

ASA Cluster for the ASA Virtual for the Private Cloud

Clustering lets you group multiple ASA Virtual's 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 the ASA Virtual clusters using:

- KVM
- VMware

Note Only routed firewall mode is supported.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 38.

- About ASA Virtual Clustering, on page 2
- Licenses for ASA Virtual Clustering, on page 8
- Requirements and Prerequisites for ASA Virtual Clustering, on page 8
- Guidelines for ASA Virtual Clustering, on page 9
- Configure the ASA Virtual Clustering Using a Day0 Configuration, on page 10
- Configure ASA Virtual Clustering after Deployment, on page 13
- Customize the Clustering Operation, on page 22
- Manage Cluster Nodes, on page 30
- Monitoring the ASA Virtual Cluster, on page 35
- Examples for ASA Virtual Clustering, on page 37
- Reference for Clustering, on page 38
- History for ASA Virtual Clustering, on page 53

About ASA Virtual Clustering

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 ASA Virtual send broadcast/multicast messages over the cluster control link.
- Management access to each firewall for configuration and monitoring. The ASA 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.



Note Layer 2 Spanned EtherChannels are not supported.

Cluster Nodes

Cluster nodes work together to accomplish the sharing of the security policy and traffic flows. This section describes the nature of each node role.

Bootstrap Configuration

On each device, you configure a minimal bootstrap configuration including the cluster name, cluster control link interface, and other cluster settings. The first node on which you enable clustering typically becomes the *control* node. When you enable clustering on subsequent nodes, they join the cluster as *data* nodes.

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 in the bootstrap configuration; the priority is set between 1 and 100, where 1 is the highest priority. All other members are data nodes. Typically, when you first create a cluster, the first node you add becomes the control node simply because it is the only node in the cluster so far.

You must perform all configuration (aside from the bootstrap configuration) on the control node only; the configuration is then replicated to the data nodes. In the case of physical assets, such as interfaces, the configuration of the control node is mirrored on all data nodes. For example, if you configure Ethernet 1/2 as the inside interface and Ethernet 1/1 as the outside interface, then these interfaces are also used on the data nodes as inside and outside interfaces.

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

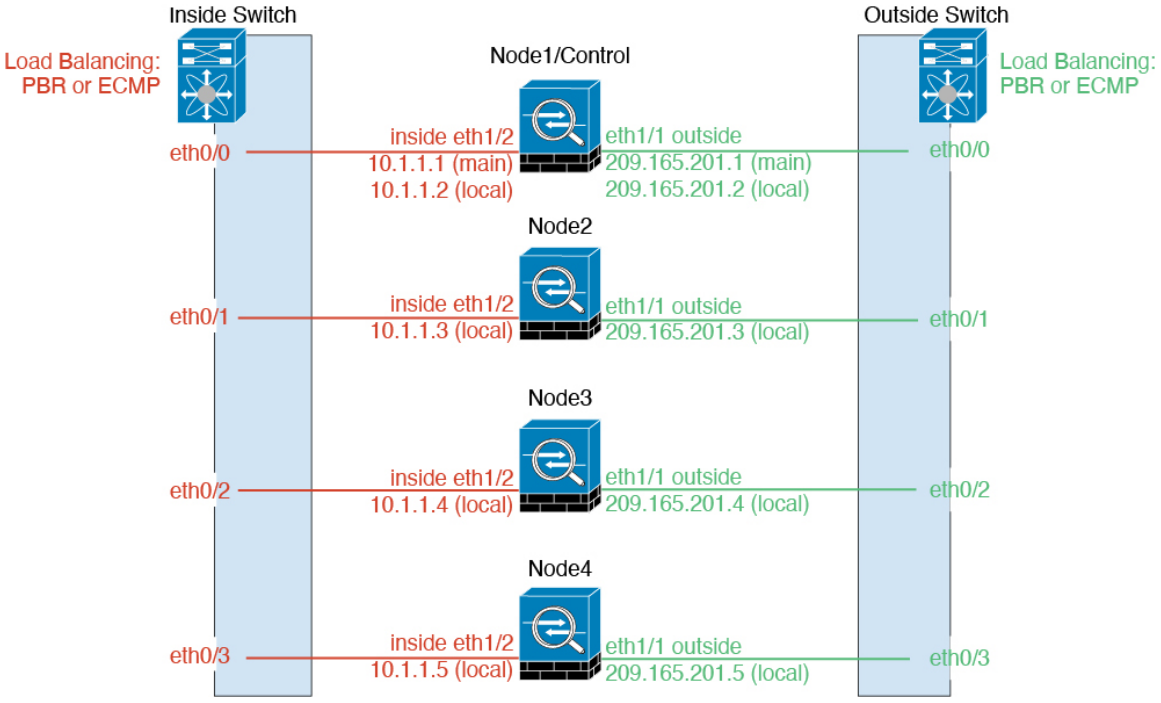
Individual Interfaces

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.

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

Load balancing must be configured separately on the upstream switch.



Note Layer 2 Spanned EtherChannels are not supported.

Policy-based routing

When using Individual interfaces, each ASA 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 ASAs 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 ASA.

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 ASA using ICMP ping. PBR can then enable or disable route maps based on reachability of a particular ASA.

Equal-cost multi-path routing

When using Individual interfaces, each ASA 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 ASA failure can cause problems; the route continues to be used, and traffic to the failed ASA 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 ASA 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 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 ASA 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, The ASA 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.

Cluster Control Link Failure

If the cluster control link line protocol goes down for a unit, then clustering is disabled; data interfaces are shut down. After you fix the cluster control link, you must manually rejoin the cluster by re-enabling clustering.

Note When the ASA Virtual becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from DHCP or the cluster IP pool. If you use a cluster IP pool, if you reload and the unit is still inactive in the cluster, then the management interface is not accessible (because it then uses the Main IP address, which is the same as the control node). You must use the console port (if available) for any further configuration.

Cluster control link MTU ping testing

Use Cluster control link MTU ping testing to

- verify MTU compatibility between cluster nodes during the join process
- use ping-based fragmentation tests to confirm the network path supports proper packet handling, and

- allow nodes to join even if tests fail while providing diagnostic messages for troubleshooting MTU mismatches.

When a node joins the cluster, the control node send a ping with a payload size of twice the MTU. This process tests the network's ability to handle packet fragmentation because the underlying IP layer will fragment packets that exceed the MTU limit.

A successful ping confirms that the network path supports proper fragmentation and that the cluster control link can reliably process traffic at the configured MTU size.

If the ping fails, view these messages:

- **show cluster history**—`Event: CCL MTU test to unit name failed`
- Console WARNING—`WARNING: Unit name is not reachable in CCL jumbo frame ICMP test, please check cluster interface and switch MTU configuration`

Even if the ping fails, the node is allowed to join the cluster. In this case, you need to resolve the MTU mismatch as soon as possible.

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.

ASA Virtual Cluster Management

One of the benefits of using ASA Virtual clustering is the ease of management. This section describes how to manage the cluster.

Management Network

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

Management Interface

Use the Management 0/0 interface for management.

Note You cannot enable dynamic routing for the management interface. You must use a static route.

You can use either static addressing or DHCP for the management IP address.

If you use static addressing, you can use a Main cluster IP address that is a fixed address for the cluster that always belongs to the current control node. For each interface, you also configure a range of addresses so that each node, including the current control node, can use a Local address from the range. The Main cluster IP address provides consistent management access to an address; when a control node changes, the Main cluster IP address moves to the new control node, so management of the cluster continues seamlessly. The Local IP address is used for routing, and is also useful for troubleshooting. For example, you can manage the cluster

by connecting to the Main cluster IP address, which is always attached to the current control node. To manage an individual member, you can connect to the Local IP address. For outbound management traffic such as TFTP or syslog, each node, including the control node, uses the Local IP address to connect to the server.

If you use DHCP, you do not use a pool of Local addresses or have a Main cluster IP address.



Note To-the-box traffic needs to be directed to the node's management IP address; to-the-box traffic is not forwarded over the cluster control link to any other node.

Control Node Management Vs. Data Node Management

All management and monitoring can take place on the control node. From the control node, you can check runtime statistics, resource usage, or other monitoring information of all nodes. You can also issue a command to all nodes in the cluster, and replicate the console messages from data nodes to the control node.

You can monitor data nodes directly if desired. Although also available from the control node, you can perform file management on data nodes (including backing up the configuration and updating images). The following functions are not available from the control node:

- Monitoring per-node cluster-specific statistics.
- Syslog monitoring per node (except for syslogs sent to the console when console replication is enabled).
- SNMP
- NetFlow

Crypto Key Replication

When you create a crypto key on the control node, the key is replicated to all data nodes. If you have an SSH session to the Main cluster IP address, you will be disconnected if the control node fails. The new control node uses the same key for SSH connections, so that you do not need to update the cached SSH host key when you reconnect to the new control node.

ASDM Connection Certificate IP Address Mismatch

By default, a self-signed certificate is used for the ASDM connection based on the Local IP address. If you connect to the Main cluster IP address using ASDM, then a warning message about a mismatched IP address might appear because the certificate uses the Local IP address, and not the Main cluster IP address. You can ignore the message and establish the ASDM connection. However, to avoid this type of warning, you can enroll a certificate that contains the Main cluster IP address and all the Local IP addresses from the IP address pool. You can then use this certificate for each cluster member. See https://www.cisco.com/c/en/us/td/docs/security/asdm/identity-cert/cert-install.html for more information.

Inter-Site Clustering

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

You can configure each cluster chassis to belong to a separate site ID. Site IDs are used to enable flow mobility using LISP inspection, director localization to improve performance and reduce round-trip time latency for

inter-site clustering for data centers, and site redundancy for connections where a backup owner of a traffic flow is always at a different site from the owner.

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

- Sizing the Data Center Interconnect—Requirements and Prerequisites for ASA Virtual Clustering, on page 8
- Inter-Site Guidelines—Guidelines for ASA Virtual Clustering, on page 9
- Configure Cluster Flow Mobility—Configure Cluster Flow Mobility, on page 27
- Enable Director Localization—Configure Basic ASA Cluster Parameters, on page 22
- Enable Site Redundancy—Configure Basic ASA Cluster Parameters, on page 22
- Inter-Site Examples—Individual Interface Routed Mode North-South Inter-Site Example, on page 37

Licenses for ASA Virtual Clustering

Each cluster node requires the same model license. We recommend using the same number of CPUs and memory for all nodes, or else peformance 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.



