

Clustering: Private Cloud

Clustering lets you group multiple Firewall Threat Defense Virtuals 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. You can deploy Firewall Threat Defense Virtual clusters in a private cloud using VMware and KVM. Only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported features and clustering, on page 36.

- About clustering: private cloud, on page 1
- Licenses for clustering: Private cloud, on page 5
- Prerequisites for clustering: Private cloud, on page 6
- Guidelines for clustering: Private cloud, on page 7
- Configure clustering: private cloud, on page 8
- Manage cluster nodes, on page 20
- Monitor cluster status, on page 31
- Reference for clustering, on page 36
- History for clustering: private cloud, on page 48

About clustering: private cloud

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 device. To act as a cluster, the firewalls need the following infrastructure:

- Isolated network for intra-cluster communication, known as the cluster control link, using VXLAN interfaces. VXLANs, which act as Layer 2 virtual networks over Layer 3 physical networks, let the Firewall Threat Defense Virtual send broadcast/multicast messages over the cluster control link.
- Management access to each firewall for configuration and monitoring. The Firewall Threat Defense Virtual deployment includes a Management 0/0 interface that you will use to manage the cluster nodes.

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 Layer 3 Individual interfaces and one of the following methods:

- Policy-Based Routing—The upstream and downstream routers perform load balancing between nodes using route maps and ACLs.
- Equal-Cost Multi-Path Routing—The upstream and downstream routers perform load balancing between nodes using equal cost static or dynamic routes.

Layer 2 Spanned EtherChannels are not supported.

Control and data node roles

When you add a cluster in the Cloud-Delivered Firewall Management Center , you choose one firewall to be the control node, and all additional firewalls to be data nodes. When you first create the cluster, the control node you specify will become the control node simply because it is the first node added to the cluster. Later, 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.

When you create the cluster and deploy the cluster configuration, the Cloud-Delivered Firewall Management Center deploys a special bootstrap configuration to each node that includes fundamental settings such as the cluster control link configuration. The bootstrap configuration lets each node join the cluster. Most bootstrap settings are defined in the cluster wizard; however, you define the cluster control link interface hardware configuration (for example, creating an EtherChannel or setting the Ethernet speed) on the control node before you create the cluster. The cluster control link settings are copied to the bootstrap configuration for each node.

All nodes in the cluster share the same policy 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. After you create the cluster, all configuration changes are made at the cluster level and are shared by all nodes.

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

Individual interfaces

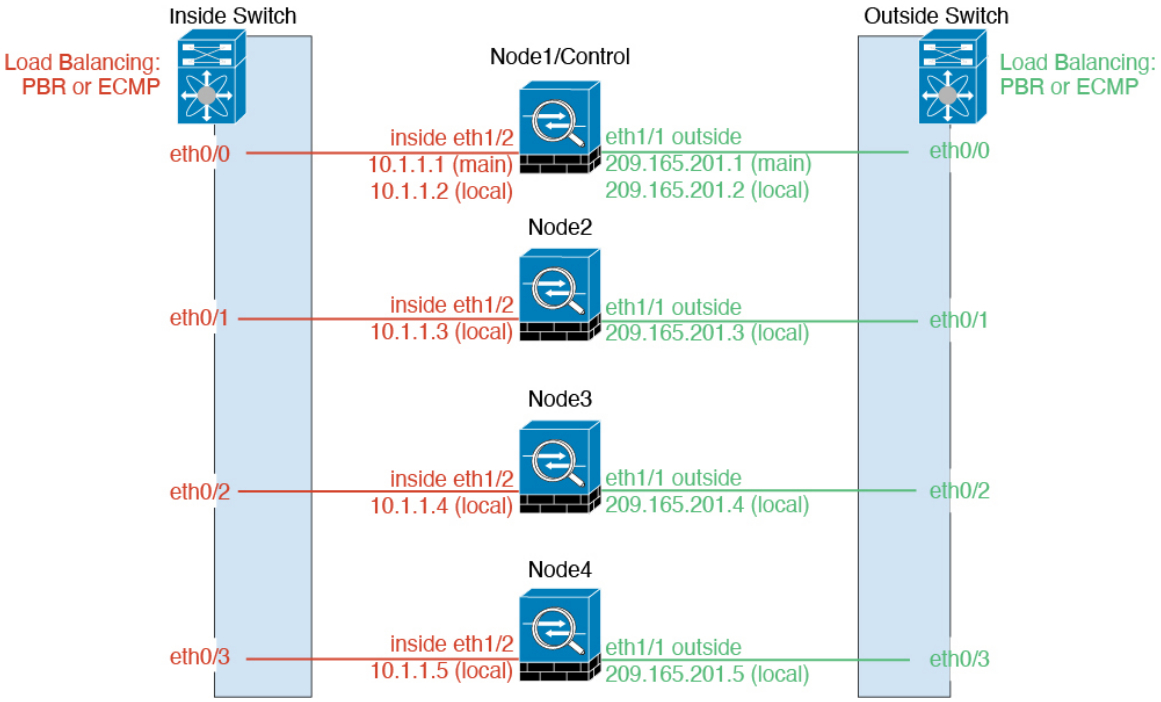
You can configure cluster interfaces as individual interfaces.

Individual interfaces are normal routed interfaces, each with their own *local IP address* used for routing. The *main cluster IP address* for each interface is a fixed address that always belongs to the control node. When the control node changes, the main cluster IP address moves to the new control node, so management of the cluster continues seamlessly.

IPS-only interfaces (inline sets and passive interfaces) are not supported as Individual interfaces.

Because interface configuration must be done only on the control node, configure a pool of IP addresses for a given interface on the cluster nodes, including one for the control node.

Load balancing must be configured separately on the upstream switch.



Layer 2 Spanned EtherChannels are not supported.

Policy-based routing

When using individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Policy-Based Routing (PBR).

We recommend this method if you are already using PBR, and want to take advantage of your existing infrastructure.

PBR makes routing decisions based on a route map and ACL. You must manually divide traffic between all Firewall Threat Defenses in a cluster. Because PBR is static, it may not achieve the optimum load balancing result at all times. To achieve the best performance, we recommend that you configure the PBR policy so that forward and return packets of a connection are directed to the same Firewall Threat Defense.

For example, if you have a Cisco router, redundancy can be achieved by using Cisco IOS PBR with Object Tracking. Cisco IOS Object Tracking monitors each Firewall Threat Defense using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular Firewall Threat Defense.

Equal-cost multi-path routing

When using Individual interfaces, each Firewall Threat Defense interface maintains its own IP address and MAC address. One method of load balancing is Equal-Cost Multi-Path (ECMP) routing.

We recommend this method if you are already using ECMP, and want to take advantage of your existing infrastructure.

ECMP routing considerations

ECMP routing can forward packets over multiple best paths that tie for top place in the routing metric. Like EtherChannel, a hash of source and destination IP addresses and/or source and destination ports can be used to send a packet to one of the next hops.

If you use static routes for ECMP routing, then the Firewall Threat Defense failure can cause problems; the route continues to be used, and traffic to the failed Firewall Threat Defense will be lost.

If you use static routes, be sure to use a static route monitoring feature such as Object Tracking. We recommend using dynamic routing protocols to add and remove routes, in which case, you must configure each Firewall Threat Defense to participate in dynamic routing.

Cluster Control Link

Each node must dedicate one interface as a VXLAN (VTEP) interface for the cluster control link. For more information about VXLAN, see Configure VXLAN interfaces .

VXLAN Tunnel Endpoint

VXLAN tunnel endpoint (VTEP) devices perform VXLAN encapsulation and decapsulation. Each VTEP has two interface types: one or more virtual interfaces called VXLAN Network Identifier (VNI) interfaces, and a regular interface called the VTEP source interface that tunnels the VNI interfaces between VTEPs. The VTEP source interface is attached to the transport IP network for VTEP-to-VTEP communication.

VTEP Source Interface

The VTEP source interface is a regular Firewall Threat Defense Virtual interface with which you plan to associate the VNI interface. You can configure one VTEP source interface to act as the cluster control link. The source interface is reserved for cluster control link use only. Each VTEP source interface 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.

VNI Interface

A VNI interface is similar to a VLAN interface. It is a virtual interface that keeps network traffic separated on a given physical interface by using tagging. You can only configure one VNI interface. Each VNI interface has an IP address on the same subnet.

Peer VTEPs

Unlike regular VXLAN for data interfaces, which allows a single VTEP peer, Firewall Threat Defense Virtual clustering allows you to configure multiple peers.

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.

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: Private cloud

Performance tier license requirements

Each Firewall Threat Defense Virtual cluster node requires the same performance tier license. We recommend using the same number of CPUs and memory for all members, or else performance will be limited on all nodes to match the least capable member. The throughput level will be replicated from the control node to each data node so they match.

Feature license assignment

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 Cloud-Delivered Firewall 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 by clicking **Edit Licenses** in **Administration** > **Licenses** > **Smart Licenses** or choosing **Devices** > **Device Management**, clicking **Edit** (✎) for the cluster, and then in the **License** area, clicking **Edit** (✎).

Note If you add the cluster before the Cloud-Delivered Firewall Management Center is licensed (and running in Evaluation mode), then when you license the Cloud-Delivered Firewall 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.

Note APEX, PLUS, and VPN-Only licenses are not supported for clustering. Before you create a cluster, remove unsupported license assignments from the devices that will become cluster nodes.

Prerequisites for clustering: Private cloud

This reference provides the complete set of requirements that must be met before configuring clustering for threat defense virtual devices in private cloud deployments.

Model requirements

- FTDv5, FTDv10, FTDv20, FTDv30, FTDv50, FTDv100
- VMware or KVM
- A maximum of 4 nodes in a cluster in a 2x2 configuration is supported. You can set up a maximum of two hosts with a maximum of two threat defense virtual instances in each host.

Hardware and software requirements

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.
- Must be the same performance tier. We recommend using the same number of CPUs and memory for all nodes, or performance will be limited on all nodes to match the least capable node.
- Must use the management interface for Cloud-Delivered Firewall Management Center communications. Data interface management is not supported.
- Must run the same version, except during upgrade. Hitless upgrade is supported.
- Must not have any deployment pending or in progress.
- Must not have any unsupported features configured on the control node: Unsupported features and clustering, on page 36.
- Must not have VPN configured on the data nodes. The control node can have site-to-site VPN configured.

Note Devices must not be in transparent mode for cluster configuration.

Cloud-Delivered Firewall Management Center Requirements

Make sure the Cloud-Delivered Firewall Management Center NTP server is set to a reliable server that is reachable by all cluster nodes to ensure proper clock sync. By default, the device uses the same NTP server as the Cloud-Delivered Firewall Management Center. If the time is not set to be the same on all cluster nodes, they can be removed automatically from the cluster.

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 154 bytes higher than the data interfaces. If the switches have an MTU mismatch, the cluster formation will fail. When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size matching the cluster control link MTU. If the initial ping fails, the node tries a ping using a smaller packet size (the MTU divided by 2, then by 4, then by 8) until a ping succeeds. A notification is generated so you can fix the MTU mismatch on connecting switches and try again.

Guidelines for clustering: Private cloud

High availability

High Availability is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional clustering guidelines

Consider these additional guidelines when configuring clustering:

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the Firewall Threat Defense Virtual, adding an additional switch to form a VSS or vPC) you should disable the health check feature and also disable interface monitoring for the interfaces that are affected by the topology changes. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the interface health check feature.
- 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.
- 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.
- We do not support VXLAN for data interfaces; only the cluster control link supports VXLAN.

Default cluster configurations

Consider these default configurations when setting up 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.

Upgrade guidelines

Configure clustering: private cloud

To configure clustering after you deploy your Firewall Threat Defense Virtuals, perform the following tasks.

Add nodes to the Management Center

This task allows you to prepare cluster nodes for clustering configuration by adding them to the management center as individual devices.

Before configuring clustering, deploy each cluster node, then add the devices as standalone units on the Cloud-Delivered Firewall Management Center.

Procedure

Step 1 Deploy each cluster node by following the procedures given in the Secure Firewall Threat Defense Virtual getting started guides.

All units in a cluster:

- Must have jumbo frame reservation enabled for the cluster control link. Do this in the Day 0 configuration when you deploy the Firewall Threat Defense Virtual by setting `"DeploymentType": "Cluster"`. Otherwise, you must restart each node to enable jumbo frames after the cluster has formed and is healthy.
- (KVM only) Must use CPU hard partitioning (CPU pinning) for all VMs on the KVM host.

