



Clustering for the Secure Firewall 3100

Clustering lets you group multiple threat defense units 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.



Note Some features are not supported when using clustering. See [Unsupported Features with Clustering, on page 36](#).

- [About Clustering for the Secure Firewall 3100, on page 1](#)
- [Licenses for Clustering, on page 2](#)
- [Requirements and Prerequisites for Clustering, on page 3](#)
- [Guidelines for Clustering, on page 4](#)
- [Configure Clustering, on page 7](#)
- [Manage Cluster Nodes, on page 22](#)
- [Monitoring the Cluster, on page 31](#)
- [Troubleshooting the Cluster, on page 33](#)
- [Examples for Clustering, on page 34](#)
- [Reference for Clustering, on page 36](#)
- [History for Clustering, on page 48](#)

About Clustering for the Secure Firewall 3100

This section describes the clustering architecture and how it works.

How the Cluster Fits into Your Network

The cluster consists of multiple firewalls acting as a single unit. To act as a cluster, the firewalls need the following infrastructure:

- Isolated, high-speed backplane network for intra-cluster communication, known as the *cluster control link*.
- Management access to each firewall for configuration and monitoring.

When you place the cluster in your network, the upstream and downstream routers need to be able to load-balance the data coming to and from the cluster using Spanned EtherChannels. Interfaces on multiple members of the cluster are grouped into a single EtherChannel; the EtherChannel performs load balancing between units.

Control and Data Node Roles

One member of the cluster is the control node. If multiple cluster nodes come online at the same time, the control node is determined by the priority setting; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. When you first create the cluster, you specify which node you want to be the control node, and it will become the control node simply because it is the first node added to the cluster.

All nodes in the cluster share the same configuration. The node that you initially specify as the control node will overwrite the configuration on the data nodes when they join the cluster, so you only need to perform initial configuration on the control node before you form the cluster.

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

Cluster Interfaces

You can configure data interfaces as Spanned EtherChannels. See [About Cluster Interfaces, on page 7](#) for more information.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link. See [Cluster Control Link, on page 7](#) for more information.

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.

Management Network

You must manage each node using the Management interface; management from a data interface is not supported with clustering.

Licenses for Clustering

You assign feature licenses to the cluster as a whole, not to individual nodes. However, each node of the cluster consumes a separate license for each feature. The clustering feature itself does not require any licenses.

When you add the control node to the management center, you can specify the feature licenses you want to use for the cluster. Before you create the cluster, it doesn't matter which licenses are assigned to the data

nodes; the license settings for the control node are replicated to each of the data nodes. You can modify licenses for the cluster in the **Devices > Device Management > Cluster > License** area.



Note If you add the cluster before the management center is licensed (and running in Evaluation mode), then when you license the management center, 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

Model Requirements

- Secure Firewall 3100—Maximum 8 units

User Roles

- Admin
- Access Admin
- Network Admin

Hardware and Software Requirements

All units in a cluster:

- Must be the same model.
- Must include the same interfaces.
- The management center access must be from the Management interface; data interface management is not supported.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported.
- Must be in the same firewall mode, routed or transparent.
- Must be in the same domain.
- Must be in the same group.
- Must not have any deployment pending or in progress.
- The control node must not have any unsupported features configured (see [Unsupported Features with Clustering, on page 36](#)).
- Data nodes must not have any VPN configured. The control node can have site-to-site VPN configured.

Switch Requirements

- Be sure to complete the switch configuration before you configure clustering. Make sure the ports connected to the cluster control link have the correct (higher) MTU configured. By default, the cluster

control link MTU is set to 100 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail.

Guidelines for Clustering

Firewall Mode

The firewall mode must match on all units.

High Availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Switches

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.
- 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:

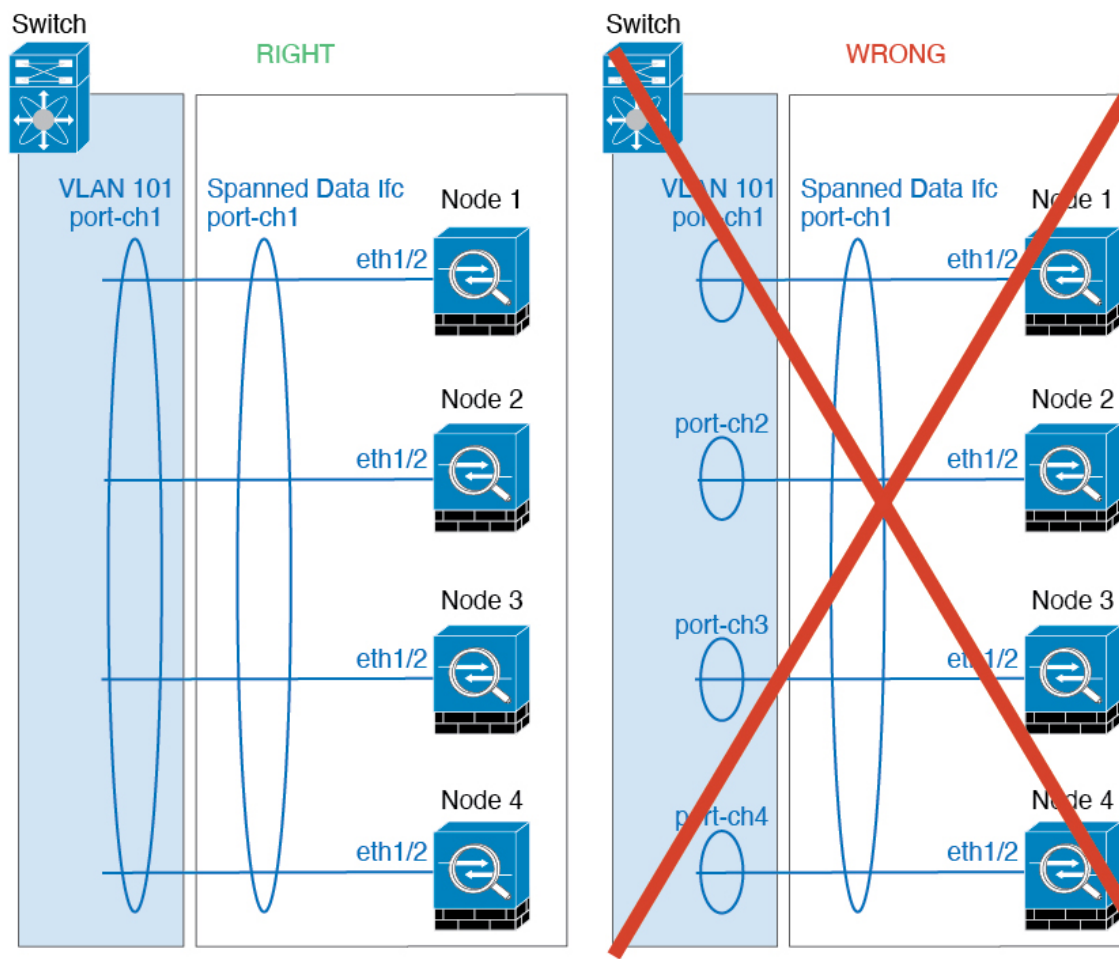
```
router(config)# port-channel id hash-distribution fixed
```

Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.

