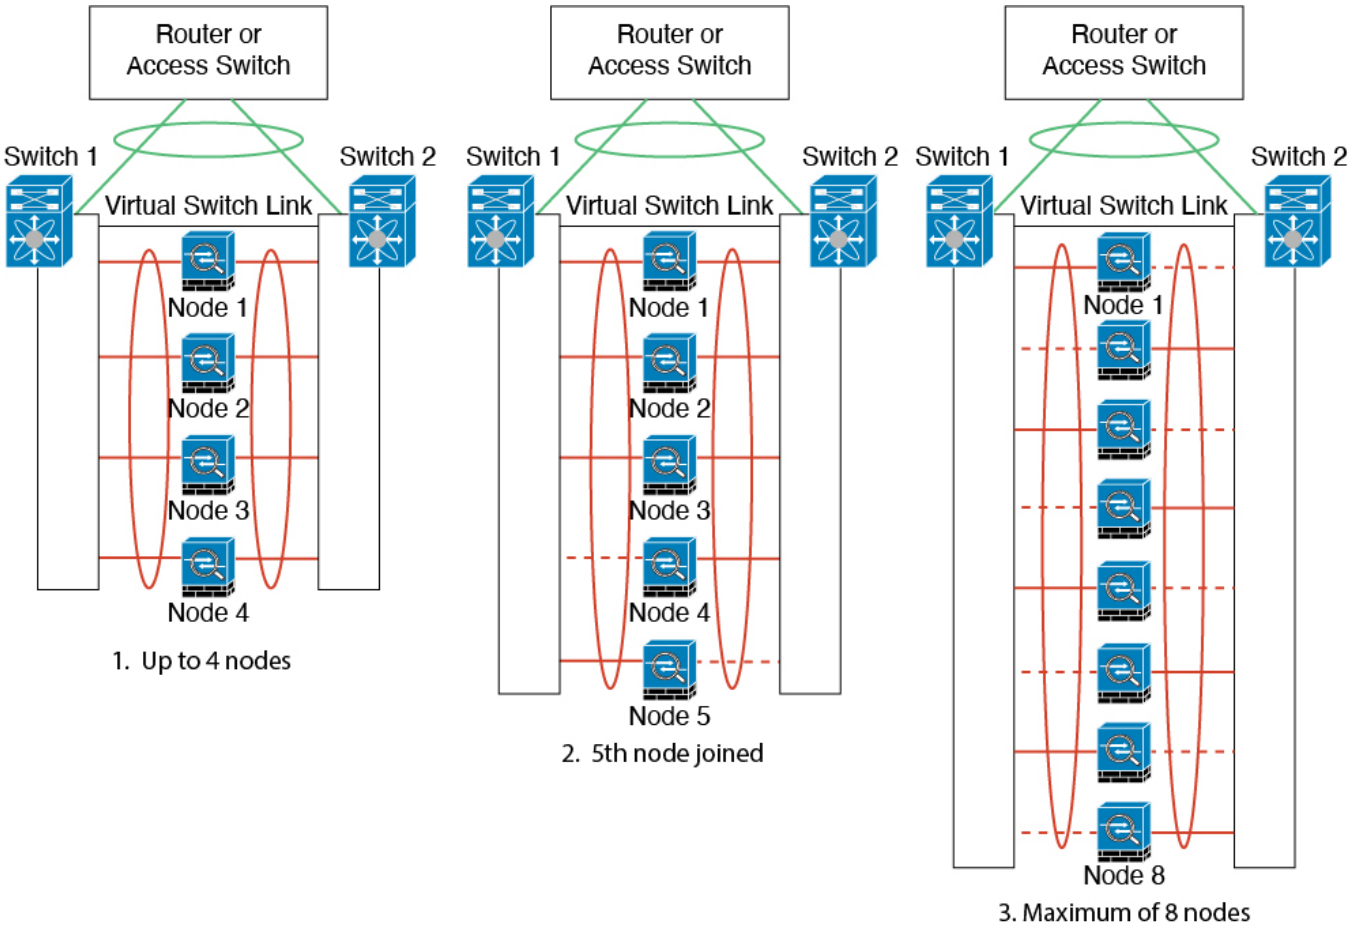


You may prefer physical separation of traffic between the inside and outside network.