Step 2 Add each node to the Cloud-Delivered Firewall Management Center as a standalone device in the same group.

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.

Create a cluster

Creating a cluster allows multiple devices to work together as a single logical unit, providing high availability and increased throughput capacity.

Form a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Before you begin

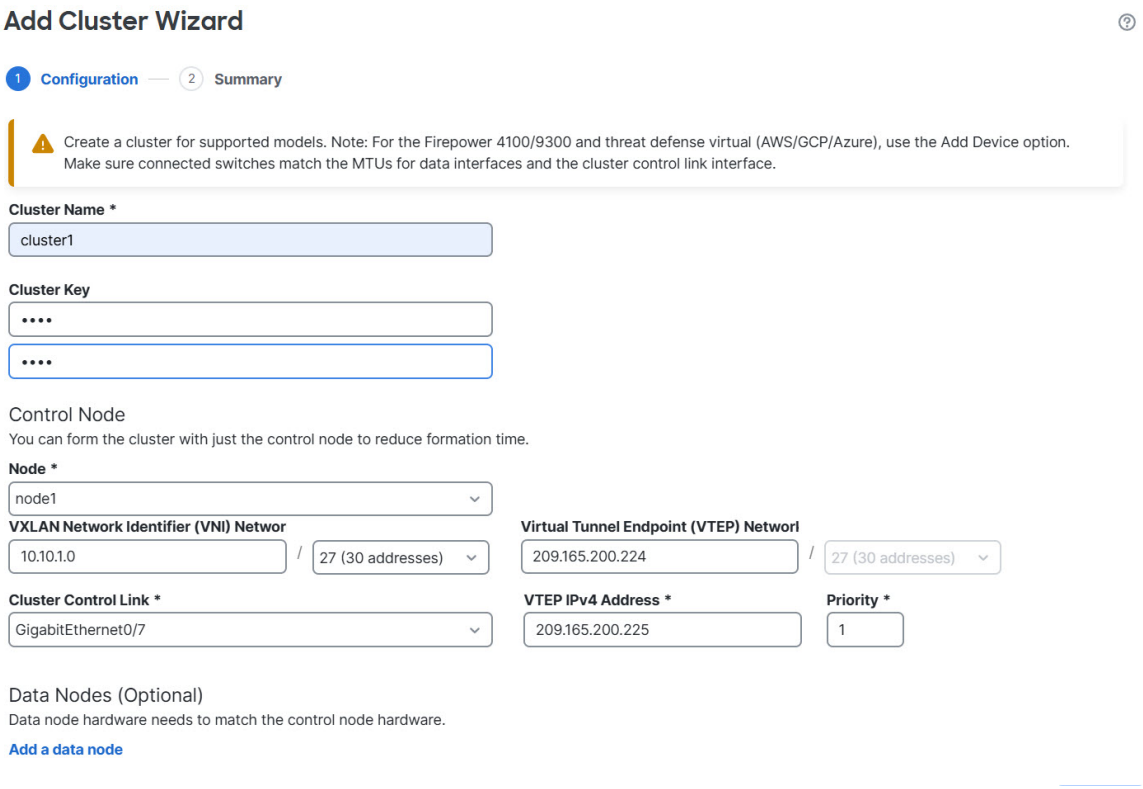
Some features are not compatible with clustering. Wait to configure until clustering is enabled. Some features will block cluster creation if they are already configured. For example, do not configure any IP addresses on interfaces, or unsupported interface types such as BVIs.

Procedure

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

The **Add Cluster Wizard** appears.

Figure 1: Add Cluster Wizard



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

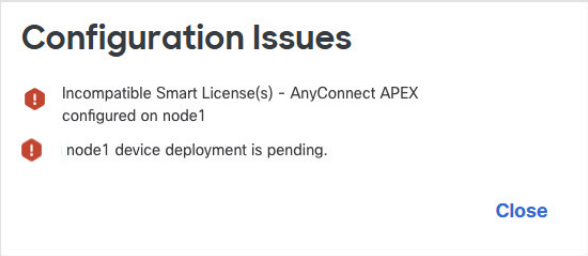
- **Cluster Name**—An ASCII string of 1 to 38 characters. The cluster name can contain only letters, digits, underscores (_), and hyphens (-).
- **Cluster Key**—An ASCII string of 1 to 63 characters. The cluster key cannot be a single digit and must not contain spaces. The **Cluster Key** value is used to generate the encryption key. This encryption does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in clear text.

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

- **Node**—Choose the device that you want to be the control node initially. When the Cloud-Delivered Firewall 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 2: Configuration Issues



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

- **VXLAN Network Identifier (VNI) Network**—Specify an IPv4 subnet for the VNI network; IPv6 is not supported for this network. Specify a **24**, **25**, **26**, or **27** subnet. An IP address will be auto-assigned to each node on this network. The VNI network is the encrypted virtual network that runs on top of the physical VTEP network.
- **Cluster Control Link**—Choose the physical interface you want to use for the cluster control link.
- **Virtual Tunnel Endpoint (VTEP) Network**—Specify an IPv4 subnet for the physical interface network; IPv6 is not supported for this network. The VTEP network is a different network than the VNI network, and it is used for the physical cluster control link.
- **VTEP IPv4 Address**—This field will be auto-populated with the first address on the VTEP network.
- **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.

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.

- **VTEP IPv4 Address**—This field will be auto-populated with the next address on the VTEP network.

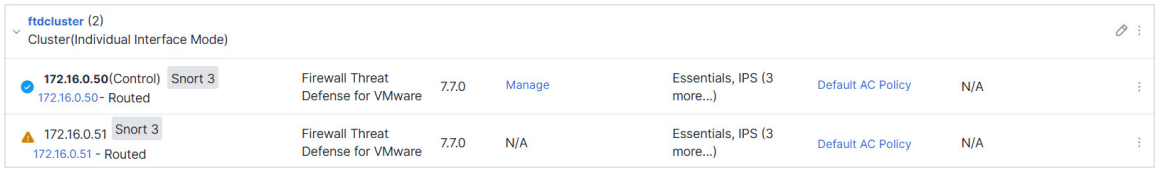
- **Priority**—Set the priority of this node for control node elections. The priority is between 1 and 100, where 1 is the highest priority.

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

The cluster bootstrap configuration is saved to the cluster nodes. The bootstrap configuration includes the VXLAN interface used for the cluster control link.

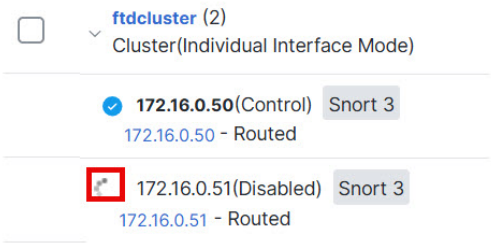
The cluster name shows on the **Devices** > **Device Management** page; expand the cluster to see the cluster nodes.

Figure 3: Cluster Management

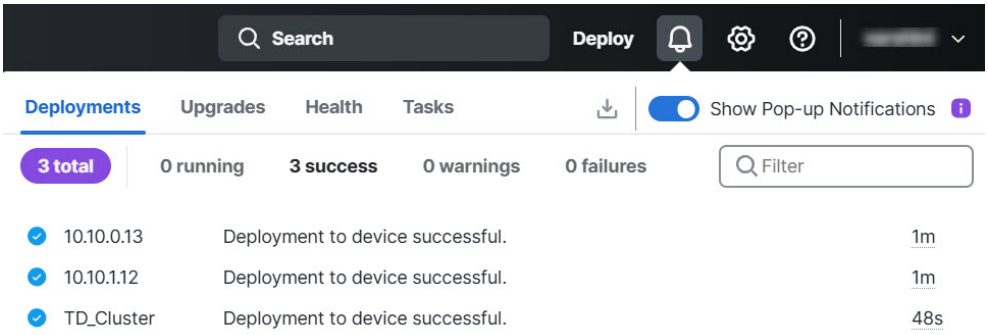


A node that is currently registering shows the loading icon.

Figure 4: Node Registration



You can monitor cluster node registration by clicking the **Notifications** icon and choosing **Tasks**. The Cloud-Delivered Firewall Management Center updates the Cluster Registration task as each node registers.

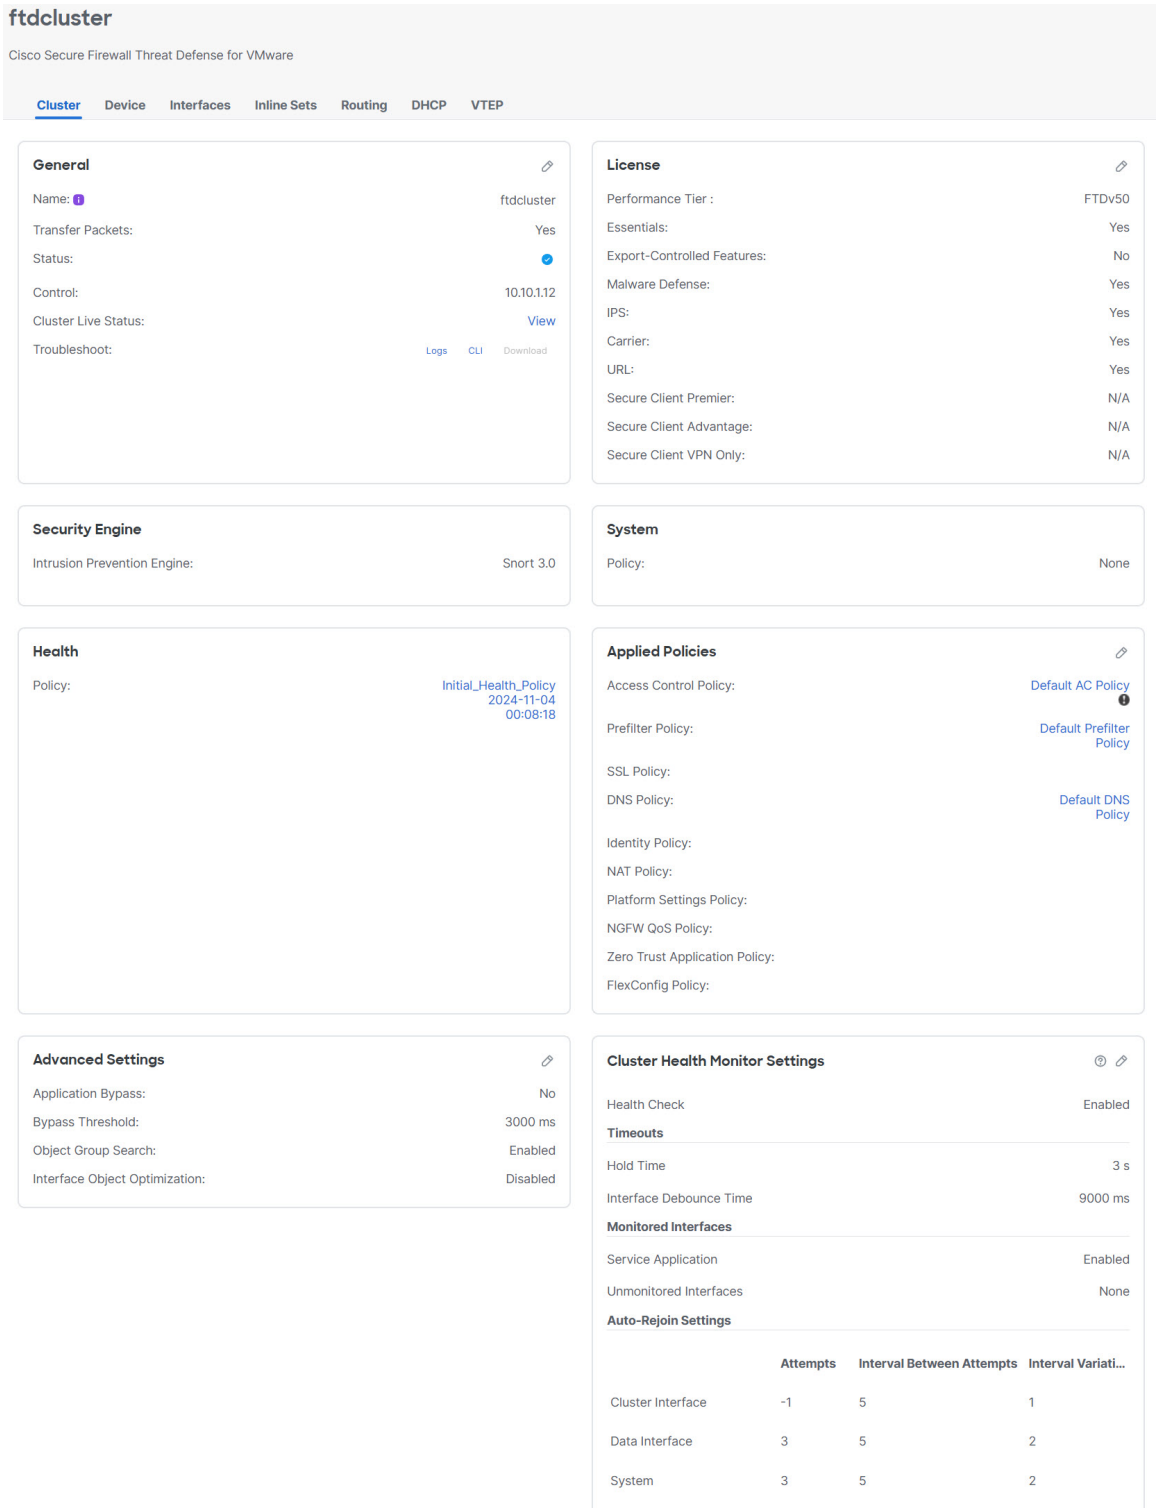


Step 6 Configure device-specific settings by clicking the **Edit** (✎) for the cluster.

