

ASA Cluster for the Secure Firewall 3100/4200

Clustering lets you group multiple ASAs together as a single logical device. A cluster provides all the convenience of a single device (management, integration into a network) while achieving the increased throughput and redundancy of multiple devices.

Note Some features are not supported when using clustering. See Unsupported Features with Clustering, on page 61.

- About ASA Clustering, on page 1
- Licenses for ASA Clustering, on page 5
- Requirements and Prerequisites for ASA Clustering, on page 6
- Guidelines for Clustering, on page 8
- Configure ASA Clustering, on page 13
- Manage Cluster Nodes, on page 40
- Monitoring the ASA Cluster, on page 47
- Examples for ASA Clustering, on page 48
- Reference for Clustering, on page 61
- History for ASA Clustering for the Secure Firewall 3100/4200, on page 75

About ASA Clustering

This section describes the clustering architecture and how it works.

How a cluster fits into your network

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

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

Infrastructure requirements

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

Cluster Members

Cluster members work together to accomplish the sharing of the security policy and traffic flows. This section describes the nature of each member 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.

Cluster interfaces

You can configure data interfaces as Spanned EtherChannels. Refer Cluster interfaces, on page 14 for more information.

Interface type restrictions

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

Cluster control links

A cluster control link is a dedicated hardware interface that each unit must allocate for cluster communication. For more information, refer Cluster control link, on page 14.

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 Cluster Management

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

Management Network

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

Management Interface

For the management interface, we recommend using one of the dedicated management interfaces. You can configure the management interfaces as Individual interfaces (for both routed and transparent modes) or as a Spanned EtherChannel interface.

We recommend using Individual interfaces for management, even if you use Spanned EtherChannels for your data interfaces. Individual interfaces let you connect directly to each unit if necessary, while a Spanned EtherChannel interface only allows remote connection to the current control unit.



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

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

For example, you can manage the cluster by connecting to the Main cluster IP address, which is always attached to the current control unit. To manage an individual member, you can connect to the Local 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.

For outbound management traffic such as TFTP or syslog, each unit, including the control unit, uses the Local IP address to connect to the server.

For a Spanned EtherChannel interface, you can only configure one IP address, and that IP address is always attached to the control unit. You cannot connect directly to a data unit using the EtherChannel interface; we recommend configuring the management interface as an Individual interface so that you can connect to each unit. Note that you can use a device-local EtherChannel for management.

Control Unit Management Vs. Data Unit 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 clustering as long as you follow the recommended guidelines.

You can configure each cluster chassis to belong to a separate site ID.

Site IDs work with site-specific MAC addresses and IP addresses. Packets egressing the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address. Site-specific MAC addresses and IP address are supported for routed mode using Spanned EtherChannels only.

Site IDs are also 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 Clustering, on page 6
- Inter-Site Guidelines—Guidelines for Clustering, on page 8
- Configure Cluster Flow Mobility—Configure Cluster Flow Mobility, on page 37
- Enable Director Localization—Configure Basic ASA Cluster Parameters, on page 30
- Enable Site Redundancy—Configure Basic ASA Cluster Parameters, on page 30
- Inter-Site Examples—Examples for Inter-Site Clustering, on page 58

Licenses for ASA Clustering

Smart Software Manager Regular and On-Prem

Each unit requires the Essentials license (enabled by default) and the same encryption license. We recommend licensing each unit with the licensing server *before* you enable clustering to avoid licensing mismatch issues, and when using the Strong Encryption license, issues with cluster control link encryption.

The clustering feature itself does not require any licenses. There is no extra cost for the Context license on data units.

The Strong Encryption license is automatically enabled for qualified customers when you apply the registration token. For the optional Strong Encryption (3DES/AES) feature license enabled in the ASA configuration, see below.

In the ASA license configuration, the Essentials license is always enabled by default on all units. You can only configure smart licensing on the control unit. The configuration is replicated to the data units, but for some licenses, they do not use the configuration; it remains in a cached state, and only the control unit requests the license. The licenses are aggregated into a single cluster license that is shared by the cluster units, and this aggregated license is also cached on the data units to be used if one of them becomes the control unit in the future. Each license type is managed as follows:

- Essentials—Each unit requests a Essentials license from the server.
- Context—Only the control unit requests the Context license from the server. The Essentials license includes 2 contexts by default and is present on all cluster members. The value from each unit's Essentials license plus the value of the Context license on the control unit are combined up to the platform limit in an aggregated cluster license. For example:
  - You have 6 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 6 units, these licenses add up to 12 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 32 contexts. Because the platform limit for one chassis is 100, the combined license allows a maximum of 100 contexts; the 32 contexts are within the limit. Therefore, you can configure up to 32 contexts on the control unit; each data unit will also have 32 contexts through configuration replication.
  - You have 3 Secure Firewall 3100s in the cluster. The Essentials license includes 2 contexts; for 3 units, these licenses add up to 6 contexts. You configure an additional 100-Context license on the control unit. Therefore, the aggregated cluster license includes 106 contexts. Because the platform limit for one unit is 100, the combined license allows a maximum of 100 contexts; the 106 contexts are over the limit. Therefore, you can only configure up to 100 contexts on the control unit; each data unit will also have 100 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 94 contexts.

- Strong Encryption (3DES)—If your Smart Account is not authorized for strong encryption, but Cisco has determined that you are allowed to use strong encryption, you can manually add a strong encryption license to your account. Only the control unit requests this license, and all units can use it due to license aggregation.

If a new control unit is elected, the new control unit continues to use the aggregated license. It also uses the cached license configuration to re-request the control unit license. When the old control unit rejoins the cluster as a data unit, it releases the control unit license entitlement. Before the data unit releases the license, the control unit's license might be in a non-compliant state if there are no available licenses in the account. The retained license is valid for 30 days, but if it is still non-compliant after the grace period, you will not be able to make configuration changes to features requiring special licenses; operation is otherwise unaffected. The new active unit sends an entitlement authorization renewal request every 35 seconds until the license is compliant. You should refrain from making configuration changes until the license requests are completely processed. If a unit leaves the cluster, the cached control configuration is removed, while the per-unit entitlements are retained. In particular, you would need to re-request the Context license on non-cluster units.

Permanent License Reservation

For permanent license reservation, you must purchase separate licenses for each chassis and enable the licenses *before* you configure clustering.

Requirements and Prerequisites for ASA Clustering

Model Requirements

- Secure Firewall 3100—Maximum 8 nodes
- Secure Firewall 4200—Maximum 8 nodes

ASA Hardware and Software Requirements

All units in a cluster:

- Must be the same model with the same DRAM. You do not have to have the same amount of flash memory.
- 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.
- Must be in the same security context mode, single or multiple.
- (Single context mode) Must be in the same firewall mode, routed or transparent.
- New cluster members must use the same SSL encryption setting (the **ssl encryption** command) as the control unit for initial cluster control link communication before configuration replication.

Switch Requirements

- Be sure to complete the switch configuration before you configure clustering on the ASAs.
- For a list of supported switches, see Cisco ASA Compatibility.

ASA Requirements

- Provide each unit with a unique IP address before you join them to the management network.
  - See the Getting Started chapter for more information about connecting to the ASA and setting the management IP address.
  - Except for the IP address used by the control unit (typically the first unit you add to the cluster), these management IP addresses are for temporary use only.
  - After a data unit joins the cluster, its management interface configuration is replaced by the one replicated from the control unit.

Sizing the Data Center Interconnect for Inter-Site Clustering

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

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

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

For example:

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

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

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

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

Other Requirements

We recommend using a terminal server to access all cluster member unit console ports. For initial setup, and ongoing management (for example, when a unit goes down), a terminal server is useful for remote management.

Guidelines for Clustering

Context Mode

The mode must match on each member unit.

Firewall Mode

For single mode, the firewall mode must match on all units.

Failover

Failover is not supported with clustering.

IPv6

The cluster control link is only supported using IPv4.

Switches

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. In addition, we do not recommend setting the cluster control link MTU between 2561 and 8362; due to block pool handling, this MTU size is not optimal for system operation.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS XR interface MTU to be 14 bytes higher than the cluster device MTU. Otherwise, OSPF adjacency peering attempts may fail unless the **mtu-ignore** option is used. Note that the cluster device MTU should match the IOS XR *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **src-dst-mixed-ip-port** (see the Cisco Nexus OS and Cisco IOS-XE **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the devices in a cluster. *Do not* change the load-balancing algorithm from the default on the cluster device.
- If you change the load-balancing algorithm of the EtherChannel on the switch, the EtherChannel interface on the switch temporarily stops forwarding traffic, and the Spanning Tree Protocol restarts. There will be a delay before traffic starts flowing again.

- Switches on the cluster control link path should not verify the L4 checksum. Redirected traffic over the cluster control link does not have a correct L4 checksum. Switches that verify the L4 checksum could cause traffic to be dropped.
- Port-channel bundling downtime should not exceed the configured keepalive interval.
- On Supervisor 2T EtherChannels, the default hash distribution algorithm is adaptive. To avoid asymmetric traffic in a VSS design, change the hash algorithm on the port-channel connected to the cluster device to fixed:

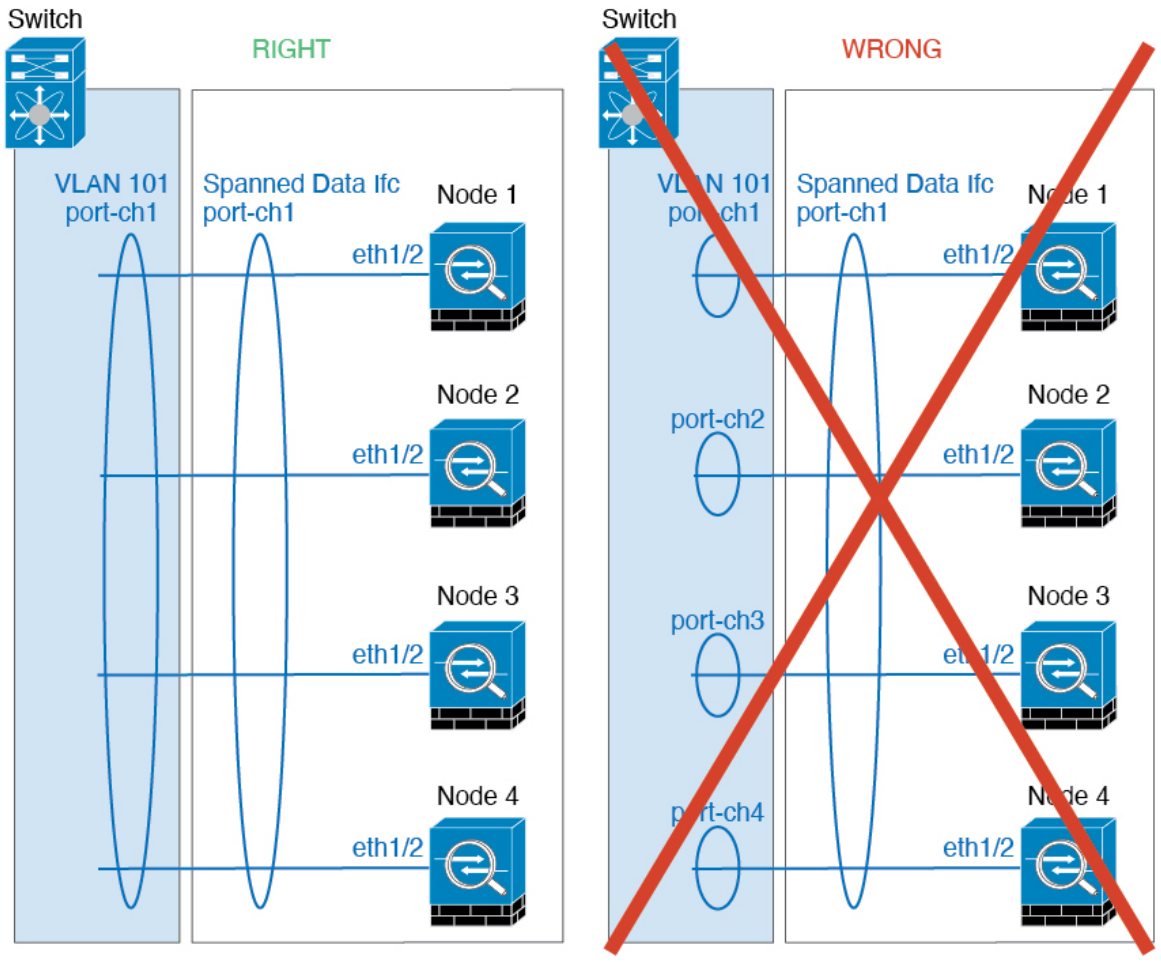
  router(config)# **port-channel** *id* **hash-distribution fixed**