As shown in the diagram above, there is one Spanned EtherChannel on the left side that connects to the inside switch, and the other on the right side to outside switch. You can also create VLAN subinterfaces on each EtherChannel if desired.

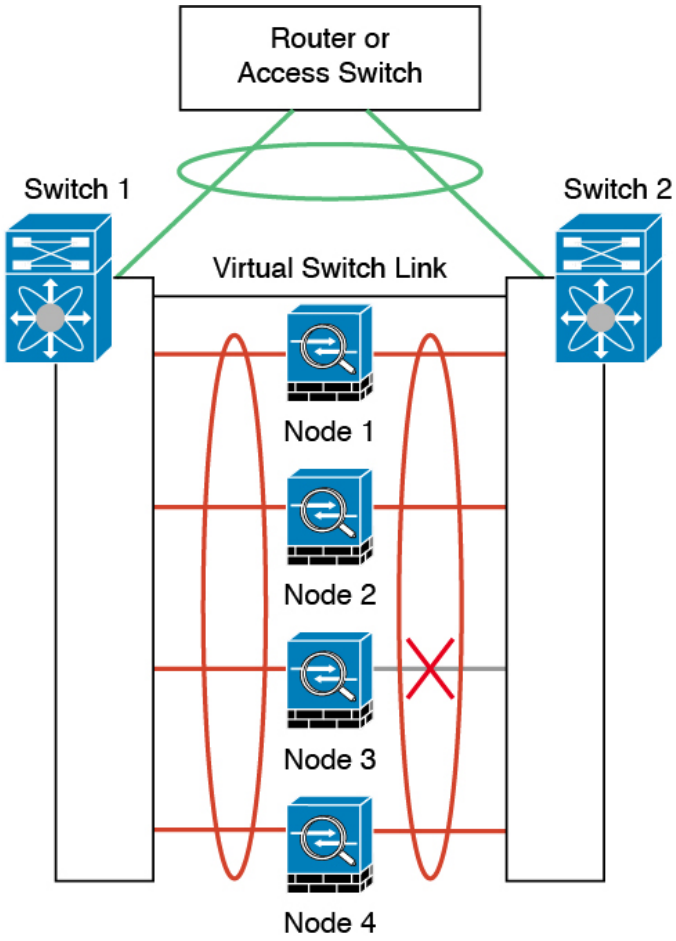
Spanned EtherChannel with Backup Links (Traditional 8 Active/8 Standby)

The maximum number of active ports in a traditional EtherChannel is limited to 8 from the switch side. If you have an 8-unit cluster, and you allocate 2 ports per unit to the EtherChannel, for a total of 16 ports total, then 8 of them have to be in standby mode. The FTD uses LACP to negotiate which links should be active or standby. If you enable multi-switch EtherChannel using VSS, vPC, StackWise, or StackWise Virtual, you can achieve inter-switch redundancy. On the FTD, all physical ports are ordered first by the slot number then by the port number. In the following figure, the lower ordered port is the “control” port (for example, Ethernet 1/1), and the other one is the “data” port (for example, Ethernet 1/2). You must guarantee symmetry in the hardware connection: all control links must terminate on one switch, and all data links must terminate on another switch if a redundant switch system is used. The following diagram shows what happens when the total number of links grows as more units join the cluster:

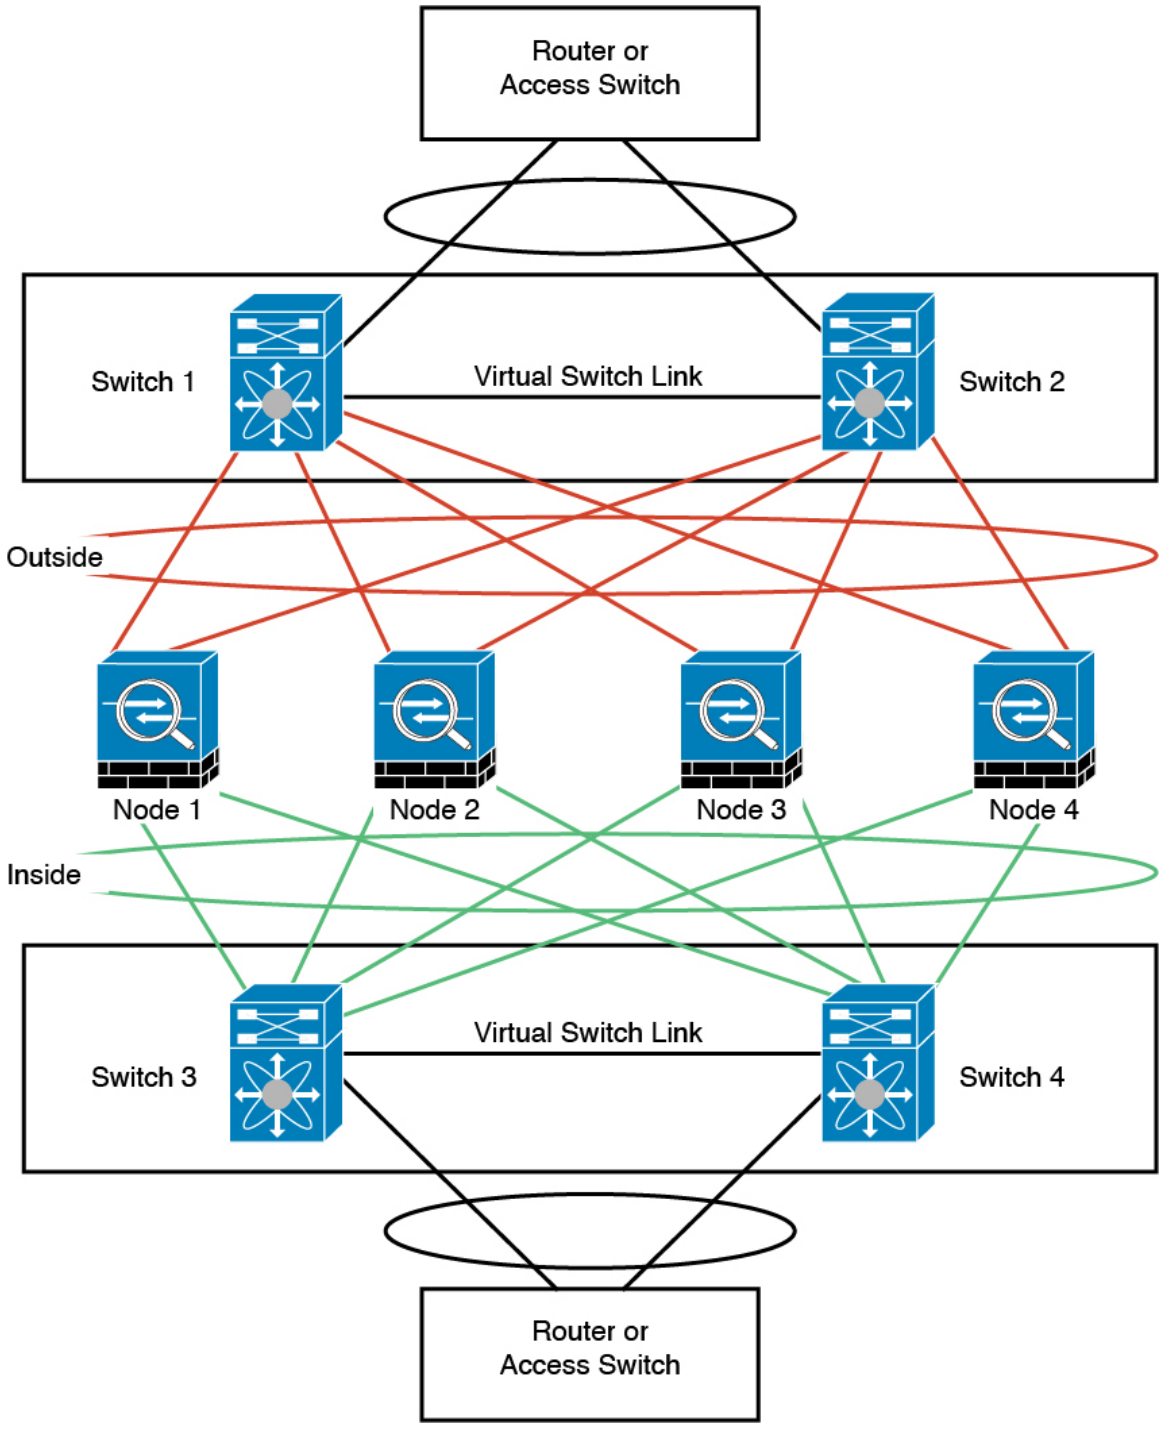


The principle is to first maximize the number of active ports in the channel, and secondly keep the number of active control ports and the number of active data ports in balance. Note that when a 5th unit joins the cluster, traffic is not balanced evenly between all units.

Link or device failure is handled with the same principle. You may end up with a less-than-perfect load balancing situation. The following figure shows a 4-unit cluster with a single link failure on one of the units.



There could be multiple EtherChannels configured in the network. The following diagram shows an EtherChannel on the inside and one on the outside. The FTD is removed from the cluster if both control and data links in one EtherChannel fail. This prevents the FTD from receiving traffic from the outside network when it has already lost connectivity to the inside network.



Reference for Clustering

This section includes more information about how clustering operates.

Firepower Threat Defense Features and Clustering

Some Firepower Threat Defense features are not supported with clustering, and some are only supported on the control unit. Other features might have caveats for proper usage.