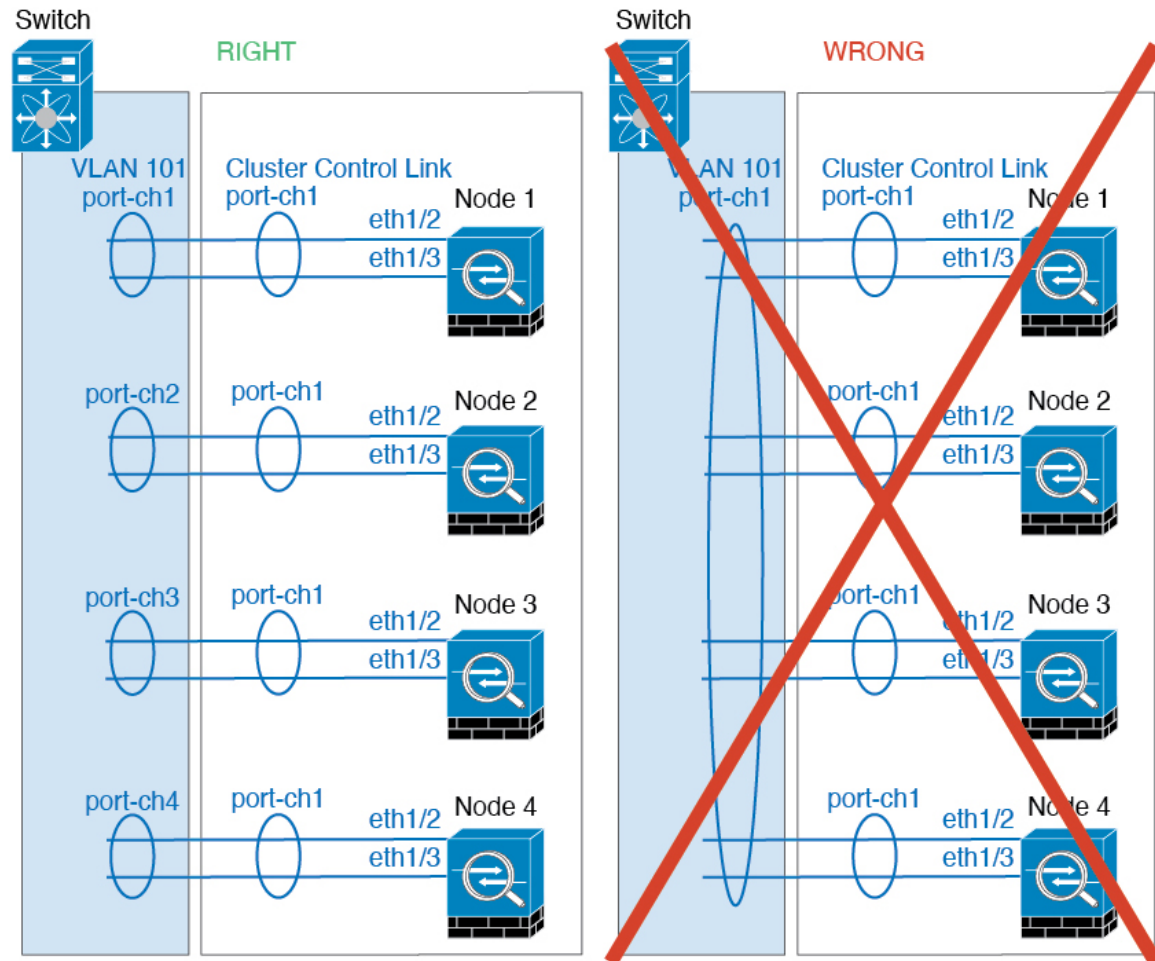
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

EtherChannels

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



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



Additional Guidelines

- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; 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, 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 ASA cluster. These messages can result in some units of the ASA cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- 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 for Clustering

- The cLACP system ID is auto-generated, and the system priority is 1 by default.
- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- The cluster auto-rejoin feature for a failed cluster control link is unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Clustering

To add a cluster to the management center, add each node to the management center as a standalone unit, configure interfaces on the unit you want to make the control node, and then form the cluster.

About Cluster Interfaces

You can configure data interfaces as Spanned EtherChannels. Each unit must also dedicate at least one hardware interface as the cluster control link.

Cluster Control Link

Each unit must dedicate at least one hardware interface as the cluster control link. We recommend using an EtherChannel for the cluster control link if available.

Cluster Control Link Traffic Overview

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

- Control node election.
- Configuration replication.
- Health monitoring.

Data traffic includes:

- State replication.
- Connection ownership queries and data packet forwarding.

Cluster Control Link Interfaces and Network

You can use any physical interface or EtherChannel for the cluster control link. You cannot use a VLAN subinterface as the cluster control link. You also cannot use the Management/Diagnostic interface.

Each cluster control link has an IP address on the same subnet. This subnet should be isolated from all other traffic, and should include only the cluster control link interfaces.



Note For a 2-member cluster, do not directly-connect the cluster control link from one node to the other node. 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. If you need to directly-connect the units (for testing purposes, for example), then you should configure and enable the cluster control link interface on both nodes before you form the cluster.

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

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.

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.



Guidelines for Maximum Throughput

To achieve maximum throughput, we recommend the following:

- Use a load-balancing hash algorithm that is “symmetric,” meaning that packets from both directions will have the same hash and will be sent to the same threat defense in the Spanned EtherChannel. We recommend using the source and destination IP address (the default) or the source and destination port as the hashing algorithm.
- Use the same type of line cards when connecting the threat defenses to the switch so that hashing algorithms applied to all packets are the same.

Load Balancing

The EtherChannel link is selected using a proprietary hash algorithm, based on source or destination IP addresses and TCP and UDP port numbers.



Note On the switch, we recommend that you use one of the following algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS or 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 ASAs in a cluster.

The number of links in the EtherChannel affects load balancing.

Symmetric load balancing is not always possible. If you configure NAT, then forward and return packets will have different IP addresses and/or ports. Return traffic will be sent to a different unit based on the hash, and the cluster will have to redirect most returning traffic to the correct unit.

EtherChannel Redundancy

The EtherChannel has built-in redundancy. It monitors the line protocol status of all links. If one link fails, traffic is re-balanced between remaining links. If all links in the EtherChannel fail on a particular unit, but other units are still active, then the unit is removed from the cluster.


Connecting to a Redundant Switch System

You can include multiple interfaces per threat defense in the Spanned EtherChannel. Multiple interfaces per threat defense are especially useful for connecting to both switches in a VSS, vPC, StackWise, or StackWise Virtual system.

Depending on your switches, you can configure up to 32 active links in the spanned EtherChannel. This feature requires both switches in the vPC to support EtherChannels with 16 active links each (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).

For switches that support 8 active links in the EtherChannel, you can configure up to 16 active links in the spanned EtherChannel when connecting to two switches in a redundant system.

The following figure shows a 16-active-link spanned EtherChannel in a 4-node cluster and an 8-node cluster.