Note If you deregister the ASA Virtual so that it is unlicensed, then it will revert to a severely rate-limited state if you reload the ASA Virtual. An unlicensed, low performing cluster node will impact the performance of the entire cluster negatively. Be sure to keep all cluster nodes licensed, or remove any unlicensed nodes.

Requirements and Prerequisites for ASA Virtual Clustering

Model Requirements

- ASAv30, ASAv50, ASAv100
- The following private cloud services:
  - KVM
  - VMware
- A maximum of 16 nodes in a cluster on *two* hosts in a 2x8 deployment configuration. We recommend you to deploy a maximum of *eight* ASAvs on each of the *two* hosts (2x8), which results in a cluster of 16 nodes.

ASA Virtual Platform and Software Requirements

All nodes in a cluster:

- Must be the same model. We recommend using the same number of CPUs and memory for all nodes, or else peformance will be limited on all nodes to match the least capable node.
- Must run the identical software except at the time of an image upgrade. Hitless upgrade is supported. Mismatched software versions can lead to poor performance, so be sure to upgrade all nodes in the same maintenance window.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the control node for initial cluster control link communication before configuration replication.

Guidelines for ASA Virtual Clustering

Failover

Failover is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Additional Guidelines

- When significant topology changes occur (such as enabling or disabling an interface on the ASA or the switch, adding an additional switch to form a VSS or vPC) you should disable the 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 interface health check feature.
- When adding a node to an existing cluster, or when reloading a node, 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.
- We do not support VXLANs for data interfaces; only the cluster control link supports VXLAN.
- It takes time to replicate changes to all the nodes in a cluster. If you make a large change, for example, adding an access control rule that uses object groups (which, when deployed, are broken out into multiple rules), the time needed to complete the change can exceed the timeout for the cluster nodes to respond with a success message. If this happens, you might see a "failed to replicate command" message. You can ignore the message.

Defaults for ASA Virtual Clustering

- 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 rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.

- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure the ASA Virtual Clustering Using a Day0 Configuration

Control Node Day0 Configuration

The following Day0 configuration for the control node includes the bootstrap configuration followed by interface configuration that will be replicated to the data nodes. Bold text shows the values you need to change for the data node Day0 configuration.



Note This configuration only includes the cluster-centric configuration. Your Day0 configuration should also include other settings for licensing, SSH access, ASDM access and more. See the getting started guide for more information about Day0 configurations.

```
!BOOTSTRAP
! Cluster interface mode
cluster interface mode individual
!
! VXLAN peer group
object-group network cluster-peers
network-object host 10.6.6.51
network-object host 10.6.6.52
network-object host 10.6.6.53
network-object host 10.6.6.54
!
! Alternate object group representation
! object-network xyz
! range 10.6.6.51 10.6.6.54
! object-group network cluster-peers
! network-object object xyz
!
! Cluster control link physical interface (VXLAN tunnel endpoint (VTEP) src interface)
interface gigabitethernet 0/7
description CCL VTEP src ifc
nve-only cluster
nameif ccl
security-level 0
ip address 10.6.6.51 255.255.255.0
no shutdown
!
! VXLAN Network Identifier (VNI) interface
interface vni1
segment-id 1
vtep-nve 1
!
! Set the CCL MTU
mtu ccl 1654
!
! Network Virtualization Endpoint (NVE) association with VTEP src interface
nve 1
encapsulation vxlan
source-interface ccl
peer-group cluster-peers
!
```

```
! Management Interface Using DHCP
interface management 0/0
nameif management
ip address dhcp setroute
no shutdown
!
! Alternate Management Using Static IP
! ip local pool mgmt_pool 10.1.1.1 10.10.10.4
! interface management 0/0
! nameif management
! ip address 10.1.1.25 255.255.255.0 cluster-pool mgmt_pool
! no shutdown
!
! Cluster Config
cluster group cluster1
local-unit A
cluster-interface vni1 ip 10.2.2.1 255.255.255.0
priority 1
enable noconfirm
!
! INTERFACES
!
ip local pool inside_pool 10.10.10.11 10.10.10.14
ip local pool outside_pool 10.11.11.11 10.11.11.14
!
interface GigabitEthernet0/1
nameif inside
security-level 100
ip address 10.10.10.10 255.255.255.0 cluster-pool inside_pool
!
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 10.11.11.10 255.255.255.0 cluster-pool outside_pool
!
!JUMBO FRAME RESERVATION for CCL MTU
jumbo-frame reservation
```

Data Node Day0 Configuration

The following Day0 configuration for the data node includes only the bootstrap configuration. Bold text shows the values you need to change from the control node Day0 configuration.

Note This configuration only includes the cluster-centric configuration. Your Day0 configuration should also include other settings for licensing, SSH access, ASDM access and more. See the getting started guide for more information about Day0 configurations.

```
!BOOTSTRAP
! Cluster interface mode
cluster interface mode individual
!
! VXLAN peer group
object-group network cluster-peers
network-object host 10.6.6.51
network-object host 10.6.6.52
network-object host 10.6.6.53
network-object host 10.6.6.54
!
```

```
! Alternate object group representation
! object-network xyz
! range 10.6.6.51 10.6.6.54
! object-group network cluster-peers
! network-object object xyz
!
! Cluster control link physical interface (VXLAN tunnel endpoint (VTEP) src interface)
interface gigabitethernet 0/7
description CCL VTEP src ifc
nve-only cluster
nameif ccl
security-level 0
ip address 10.6.6.52 255.255.255.0
no shutdown
!
! VXLAN Network Identifier (VNI) interface
interface vni1
segment-id 1
vtep-nve 1
!
! Set the CCL MTU
mtu ccl 1654
!
! Network Virtualization Endpoint (NVE) association with VTEP src interface
nve 1
encapsulation vxlan
source-interface ccl
peer-group cluster-peers
!
! Management Interface Using DHCP
interface management 0/0
nameif management
ip address dhcp setroute
no shutdown
!
! Alternate Management Using Static IP
! ip local pool mgmt_pool 10.1.1.1 10.10.10.4
! interface management 0/0
! nameif management
! ip address 10.1.1.25 255.255.255.0 cluster-pool mgmt_pool
! no shutdown
!
! Cluster Config
cluster group cluster1
local-unit B
cluster-interface vni1 ip 10.2.2.2 255.255.255.0
priority 2
enable noconfirm
!
! INTERFACES
!
ip local pool inside_pool 10.10.10.11 10.10.10.14
ip local pool outside_pool 10.11.11.11 10.11.11.14
!
interface GigabitEthernet0/1
nameif inside
security-level 100
ip address 10.10.10.10 255.255.255.0 cluster-pool inside_pool
!
interface GigabitEthernet0/0
nameif outside
security-level 0
ip address 10.11.11.10 255.255.255.0 cluster-pool outside_pool
!
```

```
!JUMBO FRAME RESERVATION for CCL MTU
jumbo-frame reservation
```

Configure ASA Virtual Clustering after Deployment

To configure clustering after you deploy your ASA Virtuals, perform the following tasks.

Back Up Your Configurations (Recommended)

When you enable clustering on a data unit, the current configuration is replaced with one synced from the active unit. If you ever want to leave the cluster entirely, it may be useful to have a backup configuration with a usable management interface configuration.

Before you begin

Perform a backup on each unit.

Procedure

Step 1 Choose **Tools** > **Backup Configurations**.

Step 2 Back up at least the running configuration. See Back Up and Restore Configurations or Other Files for a detailed procedure.

Configure Interface Settings

Configure the cluster interface mode as well as interfaces on the control node. The interface configuration will be replicated to data nodes when they join the cluster. Note that configuration of the cluster control link is covered in the bootstrap configuration procedure.

Configure the Cluster Interface Mode on the Control Node

Before you enable clustering, you need to convert the firewall to use Individual interfaces. Because clustering limits the types of interfaces you can use, this process lets you check your existing configuration for incompatible interfaces and then prevents you from configuring any unsupported interfaces.

Note If you do not add data nodes from the control node, you must set the interface mode manually on all nodes according to this section, not just the control node; if you add data nodes from the control node, ASDM sets the interface mode automatically on the data node.

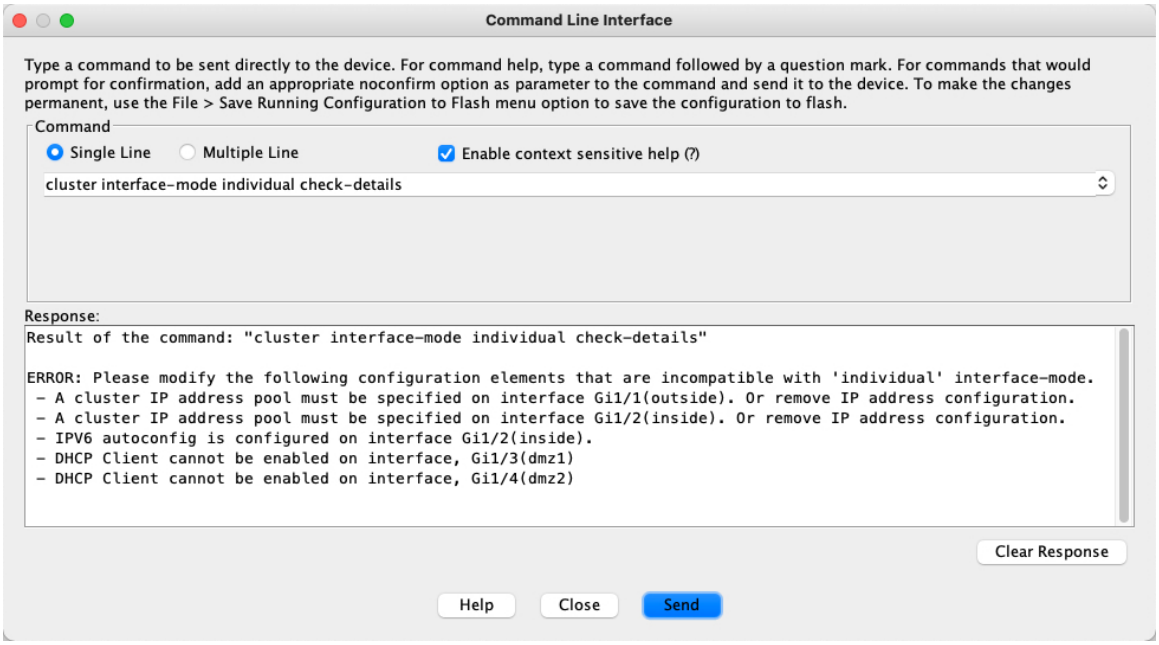
Procedure

Step 1 In ASDM on the control node, choose **Tools** > **Command Line Interface**.Show any incompatible configuration so that you can force the interface mode and fix your configuration later; the mode is not changed with this command:

cluster interface-mode individual check-details

Example:

Figure 1: Command Line Interface Output



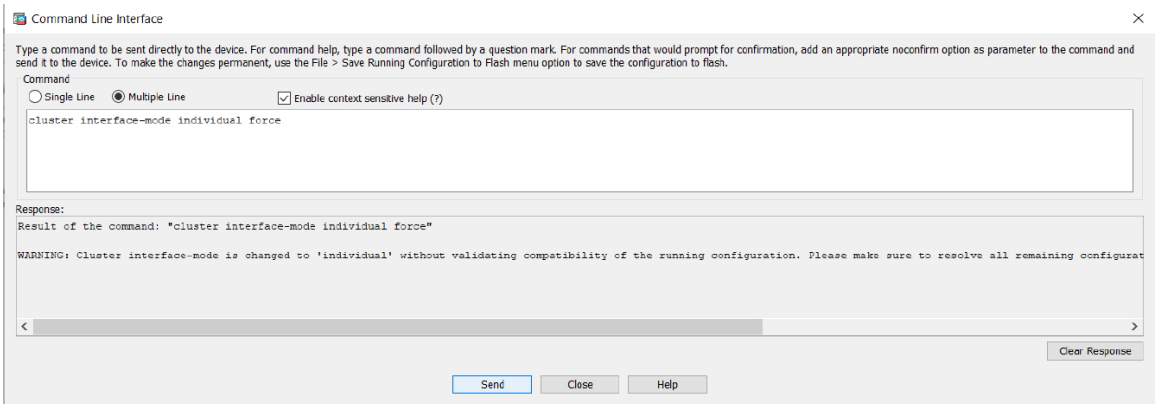
Caution
After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool or getting the IP address from DHCP), you will not be able to reconnect because cluster-incompatible interface configuration is removed. In that case, you will have to connect to the console port, if available, to fix the interface configuration.

Step 2 Set the interface mode for clustering:

cluster interface-mode individual force

Example:

Figure 2: Set the Interface Mode



There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering.

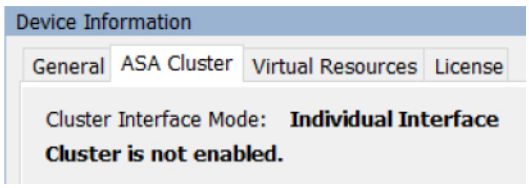
The **force** option changes the mode without checking your configuration for incompatible settings. You need to manually fix any configuration issues after you change the mode. Because any interface configuration can only be fixed after you set the mode, we recommend using the **force** option so that you can at least start from the existing configuration. You can re-run the **check-details** option after you set the mode for more guidance.

Without the **force** option, if there is any incompatible configuration, you are prompted to clear your configuration and reload, thus requiring you to connect to the console port (if available) to reconfigure your management access. If your configuration is compatible (rare), the mode is changed and the configuration is preserved. If you do not want to clear your configuration, you can exit the command by typing **n**.

To remove the interface mode, enter the **no cluster interface-mode** command.

Step 3 Quit ASDM and reload. ASDM needs to be restarted to correctly account for the cluster interface mode. After you reload, you see the ASA Cluster tab on the home page:

Figure 3: ASDM Needs Updating



Configure the Cluster Control Link on the Control Node

Configure a VXLAN interface for the cluster control link interface before you run the wizard. For more information about VXLAN and the cluster control link, see Cluster Control Link, on page 4.

Before you begin

Enable jumbo frame reservation for use with the cluster control link, so you can set the cluster control link MTU to the recommended value. Enabling jumbo frames causes the ASA to reload. See the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** screen.

Note You must enable jumbo frame reservation on each node separately.

Procedure

Step 1 Identify the VXLAN tunnel endpoint (VTEP) peer IP addresses in a network object group.

See the **Configuration** > **Firewall** > **Objects** > **Network Objects/Groups** page, and the "Objects for Access Control" chapter in the ASA firewall configuration guide for more information about network object groups.

The underlying IP network between VTEPs is independent of the cluster control link network that the VXLAN Network Identifier (VNI) interfaces use. 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.

Step 2 Configure the VTEP source interface.

The VTEP source interface is a regular ASA 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.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and edit the interface you want to use for the VTEP source interface.
b) Configure the **Interface Name**.
c) Check the **VTEP Source Interface (cluster)** check box.
d) Check **Enable Interface**.
e) Configure a static IPv4 address.

   The IP address should be included as one of the peers in the network object group.

f) Click the **Advanced** tab, and set the **MTU** to be at least 154 bytes higher than the highest MTU of the data interfaces.

   Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation, which required a reload.

   For example, when using jumbo frames, 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.

g) Click **OK**.

Step 3 Associate the VTEP source interface with the Network Virtualization Endpoint (NVE) instance.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **VXLAN**.
b) (Optional) Enter a **VXLAN Destination Port** value if you want to change from the default 4789.
c) Check the **Enable Network Virtualization Endpoint encapsulation using VXLAN** check box.
d) Choose the **VTEP Tunnel Interface** from the drop-down list.
e) Check the **Configure Packet Recipient** check box, click the **Peer Group** radio button, and choose the peer group you created.
f) Click **Apply**.

Step 4 Create the 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.

a) Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**, and click **Add** > **VNI Interface**.

b) Enter the **VNI ID**, between 1 and 10000.

This ID is just an internal interface identifier.

c) Enter the **VNI Segment ID**, between 1 and 16777215.

The segment ID is used for VXLAN tagging.

d) Check the **NVE Mapped to VTEP Interface** check box.

This setting associates the VNI interface with the VTEP source interface.

e) Click **OK**, and then **Apply**.

Configure Individual Interfaces

You must modify any interface that is currently configured with an IP address to be cluster-ready before you enable clustering. At a minimum, you may need to modify the management interface to which ASDM is currently connected when you use a static IP address for management. For other interfaces, you can configure them before or after you enable clustering; we recommend pre-configuring all of your interfaces so that the complete configuration is synced to new cluster nodes.

This section describes how to configure interfaces to be Individual interfaces compatible with clustering. 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.

For the Management interface, you can configure an IP address pool or you can use DHCP; only the Management interface supports getting an address from DHCP. To use DHCP, do not use this procedure; instead configure it as usual (see Configure General Routed Mode Interface Parameters).

Before you begin

- (Optional) Configure subinterfaces.
- For the management interface, you can use a static address or you can use DHCP. If you are using static IP addresses and connecting remotely to the management interface using ASDM, the current IP address of prospective data nodes are for temporary use.
  - Each member will be assigned an IP address from the cluster IP pool defined on the control node.
  - The cluster IP pool cannot include addresses already in use on the network, including prospective secondary IP addresses.

    For example:

    1. You configure the control node to use 10.1.1.1.
    2. Other nodes use 10.1.1.2, 10.1.1.3, and 10.1.1.4.

3. When you configure the cluster IP pool on the control node, you cannot include the .2, .3, or .4 addresses in the pool, because they are in use.
4. Instead, you need to use other IP addresses on the network, such as .5, .6, .7, and .8.

Note The pool needs as many addresses as there are members of the cluster, including the control node; the original .1 address is the main cluster IP address that belongs to the current control node.

5. After you join the cluster, the old, temporary addresses are relinquished and can be used elsewhere.

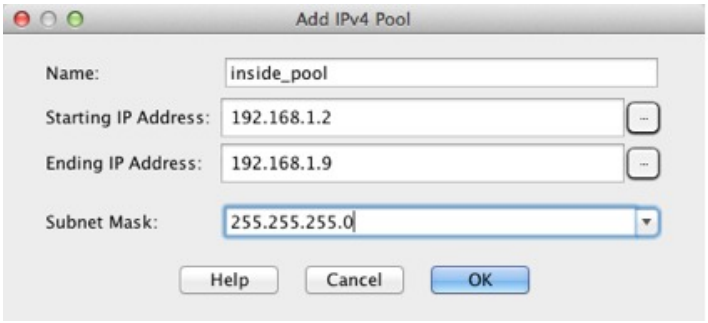
Procedure

Step 1 Choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.

Step 2 Choose the interface row, and click **Edit**. Choose **Use Static IP**. DHCP and PPPoE are not supported.

Step 3 To add the IPv4 cluster IP pool, MAC address pool, and site-specific MAC addresses, click the **Advanced** tab and set **ASA Cluster** area parameters.

a) Create a cluster IP pool by clicking the **...** button next to the **IP Address Pool** field. The valid range shown is determined by the Main IP address you set on the General tab.
b) Click **Add**.
c) Configure a range of addresses that does not include the Main cluster IP address, and that does not include any addresses currently in-use on your network. You should make the range large enough for the size of the cluster, for example, 8 addresses.



d) Click **OK** to create the new pool.
e) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool name appears in the **IP Address Pool** field.

f) (Optional) (Optional) Configure a **MAC Address Pool** if you want to manually configure MAC addresses.

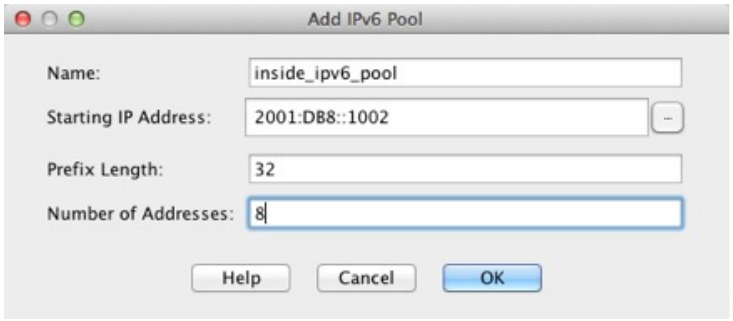
Step 4 To configure an IPv6 address, click the **IPv6** tab.

a) Check the **Enable IPv6** check box.
b) In the **Interface IPv6 Addresses** area, click **Add**.

The **Enable address autoconfiguration** option is not supported. Manually configuring the link-local address is also not supported.

The **Add IPv6 Address for Interface** dialog box appears.

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:0DB8::BA98:0:3210/48.
d) Click the **...** button to configure the cluster IP pool.
e) Click **Add**.



f) Configure the starting IP address (network prefix), prefix length, and number of addresses in the pool.
g) Click **OK** to create the new pool.
h) Select the new pool you created, and click **Assign**, and then click **OK**.

The pool appears in the **ASA Cluster IP Pool** field.

i) Click **OK**.