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

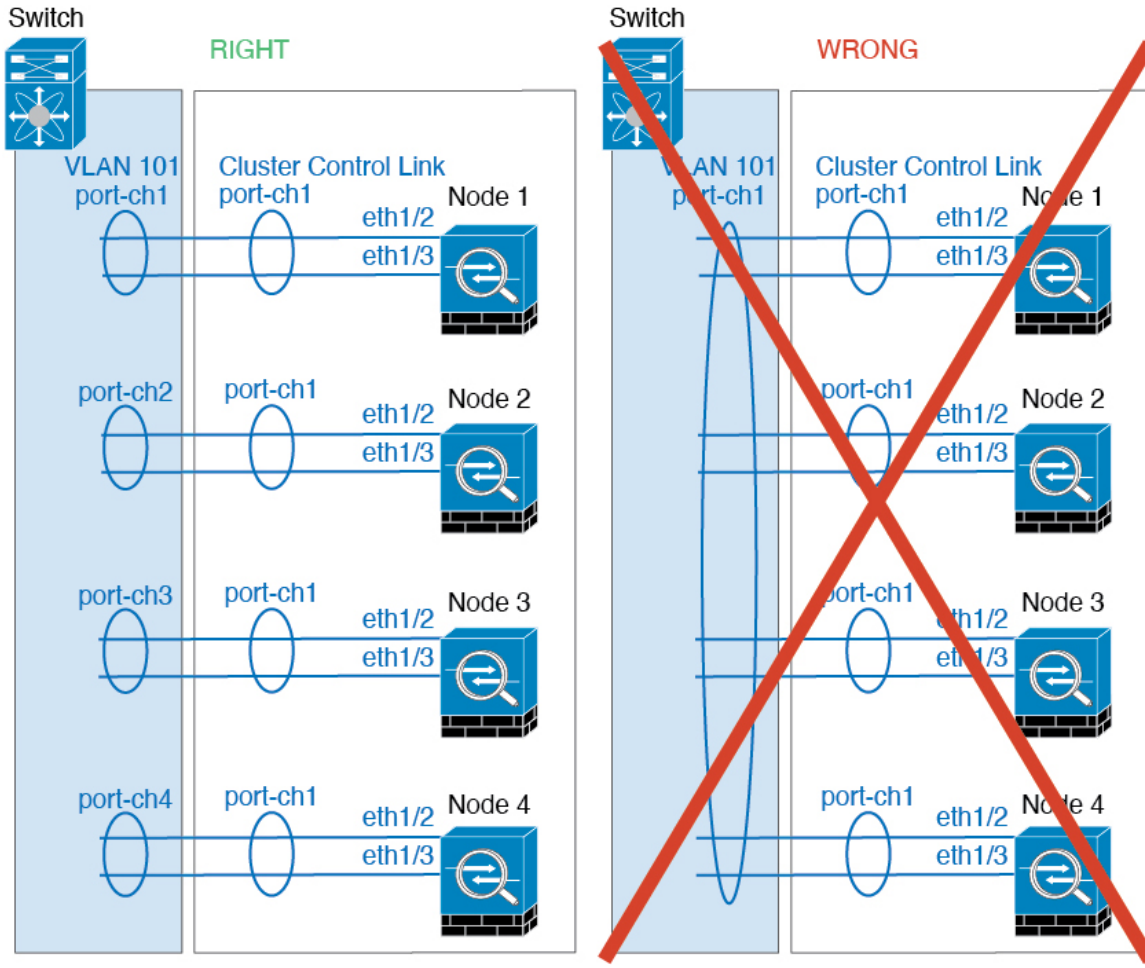
- You should disable the LACP Graceful Convergence feature on all cluster-facing EtherChannel interfaces for Cisco Nexus switches.

EtherChannels

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



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



Inter-Site Guidelines

See the following guidelines for inter-site clustering:

- The cluster control link latency must be less than 20 ms round-trip time (RTT).
- The cluster control link must be reliable, with no out-of-order or dropped packets; for example, you should use a dedicated link.
- Do not configure connection rebalancing; you do not want connections rebalanced to cluster members at a different site.
- The ASA does not encrypt forwarded data traffic on the cluster control link because it is a dedicated link, even when used on a Data Center Interconnect (DCI). If you use Overlay Transport Virtualization (OTV), or are otherwise extending the cluster control link outside of the local administrative domain, you can configure encryption on your border routers such as 802.1AE MacSec over OTV.
- The cluster implementation does not differentiate between members at multiple sites for incoming connections; therefore, connection roles for a given connection may span across sites. This is expected behavior. However, if you enable director localization, the local director role is always chosen from the same site as the connection owner (according to site ID). Also, the local director chooses a new owner

at the same site if the original owner fails (Note: if the traffic is asymmetric across sites, and there is continuous traffic from the remote site after the original owner fails, then a node from the remote site might become the new owner if it receives a data packet within the re-hosting window.).

- For director localization, the following traffic types do not support localization: NAT or PAT traffic; SCTP-inspected traffic; Fragmentation owner query.
- For UDP long-lived flows in a North-South deployment, routing loops can occur if nodes at the original flow owner site fail and then come back up, after which the flow is directed back to the original site. If the new owner at the other site doesn't have a route to the destination, it will route the flow back to the internet, causing a loop. In this case, use the **clear conn** command on the new owner to force the flow to be reestablished.
- For transparent mode, if the cluster is placed between a pair of inside and outside routers (AKA North-South insertion), you must ensure that both inside routers share a MAC address, and also that both outside routers share a MAC address. When a cluster member at site 1 forwards a connection to a member at site 2, the destination MAC address is preserved. The packet will only reach the router at site 2 if the MAC address is the same as the router at site 1.
- For transparent mode, if the cluster is placed between data networks and the gateway router at each site for firewalling between internal networks (AKA East-West insertion), then each gateway router should use a First Hop Redundancy Protocol (FHRP) such as HSRP to provide identical virtual IP and MAC address destinations at each site. The data VLANs are extended across the sites using Overlay Transport Virtualization (OTV), or something similar. You need to create filters to prevent traffic that is destined to the local gateway router from being sent over the DCI to the other site. If the gateway router becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's gateway.
- For transparent mode, if the cluster is connected to an HSRP router, you must add the router HSRP MAC address as a static MAC address table entry on the ASA (see Add a Static MAC Address for Bridge Groups). When adjacent routers use HSRP, traffic destined to the HSRP IP address will be sent to the HSRP MAC Address, but return traffic will be sourced from the MAC address of a particular router's interface in the HSRP pair. Therefore, the ASA MAC address table is typically only updated when the ASA ARP table entry for the HSRP IP address expires, and the ASA sends an ARP request and receives a reply. Because the ASA's ARP table entries expire after 14400 seconds by default, but the MAC address table entry expires after 300 seconds by default, a static MAC address entry is required to avoid MAC address table expiration traffic drops.
- For routed mode using Spanned EtherChannel, configure site-specific MAC addresses. Extend the data VLANs across the sites using OTV, or something similar. You need to create filters to prevent traffic that is destined to the global MAC address from being sent over the DCI to the other site. If the cluster becomes unreachable at one site, you need to remove any filters so traffic can successfully reach the other site's cluster nodes. Dynamic routing is not supported when an inter-site cluster acts as the first hop router for an extended segment.

Additional Guidelines

- When significant topology changes occur (such as adding or removing an EtherChannel interface, enabling or disabling an interface on the firewall or the switch, adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) 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 units, you can re-enable the interface health check feature.

- When adding a unit to an existing cluster, or when reloading a unit, there will be a temporary, limited packet/connection drop; this is expected behavior. In some cases, the dropped packets can hang your connection; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel, when the syslog server port is down and the server does not throttle ICMP error messages, then large numbers of ICMP messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.
- It takes time to replicate changes to all the units 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 units 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 Clustering

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

To configure clustering, perform the following tasks.



Note To enable or disable clustering, you must use a console connection (for CLI) or an ASDM connection.

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.

Cable the Units and Configure Interfaces

Before configuring clustering, cable the cluster control link network, management network, and data networks. Then configure your interfaces.

Cluster interfaces

You can configure data interfaces as Spanned EtherChannels.

Each unit must also dedicate at least one hardware interface as the cluster control link.

Cluster control link

A cluster control link is a dedicated hardware interface that enables communication between cluster units and supports cluster operation coordination.

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

Cluster control link traffic

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

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

Data traffic includes:

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

Cluster Control Link Interfaces and Network

You can use any data interface(s) for the cluster control link, with the following exceptions:

- You cannot use a VLAN subinterface as the cluster control link.
- You cannot use a Management *x*/*x* interface as the cluster control link, either alone or as an EtherChannel.

You can use an EtherChannel.

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

For a 2-member cluster, do not directly-connect the cluster control link from one ASA to the other ASA. If you directly connect the interfaces, then when one unit fails, the cluster control link fails, and thus the remaining healthy unit fails. If you connect the cluster control link through a switch, then the cluster control link remains up for the healthy unit.

Size the cluster control link

Cluster control link traffic is comprised mainly of state update and forwarded packets. The amount of traffic at any given time on the cluster control link varies. The amount of forwarded traffic depends on the load-balancing efficacy or whether there is a lot of traffic for centralized features.