Unsupported Features with Clustering

These features cannot be configured with clustering enabled, and the commands will be rejected.

Note To view FlexConfig features that are also not supported with clustering, for example WCCP inspection, see the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the FMC GUI.

- DHCP client, server, and proxy. DHCP relay is supported.
- High Availability
- FMC UCAPL/CC mode

Centralized Features for Clustering

The following features are only supported on the control unit, and are not scaled for the cluster.

Note Traffic for centralized features is forwarded from member units to the control unit over the cluster control link.

If you use the rebalancing feature, traffic for centralized features may be rebalanced to non-control units before the traffic is classified as a centralized feature; if this occurs, the traffic is then sent back to the control unit.

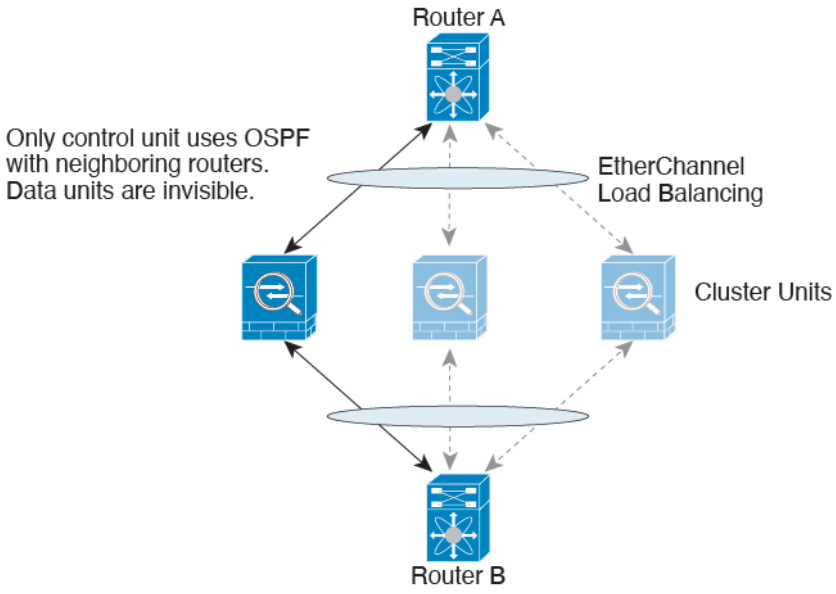
For centralized features, if the control unit fails, all connections are dropped, and you have to re-establish the connections on the new control unit.