Step 5 Click **OK** to return to the Interfaces pane.

Step 6 Click **Apply**.

Create or Join a Cluster Using the High Availability Wizard

Each node in the cluster requires a bootstrap configuration to join the cluster. Run the High Availability and Scalability wizard on one node (that will become the control node) to create the cluster, and then add data nodes to it.

Before you begin

- The VXLAN VTEP source interface you intend to use for the cluster control link interface must be in an up state on the connected switch.
- When you add a node to a running cluster, you may see temporary, limited packet/connection drops; this is expected behavior.

Procedure

Step 1 Choose **Wizards** > **High Availability and Scalability Wizard**. See select wizard guidelines in the following steps.

Step 2 On the ASA Cluster Configuration screen, configure bootstrap settings including:

- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Shared Key**—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. If enabled, ASAs in a cluster exchange load information periodically, and offload new connections from more loaded devices to less loaded devices. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.

  Note
  Do not configure connection rebalancing for inter-site topologies; you do not want connections rebalanced to cluster members at a different site.

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature. To determine node health, the ASA cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the holdtime period, the peer node is considered unresponsive or dead.

  Note
  When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you must disable the health check 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 health check.

  - **Time to Wait Before Device Considered Failed**—This value determines the amount of time between node keepalive status messages, between .3 and 45 seconds; The default is 3 seconds.

- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- **Cluster Control Link**—Specifies the cluster control link interface.
  - **MTU**—Specify the maximum transmission unit for the VTEP source interface to be at least 154 bytes higher than the highest MTU of the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation. For example, when using jumbo frames, 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. This parameter

is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
Note: If you have not pre-enabled jumbo frame reservation, you should quit the wizard, enable jumbo frames, and then restart this procedure.

Step 3 On the **Interfaces for Health Monitoring** screen, you can exempt some interfaces from monitoring for failure. You might want to disable health monitoring of non-essential interfaces, for example, the management interface.

Note
When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you must disable the health check 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 health check.

Step 4 On the **Interface Auto Rejoin settings** screen, customize the auto-rejoin settings in case of an interface or cluster control link failure. For each type, you can set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Step 5 Click **Finish**.

Step 6 The ASA scans the running configuration for incompatible commands for features that are not supported with clustering, including commands that may be present in the default configuration. Click **OK** to delete the incompatible commands. If you click **Cancel**, then clustering is not enabled.

After a period of time while ASDM enables clustering and reconnects to the ASA, the Information screen appears confirming that the ASA was added to the cluster.

Note
In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact error condition that disabled clustering; for example, the cluster control link might be down.

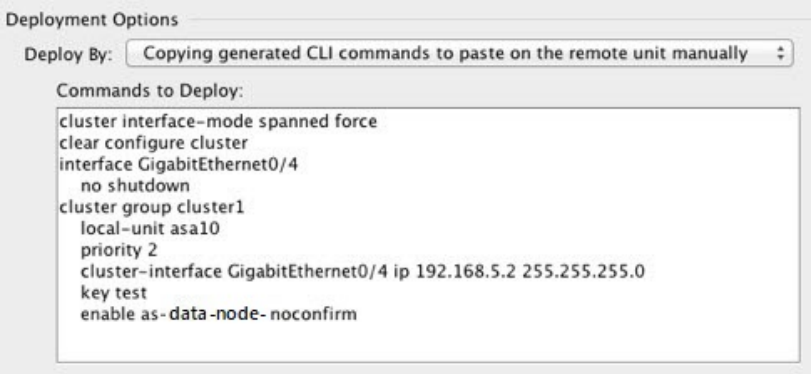
Step 7 To add a data node, click **Yes**.

If you are re-running the wizard from the control node, you can add data nodes by choosing the **Add another member to the cluster** option when you first start the wizard.

Step 8 In the **Deployment Options** area, choose one of the following **Deploy By** options:

- **Sending CLI commands to the remote unit now—**Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.

- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Customize the Clustering Operation

You can customize clustering health monitoring, TCP connection replication delay, flow mobility and other optimizations, either as part of the Day 0 configuration or after you deploy the cluster.

Perform these procedures on the control node.

Configure Basic ASA Cluster Parameters

You can customize cluster settings on the control node. If you do not use the wizard to add a node to the cluster, you can configure the cluster parameters manually. If you already enabled clustering, you can edit some cluster parameters; others that cannot be edited while clustering is enabled are grayed out. This procedure also includes advanced parameters that are not included in the wizard.

Before you begin

- If you did not use the wizard, and want to manually join the cluster, you need to pre-configure the cluster control link on each node before joining the cluster. See Configure the Cluster Control Link on the Control Node, on page 15.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.

If your device is already in the cluster, and is the control node, then this pane is on the **Cluster Configuration** tab.

Step 2 Check the **Configure ASA cluster settings** check box.

If you uncheck the check box, the settings are erased. Do not check **Participate in ASA cluster** until after you set all your parameters.

Note
After you enable clustering, do not uncheck the **Configure ASA cluster settings** check box without understanding the consequences. This action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Configure the following bootstrap parameters:

- **Cluster Name**—Names the cluster. The name must be an ASCII string from 1 to 38 characters. You can only configure one cluster per node. All members of the cluster must use the same name.
- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Site Index**—If you use inter-site clustering, set the site ID for this node so it uses a site-specific MAC address, between 1 and 8.
- (Optional) **Shared Key**—Sets an encryption key for control traffic on the cluster control link. The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the encryption key. This parameter does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear. You must configure this parameter if you also enable the password encryption service.
- (Optional) **Enable connection rebalancing for TCP traffic across all the ASAs in the cluster**—Enables connection rebalancing. This parameter is disabled by default. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. If enabled, ASAs exchange information about the connections per second periodically, and offload new connections from devices with more connections per second to less loaded devices. Existing connections are never moved. Moreover, because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal. The frequency, between 1 and 360 seconds, specifies how often the load information is exchanged. The default is 5 seconds.

  Once a connection is offloaded to a different node, it becomes an asymmetric connection.

  Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.
- **Enable cluster load monitor**—You can monitor the traffic load for cluster members, including total connection count, CPU and memory usage, and buffer drops. If the load is too high, you can choose to manually disable clustering on the node if the remaining nodes can handle the load, or adjust the load balancing on the external switch. This feature is enabled by default. You can periodically monitor the traffic load. If the load is too high, you can choose to manually disable clustering on the node.

  Set the following values:

  - **Time Interval**—Sets the time in seconds between monitoring messages, between 10 and 360 seconds. The default is 20 seconds.
  - **Number of Intervals**—Sets the number of intervals for which the ASA maintains data, between 1 and 60. The default is 30.

See **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring** to view the traffic load.

- (Optional) **Enable health monitoring of this device within the cluster**—Enables the cluster node health check feature, and determines the amount of time between node heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds. **Note**: When you are adding new nodes to the cluster, and making topology changes on the ASA or the switch, you should disable this feature temporarily until the cluster is complete, and also disable interface monitoring for the disabled interfaces (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**). You can re-enable this feature after cluster and topology changes are complete. To determine node health, the ASA cluster nodes send heartbeat messages on the cluster control link to other nodes. If a node does not receive any heartbeat messages from a peer node within the holdtime period, the peer node is considered unresponsive or dead.
- (Optional) **Debounce Time**—Configures the debounce time before the ASA considers an interface to be failed and the node is removed from the cluster. This feature allows 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 ASA waits the number of milliseconds specified before marking the interface as failed and the node is removed from the cluster. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.
- (Optional) **Replicate console output**—Enables console replication from data nodes to the control node. This feature is disabled by default. The ASA may print out some messages directly to the console for certain critical events. If you enable console replication, data nodes send the console messages to the control node so that you only need to monitor one console port for the cluster. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.
- (Optional) **Enable Clustering Flow Mobility**. See Configure LISP Inspection, on page 28.
- (Optional) **Enable Director Localization for inter-DC cluster**—To improve performance and reduce round-trip time latency for inter-site clustering for data centers, you can enable director localization. New connections are typically load-balanced and owned by cluster members within a given site. However, the ASA assigns the Director role to a member at *any* site. Director localization enables additional Director roles: a Local Director at the same site as the Owner, and a Global Director that can be at any site. Keeping the Owner and Director at the same site improves performance. Also, if the original Owner fails, the Local Director will choose a new connection Owner at the same site. The Global Director is used if a cluster member receives packets for a connection that is owned on a different site.
- (Optional) **Site Redundancy**—To protect flows from a site failure, you can enable site redundancy. If the connection backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Director localization and site redundancy are separate features; you can configure one or the other, or configure both.
- (Optional) **Enable config sync acceleration**—When a data node has the same configuration as the control node, it will skip syncing the configuration and will join faster. This feature is enabled by default. This feature is configured on each node, and is not replicated from the control node to the data node.

  Note
  Some configuration commands are not compatible with accelerated cluster joining; if these commands are present on the node, even if accelerated cluster joining is enabled, configuration syncing will always occur. You must remove the incompatible configuration for accelerated cluster joining to work. Use the **show cluster info unit-join-acceleration incompatible-config** to view incompatible configuration.

- **Enable parallel configuration replicate**—Enable the control node to sync configuration changes with data nodes in parallel. Otherwise, synching occurs sequentially, and can take more time.

- **Flow State Refresh Keepalive Interval**—Set the keepalive interval for flow state refresh messages (clu_keepalive and clu_update messages) from the flow owner to the director and backup owner, between 15 and 20 seconds. The default is 15. You may want to set the interval to be longer than the default to reduce the amount of traffic on the cluster control link.
- **Cluster Control Link**—Specifies the cluster control link interface.
  - **Interface**—Specifies the VNI interface.
  - **IP Address**—Specifies an IPv4 address for the IP address; IPv6 is not supported for this interface.
  - **Subnet Mask**—Specifies the subnet mask.
  - **MTU**—Specifies the maximum transmission unit for the VTEP source interface to be at least 154 bytes higher than the highest MTU of the data interfaces. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead (100 bytes) and VXLAN overhead (54 bytes). Set the MTU between 1554 and 9198 bytes. The default MTU is 1554 bytes. We suggest setting the cluster control link MTU to 1654 when data interfaces are set to 1500; this value requires jumbo frame reservation. For example, when using jumbo frames, 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. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes.

    Note: If you have not pre-enabled jumbo frame reservation, enable jumbo frames, and then restart this procedure.

Step 4 Check the **Participate in ASA cluster** check box to join the cluster.

Step 5 Click **Apply**.