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.

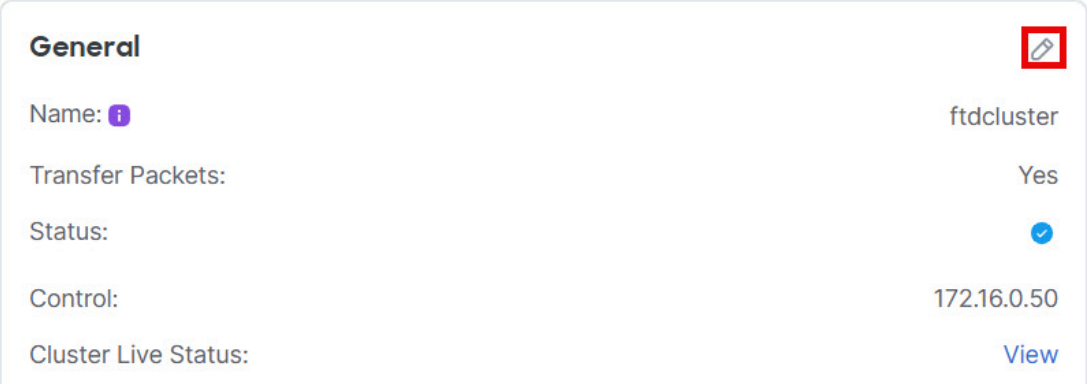
Step 7 On the **Devices** > **Device Management**, **Cluster** screen, you see **General** and other settings for the cluster.

Figure 5: Cluster Settings

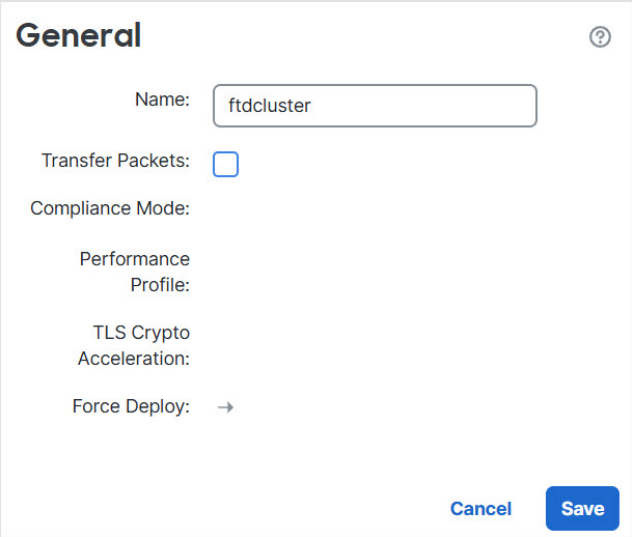


See the following cluster-specific items in the **General** area:

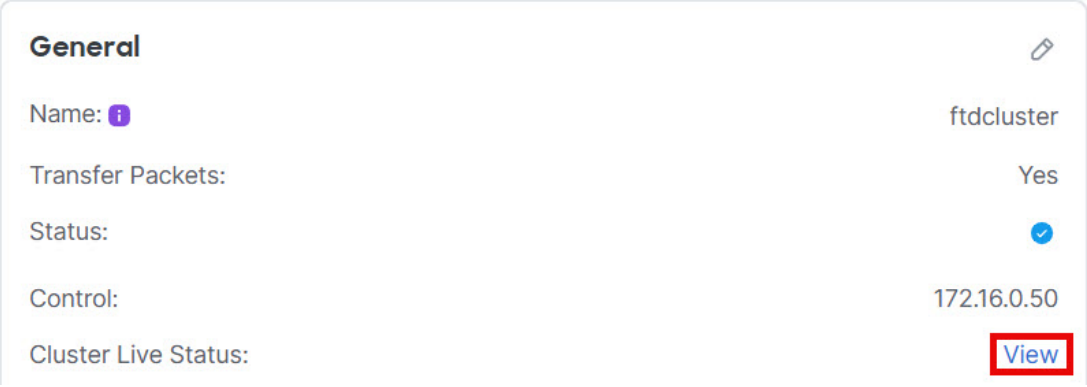
- **General > Name**—Change the cluster display name by clicking the **Edit** ( ).



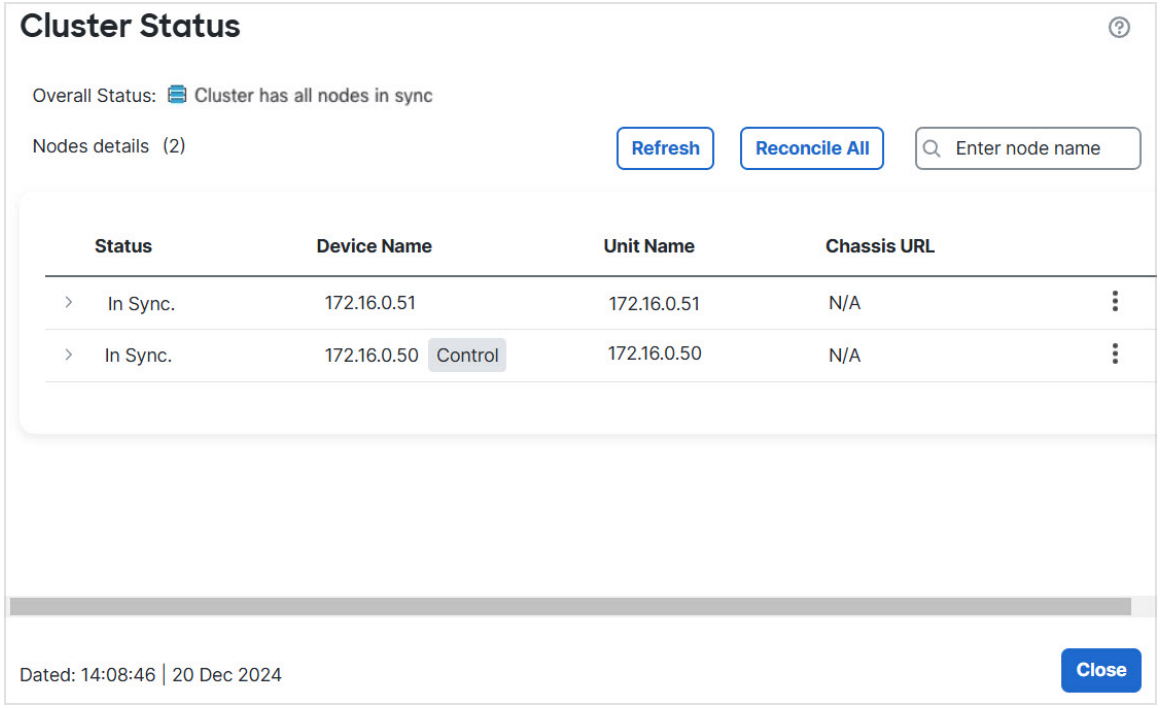
Then set the **Name** field.



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



The **Cluster Status** dialog box also lets you retry data unit registration by clicking **Reconcile All**.



Step 8 On the **Devices** > **Device Management** and then click **Add** > **Device**, you can choose each member in the cluster from the top right drop-down menu and configure the following settings.

Figure 6: Device Settings

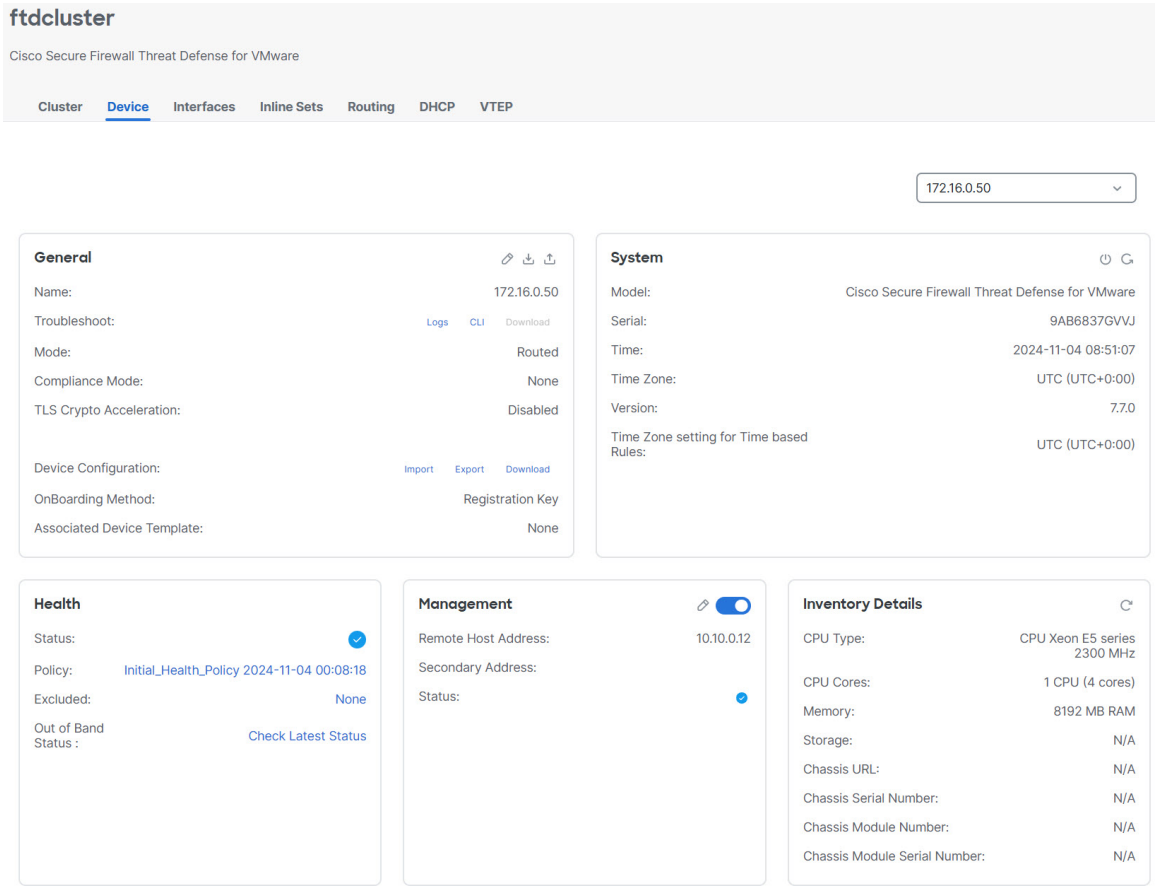
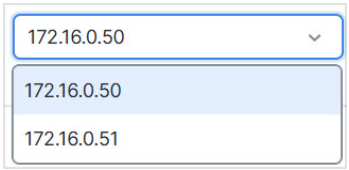
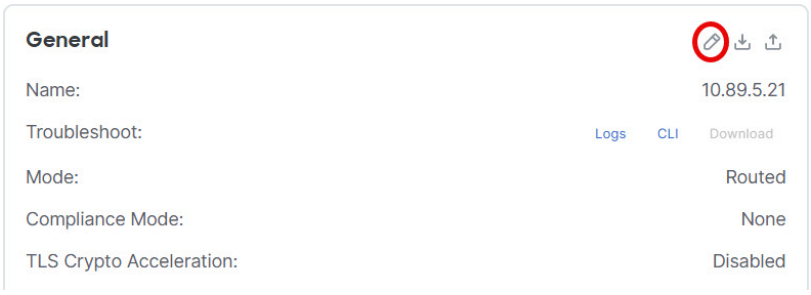