- The following application inspections:
  - DCERPC
  - NetBIOS
  - RSH
  - SUNRPC
  - TFTP
  - XDMCP
- Dynamic routing
- Static route monitoring

Dynamic Routing and Clustering

The routing process only runs on the control unit, and routes are learned through the control unit and replicated to secondaries. If a routing packet arrives at a data unit, it is redirected to the control unit.

Figure 2: Dynamic Routing



After the data units learn the routes from the control unit, each unit makes forwarding decisions independently.

The OSPF LSA database is not synchronized from the control unit to data units. If there is a control unit switchover, the neighboring router will detect a restart; the switchover is not transparent. The OSPF process picks an IP address as its router ID. Although not required, you can assign a static router ID to ensure a consistent router ID is used across the cluster. See the OSPF Non-Stop Forwarding feature to address the interruption.

Connection Settings

Connection limits are enforced cluster-wide. Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced exactly at the limit number. Each node may overestimate or underestimate the cluster-wide counter value at any given time. However, the information will get updated over time in a load-balanced cluster.

FTP and Clustering

- If FTP data channel and control channel flows are owned by different cluster members, then the data channel owner will periodically send idle timeout updates to the control channel owner and update the idle timeout value. However, if the control flow owner is reloaded, and the control flow is re-hosted, the parent/child flow relationship will not longer be maintained; the control flow idle timeout will not be updated.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different FTDs in the cluster, because the load balancing algorithm relies on IP addresses and ports, and NAT causes inbound and outbound packets to have different IP addresses and/or ports. When a packet arrives at the FTD that is not the NAT owner, it is forwarded over the cluster control link to the owner, causing large amounts of traffic on the cluster control link. Note that the receiving node does not create a forwarding flow to the owner, because the NAT owner may not end up creating a connection for the packet depending on the results of security and policy checks.

If you still want to use NAT in clustering, then consider the following guidelines:

- NAT pool address distribution for dynamic PAT—The control node evenly pre-distributes addresses across the cluster. If a member receives a connection and they have no addresses assigned, then the connection is forwarded to the control node for PAT. If a cluster member leaves the cluster (due to failure), a backup member will get the PAT IP address, and if the backup exhausts its normal PAT IP address, it can make use of the new address. Make sure to include at least as many NAT addresses as there are nodes in the cluster, plus at least one extra address, to ensure that each node receives an address, and that a failed node can get a new address if its old address is in use by the member that took over the address. Use the **show nat pool cluster** command in the device CLI to see the address allocations.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- Dynamic NAT xlates managed by the control node—The control node maintains and replicates the xlate table to data nodes. When a data node receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control node. The data node owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- No static PAT for the following inspections—
  - FTP
  - RSH
  - SQLNET
  - TFTP
  - XDMCP
  - SIP

SIP Inspection and Clustering

A control flow can be created on any node (due to load balancing); its child data flows must reside on the same node.

SNMP and Clustering

An SNMP agent polls each individual FTD by its interface Local IP address. You cannot poll consolidated data for the cluster.

You should always use the Local address, and not the Main cluster IP address for SNMP polling. If the SNMP agent polls the Main cluster IP address, if a new control node is elected, the poll to the new control node will fail.

Syslog and Clustering

- Each node in the cluster generates its own syslog messages. You can configure logging so that each node uses either the same or a different device ID in the syslog message header field. For example, the hostname configuration is replicated and shared by all nodes in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all nodes look as if they come from a single node. If you configure logging to use the local-node name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different nodes.

TLS/SSL Connections and Clustering

The decryption states of TLS/SSL connections are not synchronized, and if the connection owner fails, then the 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.

Cisco TrustSec and Clustering

Only the control node learns security group tag (SGT) information. The control node then populates the SGT to data nodes, and data nodes can make a match decision for SGT based on the security policy.

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately 80% of the maximum combined throughput.