Cable and Add Devices to the Management Center

Before configuring clustering, you need to prepare your devices. In particular, the cluster will not come up unless all nodes can communicate over the cluster control link. Therefore, before you form the cluster, the cluster control link must be ready to go.

Procedure

Step 1 Cable the cluster control link network, management network, and data networks.

Step 2 Configure the upstream and downstream equipment.

- a) For the cluster control link network, set the MTU to be at least 100 bytes higher than the data interface MTU.

By default, the data interface MTU is 1500 bytes, so the cluster control link MTU on the cluster node will be set to 1600 bytes. If you use higher MTUs on your data interfaces, increase the cluster control link MTU on connecting switches accordingly.

- b) Configure cluster control link interfaces on upstream and downstream equipment, including for an optional EtherChannel.

See [Cluster Control Link Interfaces and Network](#), on page 8 for cluster control link requirements.

- c) Configure data interfaces on upstream and downstream equipment, including Spanned EtherChannels.

See [About Cluster Interfaces](#), on page 7 for information about how to cable Spanned EtherChannels.

Step 3 Add each node to the management center as a standalone device in the same domain and group.

See [Add a Device to the Management Center](#). You can create a cluster with a single device, and then add more nodes later. The initial settings (licensing, access control policy) that you set when you add a device will be inherited by all cluster nodes from the control node. You will choose the control node when forming the cluster.

Step 4 Enable the cluster control link on the device you want to be the control node.

When you add the other nodes, they will inherit the cluster control link configuration.

Note Do not configure the name or IP addressing for the cluster control link. The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU when you form the cluster, so you do not need to set it now.

- a) On the device you want to be the control node, choose **Devices > Device Management**, and click **Edit** (✎).
- b) Click **Interfaces**.
- c) Enable the interface. If you are going to use an EtherChannel for the cluster control link, enable all members. See [Enable the Physical Interface and Configure Ethernet Settings](#).

Figure 1: Enable the Cluster Control Link Interface

The screenshot shows the 'Edit Physical Interface' configuration page. At the top, there are four tabs: 'General', 'IPv4', 'IPv6', and 'Path Monitoring'. The 'General' tab is active. Below the tabs, there is a 'Name:' label followed by an empty text input field. Below the input field, there is a checkbox labeled 'Enabled' which is checked. A red rectangular box highlights the 'Enabled' checkbox.

- d) (Optional) Add an EtherChannel. See [Configure an EtherChannel](#).

We recommend using the On mode for cluster control link member interfaces to reduce unnecessary traffic on the cluster control link (Active mode is the default). The cluster control link does not need the overhead of LACP traffic because it is an isolated, stable network. **Note:** We recommend setting *data* EtherChannels to Active mode.

- e) Click **Save** and then **Deploy** to deploy the interface changes to the control node.

Create a Cluster

Form a cluster from one or more devices in the management center.

Procedure

Step 1 Choose **Devices > Device Management**, and then choose **Add > Add Cluster**.

The **Add Cluster Wizard** appears.

Figure 2: Add Cluster Wizard

Add Cluster Wizard ×

1 Configuration — 2 Summary

▲ Create a cluster for supported models. Note: For the Firepower 4100/9300, use the Add Device option.

Cluster Name*
ftdcluster

Cluster Key

Control Node
You can form the cluster with just the control node to reduce formation time.

Node*
172.16.0.50

Cluster Control Link Network*
10.10.10.0 / 24 (254 addresses)

Cluster Control Link*
Ethernet1/7

Cluster Control Link IPv4 Address*
10.10.10.1

Priority*
1

Site ID
0

Data Nodes (Optional)
Data node hardware needs to match the control node hardware.

Node*
172.16.0.51

Cluster Control Link IPv4 Address*
10.10.10.2

Priority*
2

Site ID
0 [Remove](#)

[Add a data node](#)

Step 2 Specify a **Cluster Name** and an authentication **Cluster Key** for control traffic.

- **Cluster Name**—An ASCII string from 1 to 38 characters.
- **Cluster Key**—An ASCII string from 1 to 63 characters. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.


Step 3 For the **Control Node**, set the following:

- **Node**—Choose the device that you want to be the control node initially. When the management center forms the cluster, it will add this node to the cluster first so it will be the control node.

Note

If you see an **Error** (❗) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation. For example:

Figure 3: Configuration Issues



To resolve the above issues, remove the unsupported VPN license and deploy pending configuration changes to the device.

- **Cluster Control Link Network**—Specify an IPv4 subnet; IPv6 is not supported for this interface. Specify a **24**, **25**, **26**, or **27** subnet.
- **Cluster Control Link**—Choose the physical interface or EtherChannel you want to use for the cluster control link.

Note

The MTU of the cluster control link interface is automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes. If you want to increase the MTU, see the **Devices > Device Management > Interfaces** page.

Make sure you configure switches connected to the cluster control link to the correct (higher) MTU; otherwise, cluster formation will fail.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the first address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority. Even if you set the priority to be lower than other nodes, this node will still be the control node when the cluster is first formed.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 4

For **Data Nodes (Optional)**, click **Add a data node** to add a node to the cluster.

You can form the cluster with only the control node for faster cluster formation, or you can add all nodes now. Set the following for each data node:

- **Node**—Choose the device that you want to add.

Note

If you see an **Error** (❗) icon next to the node name, click the icon to view configuration issues. You must cancel cluster formation, resolve the issues, and then return to cluster formation.

- **Cluster Control Link IPv4 Address**—This field will be auto-populated with the next address on the cluster control link network. You can edit the host address if desired.
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.
- **Site ID**—(FlexConfig feature) Enter the site ID for this node between 1 and 8. A value of 0 disables inter-site clustering. Additional inter-site cluster customizations to enhance redundancy and stability, such as director localization, site redundancy, and cluster flow mobility, are only configurable using the FlexConfig feature.

Step 5 Click **Continue**. Review the **Summary**, and then click **Save**.

The cluster name shows on the **Devices > Device Management** page; expand the cluster to see the cluster nodes.

Figure 4: Cluster Management

Node ID	Node Type	Status	Model/Version	Policy
172.16.0.50	(Control)	Snort 3	Firewall 3120 Threat Defense 7.1.0	Default AC Policy
172.16.0.51		Snort 3	Firewall 3120 Threat Defense 7.1.0	Default AC Policy

A node that is currently registering shows the loading icon.

Figure 5: Node Registration

Node ID	Node Type	Status	Model/Version
172.16.0.50	(Control)	Snort 3	Transparent
172.16.0.51		Snort 3	Transparent

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The management center updates the Cluster Registration task as each node registers.

Task ID	Task Description	Duration
10.10.1.12	Deployment to device successful.	1m 54s
10.10.1.13	Deployment to device successful.	1m 3s
TD_Cluster	Deployment to device successful.	35s

Step 6 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

Most configuration can be applied to the cluster as a whole, and not nodes in the cluster. For example, you can change the display name per node, but you can only configure interfaces for the whole cluster.

Step 7 On the **Devices > Device Management > Cluster** screen, you see **General** and other settings for the cluster.