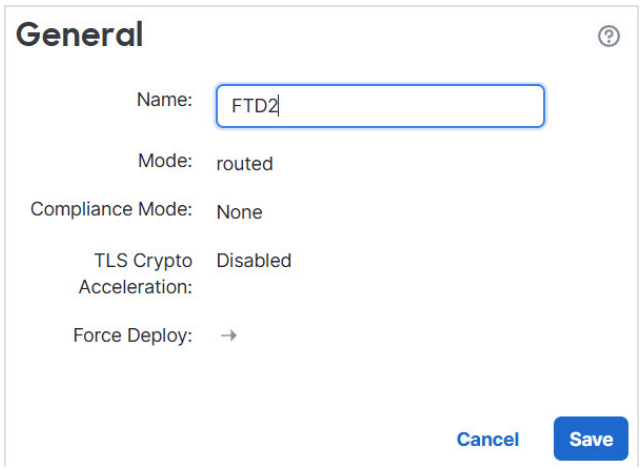
Figure 7: Choose Node



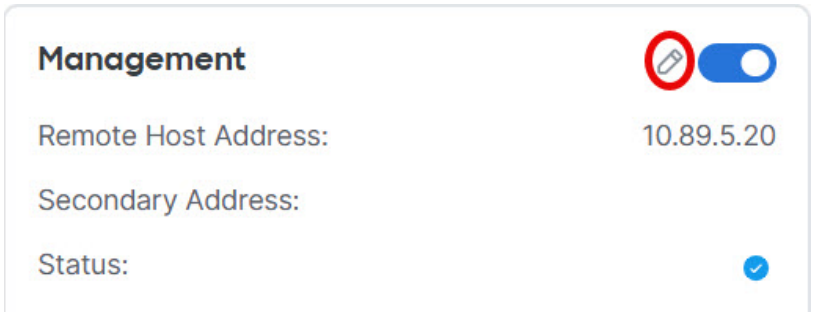
- **General > Name**—Change the cluster member display name by clicking the **Edit** (✎).



Then set the **Name** field.



- **Management > Host**—If you change the management IP address in the device configuration, you must match the new address in the Cloud-Delivered Firewall 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.



Step 9 If you deployed your cluster nodes without enabling jumbo-frame reservation, then restart all cluster nodes to enable jumbo frames, which are required for the cluster control link. See Shut down or restart the device.

If you previously enabled jumbo-frame reservation, you can skip this step.

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 (100 bytes) and VXLAN overhead (54 bytes). When you create the cluster, the MTU is set to 154 bytes higher than the highest data interface MTU (1654 by default). If you later increase the data interface MTU, be sure to also increase the cluster control link MTU. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9044, while the cluster control link can be set to 9198. See Configure the MTU.

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

Configure interfaces

Configure interfaces to be compatible with clustering by setting up individual interfaces that are normal routed interfaces, each with their own IP address taken from a pool of IP addresses.

Individual interfaces are normal routed interfaces, each with their own IP address taken from a pool of IP addresses. The Main cluster IP address is a fixed address for the cluster that always belongs to the current control node. All data interfaces must be Individual interfaces.

Note You cannot use subinterfaces.

Procedure

Step 1 Choose **Objects** > **Address Pools** to add an IPv4 address pool, an IPv6 address pool, or both. See Configure 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.

Step 2 Choose **Devices** > **Device Management**, and click **Edit** (✎) next to the cluster.

Step 3 Click **Interfaces**, and then click **Edit** (✎) for a data interface.

Step 4 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.

Step 5 From the **IPv4 Address Pool** drop-down list, choose the address pool you created.

Note If you want to manually assign a MAC address to this interface, you need to create a **MAC-address pool** using FlexConfig.

Step 6 On **IPv6** > **Basic**, from the **IPv6 Address Pool** drop-down list, choose the address pool you created.

Step 7 Configure other interface settings as normal.

Step 8 Click **Save**.

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

Configure Cluster Health Monitor Settings

This task allows you to modify cluster health monitor settings to control how cluster nodes monitor system health and automatically rejoin the cluster after failures.

The **Cluster Health Monitor Settings** section of the **Cluster** page displays the settings for monitoring cluster node health. You can monitor any port-channel ID, any single physical interface ID, and the Snort and disk-full processes. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link. It is always monitored.

When a topology change occurs (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable

the system health check feature. Also, disable interface monitoring for any interfaces that are disabled. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Procedure

Step 1 Choose **Devices** > **Device Management**.

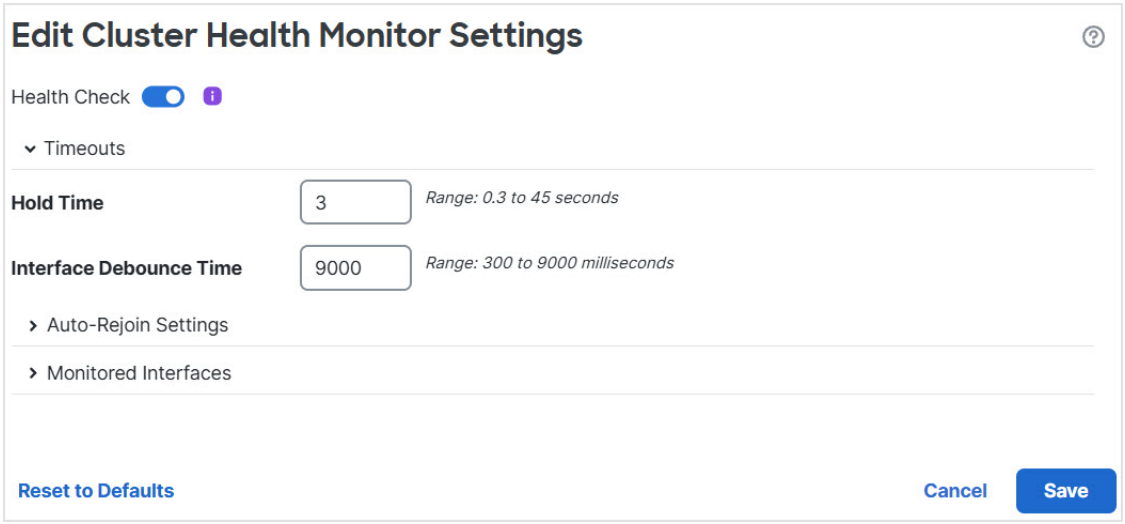
Step 2 Next to the cluster you want to modify, click **Edit** (✎).

Step 3 Click **Cluster**.

Step 4 In the **Cluster Health Monitor Settings** section, click **Edit** (✎).

Step 5 Disable the system health check by clicking the **Health Check** slider.

Figure 8: Disable the system health check



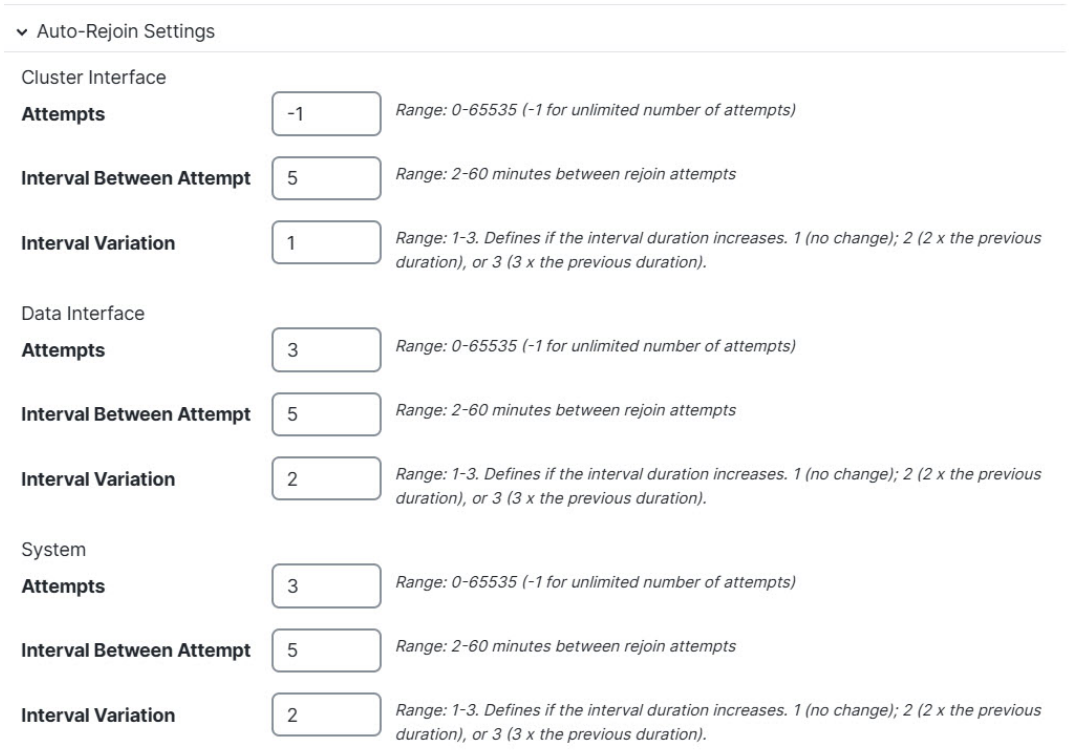
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 6 Configure the hold time and interface debounce time.

- **Hold Time**—Set the hold time to determine the amount of time between node heartbeat status messages, between 0.3 and 45 seconds. The default is 3 seconds.
- **Interface Debounce Time**—Set the debounce time between 300 and 9000 ms. The default is 500 ms. Lower values allow for faster detection of interface failures. Note that configuring a lower debounce time increases the chances of false-positives. When an interface status update occurs, the node waits the specified number of milliseconds before marking the interface as failed and removing the node from the cluster. If an EtherChannel transitions from a down state to an up state, such as after the switch reloads or enables an EtherChannel, a longer debounce time can prevent the interface from appearing failed on a cluster node if another cluster node bundles the ports more quickly.

Step 7 Customize the auto-rejoin cluster settings after a health check failure.

Figure 9: Configure Auto-Rejoin settings

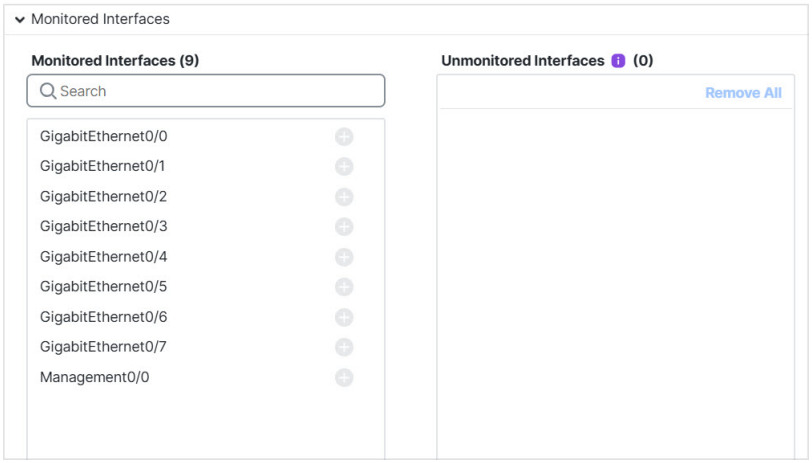


Set values for the **Cluster Interface**, **Data Interface**, and **System** (internal failures include: application sync timeout; inconsistent application statuses; and so on):

- **Attempts**—Sets the number of rejoin attempts, between -1 and 65535. **0** disables auto-rejoining. The default for the **Cluster Interface** is -1 (unlimited). The default for the **Data Interface** and **System** is 3.
- **Interval Between Attempts**—Defines the interval duration in minutes between rejoin attempts, between 2 and 60. The default value is 5 minutes. The maximum total time that the node attempts to rejoin the cluster is limited to 14400 minutes (10 days) from the time of last failure.
- **Interval Variation**—Defines if the interval duration increases. Set the value between 1 and 3: **1** (no change); **2** (2 x the previous duration), or **3** (3 x the previous duration). For example, if you set the interval duration to 5 minutes, and set the variation to 2, then the first attempt is after 5 minutes; the 2nd attempt is 10 minutes (2 x 5); the 3rd attempt 20 minutes (2 x 10), and so on. The default value is **1** for the **Cluster Interface** and **2** for the **Data Interface** and **System**.