Configure Interface Health Monitoring and Auto-Rejoin Settings

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. Health monitoring is not performed on VLAN subinterfaces. You cannot configure monitoring for the cluster control link; it is always monitored.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Interface Health Monitoring**.

Step 2 In the **Monitored Interfaces** box, select an interface, and click **Add** to move it to the **Unmonitored Interfaces** box.

Interface status messages detect link failure. If a node does not receive interface status messages within the holdtime, then the amount of time before the ASA removes a member from the cluster depends on whether the node is an established member or is joining the cluster. Health check is enabled by default for all interfaces.

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. Health monitoring is not performed on VLAN subinterfaces. You cannot configure monitoring for the cluster control link; it is always monitored.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch) you should disable the health check feature (**Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**) 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 health check feature.

Step 3 Click the **Auto Rejoin** tab to customize the auto-rejoin settings in case of an interface, system, or cluster control link failure. For each type, click **Edit** to set the following:

- **Maximum Rejoin Attempts**—Define the number of attempts at rejoining the cluster by setting **Unlimited** or a value between 0 and 65535. **0** disables auto-rejoining. The default value is **Unlimited** for the cluster-interface and **3** for the data-interface and system.
- **Rejoin Interval**—Define the interval duration in minutes between rejoin attempts by setting the interval 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**—Define if the interval duration increases by setting the interval variation 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.

Click **Restore Defaults** to restore the default settings.

Step 4 Click **Apply**.

Configure the Cluster TCP Replication Delay

Enable the cluster replication delay for TCP connections to help eliminate the “unnecessary work” related to short-lived flows by delaying the director/backup flow creation. Note that if a unit fails before the director/backup flow is created, then those flows cannot be recovered. Similarly, if traffic is rebalanced to a different unit before the flow is created, then the flow cannot be recovered. You should not enable the TCP replication delay for traffic on which you disable TCP randomization.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster Replication**.

Step 2 Click **Add** and set the following values:

- **Replication delay**—Set the seconds between 1 and 15.
- **HTTP**—Set the delay for all HTTP traffic.
- **Source Criteria**
  - **Source**—Set the source IP address.
  - **Service**—(Optional) Set the source port. Typically you set either the source or the destination port, but not both.

- **Destination Criteria**
  - **Source**—Set the destination IP address.
  - **Service**—(Optional) Set the destination port. Typically you set either the source or the destination port, but not both.

Step 3 Click **OK**.

Step 4 Click **Apply**.

Configure Inter-Site Features

For inter-site clustering, you can customize your configuration to enhance redundancy and stability.

Configure Cluster Flow Mobility

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

About LISP Inspection

You can inspect LISP traffic to enable flow mobility between sites.

About LISP

Data center virtual machine mobility such as VMware VMotion enables servers to migrate between data centers while maintaining connections to clients. To support such data center server mobility, routers need to be able to update the ingress route towards the server when it moves. Cisco Locator/ID Separation Protocol (LISP) architecture separates the device identity, or endpoint identifier (EID), from its location, or routing locator (RLOC), into two different numbering spaces, making server migration transparent to clients. For example, when a server moves to a new site and a client sends traffic to the server, the router redirects traffic to the new location.

LISP requires routers and servers in certain roles, such as the LISP egress tunnel router (ETR), ingress tunnel router (ITR), first hop routers, map resolver (MR), and map server (MS). When the first hop router for the server senses that the server is connected to a different router, it updates all of the other routers and databases so that the ITR connected to the client can intercept, encapsulate, and send traffic to the new server location.

ASA LISP Support

The ASA does not run LISP itself; it can, however, inspect LISP traffic for location changes and then use this information for seamless clustering operation. Without LISP integration, when a server moves to a new site, traffic comes to an ASA cluster member at the new site instead of to the original flow owner. The new ASA forwards traffic to the ASA at the old site, and then the old ASA has to send traffic back to the new site to reach the server. This traffic flow is sub-optimal and is known as “tromboning” or “hair-pinning.”

With LISP integration, the ASA cluster members can inspect LISP traffic passing between the first hop router and the ETR or ITR, and can then change the flow owner to be at the new site.

LISP Guidelines

- The ASA cluster members must reside between the first hop router and the ITR or ETR for the site. The ASA cluster itself cannot be the first hop router for an extended segment.

- Only fully-distributed flows are supported; centralized flows, semi-distributed flows, or flows belonging to individual nodes are not moved to new owners. Semi-distributed flows include applications, such as SIP, where all child flows are owned by the same ASA that owns the parent flow.
- The cluster only moves Layer 3 and 4 flow states; some application data might be lost.
- For short-lived flows or non-business-critical flows, moving the owner may not be worthwhile. You can control the types of traffic that are supported with this feature when you configure the inspection policy, and should limit flow mobility to essential traffic.

ASA LISP Implementation

This feature includes several inter-related configurations (all of which are described in this chapter):

1. (Optional) Limit inspected EIDs based on the host or server IP address—The first hop router might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.
2. LISP traffic inspection—The ASA inspects LISP traffic on UDP port 4342 for the EID-notify message sent between the first hop router and the ITR or ETR. The ASA maintains an EID table that correlates the EID and the site ID. For example, you should inspect LISP traffic with a source IP address of the first hop router and a destination address of the ITR or ETR. Note that LISP traffic is not assigned a director, and LISP traffic itself does not participate in cluster state sharing.
3. Service Policy to enable flow mobility on specified traffic—You should enable flow mobility on business-critical traffic. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers.
4. Site IDs—The ASA uses the site ID for each cluster node to determine the new owner.
5. Cluster-level configuration to enable flow mobility—You must also enable flow mobility at the cluster level. This on/off toggle lets you easily enable or disable flow mobility for a particular class of traffic or applications.

Configure LISP Inspection

You can inspect LISP traffic to enable flow mobility when a server moves between sites.

Before you begin

- Assign each cluster unit to a site ID according to Configure Basic ASA Cluster Parameters, on page 22.
- LISP traffic is not included in the default-inspection-traffic class, so you must configure a separate class for LISP traffic as part of this procedure.

Procedure

Step 1 (Optional) Configure a LISP inspection map to limit inspected EIDs based on IP address, and to configure the LISP pre-shared key:

a) Choose **Configuration** > **Firewall** > **Objects** > **Inspect Maps** > **LISP**.

b) Click **Add** to add a new map.
c) Enter a name (up to 40 characters) and description.
d) For the **Allowed-EID access-list**, click **Manage**.

The **ACL Manager** opens.

The first hop router or ITR/ETR might send EID-notify messages for hosts or networks the ASA cluster is not involved with, so you can limit the EIDs to only those servers or networks relevant to your cluster. For example, if the cluster is only involved with 2 sites, but LISP is running on 3 sites, you should only include EIDs for the 2 sites involved with the cluster.

e) Add an ACL with at least one ACE according to the firewall configuration guide.
f) If necessary, enter the **Validation Key**.

If you copied an encrypted key, click the **Encrypted** radio button.

g) Click **OK**.

Step 2 Add a service policy rule to configure LISP inspection:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, apply the rule to an interface or globally.

If you have an existing service policy you want to use, add a rule to that policy. By default, the ASA includes a global policy called **global_policy**. You can also create one service policy per interface if you do not want to apply the policy globally. LISP inspection is applied to traffic bidirectionally so you do not need to apply the service policy on both the source and destination interfaces; all traffic that enters or exits the interface to which you apply the rule is affected if the traffic matches the class for both directions.

d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the traffic you want to inspect. You should specify traffic between the first hop router and the ITR or ETR on UDP port 4342. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Protocol Inspection** tab.
i) Check the **LISP** check box,.
j) (Optional) Click **Configure** to choose the inspection map you created.
k) Click **Finish** to save the service policy rule.

Step 3 Add a service policy rule to enable Flow Mobility for critical traffic:

a) Choose **Configuration** > **Firewall** > **Service Policy Rules**.
b) Click **Add**.
c) On the **Service Policy** page, choose the same service policy you used for LISP inspection.
d) On the **Traffic Classification Criteria** page, click **Create a new traffic class**, and under **Traffic Match Criteria**, check **Source and Destination IP Address (uses ACL)**.
e) Click **Next**.
f) Specify the business critical traffic that you want to re-assign to the most optimal site when servers change sites. For example, you can limit flow mobility to only HTTPS traffic, and/or to traffic to specific servers. Both IPv4 and IPv6 ACLs are accepted.
g) Click **Next**.
h) On the **Rule Actions** wizard page or tab, select the **Cluster** tab.

i) Check the **Enable Cluster flow-mobility triggered by LISP EID messages** check box.
j) Click **Finish** to save the service policy rule.

Step 4 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**, and check the **Enable Clustering flow mobility** check box.

Step 5 Click **Apply**.

Manage Cluster Nodes

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

Add a New Data Node from the Control Node

You can add additional data nodes to the cluster from the control node. You can also add data nodes using the High Availability and Scalability wizard. Adding a data node from the control node has the benefit of configuring the cluster control link and setting the cluster interface mode on each data node you add.

You can alternatively log into the data node and configure clustering directly on the node. However, after you enable clustering, your ASDM session will be disconnected, and you will have to reconnect.

Before you begin

- If you want to send the bootstrap configuration over the management network, be sure the data node has an accessible IP address.

Procedure

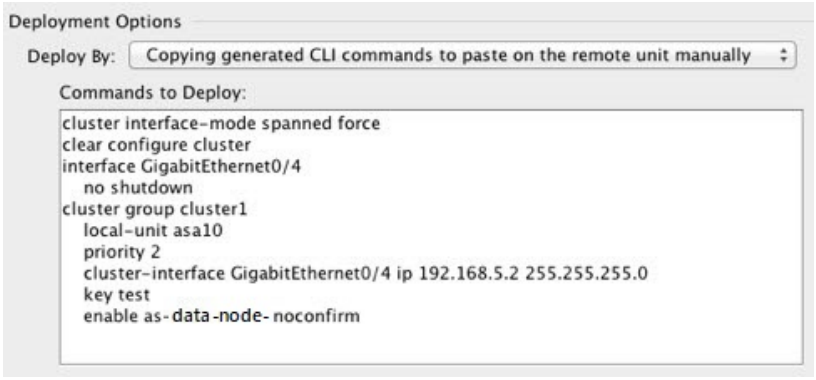
Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Click **Add**.