Figure 6: Cluster Settings

The screenshot shows the 'ftdcluster' configuration page for Cisco Secure Firewall 3100 Threat Defense. The page is divided into several sections:

- General:** Name: ftdcluster, Transfer Packets: No, Status: (green dot), Control: 172.16.0.50, Cluster Live Status: View.
- License:** Base: Yes, Export-Controlled Features: No, Malware: Yes, Threat: Yes, URL Filtering: Yes, AnyConnect Apex: N/A, AnyConnect Plus: N/A, AnyConnect VPN Only: N/A.
- Security Engine:** Intrusion Prevention Engine: Snort 3.0, with a 'Revert to Snort 2' button.
- Applied Policies:** Access Control Policy: Default AC Policy, Prefilter Policy: Default Prefilter Policy, SSL Policy: (blank), DNS Policy: Default DNS Policy, Identity Policy: (blank), NAT Policy: (blank), Platform Settings Policy: (blank), NGFW QoS Policy: (blank), FlexConfig Policy: (blank).
- Health:** Policy: Initial_Health_Policy, 2021-10-30 01:21:29.
- Advanced Settings:** Application Bypass: No, Bypass Threshold: 3000 ms, Object Group Search: Disabled, Interface Object Optimization: Disabled.

See the following cluster-specific items in the **General** area:

- **General > Name**—Change the cluster display name by clicking the **Edit** (✎).

This close-up screenshot shows the 'General' settings section. At the top right of the section is an 'Edit' icon (✎). The settings are:

- Name:** ftdcluster
- Transfer Packets:** No
- Status:** (orange triangle icon)
- Control:** 172.16.0.50
- Cluster Live Status:** View

Then set the **Name** field.

General ?

Name:

Transfer Packets:

Compliance Mode:

 TLS Crypto Acceleration:

Force Deploy: →

- **General > Cluster Live Status**—Click the **View** link to open the **Cluster Status** dialog box.

General ✎	
Name:	ftdcluster
Transfer Packets:	No
Status:	▲
Control:	172.16.0.50
Cluster Live Status:	View

The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**. You can also ping the cluster control link from a node. See [Perform a Ping on the Cluster Control Link, on page 33](#).

Cluster Status

Overall Status: Cluster has all nodes in sync

Nodes details (2)

Refresh

Reconcile All

Enter node name

	Status	Device Name	Unit Name	Chassis URL	
>	In Sync.	172.16.0.50 Control	172.16.0.50	N/A	⋮
>	In Sync.	172.16.0.51	172.16.0.51	N/A	⋮

Dated: 11:52:26 | 20 Dec 2021

Close


Step 8

On the **Devices > Device Management > Devices**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.


Figure 7: Device Settings

Figure 8: Choose Node



- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).

General	
Name:	10.89.5.21
Transfer Packets:	Yes
Mode:	routed
Compliance Mode:	None
TLS Crypto Acceleration:	Enabled

Then set the **Name** field.

General	
Name:	<input type="text" value="10.10.1.13"/>
Transfer Packets:	<input checked="" type="checkbox"/>
Mode:	routed
Compliance Mode:	None
Performance Profile:	Default
TLS Crypto Acceleration:	Disabled
Force Deploy:	→
<input type="button" value="Cancel"/> <input type="button" value="Save"/>	

- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the management center so that it can reach the device on the network. First disable the connection, edit the **Host** address in the **Management** area, then re-enable the connection.

Management	
Host:	10.89.5.20
Status:	

Configure Interfaces

Configure data interfaces as Spanned EtherChannels. You can also configure the Diagnostic interface, which is the only interface that can run as an individual interface.

Procedure

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

Step 2 Click **Interfaces**.

Step 3 Configure Spanned EtherChannel data interfaces.

- a) Configure one or more EtherChannels. See [Configure an EtherChannel](#).

You can include one or more member interfaces in the EtherChannel. Because this EtherChannel is spanned across all of the nodes, you only need one member interface per node; however, for greater throughput and redundancy, multiple members are recommended.

- b) (Optional) Configure VLAN subinterfaces on the EtherChannel. The rest of this procedure applies to the subinterfaces. See [Add a Subinterface](#).
- c) Click **Edit** (✎) for the EtherChannel interface.
- d) Configure the name, IP address, and other parameters according to [Configure Routed Mode Interfaces](#) or, for transparent mode, [Configure Bridge Group Interfaces](#).

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. By default, the cluster control link MTU is 1600 bytes. If you want to increase the MTU of data interfaces, first increase the cluster control link MTU.

- e) 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.

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

Step 4 (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) Choose **Objects > Object Management > Address Pools** to add an IPv4 and/or IPv6 address pool. See [Address Pools](#).

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.

- b) On **Devices > Device Management > Interfaces**, click **Edit** (✎) for the Diagnostic interface.
- c) 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.
- d) From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

- e) On **IPv6 > Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.
- f) Configure other interface settings as normal.

Step 5 Click **Save**.

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

Manage Cluster Nodes

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


Add a New Cluster Node

You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 Choose **Devices > Device Management**, click **More** (⋮) for the cluster, and choose **Add Nodes**.

Figure 9: Add Nodes



The **Manage Cluster Wizard** appears.

Step 2 From the **Node** menu, choose a device, adjust the IP address, priority, and Site ID if desired.

Figure 10: Manage Cluster Wizard

Manage Cluster Wizard

1 Configuration — 2 Summary

Cluster Name*
ftdcluster

Cluster Key

Control Node
You can form the cluster with just the control node to reduce formation time.

Node*
172.16.0.50

Cluster Control Link Network*
10.10.10.0 / 24 (254 addresses)

Cluster Control Link*
Ethernet1/7

Cluster Control Link IPv4 Address*
10.10.10.1

Priority*
1

Site ID
0

Data Nodes (Optional)
Data node hardware needs to match the control node hardware.

Node*
172.16.0.51

Cluster Control Link IPv4 Address*
10.10.10.2

Priority*
2

Site ID
0

Node*
Type device name

Cluster Control Link IPv4 Address*
10.10.10.3

Priority*
3

Site ID
0

[Remove](#)

[Add a data node](#)

Step 3 To add additional nodes, click **Add a data node**.

Step 4 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.


Figure 11: Node Registration

ftdcluster (2)
Cluster

172.16.0.50(Control) Snort 3
172.16.0.50 - Transparent

172.16.0.51 Snort 3
172.16.0.51 - Transparent

You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**.




Break a Node

You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

- Step 1** Choose **Devices > Device Management**, click **More** (⋮) for the node you want to break, and choose **Break Node**.

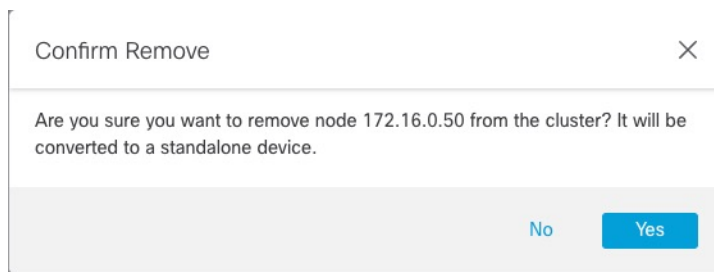
Figure 12: Break a Node



You can optionally break one or more nodes from the cluster More menu by choosing **Break Nodes**.