Step 8 Configure monitored interfaces by moving interfaces in the **Monitored Interfaces** or **Unmonitored Interfaces** window. You can also check or uncheck **Enable Service Application Monitoring** to enable or disable monitoring of the Snort and disk-full processes.

Figure 10: Configure monitored interfaces



The interface health check monitors for link failures. If all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. The time required to remove a member from the cluster depends on the type of interface and whether the node is established or is joining the cluster. By default, health check is enabled for all interfaces and the Snort and disk-full processes.

You might want to disable health monitoring of non-essential interfaces.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the node or the switch, or adding an additional switch to form a VSS or vPC or VNet) you should disable the system health check feature and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all nodes, you can re-enable the system health check feature and monitored interfaces.

Step 9 Click **Save**.

Deploy configuration changes.

Manage cluster nodes

Describes how to manage cluster nodes to maintain high availability and optimal performance across your Cisco Secure Firewall cluster deployment.

Add a new cluster node

Adding new cluster nodes expands cluster capacity and provides additional redundancy for your security infrastructure.

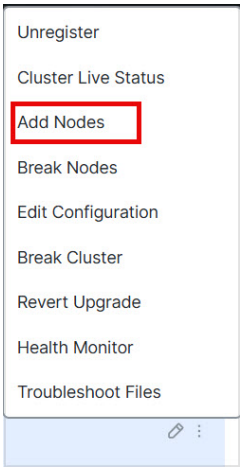
You can add one or more new cluster nodes to an existing cluster.

Procedure

Step 1 If you have enabled the new device management UI, select a cluster to open the side pane, then under **Cluster actions**, click **Add nodes**. See Device management page enhancements. Otherwise, perform the next step.

Step 2 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Add Nodes**.

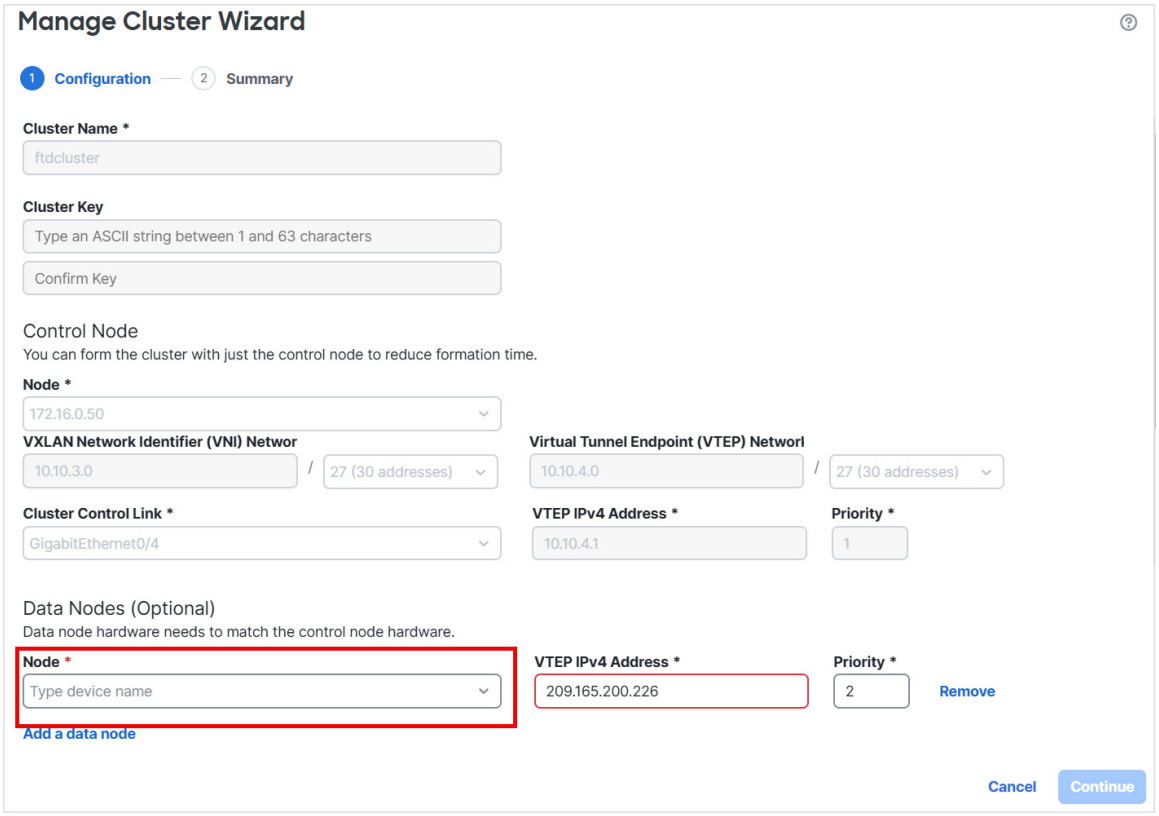
Figure 11: Add nodes



The **Manage Cluster Wizard** appears.

Step 3 From the **Node** menu, choose a device, and adjust the IP address and priority if desired.

Figure 12: Manage cluster wizard

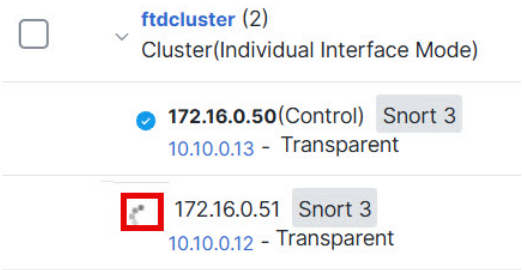


Step 4 To add additional nodes, click **Add a data node**.

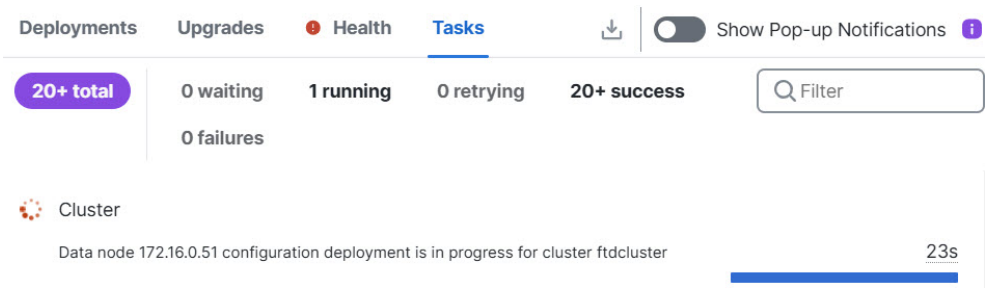
Step 5 Click **Continue**. Review the **Summary**, and then click **Save**

The node that is currently registering shows the loading icon.

Figure 13: Node registration



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



Break a node

Breaking a node removes it from the cluster configuration, converting it to a standalone device.

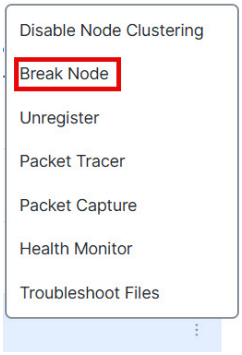
You can remove a node from the cluster so that it becomes a standalone device. You cannot break the control node unless you break the entire cluster. The data node has its configuration erased.

Procedure

Step 1 If you have enabled the new device management UI, select a cluster to open the side pane, then under **Cluster actions**, click **Break node**. See Device management page enhancements. Otherwise, perform the next step.

Step 2 Choose **Devices** > **Device Management**, click the **More** (⋮) for the node you want to break, and choose **Break Node**.

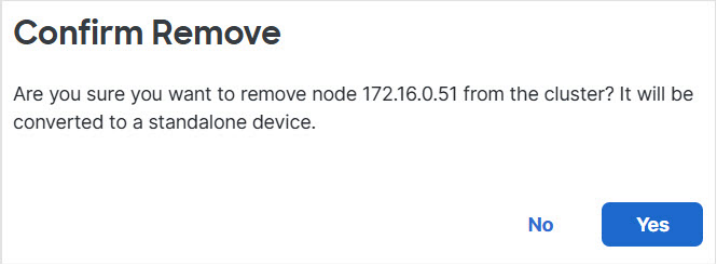
Figure 14: Break a node



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

Step 3 You are prompted to confirm the break. Click **Yes**.

Figure 15: Confirm break



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

Break the cluster

Break the cluster to convert all nodes to standalone devices when clustering is no longer needed.

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.

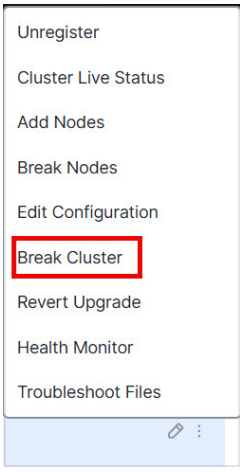
Procedure

Step 1 Make sure all cluster nodes are being managed by the Cloud-Delivered Firewall Management Center by reconciling nodes. Refer to Reconcile cluster nodes, on page 29.

Step 2 If you have enabled the new device management UI, select a cluster to open the side pane, then under **Cluster actions**, click **Break cluster**. See Device management page enhancements. Otherwise, perform the next step.

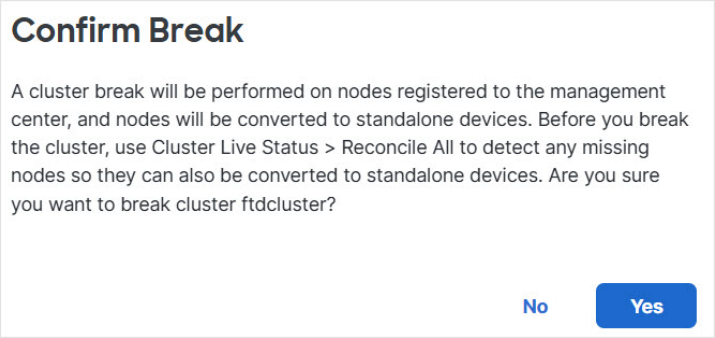
Step 3 Choose **Devices** > **Device Management**, click the **More** (⋮) for the cluster, and choose **Break Cluster**.

Figure 16: Break cluster



Step 4 You are prompted to break the cluster. Click **Yes**.

Figure 17: Confirm break



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

Disable clustering

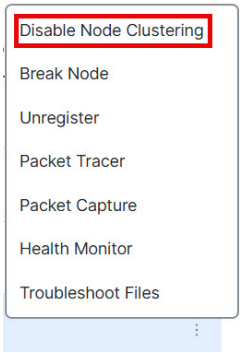
Temporarily deactivate a node in a cluster while keeping it visible in the device list.

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 Cloud-Delivered Firewall 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 18: Disable clustering



If you disable clustering on the control node, one of the data nodes will become the new control node. For centralized features, if you force a control node change, then all connections are dropped, and you have to reestablish 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 reenable clustering, refer to Rejoin the cluster, on page 26.

Rejoin the cluster

Rejoin a node to the cluster after it was removed because of failure or when clustering was manually disabled.

If a node was removed from the cluster, for example, due to 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.

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 node.

Change the control node

This task changes the control node in a cluster when you need to specify the exact unit you want to become the control node.

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. Use this procedure when you must specify the exact unit you want to become the control node.



Caution 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.