For example, for TCP throughput, the Firepower 9300 with 3 SM-40 modules can handle approximately 135 Gbps of real world firewall traffic when running alone. For 2 chassis, the maximum combined throughput will be approximately 80% of 270 Gbps (2 chassis x 135 Gbps): 216 Gbps.

Control Unit Election

Members of the cluster communicate over the cluster control link to elect a control unit as follows:

1. When you deploy the cluster, each unit broadcasts an election request every 3 seconds.
2. Any other units with a higher priority respond to the election request; the priority is set when you deploy the cluster and is not configurable.
3. If after 45 seconds, a unit does not receive a response from another unit with a higher priority, then it becomes the control unit.

Note If multiple units tie for the highest priority, the cluster unit name and then the serial number is used to determine the control unit.

4. If a unit later joins the cluster with a higher priority, it does not automatically become the control unit; the existing control unit always remains as the control unit unless it stops responding, at which point a new control unit is elected.
5. In a "split brain" scenario when there are temporarily multiple control units, then the unit with highest priority retains the role while the other units return to data unit roles.

Note You can manually force a unit to become the control unit. For centralized features, if you force a control unit change, then all connections are dropped, and you have to re-establish the connections on the new control unit.

High Availability Within the Cluster

Clustering provides high availability by monitoring chassis, unit, and interface health and by replicating connection states between units.

Chassis-Application Monitoring

Chassis-application health monitoring is always enabled. The Firepower 4100/9300 chassis supervisor checks the FTD application periodically (every second). If the FTD device is up and cannot communicate with the Firepower 4100/9300 chassis supervisor for 3 seconds, the FTD device generates a syslog message and leaves the cluster.

If the Firepower 4100/9300 chassis supervisor cannot communicate with the application after 45 seconds, it reloads the FTD device. If the FTD device cannot communicate with the supervisor, it removes itself from the cluster.

Unit Health Monitoring

Each unit periodically sends a broadcast keepalivekeepalive packet over the cluster control link. If the control node does not receive any keepalivekeepalive packets or other packets from a data node within the timeout period, then the control node removes the data node from the cluster. If the data nodes do not receive packets from the control node, then a new control node is elected from the remaining node.

If nodes cannot reach each other over the cluster control link because of a network failure and not because a node has actually failed, then the cluster may go into a "split brain" scenario where isolated data nodes will elect their own control nodes. For example, if a router fails between two cluster locations, then the original control node at location 1 will remove the location 2 data nodes from the cluster. Meanwhile, the nodes at location 2 will elect their own control node and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control node that has the higher priority will keep the control node's role. See Control Unit Election, on page 26 for more information.

Interface Monitoring

Each node monitors the link status of all hardware interfaces in use, and reports status changes to the control node. All physical interfaces are monitored by default (including the main EtherChannel for EtherChannel interfaces). Only named interfaces that are in an Up state can be monitored. For example, all member ports of an EtherChannel must fail before a *named* EtherChannel is removed from the cluster.

If a monitored interface fails on a particular node, but it is active on other nodes, then the node is removed from the cluster. The amount of time before the FTD device removes a node from the cluster depends on whether the node is an established member or is joining the cluster. The FTD device does not monitor interfaces for the first 90 seconds that a node joins the cluster. Interface status changes during this time will not cause the FTD device to be removed from the cluster. For an established member, the node is removed after 500 ms.

Decorator Application Monitoring

When you install a decorator application on an interface, such as the Radware DefensePro application, then both the FTD device and the decorator application must be operational to remain in the cluster. The unit does not join the cluster until both applications are operational. Once in the cluster, the unit monitors the decorator application health every 3 seconds. If the decorator application is down, the unit is removed from the cluster.

Status After Failure

When a node in the cluster fails, the connections hosted by that node are seamlessly transferred to other nodes; state information for traffic flows is shared over the control node's cluster control link.

If the control node fails, then another member of the cluster with the highest priority (lowest number) becomes the control node.

The FTD automatically tries to rejoin the cluster, depending on the failure event.