For example:

- NAT results in poor load balancing of connections, and the need to rebalance all returning traffic to the correct units.
- AAA for network access is a centralized feature, so all traffic is forwarded to the control unit.
- When membership changes, the cluster needs to rebalance a large number of connections, thus temporarily using a large amount of cluster control link bandwidth.

A higher-bandwidth cluster control link helps the cluster to converge faster when there are membership changes and prevents throughput bottlenecks.



Note If your cluster has large amounts of asymmetric (rebalanced) traffic, then you should increase the cluster control link size.

Cluster control link 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:

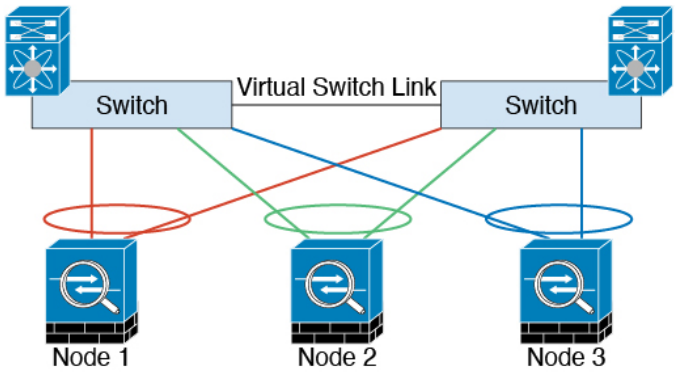
- **show cluster history**—`Event: CCL MTU test to unit name failed`
- Console WARNING—`WARNING: Unit name is not reachable in CCL jumbo frame ICMP test, please check cluster interface and switch MTU configuration`

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.

Cluster control link redundancy

We recommend using an EtherChannel for the cluster control link, so that you can pass traffic on multiple links in the EtherChannel while still achieving redundancy.

The diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS), Virtual Port Channel (vPC), StackWise, or StackWise Virtual environment. All links in the EtherChannel are active. When the switch is part of a redundant system, then you can connect firewall interfaces within the same EtherChannel to separate switches in the redundant system. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch. Note that this EtherChannel is device-local, not a Spanned EtherChannel.



Cluster control link reliability

To ensure cluster control link functionality, be sure the round-trip time (RTT) between units is less than 20 ms. This maximum latency enhances compatibility with cluster members installed at different geographical sites. To check your latency, perform a ping on the cluster control link between units.

The cluster control link must be reliable, with no out-of-order or dropped packets; for example, for inter-site deployment, you should use a dedicated link.

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.

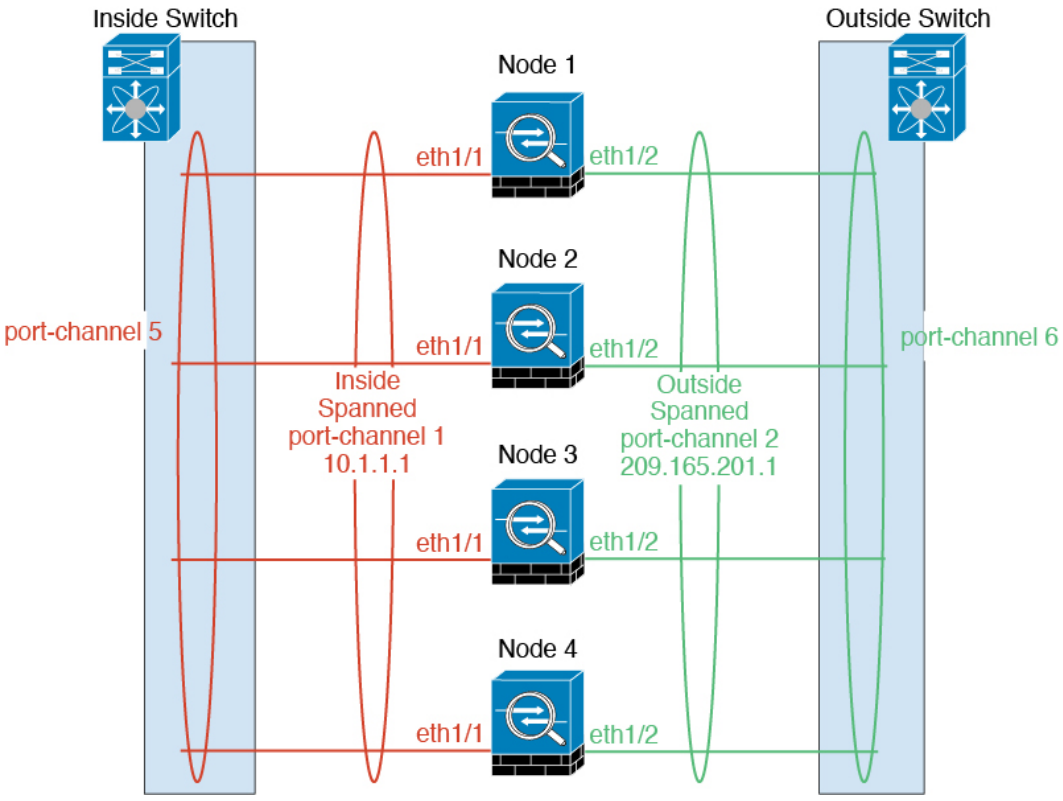
Note When the ASA becomes inactive, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit is still inactive in the cluster,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

Spanned EtherChannels

You can group one or more interfaces per chassis into an EtherChannel that spans all chassis in the cluster. The EtherChannel aggregates the traffic across all the available active interfaces in the channel.

A Spanned EtherChannel can be configured in both routed and transparent firewall modes. In routed mode, the EtherChannel is configured as a routed interface with a single IP address. In transparent mode, the IP address is assigned to the BVI, not to the bridge group member interface.

The EtherChannel inherently provides load balancing as part of basic operation.



Guidelines for maximum throughput

To achieve maximum throughput, follow these guidelines:

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

Load balancing

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

Note On the switch, we recommend that you use one of the following algorithms: **source-dest-IP** or **source-dest-IP-port** (see the Cisco Nexus OS or Cisco IOS **port-channel load-balance** command). Do not use a **vlan** keyword in the load-balance algorithm because it can cause unevenly distributed traffic to the nodes in a cluster.

The number of links in the EtherChannel affects load balancing.

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

EtherChannel redundancy

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

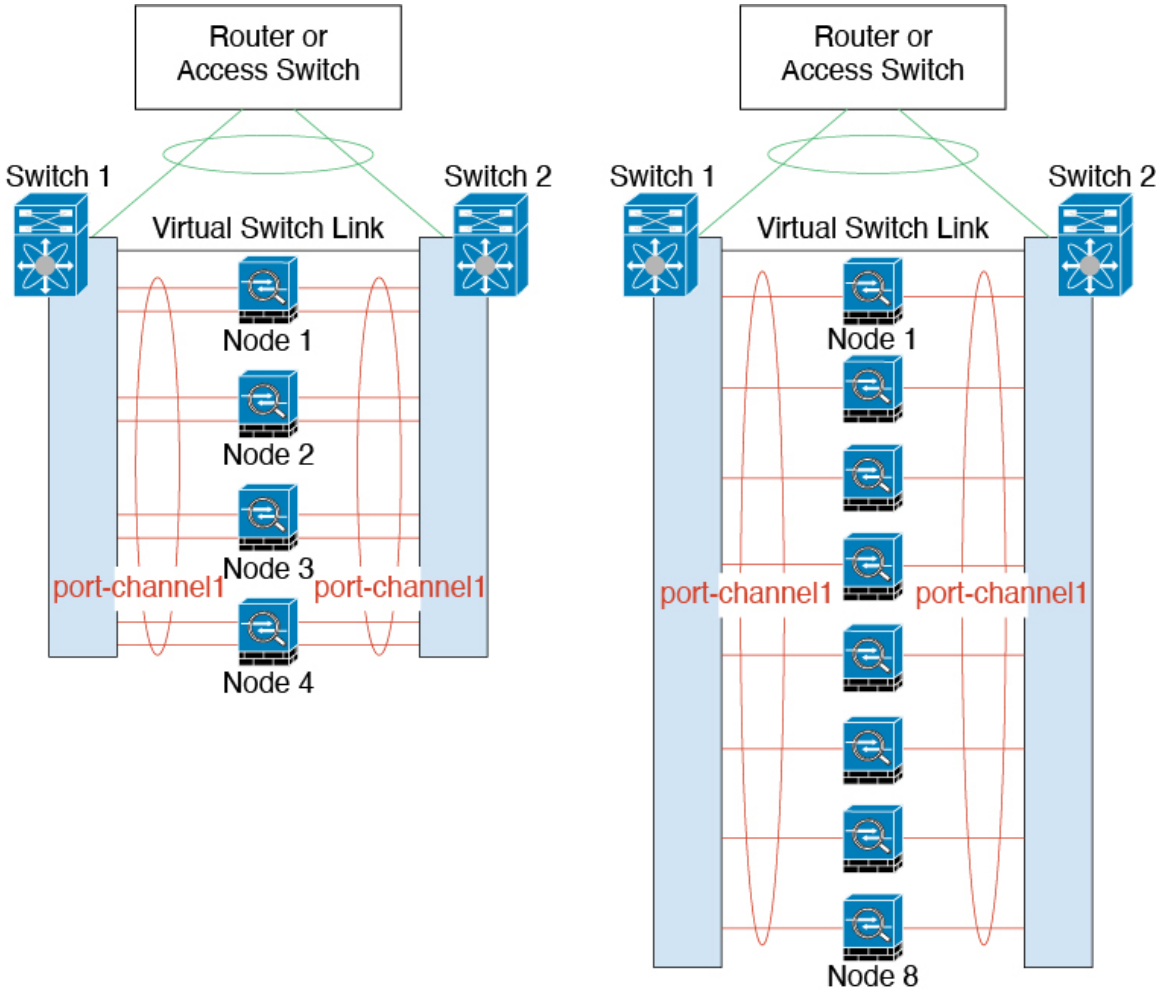
Connecting to a redundant switch system

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

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

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

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



Cable the Cluster Units and Configure Upstream and Downstream Equipment

Before configuring clustering, cable the cluster control link network, management network, and data networks.

Procedure

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

Note
At a minimum, an active cluster control link network is required before you configure the nodes to join the cluster.

Step 2 You should also configure the upstream and downstream equipment. For example, if you use EtherChannels, then you should configure the upstream and downstream equipment for the EtherChannels.

Configure the Cluster Interface Mode on the Control Unit

Before you enable clustering, you need to convert the firewall to use Spanned EtherChannels. 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 units from the control unit, you must set the interface mode manually on all units according to this section, not just the control unit; if you add secondaries from the control unit, ASDM sets the interface mode automatically on the data unit.