- Step 2** You are prompted to confirm the break; click **Yes**.

Figure 13: Confirm Break



You can monitor the cluster node break by clicking the **Notifications** icon and choosing **Tasks**.


Break the Cluster

You can break the cluster and convert all nodes to standalone devices. The control node retains the interface and security policy configuration, while data nodes have their configuration erased.

Procedure

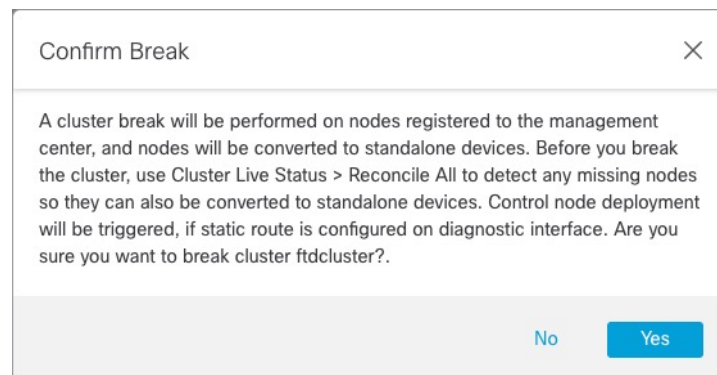
- Step 1** Make sure all cluster nodes are being managed by the management center by reconciling nodes. See [Reconcile Cluster Nodes](#), on page 29.
- Step 2** Choose **Devices > Device Management**, click **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 14: Break Cluster



- Step 3** You are prompted to break the cluster; click **Yes**.

Figure 15: Confirm Break



You can monitor the cluster break by clicking the **Notifications** icon and choosing **Tasks**.


Disable Clustering

You may want to deactivate a node in preparation for deleting the node, or temporarily for maintenance. This procedure is meant to temporarily deactivate a node; the node will still appear in the management center device list. When a node becomes inactive, all data interfaces are shut down.

Procedure

- Step 1** For the unit you want to disable, choose **Devices > Device Management**, click **More** (⋮), and choose **Disable Node Clustering**.

Figure 16: Disable Clustering



If you disable clustering on the control node, one of the data nodes will become the new control node. Note that for centralized features, if you force a control node change, then all connections are dropped, and you have to re-establish the connections on the new control node. You cannot disable clustering on the control node if it is the only node in the cluster.

- Step 2** Confirm that you want to disable clustering on the node.
The node will show **(Disabled)** next to its name in the **Devices > Device Management** list.
- Step 3** To reenabling clustering, see [Rejoin the Cluster, on page 26](#).
-

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually disabled clustering, you must manually rejoin the cluster. Make sure the failure is resolved before you try to rejoin the cluster. See [Rejoining the Cluster, on page 43](#) for more information about why a node can be removed from a cluster.

Procedure

- Step 1** For the unit you want to reactivate, choose **Devices > Device Management**, click **More** (⋮), and choose **Enable Node Clustering**.
- Step 2** Confirm that you want to enable clustering on the unit.
-

Change the Control Node



Caution The best method to change the control node is to disable clustering on the control node, wait for a new control election, and then re-enable clustering. If you must specify the *exact* unit you want to become the control node, use the procedure in this section. Note that for centralized features, if you force a control node change using either method, then all connections are dropped, and you have to re-establish the connections on the new control node.

To change the control node, perform the following steps.

Procedure

Step 1 Open the **Cluster Status** dialog box by choosing **Devices > Device Management > More (⋮) > Cluster Live Status**.

Figure 17: Cluster Status

Cluster Status ?

Overall Status: Cluster has all nodes in sync

Nodes details (2) Refresh Reconcile All 🔍 Enter node name

	Status	Device Name	Unit Name	Chassis URL	
>	In Sync.	172.16.0.50 Control	172.16.0.50	N/A	⋮
>	In Sync.	172.16.0.51	172.16.0.51	N/A	⋮

Dated: 11:52:26 | 20 Dec 2021 Close

Step 2 For the unit you want to become the control unit, choose **More (⋮) > Change Role to Control**.

Step 3 You are prompted to confirm the role change. Check the checkbox, and click **OK**.

Edit the Cluster Configuration


You can edit the cluster configuration. If you change the cluster key, cluster control link interface, or cluster control link network, the cluster will be broken and reformed automatically. Until the cluster is reformed, you

may experience traffic disruption. If you change the cluster control link IP address for a node, node priority, or site ID, only the affected nodes are broken and readded to the cluster.

Procedure

Step 1 Choose **Devices > Device Management**, click **More** (⋮) for the cluster, and choose **Edit Configuration**.

Figure 18: Edit Configuration



The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 19: Manage Cluster Wizard

The screenshot shows the 'Manage Cluster Wizard' Configuration step. It includes a warning message about editing the bootstrap configuration. The form contains the following fields and sections:

- Cluster Name***: ftd_cluster
- Cluster Key**: Two masked input fields (highlighted in red).
- Control Node**: You can form the cluster with just the control node to reduce formation time.
 - Node***: 172.16.0.51
 - Cluster Control Link Network***: 10.10.10.0 / 24 (254 addresses) (highlighted in red).
 - Cluster Control Link IPv4 Address***: 10.10.10.2
 - Priority***: 2
 - Site ID**: 0
- Data Nodes (Optional)**: Data node hardware needs to match the control node hardware.
 - Node***: 172.16.0.50
 - Cluster Control Link IPv4 Address***: 10.10.10.1
 - Priority***: 1
 - Site ID**: 0

The 'Cluster Control Link IPv4 Address*', 'Priority*', and 'Site ID' fields for both control and data nodes are highlighted in green.

If the cluster control link is an EtherChannel, you can edit the interface membership and LACP configuration by clicking **Edit** (✎) next to the interface drop-down menu.

Step 3 Click **Continue**. Review the **Summary**, and then click **Save**


Reconcile Cluster Nodes

If a cluster node fails to register, you can reconcile the cluster membership from the device to the management center. For example, a data node might fail to register if the management center is occupied with certain processes, or if there is a network issue.

Procedure

- Step 1** Choose **Devices > Device Management > More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 20: Cluster Live Status



- Step 2** Click **Reconcile All**.

Figure 21: Reconcile All

The screenshot shows the 'Cluster Status' dialog box. At the top, it says 'Cluster Status' with a help icon. Below that, the overall status is 'Cluster has all nodes in sync'. Under 'Nodes details (2)', there are two buttons: 'Refresh' and 'Reconcile All' (circled in red). To the right of the buttons is a search input field labeled 'Enter node name'. Below the buttons is a table with the following data:

	Status	Device Name	Unit Name	Chassis URL	
>	In Sync.	172.16.0.50 Control	172.16.0.50	N/A	⋮
>	In Sync.	172.16.0.51	172.16.0.51	N/A	⋮

At the bottom of the dialog, it says 'Dated: 11:52:26 | 20 Dec 2021' and a 'Close' button.

For more information about the cluster status, see [Monitoring the Cluster, on page 31](#).