Note When the FTD becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management/Diagnostic interface can send and receive traffic.

Rejoining the Cluster

After a cluster member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The FTD automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firepower Threat Defense automatically tries to rejoin at 5 minutes, then at 10 minutes, and finally at 20 minutes. If the join is not successful after 20 minutes, then the Firepower Threat Defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- Failed unit—If the unit was removed from the cluster because of a unit health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the unit

will rejoin the cluster when it starts up again as long as the cluster control link is up. The Firepower Threat Defense application attempts to rejoin the cluster every 5 seconds.

- Failed Chassis-Application Communication—When the Firepower Threat Defense application detects that the chassis-application health has recovered, it tries to rejoin the cluster automatically.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. After you resolve the problem, you must manually rejoin the cluster by re-enabling clustering.
- Failed configuration deployment—If you deploy a new configuration from FMC, and the deployment fails on some cluster members but succeeds on others, then the units that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control unit, then the deployment is rolled back, and no members are removed.

Data Path Connection State Replication

Every connection has one owner and at least one backup owner in the cluster. The backup owner does not take over the connection in the event of a failure; instead, it stores TCP/UDP state information, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner is usually also the director.

Some traffic requires state information above the TCP or UDP layer. See the following table for clustering support or lack of support for this kind of traffic.

Table 1: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	Transparent mode only.
MAC address table	Yes	Transparent mode only.
User Identity	Yes	—
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—

How the Cluster Manages Connections

Connections can be load-balanced to multiple nodes of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the node that initially receives the connection. The owner maintains the TCP state and processes packets. A connection has only one owner. If the original owner fails, then when new nodes receive packets from the connection, the director chooses a new owner from those nodes.

- Backup owner—The node that stores TCP/UDP state information received from the owner, so that the connection can be seamlessly transferred to a new owner in case of a failure. The backup owner does not take over the connection in the event of a failure. If the owner becomes unavailable, then the first node to receive packets from the connection (based on load balancing) contacts the backup owner for the relevant state information so it can become the new owner.

  As long as the director (see below) is not the same node as the owner, then the director is also the backup owner. If the owner chooses itself as the director, then a separate backup owner is chosen.

- Director—The node that handles owner lookup requests from forwarders. When the owner receives a new connection, it chooses a director based on a hash of the source/destination IP address and ports (see below for ICMP hash details), and sends a message to the director to register the new connection. If packets arrive at any node other than the owner, the node queries the director about which node is the owner so it can forward the packets. A connection has only one director. If a director fails, the owner chooses a new director.

  As long as the director is not the same node as the owner, then the director is also the backup owner (see above). If the owner chooses itself as the director, then a separate backup owner is chosen.

  ICMP/ICMPv6 hash details:

  - For Echo packets, the source port is the ICMP identifier, and the destination port is 0.
  - For Reply packets, the source port is 0, and the destination port is the ICMP identifier.
  - For other packets, both source and destination ports are 0.

- Forwarder—A node that forwards packets to the owner. If a forwarder receives a packet for a connection it does not own, it queries the director for the owner, and then establishes a flow to the owner for any other packets it receives for this connection. The director can also be a forwarder. Note that if a forwarder receives the SYN-ACK packet, it can derive the owner directly from a SYN cookie in the packet, so it does not need to query the director. (If you disable TCP sequence randomization, the SYN cookie is not used; a query to the director is required.) For short-lived flows such as DNS and ICMP, instead of querying, the forwarder immediately sends the packet to the director, which then sends them to the owner. A connection can have multiple forwarders; the most efficient throughput is achieved by a good load-balancing method where there are no forwarders and all packets of a connection are received by the owner.



Note We do not recommend disabling TCP sequence randomization when using clustering. There is a small chance that some TCP sessions won't be established, because the SYN/ACK packet might be dropped.