Before you begin

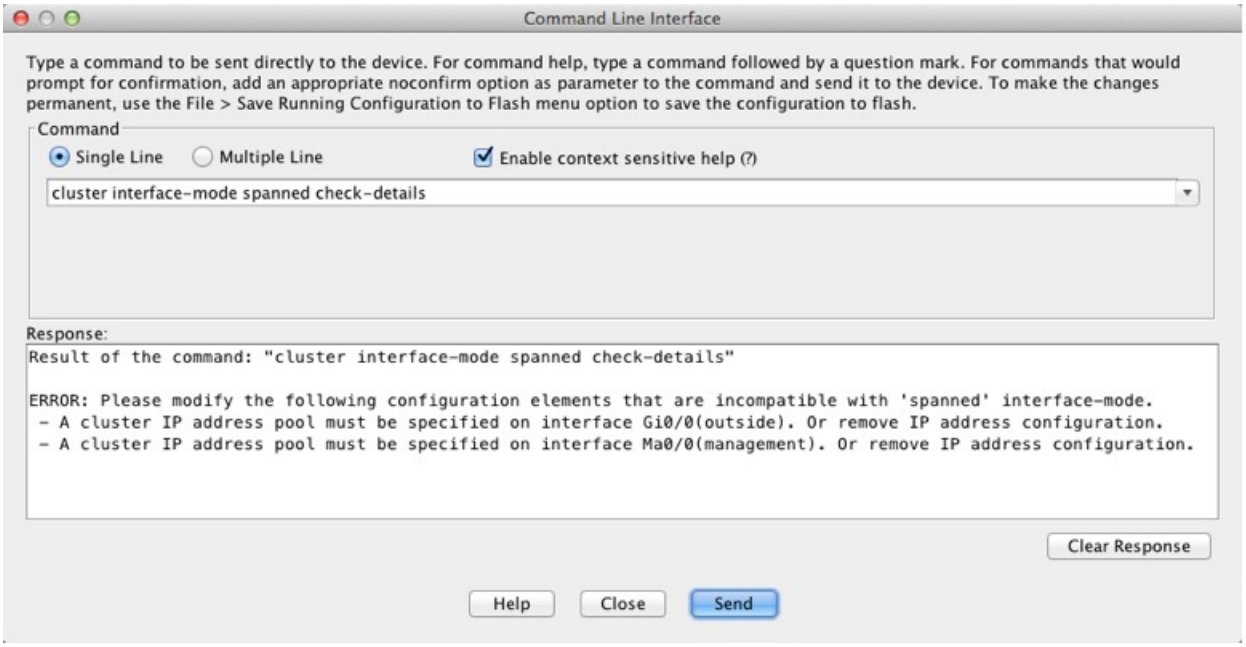
- You can always configure the management-only interface as an Individual interface (recommended). The management interface can be an Individual interface even in transparent firewall mode.
- If you configure the management interface as an Individual interface, you cannot enable dynamic routing for the management interface. You must use a static route.

Procedure

Step 1 In ASDM on the control unit, 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 spanned check-details

Example:

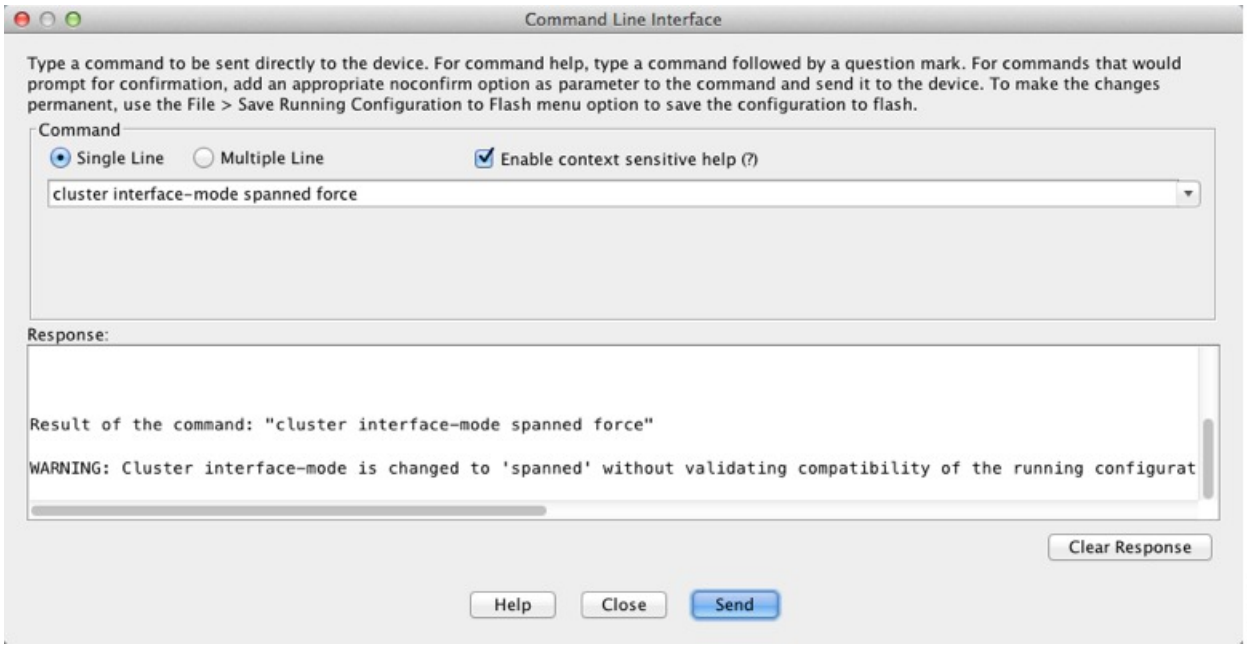


Caution
After you set the interface mode, you can continue to connect to the interface; however, if you reload the ASA before you configure your management interface to comply with clustering requirements (for example, adding a cluster IP pool), you will not be able to reconnect because cluster-incompatible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

Step 2 Set the interface mode for clustering:

cluster interface-mode spanned force

Example:



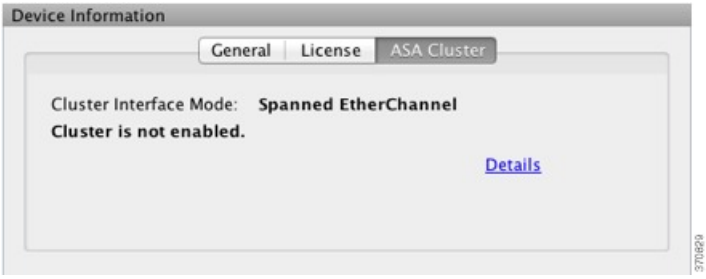
There is no default setting; you must explicitly choose the mode. If you have not set the mode, you cannot enable clustering.

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 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:



(Recommended; Required in Multiple Context Mode) Configure Interfaces on the Control Node

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 must modify the management interface to which ASDM is currently connected. 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 members. In multiple context mode, you must use the procedures in this section to fix existing interfaces or to configure new interfaces. However, in single mode, you can skip this section and configure common interface parameters within the High Availability and Scalability wizard (see Create or Join a Cluster Using the High Availability Wizard, on page 27). Note that advanced interface settings such as creating EtherChannels for Individual interfaces are not available in the wizard.

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

Configure the Management Interface as an Individual Interface

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 control node.

We recommend configuring the management interface as an Individual interface. Individual management interfaces let you connect directly to each unit if necessary, while a Spanned EtherChannel interface only allows connection to the control node.

Before you begin

- For multiple context mode, perform this procedure in each context. If you are not already in the context configuration mode in the Configuration > Device List pane, double-click the context name under the active device IP address.
- (Optional) Configure the interface as a device-local EtherChannel interface, and/or configure subinterfaces.
  - For an EtherChannel, this EtherChannel is local to the unit, and is not a Spanned EtherChannel.
- If you are connecting remotely to the management interface using ASDM, the current IP address of prospective secondary units are for temporary use.
  - Each member will be assigned an IP address from the cluster IP pool defined on the primary unit.
  - The cluster IP pool cannot include addresses already in use on the network, including prospective secondary IP addresses.

    For example:

1. You configure the primary unit to use 10.1.1.1.
2. Other units use 10.1.1.2, 10.1.1.3, and 10.1.1.4.
3. When you configure the cluster IP pool on the primary unit, 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 primary unit; the original .1 address is the main cluster IP address that belongs to the current primary unit.

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

Procedure

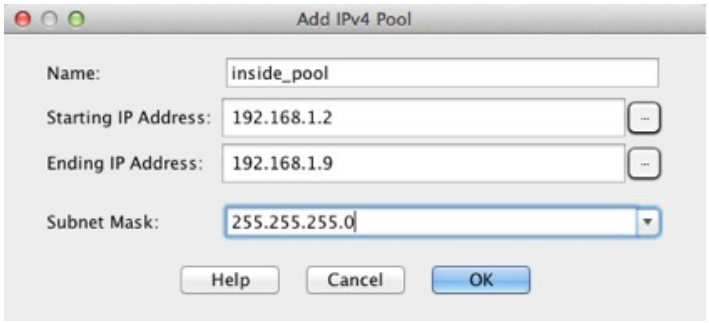
Step 1 Choose the **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces** pane.

Step 2 Choose the interface row, and click **Edit**. Set the interface parameters. See the following guidelines:

- **Dedicate this interface to management only**—Sets the interface to management-only mode so that it does not pass through traffic. By default, Management type interfaces are configured as management-only. In transparent mode, this command is always enabled for a Management type interface.
- **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.

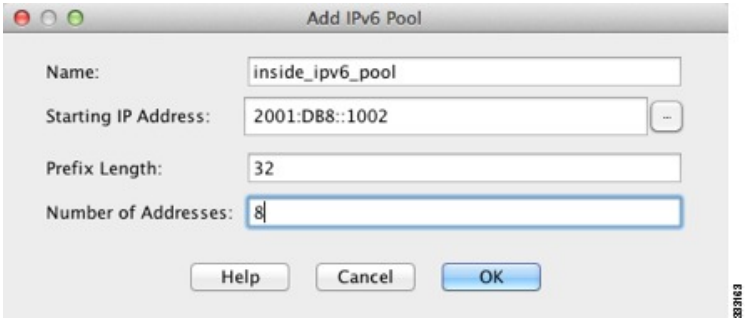
Step 4 To configure an IPv6 address, click the **IPv6** tab.

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**.

Configure Spanned EtherChannels

A Spanned EtherChannel spans all ASAs in the cluster, and provides load balancing as part of the EtherChannel operation.

Before you begin

- You must be in Spanned EtherChannel interface mode.

- For multiple context mode, start this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.
- For transparent mode, configure the bridge group. See Configure the Bridge Virtual Interface (BVI).
- When using Spanned EtherChannels, the port-channel interface will not come up until clustering is fully enabled. This requirement prevents traffic from being forwarded to a unit that is not an active unit in the cluster.

Procedure

Step 1 Depending on your context mode:

- For single mode, choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane.
- For multiple mode in the System execution space, choose the **Configuration > Context Management > Interfaces** pane.

Step 2 Choose **Add** > **EtherChannel Interface**.

The **Add EtherChannel Interface** dialog box appears.