Step 3 Configure the following parameters:

- **Member Name**—Names this member of the cluster with a unique ASCII string from 1 to 38 characters.
- **Member Priority**—Sets the priority of this node for control node elections, between 1 and 100, where 1 is the highest priority.
- **Cluster Control Link > IP Address**—Specifies a unique IP address for this member for the cluster control link, on the same network as the control node cluster control link.
- In the **Deployment Options** area, choose one of the following **Deploy By** options:
  - **Sending CLI commands to the remote unit now**—Send the bootstrap configuration to the data node (temporary) management IP address. Enter the data node management IP address, username, and password.

- **Copying generated CLI commands to paste on the remote unit manually**—Generates the commands so that you can cut and paste them at the data node CLI or using the CLI tool in ASDM. In the Commands to Deploy box, select and copy the generated commands for later use.



Step 4 Click **OK**, then **Apply**.

Become an Inactive Node

To become an inactive member of the cluster, disable clustering on the node while leaving the clustering configuration intact.

Note When an ASA becomes inactive (either manually or through a health check failure), all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering; or you can remove the node altogether from the cluster. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, you saved the configuration with clustering disabled), then the management interface is disabled. You must use the console port for any further configuration.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**.

Step 2 Uncheck the **Participate in ASA cluster** check box.

Note
Do not uncheck the **Configure ASA cluster settings** check box; this action clears all cluster configuration, and also shuts down all interfaces including the management interface to which ASDM is connected. To restore connectivity in this case, you need to access the CLI at the console port.

Step 3 Click **Apply**.

Deactivate a Data Node from the Control Node

To deactivate a data node, perform the following steps.

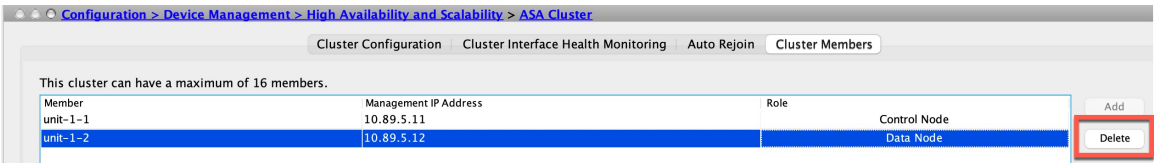
Note When an ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. To resume traffic flow, re-enable clustering. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster (for example, if you saved the configuration with clustering disabled), the management interface is disabled. You must use the console port for any further configuration.

Procedure

Step 1 Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Members**.

Step 2 Select the data node that you want to remove, and click **Delete**.

Figure 4: Delete Node



The data node bootstrap configuration remains intact, so that you can later re-add the data node without losing your configuration.

Step 3 Click **Apply**.

Rejoin the Cluster

If a node was removed from the cluster, for example for a failed interface or if you manually deactivated a member, you must manually rejoin the cluster.

Procedure

Step 1 If you still have ASDM access, you can reenable clustering in ASDM by connecting ASDM to the node you want to reenable.

You cannot reenable clustering for a data node from the control node unless you add it as a new member.

a) Choose **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**.
b) Check the **Participate in ASA cluster** check box.
c) Click **Apply**.

Step 2 If you cannot use ASDM: At the console, enter cluster configuration mode:

cluster group *name*

Example:

```
ciscoasa(config)# cluster group pod1
```

Step 3 Enable clustering.

enable

Leave the Cluster

If you want to leave the cluster altogether, you need to remove the entire cluster bootstrap configuration. Because the current configuration on each node is the same (synced from the active unit), leaving the cluster also means either restoring a pre-clustering configuration from backup, or clearing your configuration and starting over to avoid IP address conflicts.

Procedure

Step 1 For a data node, disable clustering:

cluster group *cluster_name*
no enable

Example:

```
ciscoasa(config)# cluster group cluster1
ciscoasa(cfg-cluster)# no enable
```

You cannot make configuration changes while clustering is enabled on a data node.

Step 2 Clear the cluster configuration:

clear configure cluster

The ASA shuts down all interfaces including the management interface and cluster control link.

Step 3 Disable cluster interface mode:

no cluster interface-mode

The mode is not stored in the configuration and must be reset manually.

Step 4 If you have a backup configuration, copy the backup configuration to the running configuration:

copy *backup_cfg* **running-config**

Example:

```
ciscoasa(config)# copy backup_cluster.cfg running-config

Source filename [backup_cluster.cfg]?

Destination filename [running-config]?
ciscoasa(config)#
```

Step 5 Save the configuration to startup:

write memory

Step 6 If you do not have a backup configuration, reconfigure management access. Be sure to change the interface IP addresses, and restore the correct hostname, for example.

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 node you want to become the control node, use the procedure in this section. Note, however, that for centralized features, if you force a control node change using this procedure, 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 Choose **Monitoring** > **ASA Cluster** > **Cluster Summary**.

Step 2 From the drop-down list, choose a data node to become control, and click the button to make it the control node.

Step 3 You are prompted to confirm the control node change. Click **Yes**.

Step 4 Quit ASDM, and reconnect using the Main cluster IP address.

Execute a Command Cluster-Wide

To send a command to all nodes in the cluster, or to a specific node, perform the following steps. Sending a **show** command to all nodes collects all output and displays it on the console of the current node. Other commands, such as **capture** and **copy**, can also take advantage of cluster-wide execution.

Before you begin

Perform this procedure at the Command Line Interface tool: choose **Tools > Command Line Interface**.

Procedure

Send a command to all nodes, or if you specify the node name, a specific node:

cluster exec [**unit** *node_name*] *command*

Example:

```
ciscoasa# cluster exec show xlate
```

To view node names, enter **cluster exec unit ?** (to see all names except the current node), or enter the **show cluster info** command.

Examples

To copy the same capture file from all nodes in the cluster at the same time to a TFTP server, enter the following command on the control node:

```
ciscoasa# cluster exec copy /pcap capture: tftp://10.1.1.56/capture1.pcap
```

Multiple PCAP files, one from each node, are copied to the TFTP server. The destination capture file name is automatically attached with the node name, such as capture1_asa1.pcap, capture1_asa2.pcap, and so on. In this example, asa1 and asa2 are cluster node names.

Monitoring the ASA Virtual Cluster

You can monitor and troubleshoot cluster status and connections.

Monitoring Cluster Status

See the following screens for monitoring cluster status:

- **Monitoring > ASA Cluster > Cluster Summary**

  This pane shows cluster information about the node to which you are connected, as well as other nodes in the cluster. You can also change the primary node from this pane.

- **Cluster Dashboard**

  On the home page on the primary node, you can monitor the cluster using the Cluster Dashboard and the Cluster Firewall Dashboard.

Capturing Packets Cluster-Wide

See the following screen for capturing packets in a cluster:

Wizards > Packet Capture Wizard

To support cluster-wide troubleshooting, you can enable capture of cluster-specific traffic on the control node, which is then automatically enabled on all of the data nodes in the cluster.

Monitoring Cluster Resources

See the following screens for monitoring cluster resources:

- **Monitoring > ASA Cluster > System Resources Graphs > CPU**

  This pane lets you create graphs or tables showing the CPU utilization across the cluster nodes.

- **Monitoring > ASA Cluster > System Resources Graphs > Memory**. This pane lets you create graphs or tables showing the Free Memory and Used Memory across the cluster nodes.

Monitoring Cluster Traffic

See the following screens for monitoring cluster traffic:

- **Monitoring > ASA Cluster > Traffic Graphs > Connections**.

  This pane lets you create graphs or tables showing the Connections across the cluster members.

- **Monitoring > ASA Cluster > Traffic Graphs > Throughput**.

  This pane lets you create graphs or tables showing the traffic throughput across the cluster members.

- **Monitoring** > **ASA Cluster** > **Cluster Load-Monitoring**

  This section includes the **Load Monitor-Information** and **Load-Monitor Details** panes. **Load Monitor-Information** shows the traffic load for cluster members for the last interval and also the average over total number of intervals configured (30 by default). Use the **Load-Monitor Details** pane to view the value for each measure at each interval.

Monitoring the Cluster Control Link

See the following screen for monitoring cluster status:

Monitoring > Properties > System Resources Graphs > Cluster Control Link.

This pane lets you create graphs or tables showing the cluster control link Receival and Transmittal capacity utilization.

Monitoring Cluster Routing

See the following screen for cluster routing:

- **Monitoring** > **Routing** > **LISP-EID Table**

  Shows the ASA EID table showing EIDs and site IDs.

Configuring Logging for Clustering

See the followingscreen for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

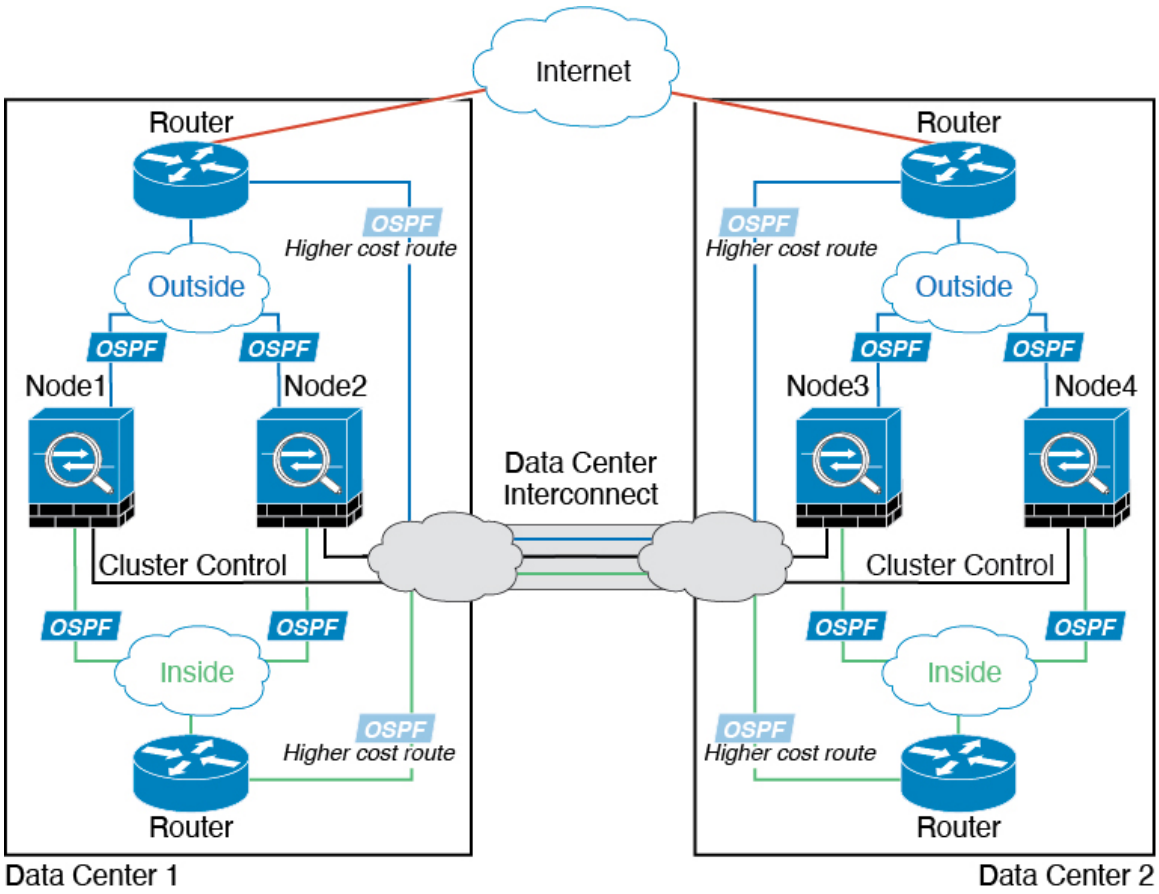
Each node in the cluster generates syslog messages independently. You can generate syslog messages with identical or different device IDs to make messages appear to come from the same or different nodes in the cluster.