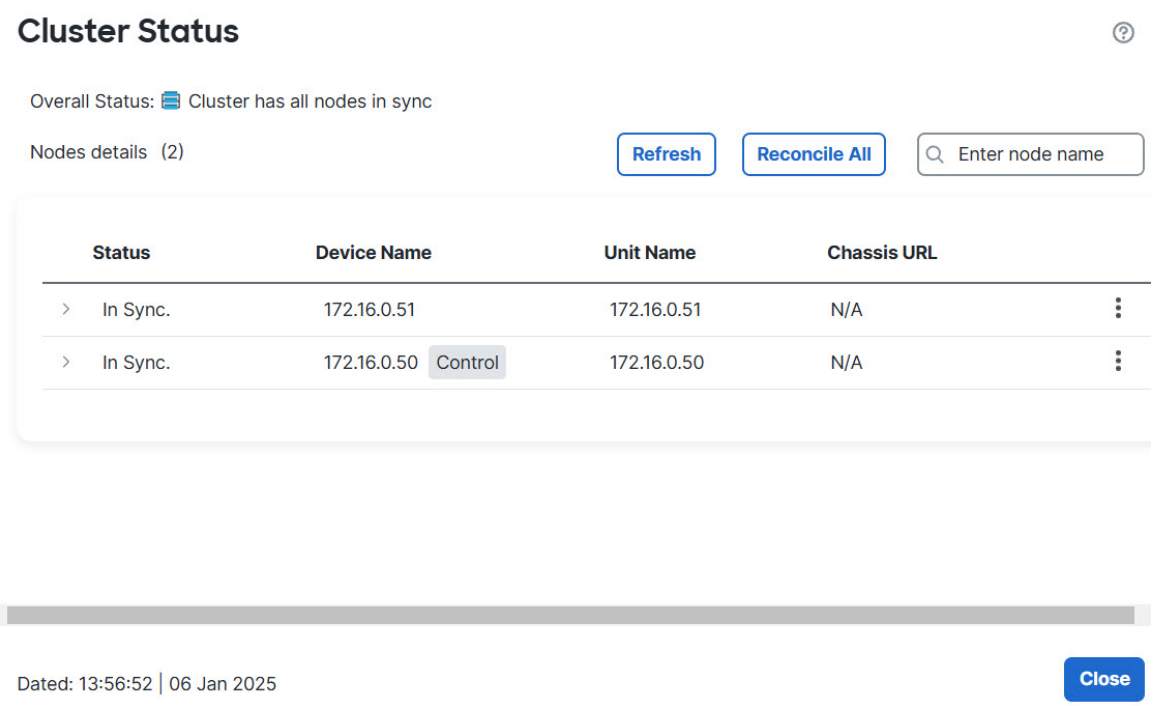
Before you begin

Follow these steps to change the control node:

Procedure

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

Figure 19: Cluster status



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. Select the checkbox and click **OK**.

Edit the cluster configuration

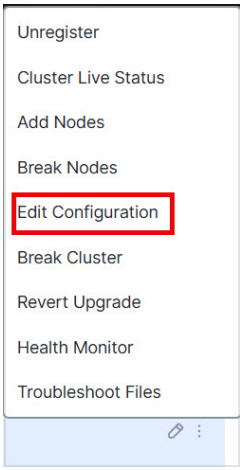
Edit the cluster configuration to modify cluster settings and parameters as needed for your network environment.

You can edit the cluster configuration. If you change any values other than the VTEP IP address for a node or node priority, the cluster will be broken and reformed automatically. Until the cluster is reformed, you may experience traffic disruption. If you change the VTEP IP address for a node or node priority, only the affected nodes are broken and readded to the cluster.

Procedure

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

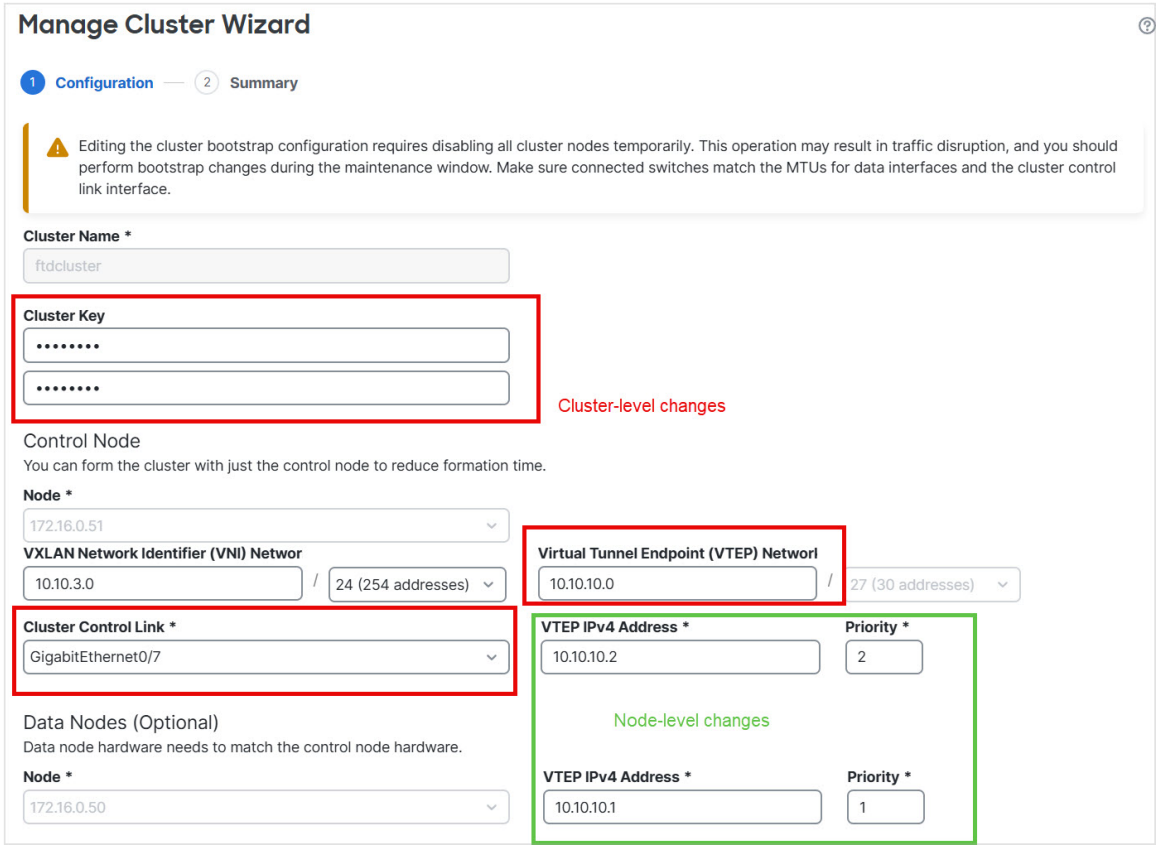
Figure 20: Edit configuration



The **Manage Cluster Wizard** appears.

Step 2 Update the cluster configuration.

Figure 21: Manage cluster wizard



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

Reconcile cluster nodes

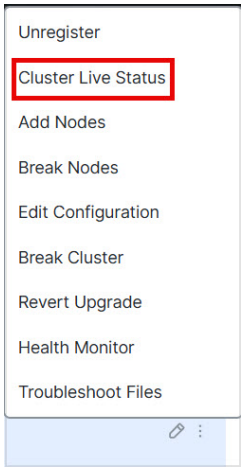
Use this procedure to reconcile cluster membership when a cluster node fails to register with the Cloud-Delivered Firewall Management Center.

If a cluster node fails to register, you can reconcile the cluster membership from the device to the Cloud-Delivered Firewall Management Center. For example, a data node might fail to register if the Cloud-Delivered Firewall Management Center is occupied with certain processes, or if there is a network issue.

Procedure

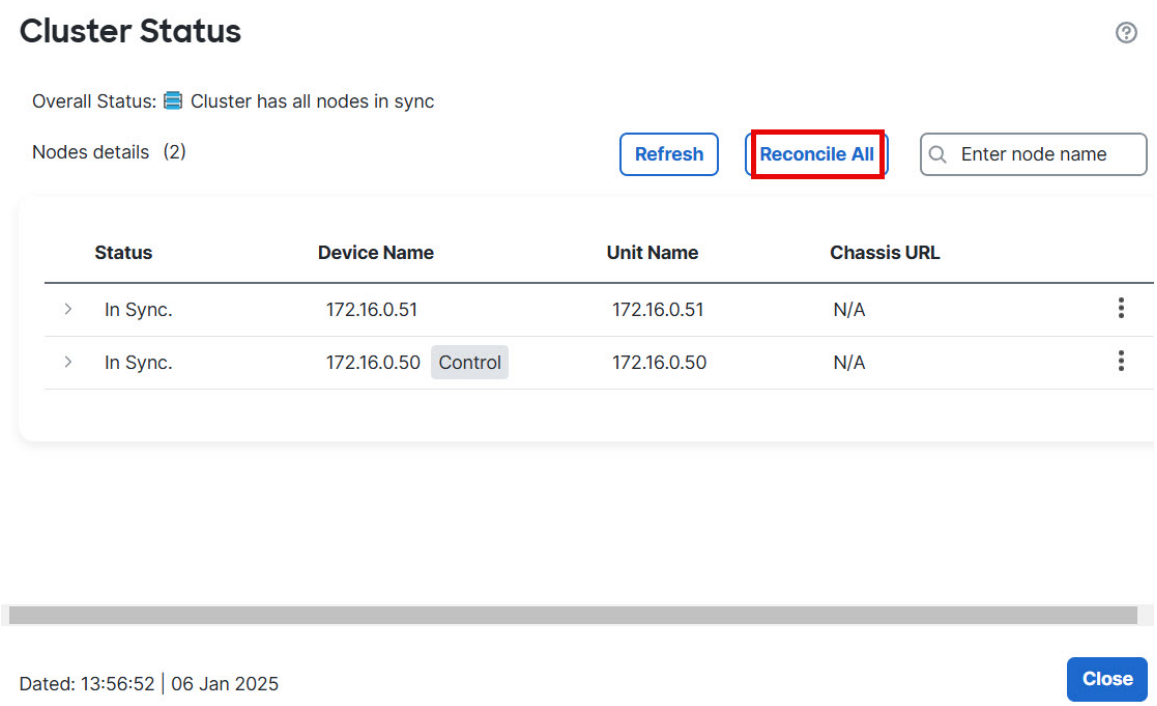
Step 1 Choose **Devices** > **Device Management**, **More** (⋮) for the cluster, and then choose **Cluster Live Status** to open the **Cluster Status** dialog box.

Figure 22: Cluster live status



Step 2 Click **Reconcile All**.

Figure 23: Reconcile all



Delete the cluster or nodes from the Management Center

Delete clusters or nodes from Security Cloud Control to manage cluster membership and device visibility.

You can delete the cluster from the Security Cloud Control, which keeps the cluster intact. You might want to delete the cluster if you want to add the cluster to a new Security Cloud Control.

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

Procedure

Step 1 Log into Security Cloud Control and click **Security Devices**.

Step 2 Click the **FTD** tab and locate the cluster you want. Select it so the device row is highlighted.

Step 3 Do the following:

- To delete a node within the cluster, in the **Cluster** pane on the right, click the delete icon next to the device you want to delete.
- To delete the cluster, in the **Device Actions** pane on the right, click **Remove**.

Step 4 When prompted, select **OK** to confirm the removal of the selected device.

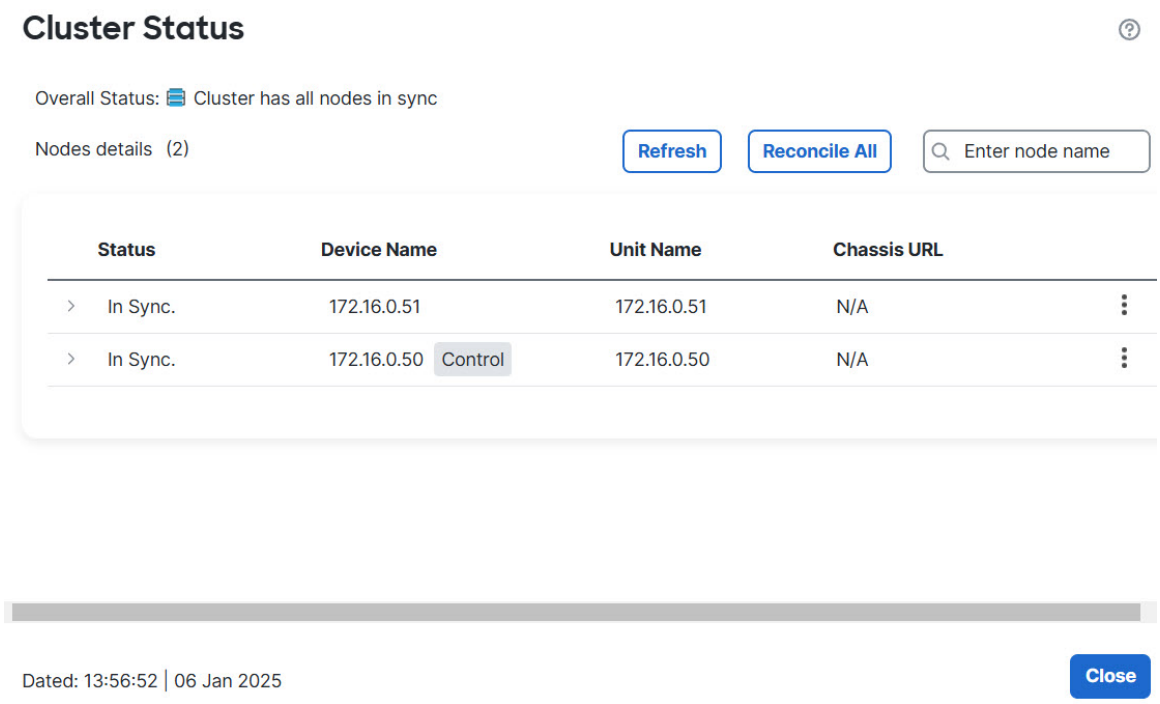
Monitor cluster status

You can monitor the cluster in the Cloud-Delivered Firewall Management Center and at the Firewall Threat Defense CLI.