Step 3 Enable the following:

- **Port Channel ID**
- **Enable Interface** (checked by default)
- **Members in Group**—In the **Members in Group** list, you need to add at least one interface. Multiple interfaces in the EtherChannel per unit are useful for connecting to switches in a VSS, vPC, StackWise, or StackWise Virtual.

  Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. ASDM does not prevent you from adding non-matching interfaces.

The rest of the fields on this screen are described later in this procedure.

Step 4 To configure the MAC address and optional parameters, click the **Advanced** tab.

- In the **MAC Address Cloning** area, set a manual global MAC address for the EtherChannel. Do not set the Standby MAC Address; it is ignored. You must configure a MAC address for a Spanned EtherChannel to avoid potential network connectivity problems. With a manually-configured MAC address, the MAC address stays with the current control unit. If you do not configure a MAC address, then if the control unit changes, the new control unit uses a new MAC address for the interface, which can cause a temporary network outage.

  In multiple context mode, if you share an interface between contexts, you should instead enable auto-generation of MAC addresses so you do not need to set the MAC address manually. Note that you must manually configure the MAC address using this command for *non-shared* interfaces.
- (Routed mode) In the **ASA Cluster** area, for inter-site clustering set **Site specific MAC Addresses** and IP addresses for a site by clicking **Add** and specifying a MAC address and IP address for the site ID (1

through 8). Repeat for up to 8 sites. The site-specific IP addresses must be on the same subnet as the global IP address. The site-specific MAC address and IP address used by a unit depends on the site ID you specify in each unit's bootstrap configuration.

Step 5 (Optional) Configure VLAN subinterfaces on this EtherChannel. The rest of this procedure applies to the subinterfaces.

Step 6 (Multiple context mode) Before you complete this procedure, you need to allocate interfaces to contexts.

a) Click **OK** to accept your changes.
b) Allocate interfaces.
c) Change to the context that you want to configure: in the **Device List** pane, double-click the context name under the active device IP address.
d) Choose the **Configuration > Device Setup > Interface Settings > Interfaces** pane, select the port-channel interface that you want to customize, and click **Edit**.

The **Edit Interface** dialog box appears.

Step 7 Click the **General** tab.

Step 8 (Transparent Mode) From the **Bridge Group** drop-down list, choose the bridge group to which you want to assign this interface.

Step 9 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 10 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 11 (Routed Mode) For an IPv4 address, click the **Use Static IP** radio button and enter the IP address and mask. DHCP and PPPoE are not supported. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses. For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

Step 12 (Routed Mode) To configure an IPv6 address, click the **IPv6** tab.

For transparent mode, you configure the IP address for the bridge group interface, not the EtherChannel interface.

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

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

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

c) In the **Address/Prefix Length** field, enter the global IPv6 address and the IPv6 prefix length. For example, 2001:DB8::BA98:0:3210/64.
d) (Optional) To use the Modified EUI-64 interface ID as the host address, check the **EUI-64** check box. In this case, just enter the prefix in the **Address/Prefix Length** field.
e) Click **OK**.

Step 13 Click **OK** to return to the **Interfaces** screen.

Step 14 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.

Note The wizard does not let you change the default cLACP system ID and priority, which can only be configured while clustering is disabled. So either use the wizard and then disable clustering afterwards, or see Configure Basic ASA Cluster Parameters, on page 30 to configure the cluster manually.

Before you begin

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- The interfaces 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 **Interfaces** screen, you cannot create new EtherChannels from this screen (except for the cluster control link).

Step 3 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, adding an additional switch to form a VSS or vPC) 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) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one node in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends heartbeat messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the heartbeat messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

- (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**—Specifies the maximum transmission node for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces, between 1400 and 9198 bytes, but not between 2561 and 8362. Due to block pool handling, this MTU size is not optimal for system operation. The default MTU is 1500 bytes.We suggest setting the MTU to the maximum. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9098, while the cluster control link can be set to 9198.

Step 4 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, adding an additional switch to form a VSS or vPC) 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 5 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 6 Click **Finish**.

Step 7 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.

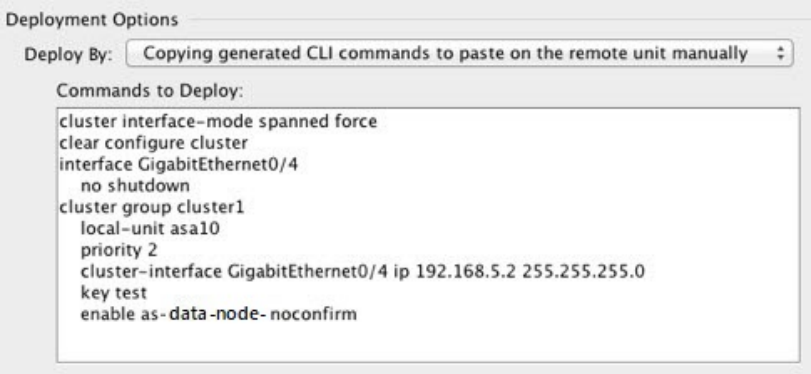
Note
In some cases, there might be an error when joining the cluster after you finish the wizard. If ASDM was disconnected, ASDM will not receive any subsequent errors from the ASA. If clustering remains disabled after you reconnect ASDM, you should connect to the ASA console port to determine the exact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

Step 8 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 9 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.

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.

For parameters that can only be changed while clustering is disabled, you need to disable clustering on all nodes and then re-enable clustering, which requires console CLI access.

Before you begin

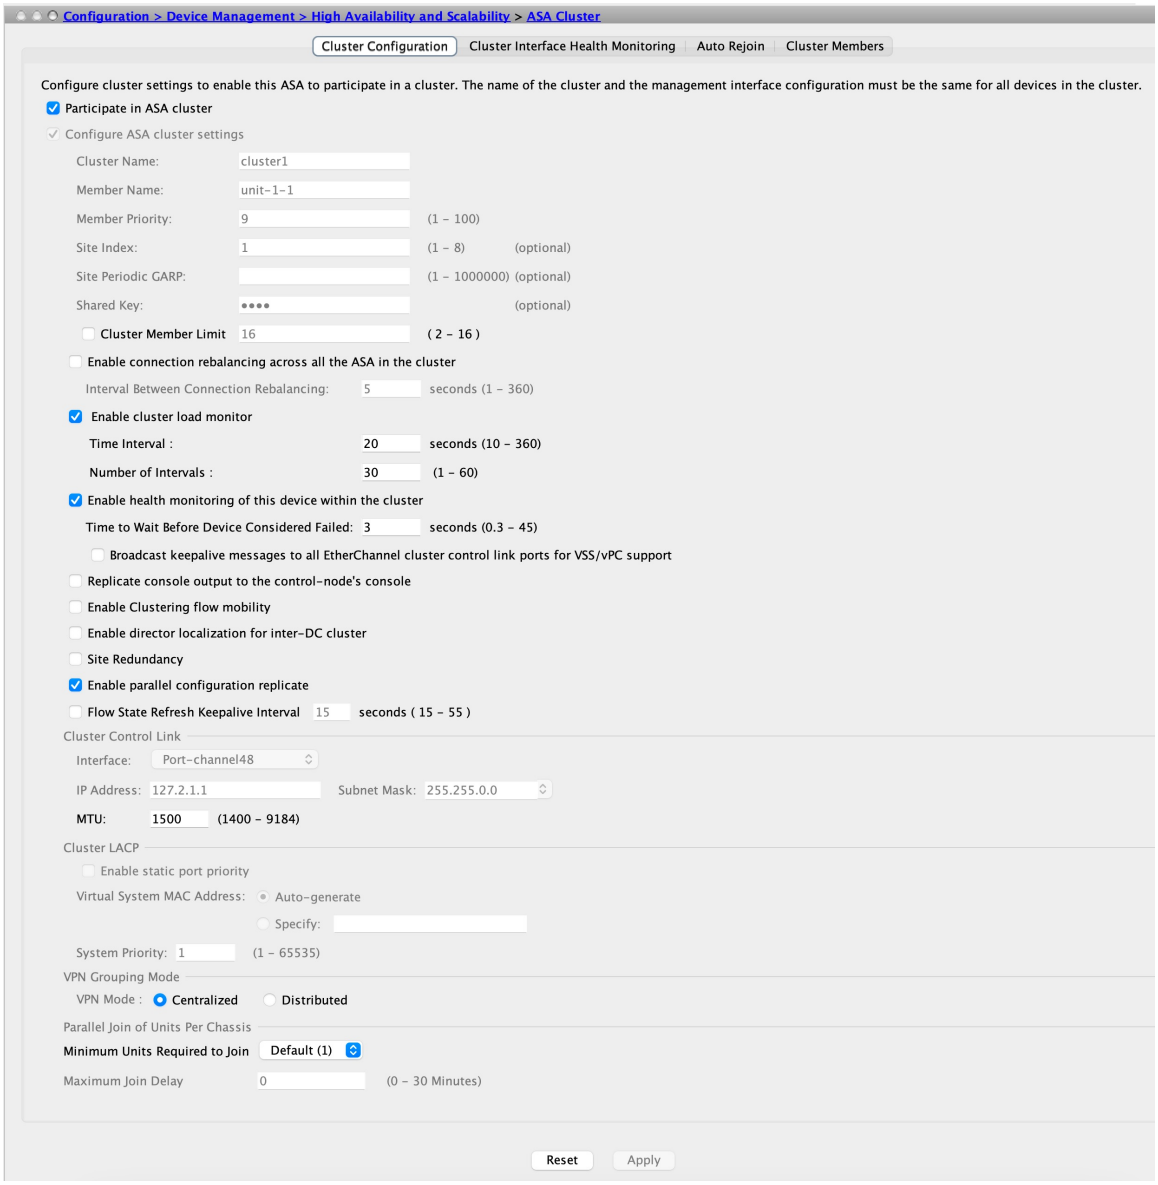
- If you did not use the wizard and want to manually join the cluster, you need to pre-configure the cluster-control-link interfaces on each node before joining the cluster. For a single interface, you must enable it; do not configure any other settings. For an EtherChannel interface, enable it and set the EtherChannel mode to On.
- For multiple context mode, complete this procedure in the system execution space on the control node. If you are not already in the System configuration mode, in the **Configuration** > **Device List** pane, double-click **System** under the active device IP address.

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.

Figure 1: Cluster Configuration

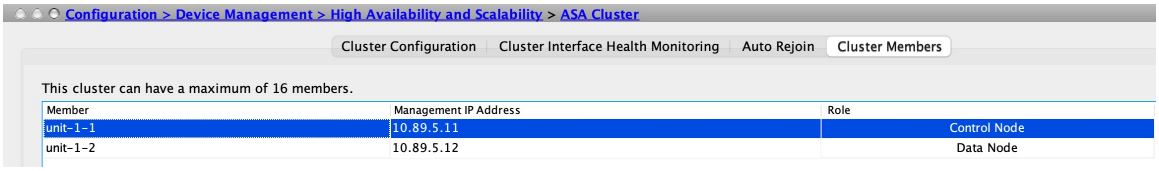


Step 2 If clustering was already enabled, and you want to configure grayed-out settings, then disable clustering on all nodes.