Examples for ASA Virtual Clustering

These examples include all cluster-related ASA configuration for typical deployments.

Individual Interface Routed Mode North-South Inter-Site Example

The following example shows 2 ASA cluster nodes at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster nodes are connected by the cluster control link over the DCI. The inside and outside routers at each data center use OSPF and PBR or ECMP to load balance the traffic between cluster members. By assigning a higher cost route across the DCI, traffic stays within each data center unless all ASA cluster nodes at a given site go down. In the event of a failure of all cluster nodes at one site, traffic goes from each router over the DCI to the ASA cluster nodes at the other site.



Reference for Clustering

This section includes more information about how clustering operates.

ASA Features and Clustering

Some ASA features are not supported with ASA clustering, and some are only supported on the control node. 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.

- Unified Communication features that rely on TLS Proxy
- Remote access VPN (SSL VPN and IPsec VPN)
- Virtual Tunnel Interfaces (VTIs)
- The following application inspections:
  - CTIQBE

    - H323, H225, and RAS
    - IPsec passthrough
    - MGCP
    - MMP
    - RTSP
    - SCCP (Skinny)
    - WAAS
    - WCCP
- Botnet Traffic Filter
- Auto Update Server
- DHCP client, server, and proxy. DHCP relay is supported.
- VPN load balancing
- Failover on Azure
- Integrated Routing and Bridging
- FIPS mode

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.

- These application inspections are supported:
    - DCERPC and ESMTP
    - IM
    - NetBIOS and PPTP
    - RADIUS
    - RSH
    - SNMP

  - SQLNET and SUNRPC
  - TFTP
  - XDMCP
- Static route monitoring
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services
- Site-to-site VPN
- Multicast routing

Features Applied to Individual Nodes

These features are applied to each ASA node, instead of the cluster as a whole or to the control node.

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each node independently. For example, if you configure policing on output, then the conform rate and conform burst values are enforced on traffic exiting a particular ASA. In a cluster with 3 nodes and with traffic evenly distributed, the conform rate actually becomes 3 times the rate for the cluster.
- Threat detection—Threat detection works on each node independently; for example, the top statistics is node-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all nodes, and one node will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each node based on local usage.
- LISP traffic—LISP traffic on UDP port 4342 is inspected by each receiving node, but is not assigned a director. Each node adds to the EID table that is shared across the cluster, but the LISP traffic itself does not participate in cluster state sharing.

AAA for Network Access and Clustering

AAA for network access consists of three components: authentication, authorization, and accounting. Authentication and authorization are implemented as centralized features on the clustering control node with replication of the data structures to the cluster data nodes. If a control node is elected, the new control node will have all the information it needs to continue uninterrupted operation of the established authenticated users and their associated authorizations. Idle and absolute timeouts for user authentications are preserved when a control node change occurs.

Accounting is implemented as a distributed feature in a cluster. Accounting is done on a per-flow basis, so the cluster node owning a flow will send accounting start and stop messages to the AAA server when accounting is configured for a flow.

Connection Settings and Clustering

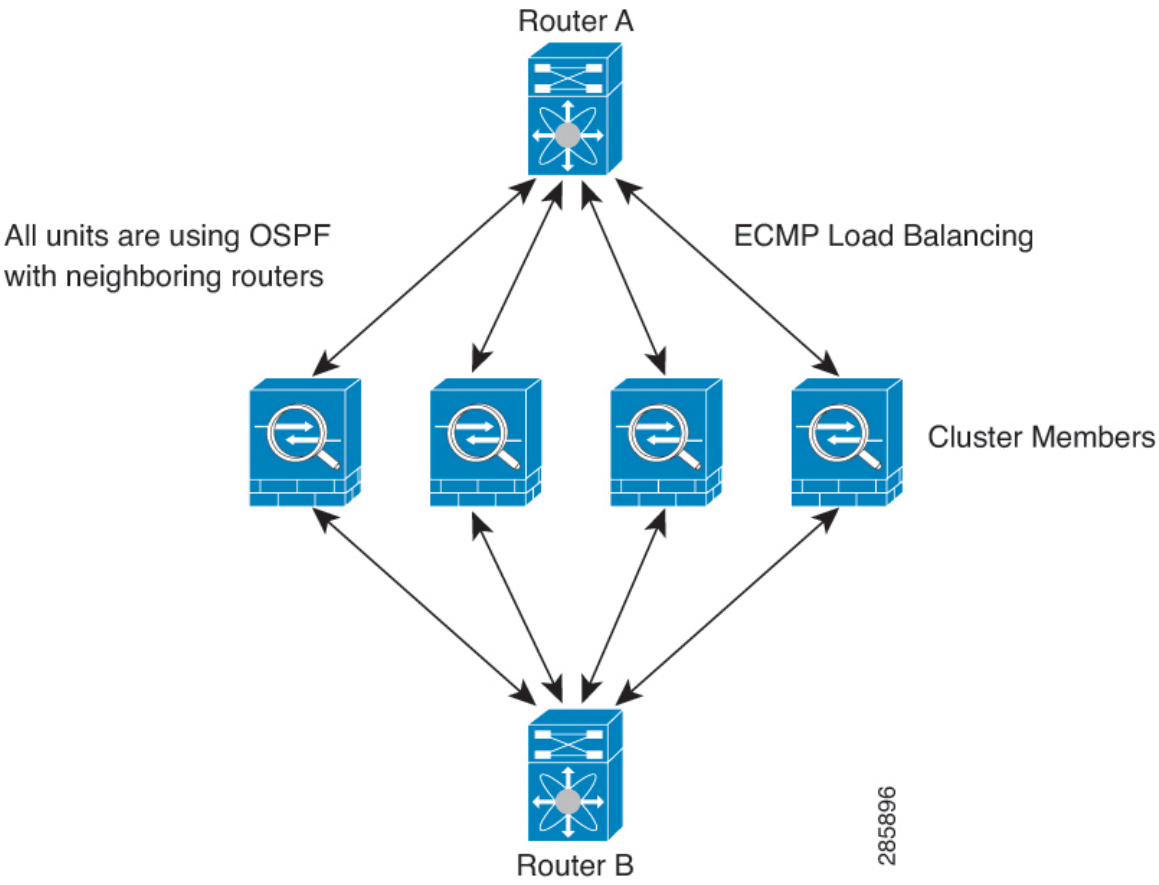
Connection limits are enforced cluster-wide (see **Configuration** > **Firewall** > **Service Policy** page). Each node has an estimate of the cluster-wide counter values based on broadcast messages. Due to efficiency considerations, the configured connection limit across the cluster might not be enforced 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

In Individual interface mode, each node runs the routing protocol as a standalone router, and routes are learned by each node independently.

Figure 5: Dynamic Routing in Individual Interface Mode



In the above diagram, 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 talking to external routers.

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

EIGRP does not form neighbor relationships with cluster peers in individual interface mode.

Note If the cluster has multiple adjacencies to the same router for redundancy purposes, asymmetric routing can lead to unacceptable traffic loss. To avoid asymmetric routing, group all of these node interfaces into the same traffic zone. See Configure a Traffic Zone.

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.
- If you use AAA for FTP access, then the control channel flow is centralized on the control node.

ICMP Inspection and Clustering

The flow of ICMP and ICMP error packets through the cluster varies depending on whether ICMP/ICMP error inspection is enabled. Without ICMP inspection, ICMP is a one-direction flow, and there is no director flow support. With ICMP inspection, the ICMP flow becomes two-directional and is backed up by a director/backup flow. One difference for an inspected ICMP flow is in the director handling of a forwarded packet: the director will forward the ICMP echo reply packet to the flow owner instead of returning the packet to the forwarder.

Multicast Routing and Clustering

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 ASAs 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 ASA 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 needs a static route or PBR with Object Tracking for the mapped addresses that points to the Main cluster IP address. This is not an issue for a Spanned EtherChannel, because there is only one IP address associated with the cluster interface.
- 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.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each data node to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the control node. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate, whereas ICMP and all other UDP traffic uses multi-session. You can configure per-session NAT rules to change these defaults for TCP and UDP, but you cannot configure per-session PAT for ICMP. For example, with an increasing use of Quic protocol over UDP/443 as a best performance alternative compared to HTTPS TLS over TCP/443, you should enable per-session PAT for UDP/443. For traffic that benefits from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT for the associated TCP ports (the UDP ports for those H.323 and SIP are already multi-session by default). For more information about per-session PAT, see the firewall configuration guide.
- No static PAT for the following inspections—
  - FTP
  - PPTP

- 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.

SCTP and Clustering

An SCTP association can be created on any node (due to load balancing); its multi-homing connections must reside on the same node.

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.

TLS Proxy configuration is not supported.

SNMP and Clustering

An SNMP agent polls each individual ASA by its 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 re-add them on the control node to force the users to replicate to the new node, or directly on the data node.

STUN and Clustering

STUN inspection is supported in failover and cluster modes, as pinholes are replicated. However, the transaction ID is not replicated among nodes. In the case where a node fails after receiving a STUN Request and another node received the STUN Response, the STUN Response will be dropped.

Syslog and NetFlow and Clustering

- Syslog—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.

- NetFlow—Each node in the cluster generates its own NetFlow stream. The NetFlow collector can only treat each ASA as a separate NetFlow exporter.

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.

For connections to an Individual interface when using PBR or 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.