Monitor cluster status through these methods:

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

Figure 24: Cluster status



The Control node has a graphic indicator identifying its role.

Cluster member **Status** includes these states:

- In Sync—The node is registered with the Cloud-Delivered Firewall Management Center.
- Pending Registration—The node is part of the cluster, but has not yet registered with the Cloud-Delivered Firewall 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 Cloud-Delivered Firewall 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 Cloud-Delivered Firewall Management Center.

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

Figure 25: Node summary

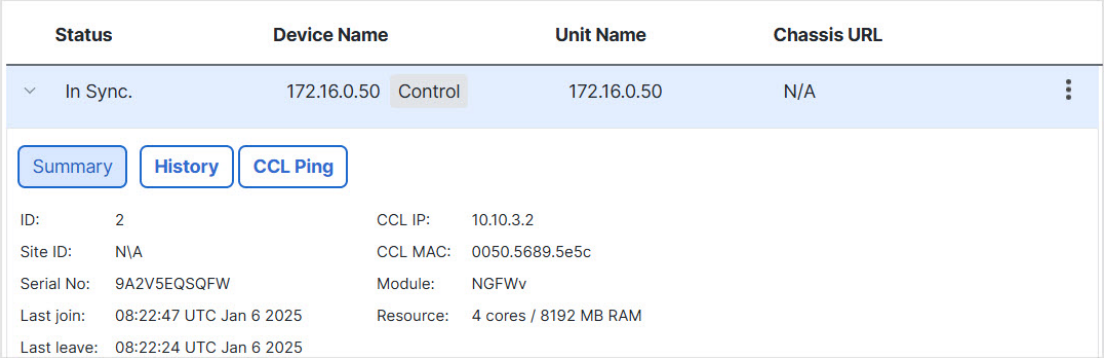
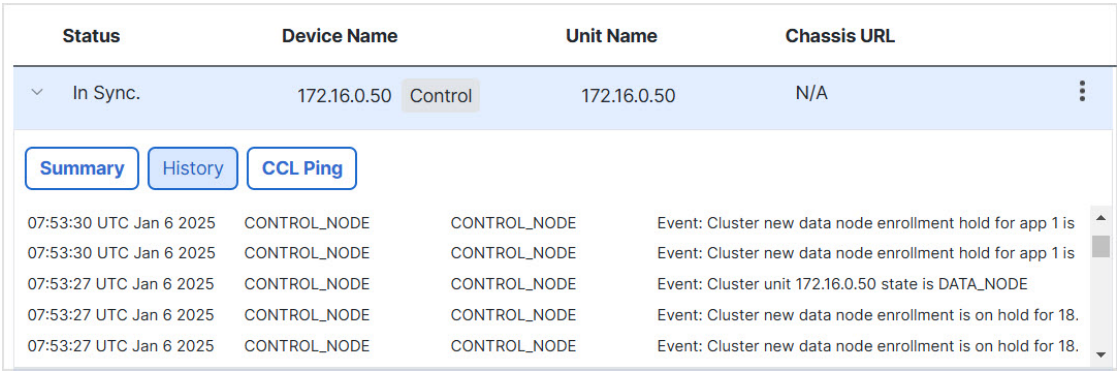


Figure 26: Node history



- **System** (⚙) > **Tasks** page.

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

- **Devices** > **Device Management** and then click **Add** > **Device** *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.

Cluster health monitor

When a Firewall Threat Defense is the control node of a cluster, the Cloud-Delivered Firewall Management Center collects various metrics periodically from the device metric data collector.

The cluster health monitor is comprised of these components:

- Overview dashboard: Displays information about the cluster topology, cluster statistics, and metric charts:
  - The topology section displays a cluster's live status, the health of individual threat defense, threat defense node type (control node or data node), and the status of the device. The status of the device could be *Disabled* (when the device leaves the cluster), *Added out of box* (in a public cloud cluster, the additional nodes that do not belong to the Cloud-Delivered Firewall Management Center), or *Normal* (ideal state of the node).
  - The cluster statistics section displays current metrics of the cluster with respect to the CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.



Note The CPU and memory metrics display the individual average of the data plane and snort usage.

  - The metric charts, namely, CPU Usage, Memory Usage, Throughput, and Connections, diagrammatically display the statistics of the cluster over the specified time period.

- Load Distribution dashboard: Displays load distribution across the cluster nodes in two widgets:
  - The Distribution widget displays the average packet and connection distribution over the time range across the cluster nodes. This data depicts how the load is being distributed by the nodes. Using this widget, you can easily identify any abnormalities in the load distribution and rectify it.
  - The Node Statistics widget displays the node level metrics in table format. It displays metric data on CPU usage, memory usage, input rate, output rate, active connections, and NAT translations across the cluster nodes. This table view enables you to correlate data and easily identify any discrepancies.
- Member Performance dashboard: Displays current metrics of the cluster nodes. You can use the selector to filter the nodes and view the details of a specific node. The metric data include CPU usage, memory usage, input rate, output rate, active connections, and NAT translations.
- CCL dashboard: Displays, graphically, the cluster control link data namely, the input, and output rate.
- Troubleshooting and Links: Provides convenient links to frequently used troubleshooting topics and procedures.
- Time range: An adjustable time window to constrain the information that appears in the various cluster metrics dashboards and widgets.
- Custom Dashboard: Displays data on both cluster-wide metrics and node-level metrics. However, node selection only applies for the threat defense metrics and not for the entire cluster to which the node belongs.

View the cluster health monitor

The cluster health monitor provides a detailed view of the health status of a cluster and its nodes. This cluster health monitor provides health status and trends of the cluster in an array of dashboards.

Before you begin

Ensure you have created a cluster from one or more devices in the Cloud-Delivered Firewall Management Center.

Follow these steps to view the cluster health monitor:

Procedure

Step 1 Choose **Troubleshooting** > **Health** > **Monitor**.

Use the Monitoring navigation pane to access node-specific health monitors.

Step 2 In the device list, click **Expand**(>) and **Collapse** (V) to expand and collapse the list of managed cluster devices.

Step 3 Click on the cluster name to view the cluster health statistics.

The cluster monitor reports health and performance metrics in several predefined dashboards by default. The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including its nodes, CPU, memory, input and output rates, connection statistics, and NAT translation information.
- Load Distribution: Traffic and packet distribution across the cluster nodes.
- Member Performance: Node-level statistics on CPU usage, memory usage, input throughput, output throughput, active connection, and NAT translation.
- CCL: Interface status and aggregate traffic statistics.

You can navigate through the various metrics dashboards by clicking on the labels. For a comprehensive list of the supported cluster metrics, refer to Cisco Secure Firewall Threat Defense Health Metrics.

Step 4 Configure the time range from the drop-down in the upper-right corner.

The time range can reflect a period as short as the last hour (the default) or as long as two weeks. Select **Custom** from the drop-down to configure a custom start and end date.

Click the refresh icon to set auto refresh to 5 minutes or to toggle off auto refresh.

Step 5 Click on deployment icon for a deployment overlay on the trend graph, with respect to the selected time range.

The deployment icon indicates the number of deployments during the selected time-range. A vertical band indicates the deployment start and end time. For multiple deployments, multiple bands/lines appear. Click on the icon on top of the dotted line to view the deployment details.

Step 6 View the **Health Alerts** for the node in the alert notification at the top of page, directly to the right of the device name.

Hover your pointer over the **Health Alerts** to view the health summary of the node. The popup window shows a truncated summary of the top five health alerts. Click on the popup to open a detailed view of the health alert summary.

Step 7 View device monitor health and performance metrics in several predefined dashboards.

The metrics dashboards include:

- Overview: Highlights key metrics from the other predefined dashboards, including CPU, memory, interfaces, connection statistics; plus disk usage and critical process information.
- CPU: CPU utilization, including the CPU usage by process and by physical cores.
- Memory: Device memory utilization, including data plane and Snort memory usage.
- Interfaces: Interface status and aggregate traffic statistics.
- Connections: Connection statistics (such as elephant flows, active connections, peak connections, and so on) and NAT translation counts.
- Snort: Statistics that are related to the Snort process.
- ASP drops: Statistics related to the dropped packets against various reasons.

You can navigate through the various metrics dashboards by clicking on the labels. Refer to Cisco Secure Firewall Threat Defense Health Metrics for a comprehensive list of the supported device metrics.

Step 8 Click the plus sign **Add New Dashboard**(+) in the upper right corner of the health monitor to create a custom dashboard by building your own variable set from the available metric groups.

For cluster-wide dashboard, choose Cluster metric group, and then choose the metric.

Cluster metrics

The cluster health monitor tracks statistics that are related to a cluster and its nodes, and aggregate of load distribution, performance, and CCL traffic statistics.

Table 1: Cluster Metrics

Metric	Description	Format
CPU	Average of CPU metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Memory	Average of memory metrics on the nodes of a cluster (individually for data plane and snort).	percentage
Data Throughput	Incoming and outgoing data traffic statistics for a cluster.	bytes
CCL Throughput	Incoming and outgoing CCL traffic statistics for a cluster.	bytes
Connections	Count of active connections in a cluster.	number

Metric	Description	Format
NAT Translations	Count of NAT translations for a cluster.	number
Distribution	Connection distribution count in the cluster for every second.	number
Packets	Packet distribution count in the cluster for every second.	number

Reference for clustering

This section includes more information about how clustering operates.

Threat Defense features and clustering

Some Firewall 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 and clustering

Centralized features for clustering

These 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, refer to the ASA general operations configuration guide. FlexConfig lets you configure many ASA features that are not present in the Cloud-Delivered Firewall Management Center GUI. For more information, refer to FlexConfig policies.

- These application inspections are supported:
  - DCERPC and ESMTP
  - NetBIOS and PPTP

  - RSH
  - SQLNET and 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.

Dynamic routing and clustering

Dynamic routing and clustering are network configuration approaches that

- enable each node to run the routing protocol as a standalone router in Individual interface mode,
- allow routes to be learned by each node independently, and
- require separate router IDs for each node when communicating with external routers.