When an ASA 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 node received from the cluster IP pool, not the main IP address.

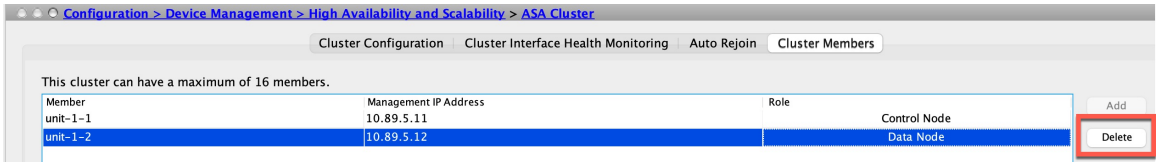
a) You can view the local IP address of the control node by clicking the **Cluster Members** tab. You will need to reconnect to this address after you disable clustering, so note the control node IP address.

Figure 2: Management IP Address of the Control Node



b) On the **Cluster Members** tab, select each data node and then click **Delete**.

Figure 3: Delete Node



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

c) On the **Cluster Configuration** tab, uncheck the **Participate in ASA cluster** check box, and click **Apply**.

You will be disconnected from ASDM.

d) Reconnect ASDM to the control node's local IP address.

e) Continue to make your changes according to the rest of this procedure. To set each data node to rejoin the cluster after finishing this procedure, you can either add each node as a new node, or you can go to the console port of each data node and re-enable clustering at the CLI.

Example:

```
node2(config)# cluster group cluster1
node2(cfg-cluster)# enable
```

Step 3 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 have 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 4 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) **Site Periodic GARP**—The ASA generates gratuitous ARP (GARP) packets to keep the switching infrastructure up to date: the highest priority member at each site periodically generates GARP traffic for the global MAC/IP addresses. GARP is enabled by default when you set the site ID for each node and the site MAC and IP address for each Spanned EtherChannel. Set the GARP interval between 1 and 1000000 seconds. The default is 290 seconds.

  When using per-site MAC and IP addresses, packets sourced from the cluster use a site-specific MAC address and IP address, while packets received by the cluster use a global MAC address and IP address. If traffic is not generated from the global MAC address periodically, you could experience a MAC address timeout on your switches for the global MAC address. After a timeout, traffic destined for the global MAC address will be flooded across the entire switching infrastructure, which can cause performance and security concerns.
- (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) **Broadcast keepalive messages to all EtherChannel cluster control link ports for VSS/vPC support**—If you configure the cluster control link as an EtherChannel (recommended), and it is connected to a VSS or vPC pair, then you might need to enable this option. For some switches, when one node in the VSS/vPC is shutting down or booting up, EtherChannel member interfaces connected to that switch may appear to be Up to the ASA, but they are not passing traffic on the switch side. The ASA can be erroneously removed from the cluster if you set the ASA holdtime timeout to a low value (such as .8 seconds), and the ASA sends heartbeat messages on one of these EtherChannel interfaces. When you enable this option, the ASA floods the heartbeat messages on all EtherChannel interfaces in the cluster control link to ensure that at least one of the switches can receive them.

- (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. In the case of an EtherChannel that transitions from a down state to an up state (for example, the switch reloaded, or the switch enabled an EtherChannel), a longer debounce time can prevent the interface from appearing to be failed on a cluster node just because another cluster node was faster at bundling the ports. 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 39.
- (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. This interface cannot have a name configured; available interfaces are shown in the drop-down list.
  - **Interface**—Specifies the interface ID, preferably an EtherChannel. Subinterfaces and Management type interfaces are not allowed.
  - **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 node for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces, between 1400 and 9198 bytes. The default MTU is 1500 bytes. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead. We suggest setting the cluster control link MTU to the maximum. For example, because the maximum MTU is 9198 bytes, then the highest data interface MTU can be 9098, while the cluster control link can be set to 9198.
- (Optional) **Cluster LACP**—When using Spanned EtherChannels, the ASA uses cLACP to negotiate the EtherChannel with the neighbor switch. ASAs in a cluster collaborate in cLACP negotiation so that they appear as a single (virtual) device to the switch.
  - **Virtual System MAC Address**—Sets the cLACP system ID, which is in the format of a MAC address. All ASAs use the same system ID: auto-generated by the control node (the default) and replicated to all secondaries; or manually specified in the form *H.H.H*, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. However, you can only change this value if you disable clustering.
  - **System Priority**—Sets the system priority, between 1 and 65535. The priority is used to decide which node is in charge of making a bundling decision. By default, the ASA uses priority 1, which is the highest priority. The priority needs to be higher than the priority on the switch. This parameter is not part of the bootstrap configuration, and is replicated from the control node to the data nodes. However, you can only change this value if you disable clustering.

Step 5 Check the **Participate in ASA cluster** check box to join the cluster.

Step 6 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. You can monitor any port-channel ID, redundant ID, or single physical interface ID. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

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 all physical ports for a given logical interface fail on a particular node, but there are active ports under the same logical interface on other nodes, then the node is removed from the cluster. 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 the type of interface and 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. You can specify any port-channel ID, redundant ID, or single physical interface ID. Health monitoring is not performed on VLAN subinterfaces or virtual interfaces such as VNIs or BVIs. You cannot configure monitoring for the cluster control link; it is always monitored.

When any topology changes occur (such as adding or removing a data interface, enabling or disabling an interface on the ASA or the switch, or adding an additional switch to form a VSS, vPC, StackWise, or StackWise Virtual) 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 30.
- 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

- For multiple context mode, complete this procedure in the system execution space. If you are not already in the System configuration mode, in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- If you want to send the bootstrap configuration over the management network, be sure the data node has an accessible IP address.

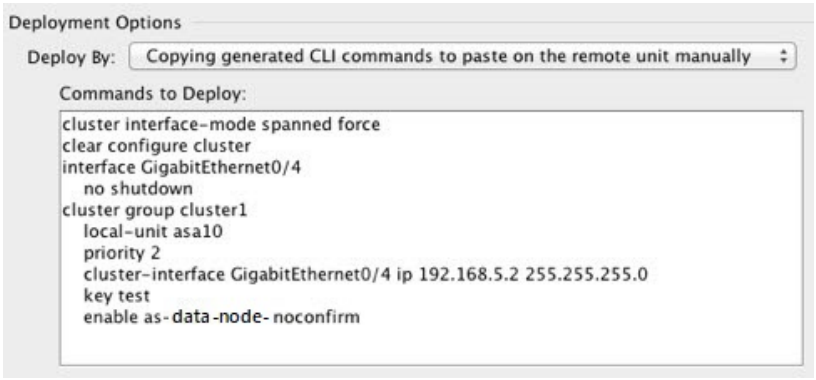
Procedure

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.

Before you begin

- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

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.

Before you begin

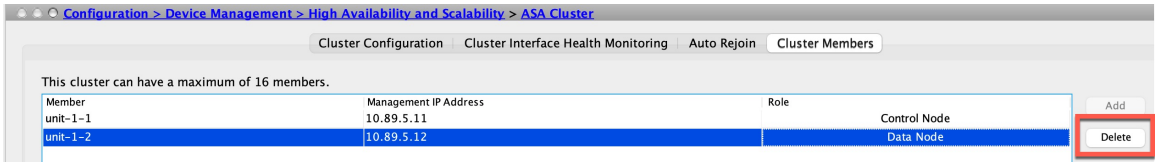
For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.

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.

Before you begin

- You must use the console port to reenable clustering. Other interfaces are shut down. The exception is if you manually disabled clustering in ASDM, then you can reenable clustering in ASDM if you did not save the configuration and reload. After reloading, the management interface is disabled, so console access is the only method to reenable clustering.
- For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the **Configuration > Device List** pane, double-click **System** under the active device IP address.
- Make sure the failure is resolved before you try to 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.

Before you begin

You must use the console port; when you remove the cluster configuration, all interfaces are shut down, including the management interface and cluster control link.

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.

Before you begin

For multiple context mode, perform this procedure in the system execution space. If you are not already in the System configuration mode in the Configuration > Device List pane, double-click **System** under the active device IP address.

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.

The following sample output for the **cluster exec show port-channel** summary command shows EtherChannel information for each node in the cluster:

```
ciscoasa# cluster exec show port-channel summary
control node(LOCAL):**************************************************************
 Number of channel-groups in use: 2
Group  Port-channel  Protocol  Span-cluster  Ports
------+-------------+-----------+-----------------------------------------------
1         Po1             LACP       Yes  Gi0/0(P)
2         Po2             LACP       Yes  Gi0/1(P)
 slave:***********************************************************************
 Number of channel-groups in use: 2
```

```
Group  Port-channel  Protocol  Span-cluster   Ports
------+-------------+-----------+-----------------------------------------------
1         Po1             LACP       Yes   Gi0/0(P)
2         Po2             LACP       Yes   Gi0/1(P)
```

Monitoring the ASA 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 Clustering

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

Sample ASA and Switch Configuration

The following sample configurations connect the following interfaces between the ASA and the switch:

ASA Interface	Switch Interface
Ethernet 1/2	GigabitEthernet 1/0/15
Ethernet 1/3	GigabitEthernet 1/0/16

ASA Interface	Switch Interface
Ethernet 1/4	GigabitEthernet 1/0/17
Ethernet 1/5	GigabitEthernet 1/0/18

ASA Configuration

Interface Mode on Each Unit

```
cluster interface-mode spanned force
```

ASA1 Control Unit Bootstrap Configuration

```
interface Ethernet1/6
 channel-group 1 mode on
 no shutdown
!
interface Ethernet1/7
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
!
cluster group Moya
 local-unit A
 cluster-interface Port-channel1 ip 10.0.0.1 255.255.255.0
 priority 10
 key emphyri0
 enable noconfirm
```

ASA2 Data Unit Bootstrap Configuration

```
interface Ethernet1/6
 channel-group 1 mode on
 no shutdown
!
interface Ethernet1/7
 channel-group 1 mode on
 no shutdown
!
interface Port-channel1
 description Clustering Interface
!
cluster group Moya
 local-unit B
 cluster-interface Port-channel1 ip 10.0.0.2 255.255.255.0
 priority 11
 key emphyri0
 enable as-data-node
```

Control Unit Interface Configuration

```
ip local pool mgmt-pool 10.53.195.231-10.53.195.232

interface Ethernet1/2
 channel-group 10 mode active
 no shutdown
!
interface Ethernet1/3
 channel-group 10 mode active
 no shutdown
!
interface Ethernet1/4
 channel-group 11 mode active
 no shutdown
!
interface Ethernet1/5
 channel-group 11 mode active
 no shutdown
!
interface Management1/1
 management-only
 nameif management
 ip address 10.53.195.230 cluster-pool mgmt-pool
 security-level 100
 no shutdown
!
interface Port-channel10
 mac-address aaaa.bbbb.cccc
 nameif inside
 security-level 100
 ip address 209.165.200.225 255.255.255.224
!
interface Port-channel11
 mac-address aaaa.dddd.cccc
 nameif outside
 security-level 0
 ip address 209.165.201.1 255.255.255.224
```