Delete (Unregister) the Cluster or Nodes and Register to a New Management Center

You can unregister the cluster from the management center, which keeps the cluster intact. You might want to unregister the cluster if you want to add the cluster to a new management center.

You can also unregister a node from the management center without breaking the node from the cluster. Although the node is not visible in the management center, it is still part of the cluster, and it will continue to pass traffic and could even become the control node. You cannot unregister the current control node. You might want to unregister the node if it is no longer reachable from the management center, but you still want to keep it as part of the cluster while you troubleshoot management connectivity.

Unregistering a cluster:

- Severs all communication between the management center and the cluster.
- Removes the cluster from the **Device Management** page.
- Returns the cluster to local time management if the cluster's platform settings policy is configured to receive time from the management center using NTP.
- Leaves the configuration intact, so the cluster continues to process traffic.

Policies, such as NAT and VPN, ACLs, and the interface configurations remain intact.

Registering the cluster again to the same or a different management center causes the configuration to be removed, so the cluster will stop processing traffic at that point; the cluster configuration remains intact so you can add the cluster as a whole. You can choose an access control policy at registration, but you will have to re-apply other policies after registration and then deploy the configuration before it will process traffic again.


Before you begin

This procedure requires CLI access to one of the nodes.

Procedure

Step 1 Choose **Devices > Device Management**, click **More** (⋮) for the cluster or node, and choose **Delete**.

Figure 22: Delete Cluster or Node



- Step 2** You are prompted to delete the cluster or node; click **Yes**.
- Step 3** You can register the cluster to a new (or the same) management center by adding one of the cluster members as a new device.
- You only need to add one of the cluster nodes as a device, and the rest of the cluster nodes will be discovered.
- Connect to one cluster node's CLI, and identify the new management center using the **configure manager add** command. See [Modify Threat Defense Management Interfaces at the CLI](#).
 - Choose **Devices > Device Management**, and then click **Add > Device**.
- Step 4** To re-add an unregistered node, see [Reconcile Cluster Nodes, on page 29](#).

Monitoring the Cluster

You can monitor the cluster in the management center and at the threat defense CLI.

- Cluster Status** dialog box, which is available from the **Devices > Device Management > More** (⋮) icon or from the **Devices > Device Management > Cluster page > General area > Cluster Live Status** link.

Figure 23: Cluster Status

The screenshot shows the 'Cluster Status' dialog box. At the top, it says 'Cluster Status' with a help icon. Below that, the 'Overall Status' is 'Cluster has all nodes in sync'. There are two buttons: 'Refresh' and 'Reconcile All', and a search input field with the placeholder 'Enter node name'. Under 'Nodes details (2)', there is a table with the following data:

	Status	Device Name	Unit Name	Chassis URL	
>	In Sync.	172.16.0.50 Control	172.16.0.50	N/A	⋮
>	In Sync.	172.16.0.51	172.16.0.51	N/A	⋮

At the bottom, it shows 'Dated: 11:52:26 | 20 Dec 2021' and a 'Close' button.

The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes the following states:

- In Sync.**—The node is registered with the management center.
- Pending Registration**—The node is part of the cluster, but has not yet registered with the management center. If a node fails to register, you can retry registration by clicking **Reconcile All**.

- Clustering is disabled—The node is registered with the management center, but is an inactive member of the cluster. The clustering configuration remains intact if you intend to later re-enable it, or you can delete the node from the cluster.
- Joining cluster...—The node is joining the cluster on the chassis, but has not completed joining. After it joins, it will register with the management center.

For each node, you can view the **Summary** or the **History**.

Figure 24: Node Summary

Status	Device Name	Unit Name	Chassis URL
In Sync.	172.16.0.50 Control	172.16.0.50	N/A

Summary		History	
ID:	0	CCL IP:	10.10.10.1
Site ID:	N/A	CCL MAC:	6c13.d509.4d9a
Serial No:	FJZ2512139M	Module:	N/A
Last join:	05:41:26 UTC Dec 17 2021	Resource:	N/A
Last leave:	N/A		

Figure 25: Node History

Status	Device Name	Unit Name	Chassis URL
In Sync.	172.16.0.50 Control	172.16.0.50	N/A

Summary		History	
Timestamp	From State	To State	Event
05:56:31 UTC Dec 17 2021	MASTER	MASTER	Event: Cluster new slave enrollment hold for app 1 is relea...
05:56:31 UTC Dec 17 2021	MASTER	MASTER	Event: Cluster new slave enrollment hold for app 1 is relea...
05:56:29 UTC Dec 17 2021	MASTER	MASTER	Event: Cluster new slave enrollment is on hold for app 1 fo...
05:56:29 UTC Dec 17 2021	MASTER	MASTER	Event: Cluster new slave enrollment is on hold for app 1 fo...

- **System** (⚙) > **Tasks** page.

The **Tasks** page shows updates of the Cluster Registration task as each node registers.

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

When you expand the cluster on the devices listing page, you can see all member nodes, including the control node shown with its role next to the IP address. For nodes that are still registering, you can see the loading icon.

- **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.

Troubleshooting the Cluster

You can use the **CCL Ping** tool to make sure the cluster control link is operating correctly.

Perform a Ping on the Cluster Control Link

You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs.

Procedure

- Step 1** Choose **Devices > Device Management**, click the **More** (⋮) icon next to the cluster, and choose **> Cluster Live Status**.

Figure 26: Cluster Status

Cluster Status

Overall Status: Cluster has all nodes in sync

Nodes details (2) Refresh Reconcile All

	Status	Device Name	Unit Name	Chassis URL	
>	In Sync.	172.16.0.50 Control	172.16.0.50	N/A	⋮
>	In Sync.	172.16.0.51	172.16.0.51	N/A	⋮

Dated: 11:52:26 | 20 Dec 2021 Close

- Step 2** Expand one of the nodes, and click **CCL Ping**.

Figure 27: CCL Ping

Cluster Status ?

Overall Status: ❌ Clustering is disabled for 1 node(s)

Nodes details (3) Refresh Reconcile All

Status	Device Name	Unit Name	Chassis URL	
In Sync.	10.10.43.21	Control	10.10.43.21	N/A
<div style="display: flex; justify-content: space-between; border-bottom: 1px solid #ccc; padding: 5px 0;"> Summary History CCL Ping </div> <pre> ping 10.10.3.2 size 1654 Sending 5, 1654-byte ICMP Echos to 10.10.3.2, timeout is 2 seconds: ????? Success rate is 0 percent (0/5) ----- ping 10.10.3.1 size 1654 </pre>				
> Clustering is disabled	10.10.43.22		10.10.43.22	N/A


Dated: 18:38:41 | 01 Mar 2023 Close

The node sends a ping on the cluster control link to every other node using a packet size that matches the maximum MTU.

Examples for Clustering

These examples include examples for typical deployments.