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 3 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 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, 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 ASA Virtual Cluster

The ASA Virtual 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.

When you enable health monitoring, all physical interfaces are monitored by default; you can optionally disable monitoring per interface. Only named interfaces can be monitored.

A node is removed from the cluster if its monitored interfaces fail. The amount of time before the ASA removes a member from the cluster depends on whether the node is an established member or is joining the cluster. The ASA 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 ASA to be removed from the cluster. The node is removed after 500 ms, regardless of the node 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 ASA automatically tries to rejoin the cluster, depending on the failure event.



Note When the ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the node received from the cluster IP pool. However if you reload, and the node is still inactive in the cluster, the management interface is disabled. You must use the console port for any further configuration.

Rejoining the Cluster

After a cluster node 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 ASA automatically tries to rejoin every 5 minutes, indefinitely. This behavior is configurable.
- Failed data interface—The ASA 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 ASA disables clustering. After you resolve the problem with the data interface, you have to manually enable clustering. This behavior is configurable.
- 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 and clustering is still enabled. The ASA attempts to rejoin the cluster every 5 seconds.
- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. A node will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. This behavior is configurable.

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	Includes AAA rules (uauth).
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	**No**	—
Distributed VPN (Site-to-Site) for Firepower 4100/9300	Yes	Backup session becomes the active session, then a new backup session is created.

How the ASA Virtual 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.

  If you enable director localization for inter-site clustering, then there are two backup owner roles: the local backup and the global backup. The owner always chooses a local backup at the same site as itself

(based on site ID). The global backup can be at any site, and might even be the same node as the local backup. The owner sends connection state information to both backups.

If you enable site redundancy, and the backup owner is at the same site as the owner, then an additional backup owner will be chosen from another site to protect flows from a site failure. Chassis backup and site backup are independent, so in some cases a flow will have both a chassis backup and a site backup.

- 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.

  If you enable director localization for inter-site clustering, the cluster uses two director roles: the local director and the global director. For a new-flow packet, the receiving member selects the local director from members at the receiving member's site. The global director can be at any site and might be the same member as the local director. If the original owner fails and localization applies, the local director chooses a new connection owner at the local director's site.

  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. If you enable director localization, the forwarder queries the local director when localization applies. The forwarder queries the global director if the local director does not know the owner, if the connection is owned at another site, or if localization is bypassed. A cross-site director query is expected behavior in these cases. 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

When a connection uses Port Address Translation (PAT), then the PAT type (per-session or multi-session) influences which member of the cluster becomes the owner of a new connection:

- Per-session PAT—The owner is the node that receives the initial packet in the connection.

  By default, TCP and DNS UDP traffic use per-session PAT.

- Multi-session PAT—The owner is always the control node. If a multi-session PAT connection is initially received by a data node, then the data node forwards the connection to the control node.

  By default, UDP (except for DNS UDP) and ICMP traffic use multi-session PAT, so these connections are always owned by the control node.

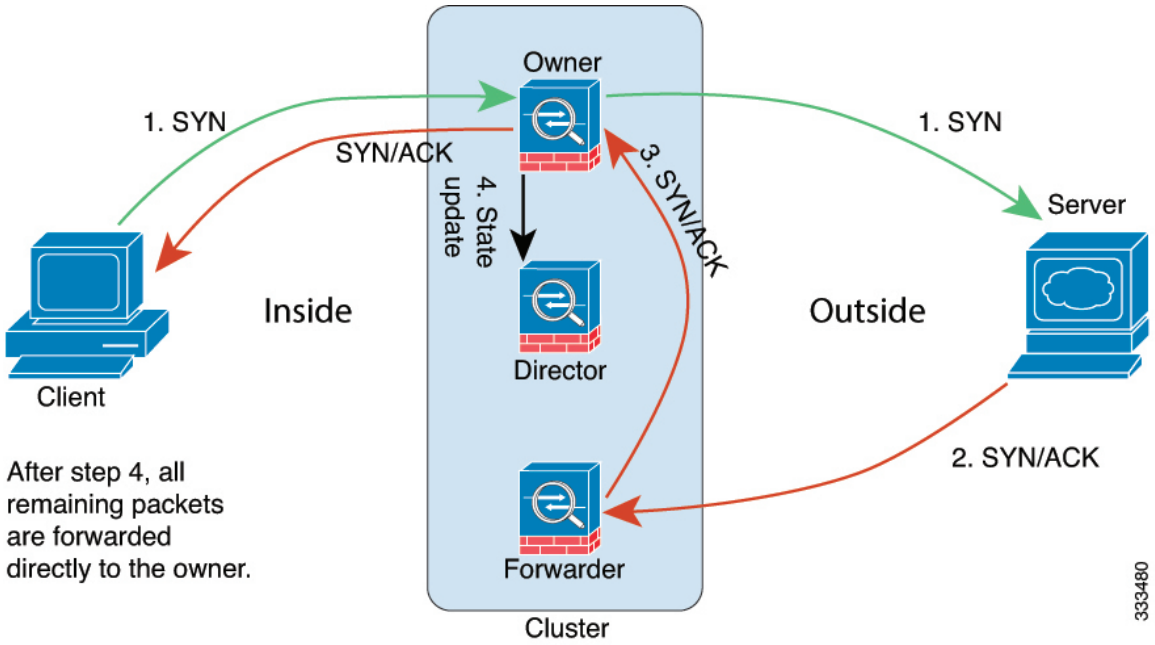
You can change the per-session PAT defaults for TCP and UDP so connections for these protocols are handled per-session or multi-session depending on the configuration. For ICMP, you cannot change from the default multi-session PAT. For more information about per-session PAT, see the firewall configuration guide.

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. For best performance, proper external load balancing is required for both directions of a flow to arrive at the same node, and for flows to be distributed evenly between nodes. 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.

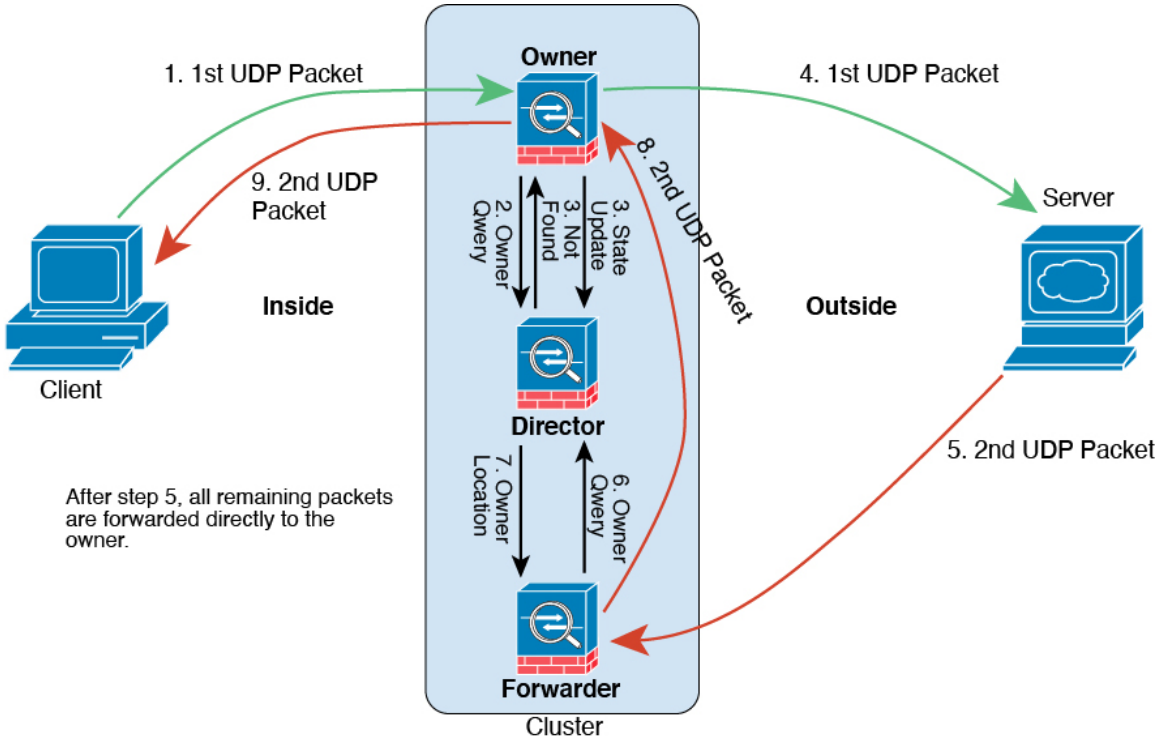


1. The SYN packet originates from the client and is delivered to one ASA (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and is delivered to a different ASA (based on the load balancing method). This ASA 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 6: ICMP and UDP Data Flow*



The first UDP packet originates from the client and is delivered to one ASA (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.

Rebalancing New TCP Connections Across the Cluster

If the load balancing capabilities of the upstream or downstream routers result in unbalanced flow distribution, you can configure new connection rebalancing so nodes with higher new connections per second will redirect new TCP flows to other nodes. No existing flows will be moved to other nodes.

Because this command only rebalances based on connections per second, the total number of established connections on each node is not considered, and the total number of connections may not be equal.

Once a connection is offloaded to a different node, it becomes an asymmetric connection.

Do not configure connection rebalancing for inter-site topologies; you do not want new connections rebalanced to cluster members at a different site.

History for ASA Virtual Clustering

Feature Name	Version	Feature Information
Configurable cluster keepalive interval for flow status	9.20(1)	The flow owner sends keepalives (clu_keepalive messages) and updates (clu_update messages) to the director and backup owner to refresh the flow state. You can now set the keepalive interval. The default is 15 seconds, and you can set the interval between 15 and 55 seconds. You may want to set the interval to be longer to reduce the amount of traffic on the cluster control link. New/Modified screens: **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster** > **Cluster Configuration**
Removal of biased language	9.19(1)	Commands, command output, and syslog messages that contained the terms "Master" and "Slave" have been changed to "Control" and "Data." New/Modified commands: **cluster control-node**, **enable as-data-node**, **prompt**, **show cluster history**, **show cluster info**
ASAv30, ASAv50, and ASAv100 clustering for VMware and KVM	9.17(1)	The ASA Virtual clustering lets you group up to 16 ASA 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. The ASA Virtual clustering supports Individual Interface mode in routed firewall mode; Spanned EtherChannels are not supported. The ASA Virtual uses a VXLAN virtual interface (VNI) for the cluster control link. New/Modified screens: • **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** • **Configuration** > **Device Management** > **High Availability and Scalability** > **ASA Cluster**