Cisco IOS Switch Configuration

```
interface GigabitEthernet1/0/15
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/16
 switchport access vlan 201
 switchport mode access
 spanning-tree portfast
 channel-group 10 mode active
!
interface GigabitEthernet1/0/17
 switchport access vlan 401
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active
!
interface GigabitEthernet1/0/18
 switchport access vlan 401
```

```
 switchport mode access
 spanning-tree portfast
 channel-group 11 mode active

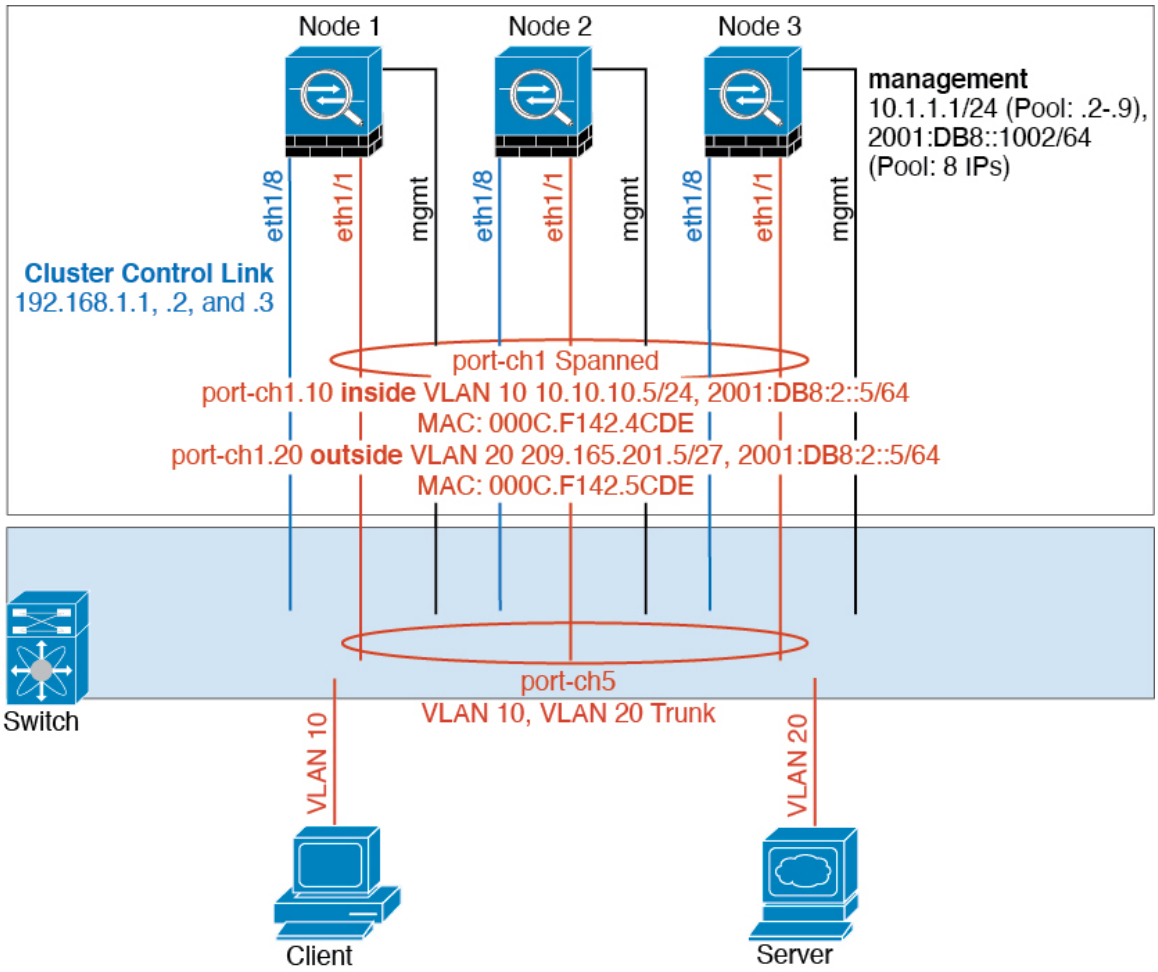
interface Port-channel10
 switchport access vlan 201
 switchport mode access

interface Port-channel11
 switchport access vlan 401
 switchport mode access
```

Firewall on a stick

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

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



Interface mode on each unit

```
cluster interface-mode spanned force
```

Unit 1 control unit bootstrap configuration

```
interface ethernet1/8
no shutdown
description CCL

cluster group cluster1
local-unit asa1
cluster-interface ethernet1/8 ip 192.168.1.1 255.255.255.0
priority 1
key chuntheunavoidable
enable noconfirm
```

Unit 2 data unit bootstrap configuration

```
interface ethernet1/8
no shutdown
description CCL

cluster group cluster1
local-unit asa2
cluster-interface ethernet1/8 ip 192.168.1.2 255.255.255.0
priority 2
key chuntheunavoidable
enable as-data-node
```

Unit 3 data unit bootstrap configuration

```
interface ethernet1/8
no shutdown
description CCL

cluster group cluster1
local-unit asa3
cluster-interface ethernet1/8 ip 192.168.1.3 255.255.255.0
priority 3
key chuntheunavoidable
enable as-data-node
```

Control unit interface configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

interface management 1/1
nameif management
ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
security-level 100
management-only
```

```
no shutdown

interface ethernet1/1
channel-group 1 mode active
no shutdown

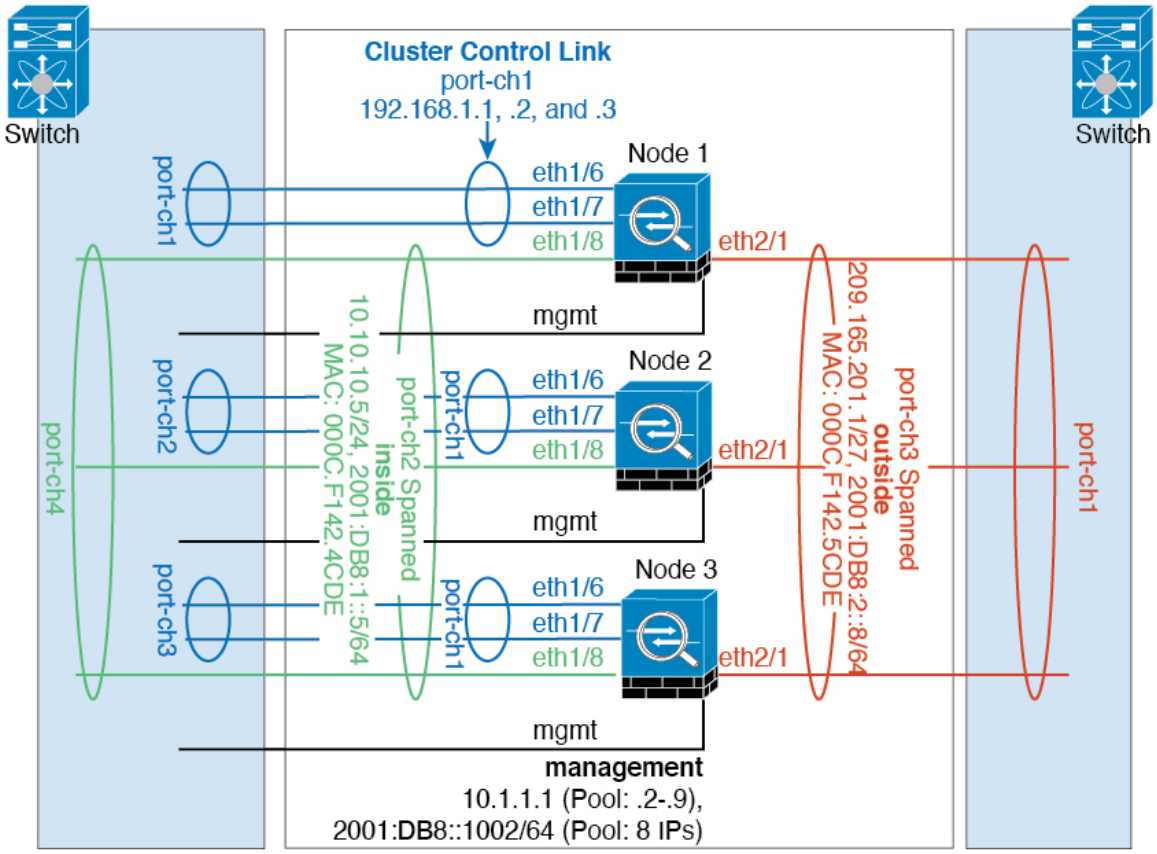
interface port-channel 1

interface port-channel 1.10
vlan 10
nameif inside
ip address 10.10.10.5 255.255.255.0
ipv6 address 2001:DB8:1::5/64
mac-address 000C.F142.4CDE

interface port-channel 1.20
vlan 20
nameif outside
ip address 209.165.201.1 255.255.255.224
ipv6 address 2001:DB8:2::8/64
mac-address 000C.F142.5CDE
```

Traffic segregation

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



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

Interface mode on each unit

```
cluster interface-mode spanned force
```

Unit 1 control unit bootstrap configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa1
 cluster-interface port-channel1 ip 192.168.1.1 255.255.255.0
 priority 1
 key chuntheunavoidable
 enable noconfirm
```

Unit 2 data unit bootstrap configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa2
 cluster-interface port-channel1 ip 192.168.1.2 255.255.255.0
 priority 2
 key chuntheunavoidable
 enable as-data-node
```

Unit 3 data unit bootstrap configuration

```
interface ethernet 1/6
 channel-group 1 mode on
 no shutdown

interface ethernet 1/7
```

```
 channel-group 1 mode on
 no shutdown

interface port-channel 1
 description CCL

cluster group cluster1
 local-unit asa3
 cluster-interface port-channel1 ip 192.168.1.3 255.255.255.0
 priority 3
 key chuntheunavoidable
 enable as-data-node
```

Control unit interface configuration

```
ip local pool mgmt 10.1.1.2-10.1.1.9
ipv6 local pool mgmtipv6 2001:DB8::1002/64 8

interface management 1/1
 nameif management
 ip address 10.1.1.1 255.255.255.0 cluster-pool mgmt
 ipv6 address 2001:DB8::1001/32 cluster-pool mgmtipv6
 security-level 100
 management-only
 no shutdown

interface ethernet 1/8
 channel-group 2 mode active
 no shutdown

interface port-channel 2
 nameif inside
 ip address 10.10.10.5 255.255.255.0
 ipv6 address 2001:DB8:1::5/64
 mac-address 000C.F142.4CDE

interface ethernet 2/1
 channel-group 3 mode active
 no shutdown

interface port-channel 3
 nameif outside
 ip address 209.165.201.1 255.255.255.224
 ipv6 address 2001:DB8:2::8/64
 mac-address 000C.F142.5CDE
```

OTV Configuration for Routed Mode Inter-Site Clustering

The success of inter-site clustering for routed mode with Spanned EtherChannels depends on the proper configuration and monitoring of OTV. OTV plays a major role by forwarding the packets across the DCI. OTV forwards unicast packets across the DCI only when it learns the MAC address in its forwarding table. If the MAC address is not learned in the OTV forwarding table, it will drop the unicast packets.

Sample OTV Configuration