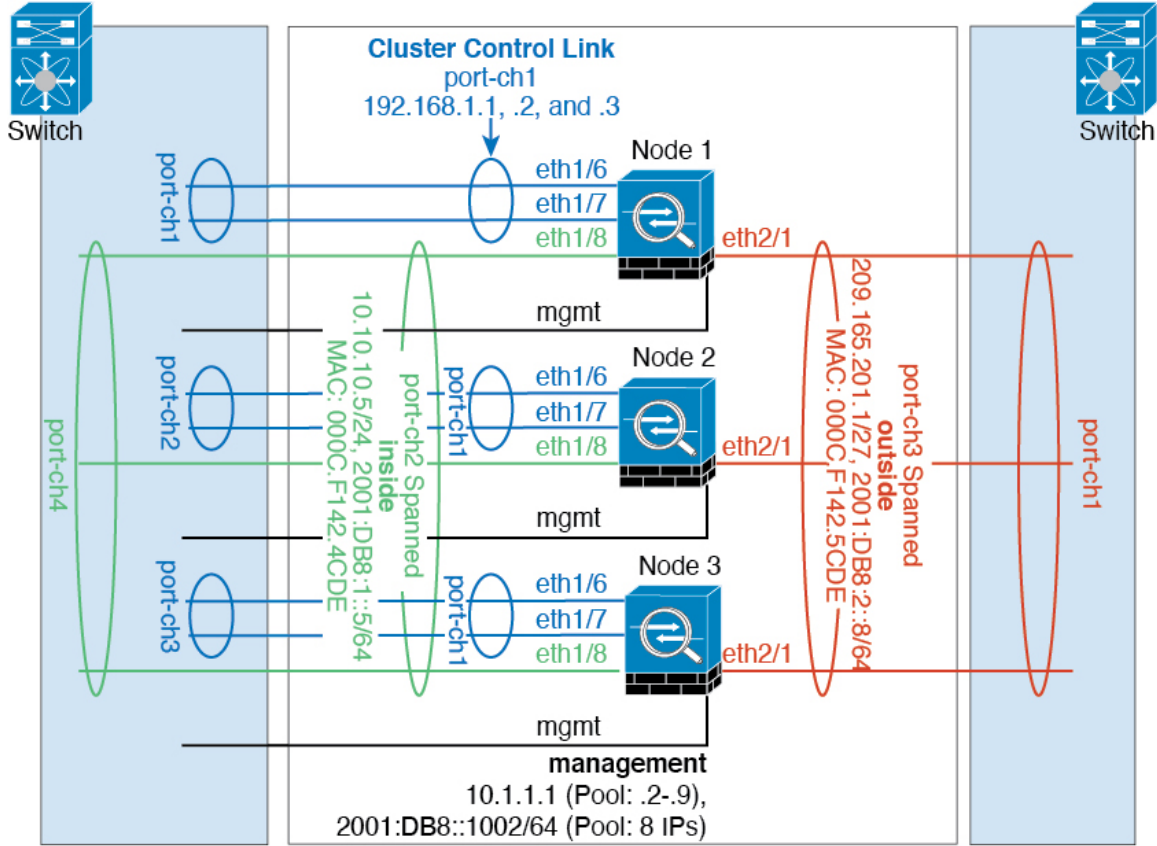
Firewall on a Stick



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. This 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.

Reference for Clustering

This section includes more information about how clustering operates.

Threat Defense Features and Clustering

Some 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 management center GUI. See [FlexConfig Policies](#).

- Remote access VPN (SSL VPN and IPsec VPN)
- DHCP client, server, and proxy. DHCP relay is supported.
- Virtual Tunnel Interfaces (VTIs)
- High Availability
- Integrated Routing and Bridging
- Management Center UCAPL/CC mode

Centralized Features for Clustering

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



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

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

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



Note To view FlexConfig features that are also centralized with clustering, for example RADIUS inspection, see the [ASA general operations configuration guide](#). FlexConfig lets you configure many ASA features that are not present in the management center GUI. See [FlexConfig Policies](#).

- The following application inspections:
 - DCERPC
 - ESMTP
 - NetBIOS
 - PPTP
 - RSH
 - SQLNET
 - SUNRPC

- TFTP
- XDMCP
- Static route monitoring
- Site-to-site VPN
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

Connection Settings and Clustering

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.

Multicast Routing in Individual Interface Mode

In Individual interface mode, units do not act independently with multicast. All data and routing packets are processed and forwarded by the control unit, thus avoiding packet replication.

NAT and Clustering

NAT can affect the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different threat defenses 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 threat defense 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:

- PAT with Port Block Allocation—See the following guidelines for this feature:
 - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each node individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 nodes, then it can get allocated 3 blocks with 1 in each node.


- Port blocks created on the backup node from the backup pools are not accounted for when enforcing the maximum-per-host limit.
 - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster nodes.
 - When operating in a cluster, you cannot simply change the block allocation size. The new size is effective only after you reload each device in the cluster. To avoid having to reload each device, we recommend that you delete all block allocation rules and clear all xlates related to those rules. You can then change the block size and recreate the block allocation rules.
- NAT pool address distribution for dynamic PAT—When you configure a PAT pool, the cluster divides each IP address in the pool into port blocks. By default, each block is 512 ports, but if you configure port block allocation rules, your block setting is used instead. These blocks are distributed evenly among the nodes in the cluster, so that each node has one or more blocks for each IP address in the PAT pool. Thus, you could have as few as one IP address in a PAT pool for a cluster, if that is sufficient for the number of PAT'ed connections you expect. Port blocks cover the 1024-65535 port range, unless you configure the option to include the reserved ports, 1-1023, on the PAT pool NAT rule.
 - 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.
 - No extended PAT—Extended PAT 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

- If you have an extremely large number of NAT rules, over ten thousand, you should enable the transactional commit model using the `asp rule-engine transactional-commit nat` command in the device CLI. Otherwise, the node might not be able to join the cluster.

Dynamic Routing

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

Figure 28: Dynamic Routing in Spanned EtherChannel Mode



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

The OSPF LSA database is not synchronized from the control node to data nodes. If there is a control node 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.

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 threat defense by its Diagnostic 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.

When using SNMPv3 with clustering, if you add a new cluster node after the initial cluster formation, then SNMPv3 users are not replicated to the new node. You must remove the users, and re-add them, and then reapply your configuration to force the users to replicate to the new node.

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.

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.

VPN and Clustering

Site-to-site VPN is a centralized feature; only the control node supports VPN connections.



Note Remote access VPN is not supported with clustering.

VPN functionality is limited to the control node and does not take advantage of the cluster high availability capabilities. If the control node fails, all existing VPN connections are lost, and VPN users will see a disruption in service. When a new control node is elected, you must reestablish the VPN connections.

When you connect a VPN tunnel to a Spanned EtherChannel address, connections are automatically forwarded to the control node.

VPN-related keys and certificates are replicated to all nodes.

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, if your model can handle approximately 10 Gbps of traffic when running alone, then for a cluster of 8 units, the maximum combined throughput will be approximately 80% of 80 Gbps (8 units x 10 Gbps): 64 Gbps.

Control Node Election

Nodes of the cluster communicate over the cluster control link to elect a control node as follows:

1. When you enable clustering for a node (or when it first starts up with clustering already enabled), it broadcasts an election request every 3 seconds.
2. Any other nodes with a higher priority respond to the election request; the priority is set between 1 and 100, where 1 is the highest priority.
3. If after 45 seconds, a node does not receive a response from another node with a higher priority, then it becomes the control node.