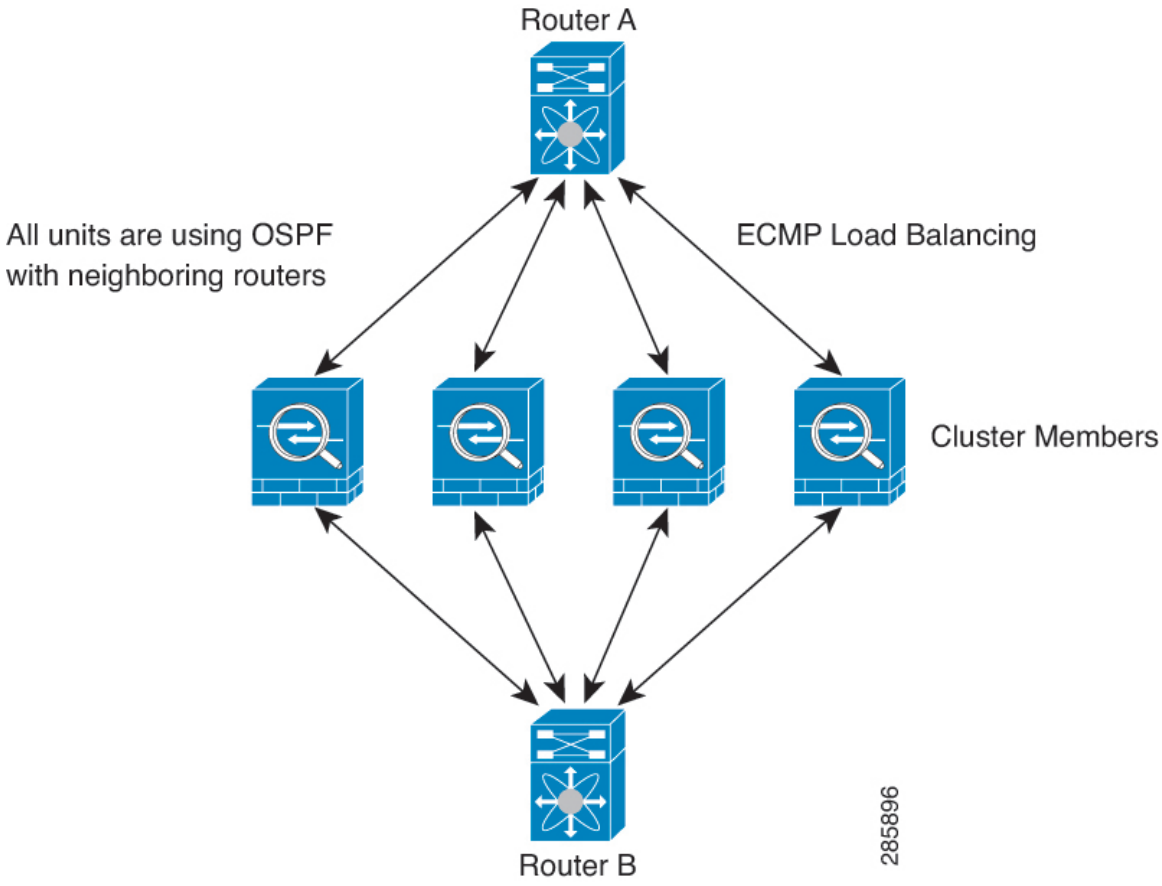
Routing behavior and load balancing

In Individual interface mode, ECMP is used to load balance traffic between multiple equal-cost paths. External routers can discover multiple paths to destinations through different cluster nodes.

Router A learns that there are 4 equal-cost paths to Router B, each through a node. ECMP is used to load balance traffic between the 4 paths. Each node picks a different router ID when communicating with external routers.

You must configure a cluster pool for the router ID so that each node has a separate router ID.

Figure 27: Dynamic routing in individual interface mode



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

Provides information about NAT behavior and limitations in clustering deployments.

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

- No Proxy ARP—For individual interfaces, a proxy ARP reply is never sent for mapped addresses. This prevents the adjacent router from maintaining a peer relationship with an ASA that may no longer be in the cluster. The upstream router requires a static route or policy-based routing (PBR) with Object Tracking for the mapped addresses, which must point to the main cluster IP address.
- No interface PAT on an individual interface—Interface PAT is not supported for Individual interfaces.
- 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.

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

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 redeploy 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

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.

For connections to an individual interface when using Policy-Based Routing (PBR) or Equal-Cost Multi-Path (ECMP), you must always connect to the Main cluster IP address, not a Local address.

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

Note Remote access VPN is not supported with clustering.

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.

Throughput calculation

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

The key components involved in the control node election process are:

- Cluster nodes: Communicate over the cluster control link and participate in the election process
- Election requests: Broadcast every three seconds by nodes when clustering is enabled
- Priority values: Set between 1 and 100 where 1 is the highest priority
- Control node: The elected node that coordinates cluster operations

These stages describe how cluster nodes elect a control node:

1. When clustering is enabled for a node (or when it first starts up with clustering already enabled), the node broadcasts an election request every three 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 where multiple control nodes exist temporarily, the node with the 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 configurable 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.

Split brain scenario

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.

Interface monitoring

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

All physical interfaces are monitored; only named interfaces can be monitored. You can optionally disable monitoring per interface.

A node is removed from the cluster if its monitored interfaces fail. The node is removed after 500 milliseconds.

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 Firewall Threat Defense automatically tries to rejoin the cluster, depending on the failure event.

Note When the Firewall Threat Defense becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the Management 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 Firewall Threat Defense automatically tries to rejoin every 5 minutes, indefinitely.
- Failed data interface—The Firewall 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 Firewall 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 Firewall 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 Cloud-Delivered Firewall 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 2: 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	—

Traffic	State Support	Notes
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.

Port Address Translation Connections

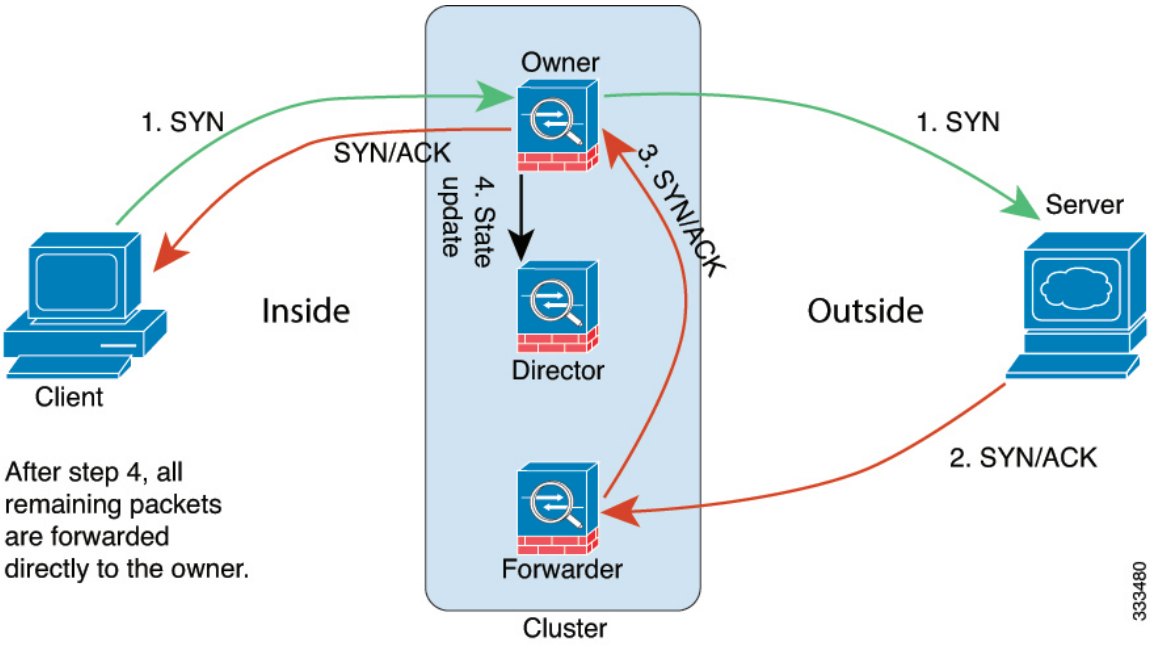
New connection ownership

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.

Traffic redirection is not supported in this release. When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. All the subsequent packets for the same connection should arrive the same node. If any connection packets arrive at a different node, they will be dropped. If a reverse flow arrives at a different node, it will be dropped as well. For centralized features, if the connections do not arrive on the control node, they will be dropped.

Sample data flow for TCP

The following example shows the establishment of a new connection.

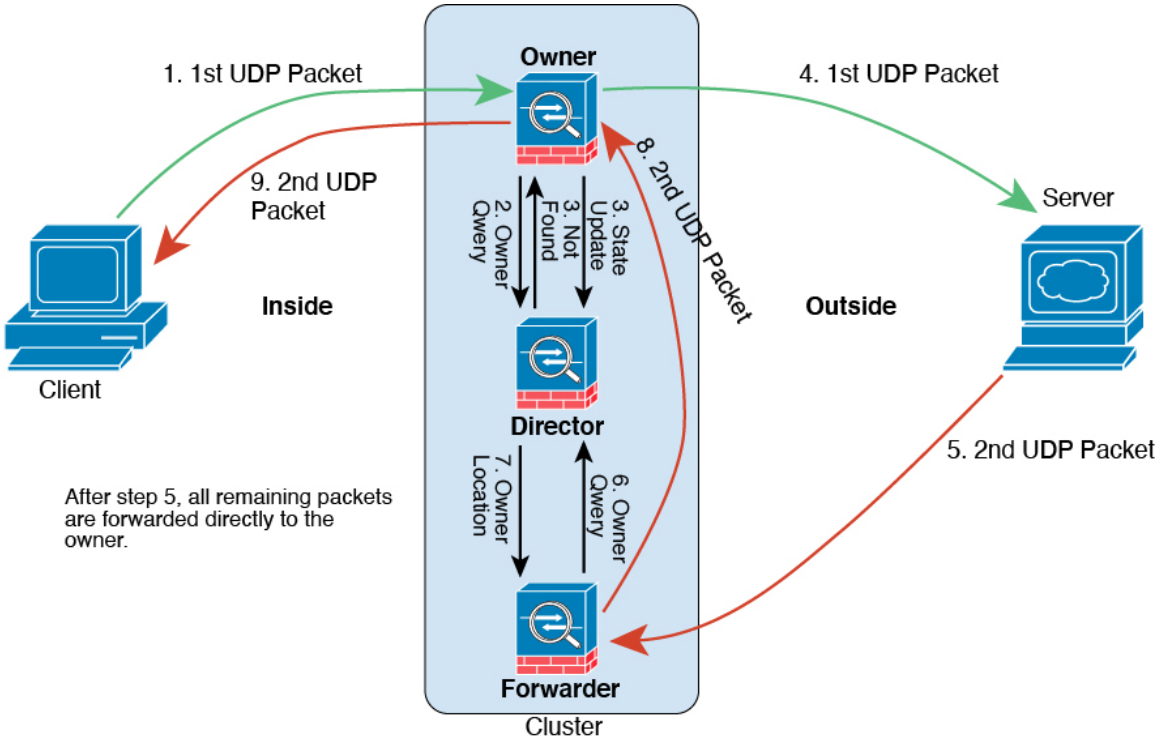


1. The SYN packet originates from the client and is delivered to one Firewall Threat Defense (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different Firewall Threat Defense (based on the load balancing method). This Firewall 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 28: ICMP and UDP Data Flow*



The first UDP packet originates from the client and is delivered to one Firewall 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: private cloud

This reference provides the feature history for clustering functionality in private cloud environments, including version information, platform compatibility, and detailed descriptions of enhancements and new capabilities.

Feature	Minimum Cloud-Delivered Firewall Management Center	Minimum Firewall Threat Defense	Details
MTU ping test from data node on node join	20241030	7.6.0	When a node joins the cluster, it checks MTU compatibility by sending a ping to the control node with a packet size twice the cluster control link MTU. Previously, only the control node sent a ping. If the ping fails, a notification is generated so you can fix the MTU mismatch on connecting switches and try again. Added/modified commands: **show cluster history**. See: Clustering for the Secure Firewall 3100/4200, Clustering for Threat Defense Virtual in a Private Cloud, Clustering for Threat Defense Virtual in a Public Cloud, Clustering for the Firepower 4100/9300
Cluster health monitor settings	20221213	Any	You can now edit cluster health monitor settings. New/Modified screens: **Devices** > **Device Management** > **Cluster** > **Cluster Health Monitor Settings** **Note** If you previously configured these settings using FlexConfig, be sure to remove the FlexConfig configuration before you deploy. Otherwise the FlexConfig configuration will overwrite the management center configuration.
Cluster health monitor dashboard	20221213	Any	You can now view cluster health on the cluster health monitor dashboard. New/Modified screens: **System** > **Health** > **Monitor**
Clustering for the Firewall Threat Defense Virtual on VMware and KVM	Any	7.2.0	The Firewall Threat Defense Virtual supports Individual interface clustering for up to 4 nodes on VMware and KVM. New/Modified screens: • **Devices** > **Device Management** > **Add Cluster** • **Devices** > **Device Management** > **More** menu • **Devices** > **Device Management** > **Cluster** Supported platforms: Firewall Threat Defense Virtual on VMware and KVM