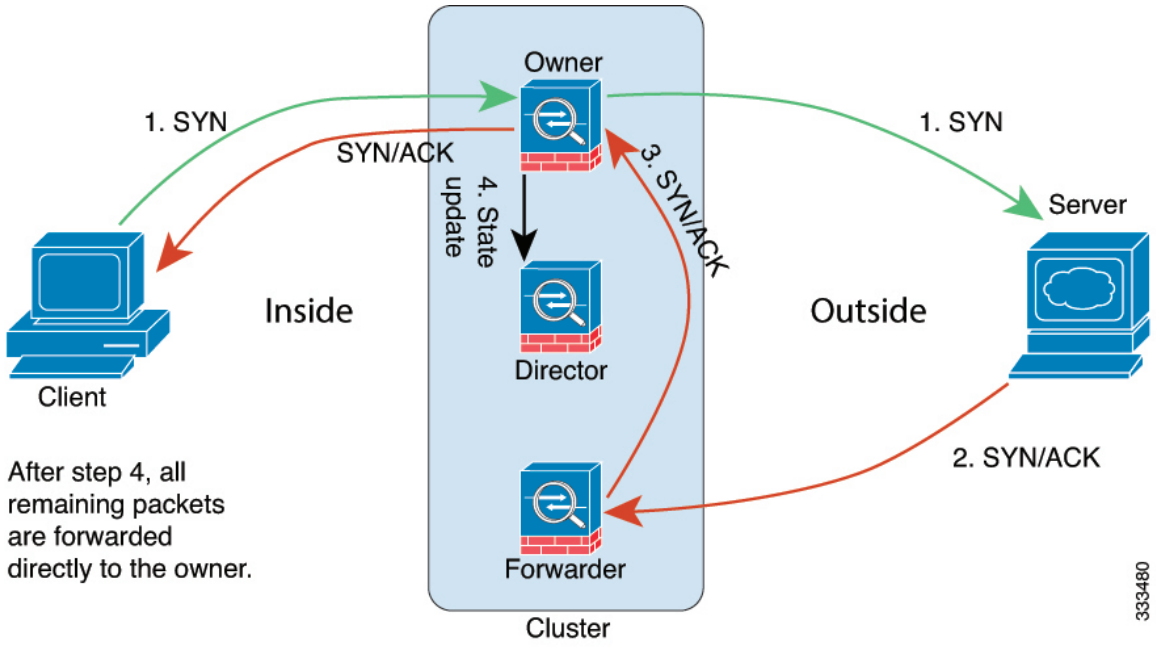
- Fragment Owner—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may be load-balanced to different cluster nodes, because only the first fragment includes the 5-tuple used in the switch load balance hash. Other fragments do not contain the source and destination ports and may be load-balanced to other cluster nodes. The fragment owner temporarily reassembles the packet so it can determine the director based on a hash of the source/destination IP address and ports. If it is a new connection, the fragment owner will register to be the connection owner. If it is an existing connection, the fragment owner forwards all fragments to the provided connection owner over the cluster control link. The connection owner will then reassemble all fragments.

New Connection Ownership

When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. If a reverse flow arrives at a different node, it is redirected back to the original node.

Sample Data Flow for TCP

The following example shows the establishment of a new connection.



1. The SYN packet originates from the client and is delivered to one FTD (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different FTD (based on the load balancing method). This FTD is the forwarder.
3. Because the forwarder does not own the connection, it decodes owner information from the SYN cookie, creates a forwarding flow to the owner, and forwards the SYN-ACK to the owner.
4. The owner sends a state update to the director, and forwards the SYN-ACK to the client.
5. The director receives the state update from the owner, creates a flow to the owner, and records the TCP state information as well as the owner. The director acts as the backup owner for the connection.
6. Any subsequent packets delivered to the forwarder will be forwarded to the owner.
7. If packets are delivered to any additional nodes, it will query the director for the owner and establish a flow.
8. Any state change for the flow results in a state update from the owner to the director.

History for Clustering

Feature	Version	Details
Intra-chassis Clustering for the Firepower 9300	6.0.1	You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster. New/Modified screens: **Devices** > **Device Management** > **Add** > **Add Cluster** **Devices** > **Device Management** > **Cluster** Supported platforms: Firepower Threat Defense on the Firepower 9300