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

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



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

High Availability Within the Cluster

Clustering provides high availability by monitoring node and interface health and by replicating connection states between nodes.

Node Health Monitoring

Each node periodically sends a broadcast heartbeat packet over the cluster control link. If the control node does not receive any heartbeat 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 nodes.

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 Node Election, on page 42](#) for more information.

Interface Monitoring

Each node monitors the link status of all named hardware interfaces in use, and reports status changes to the control node.

- **Spanned EtherChannel**—Uses cluster Link Aggregation Control Protocol (cLACP). Each node monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel. The status is reported to the control node.

All physical interfaces (including the main EtherChannel) are monitored by default. Only named interfaces can be monitored. For example, the named EtherChannel must fail to be considered failed, which means all member ports of an EtherChannel must fail to trigger cluster removal.

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the threat defense removes a member from the cluster depends on whether the node is an established member or is joining the cluster. If the interface is down on an established member, then the threat defense removes the member after 9 seconds. The threat defense 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 threat defense to be removed from the cluster. For non-EtherChannels, the node is removed after 500 ms, regardless of the member state.

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 threat defense automatically tries to rejoin the cluster, depending on the failure event.



Note When the threat defense 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 threat defense automatically tries to rejoin every 5 minutes, indefinitely.
- **Failed data interface**—The 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 threat defense application disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering.
- **Failed node**—If the node was removed from the cluster because of a node health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the node will rejoin the cluster when it starts up again as long as the cluster control link is up. The threat defense application attempts to rejoin the cluster every 5 seconds.

- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on.
- Failed configuration deployment—If you deploy a new configuration from management center, and the deployment fails on some cluster members but succeeds on others, then the nodes that failed are removed from the cluster. You must manually rejoin the cluster by re-enabling clustering. If the deployment fails on the control node, then the deployment is rolled back, and no members are removed. If the deployment fails on all data nodes, 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	—
MAC address table	Yes	—
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.

For clustering on the Firepower 9300, which can include up to 3 cluster nodes in one chassis, if the backup owner is on the same chassis as the owner, then an additional backup owner will be chosen from another chassis to protect flows from a chassis failure.

- **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.

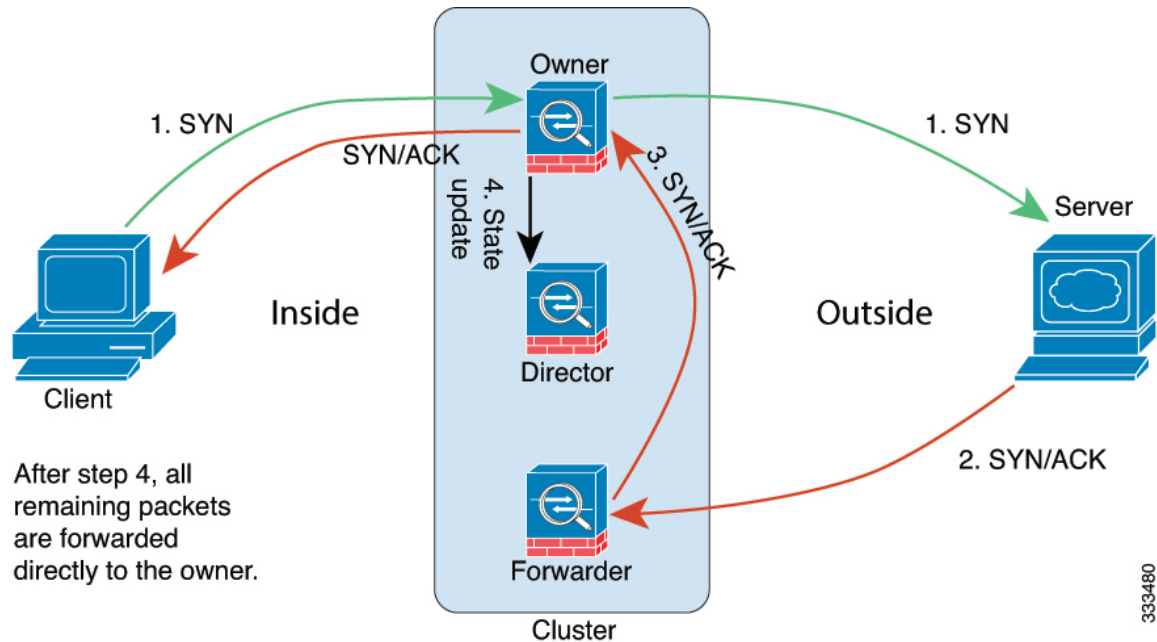
- **Fragment Owner**—For fragmented packets, cluster nodes that receive a fragment determine a fragment owner using a hash of the fragment source IP address, destination IP address, and the packet ID. All fragments are then forwarded to the fragment owner over the cluster control link. Fragments may 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 threat defense (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 threat defense (based on the load balancing method). This threat defense 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.

Sample Data Flow for ICMP and UDP

The following example shows the establishment of a new connection.

1. *Figure 29: ICMP and UDP Data Flow*



The first UDP packet originates from the client and is delivered to one threat defense (based on the load balancing method).

2. The node that received the first packet queries the director node that is chosen based on a hash of the source/destination IP address and ports.
3. The director finds no existing flow, creates a director flow and forwards the packet back to the previous node. In other words, the director has elected an owner for this flow.
4. The owner creates the flow, sends a state update to the director, and forwards the packet to the server.
5. The second UDP packet originates from the server and is delivered to the forwarder.
6. The forwarder queries the director for ownership information. For short-lived flows such as DNS, instead of querying, the forwarder immediately sends the packet to the director, which then sends it to the owner.
7. The director replies to the forwarder with ownership information.
8. The forwarder creates a forwarding flow to record owner information and forwards the packet to the owner.
9. The owner forwards the packet to the client.

History for Clustering

Feature	Minimum Management Center	Minimum Threat Defense	Details
Cluster control link ping tool.	7.2.6	Any	<p>You can check to make sure all the cluster nodes can reach each other over the cluster control link by performing a ping. One major cause for the failure of a node to join the cluster is an incorrect cluster control link configuration; for example, the cluster control link MTU may be set higher than the connecting switch MTUs.</p> <p>New/modified screens: Devices > Device Management > More (⚙) > Cluster Live Status</p> <p>Other version restrictions: Not supported with management center Version 7.3.x or 7.4.0.</p>
Automatic configuration of the cluster control link MTU	7.2.0	7.2.0	The MTU of the cluster control link interface is now automatically set to 100 bytes more than the highest data interface MTU; by default, the MTU is 1600 bytes.
Clustering for the Secure Firewall 3100	7.1.0	7.1.0	<p>The Secure Firewall 3100 supports Spanned EtherChannel clustering for up to 8 nodes.</p> <p>New/Modified screens:</p> <ul style="list-style-type: none"> • Devices > Device Management > Add Cluster • Devices > Device Management > More menu • Devices > Device Management > Cluster <p>Supported platforms: Secure Firewall 3100</p>