```
//Sample OTV config:
//3151 – Inside VLAN, 3152 – Outside VLAN, 202 – CCL VLAN
```

```
//aaaa.1111.1234 – ASA inside interface global vMAC
//0050.56A8.3D22 – Server MAC

feature ospf
feature otv

mac access-list ALL_MACs
  10 permit any any
mac access-list HSRP_VMAC
  10 permit aaaa.1111.1234 0000.0000.0000 any
  20 permit aaaa.2222.1234 0000.0000.0000 any
  30 permit any aaaa.1111.1234 0000.0000.0000
  40 permit any aaaa.2222.1234 0000.0000.0000
vlan access-map Local 10
  match mac address HSRP_VMAC
  action drop
vlan access-map Local 20
  match mac address ALL_MACs
  action forward
vlan filter Local vlan-list 3151-3152


//To block global MAC with ARP inspection:
arp access-list HSRP_VMAC_ARP
  10 deny aaaa.1111.1234 0000.0000.0000 any
  20 deny aaaa.2222.1234 0000.0000.0000 any
  30 deny any aaaa.1111.1234 0000.0000.0000
  40 deny any aaaa.2222.1234 0000.0000.0000
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP 3151-3152

no ip igmp snooping optimise-multicast-flood
vlan 1,202,1111,2222,3151-3152

otv site-vlan 2222
mac-list GMAC_DENY seq 10 deny aaaa.aaaa.aaaa ffff.ffff.ffff
mac-list GMAC_DENY seq 20 deny aaaa.bbbb.bbbb ffff.ffff.ffff
mac-list GMAC_DENY seq 30 permit 0000.0000.0000 0000.0000.0000
route-map stop-GMAC permit 10
  match mac-list GMAC_DENY

interface Overlay1
  otv join-interface Ethernet8/1
  otv control-group 239.1.1.1
  otv data-group 232.1.1.0/28
  otv extend-vlan 202, 3151
  otv arp-nd timeout 60
  no shutdown

interface Ethernet8/1
  description uplink_to_OTV_cloud
  mtu 9198
  ip address 10.4.0.18/24
  ip igmp version 3
  no shutdown

interface Ethernet8/2

interface Ethernet8/3
  description back_to_default_vdc_e6/39
  switchport
    switchport mode trunk
    switchport trunk allowed vlan 202,2222,3151-3152
  mac packet-classify
```

```
  no shutdown

otv-isis default
  vpn Overlay1
    redistribute filter route-map stop-GMAC
otv site-identifier 0x2
//OTV flood not required for ARP inspection:
otv flood mac 0050.56A8.3D22 vlan 3151
```

OTV Filter Modifications Required Because of Site Failure

If a site goes down, the filters need to be removed from OTV because you do not want to block the global MAC address anymore. There are some additional configurations required.

You need to add a static entry for the ASA global MAC address on the OTV switch in the site that is functional. This entry will let the OTV at the other end add these entries on the overlay interface. This step is required because if the server and client already have the ARP entry for the ASA, which is the case for existing connections, then they will not send the ARP again. Therefore, OTV will not get a chance to learn the ASA global MAC address in its forwarding table. Because OTV does not have the global MAC address in its forwarding table, and per OTV design it will not flood unicast packets over the overlay interface, then it will drop the unicast packets to the global MAC address from the server, and the existing connections will break.

```
//OTV filter configs when one of the sites is down

mac-list GMAC_A seq 10 permit 0000.0000.0000 0000.0000.0000
route-map a-GMAC permit 10
  match mac-list GMAC_A

otv-isis default
  vpn Overlay1
    redistribute filter route-map a-GMAC

no vlan filter Local vlan-list 3151

//For ARP inspection, allow global MAC:
arp access-list HSRP_VMAC_ARP_Allow
  50 permit ip any mac
ip arp inspection filter HSRP_VMAC_ARP_Allow 3151-3152

mac address-table static aaaa.1111.1234 vlan 3151 interface Ethernet8/3
//Static entry required only in the OTV in the functioning Site
```

When the other site is restored, you need to add the filters back again and remove this static entry on the OTV. It is very important to clear the dynamic MAC address table on both the OTVs to clear the overlay entry for the global MAC address.

MAC Address Table Clearing

When a site goes down, and a static entry for the global MAC address is added to OTV, you need to let the other OTV learn the global MAC address on the overlay interface. After the other site comes up, these entries should be cleared. Make sure to clear the mac address table to make sure OTV does not have these entries in its forwarding table.

```
cluster-N7k6-OTV# show mac address-table
Legend:
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
```

```
age - seconds since last seen,+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False
VLAN MAC Address Type age Secure NTFY Ports/SWID.SSID.LID
---------+-----------------+--------+---------+------+----+------------------
G -    d867.d900.2e42 static  -  F F sup-eth1(R)
O 202  885a.92f6.44a5 dynamic -  F F Overlay1
* 202  885a.92f6.4b8f dynamic 5  F F Eth8/3
O 3151 0050.5660.9412 dynamic -  F F Overlay1
* 3151 aaaa.1111.1234 dynamic 50 F F Eth8/3
```

OTV ARP Cache Monitoring

OTV maintains an ARP cache to proxy ARP for IP addresses that it learned across the OTV interface.

```
cluster-N7k6-OTV# show otv arp-nd-cache
OTV ARP/ND L3->L2 Address Mapping Cache

Overlay Interface Overlay1
VLAN MAC Address Layer-3 Address Age Expires In
3151 0050.5660.9412 10.0.0.2 1w0d 00:00:31
cluster-N7k6-OTV#
```

Examples for Inter-Site Clustering

The following examples show supported cluster deployments.

Spanned EtherChannel Routed Mode Example with Site-Specific MAC and IP Addresses

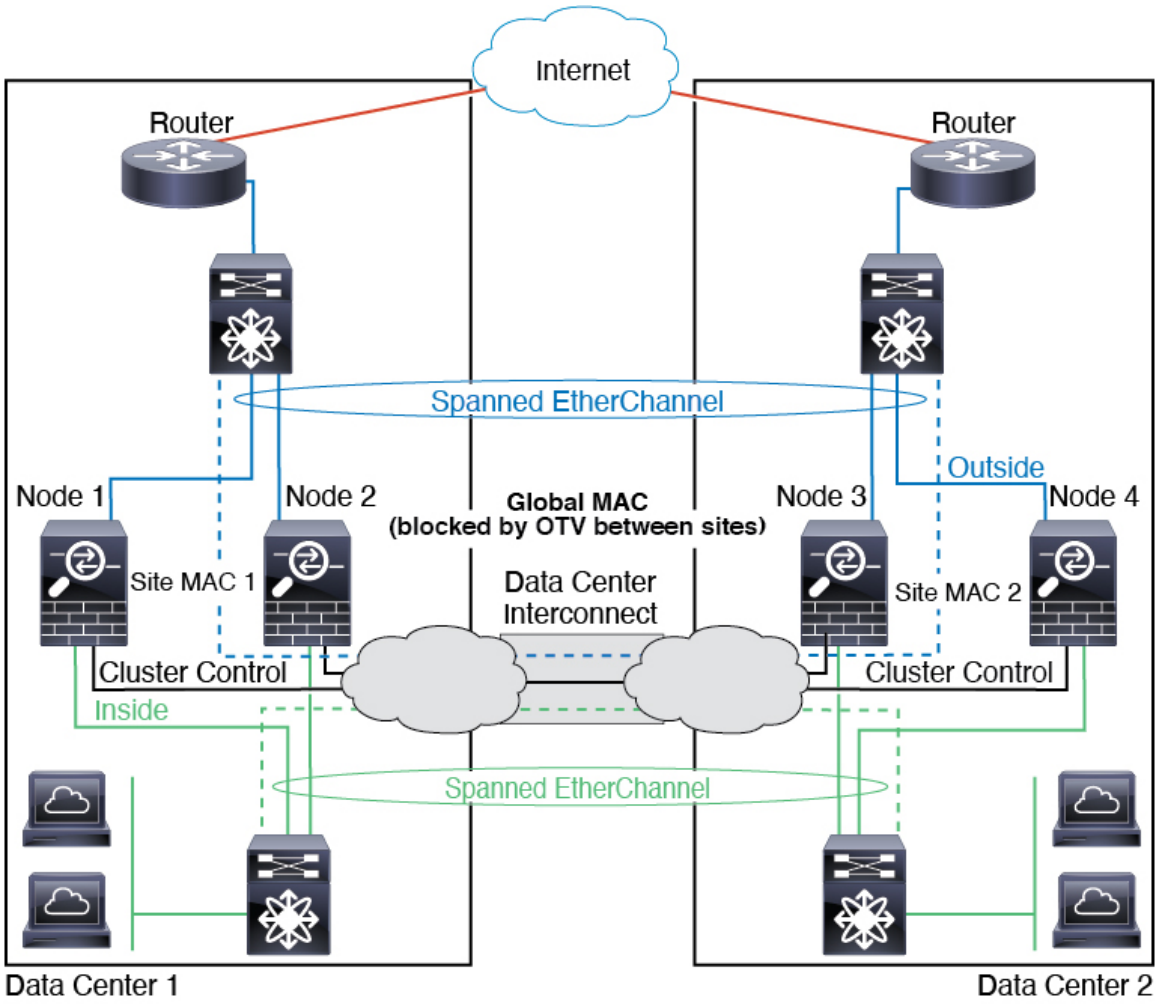
The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and an inside network at each site (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the inside and outside networks. Each EtherChannel is spanned across all chassis in the cluster.

The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters blocking the global MAC address to prevent traffic from traversing the DCI to the other site when the traffic is destined for the cluster. If the cluster nodes at one site become unreachable, you must remove the filters so traffic can be sent to the other site's cluster nodes. You should use VACLs to filter the global MAC address.For some switches, such as Nexus with the F3-series line card, you must also use ARP inspection to block ARP packets from the global MAC address. ARP inspection requires you to set both the site MAC address and the site IP address on the ASA. If you only configure the site MAC address be sure to disable ARP inspection.

The cluster acts as the gateway for the inside networks. The global virtual MAC, which is shared across all cluster nodes, is used only to receive packets. Outgoing packets use a site-specific MAC address from each DC cluster. This feature prevents the switches from learning the same global MAC address from both sites on two different ports, which causes MAC flapping; instead, they only learn the site MAC address.

In this scenario:

- All egress packets sent from the cluster use the site MAC address and are localized at the data center.
- All ingress packets to the cluster are sent using the global MAC address, so they can be received by any of the nodes at both sites; filters at the OTV localize the traffic within the data center.



Spanned EtherChannel Transparent Mode North-South Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between inside and outside routers (North-South insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

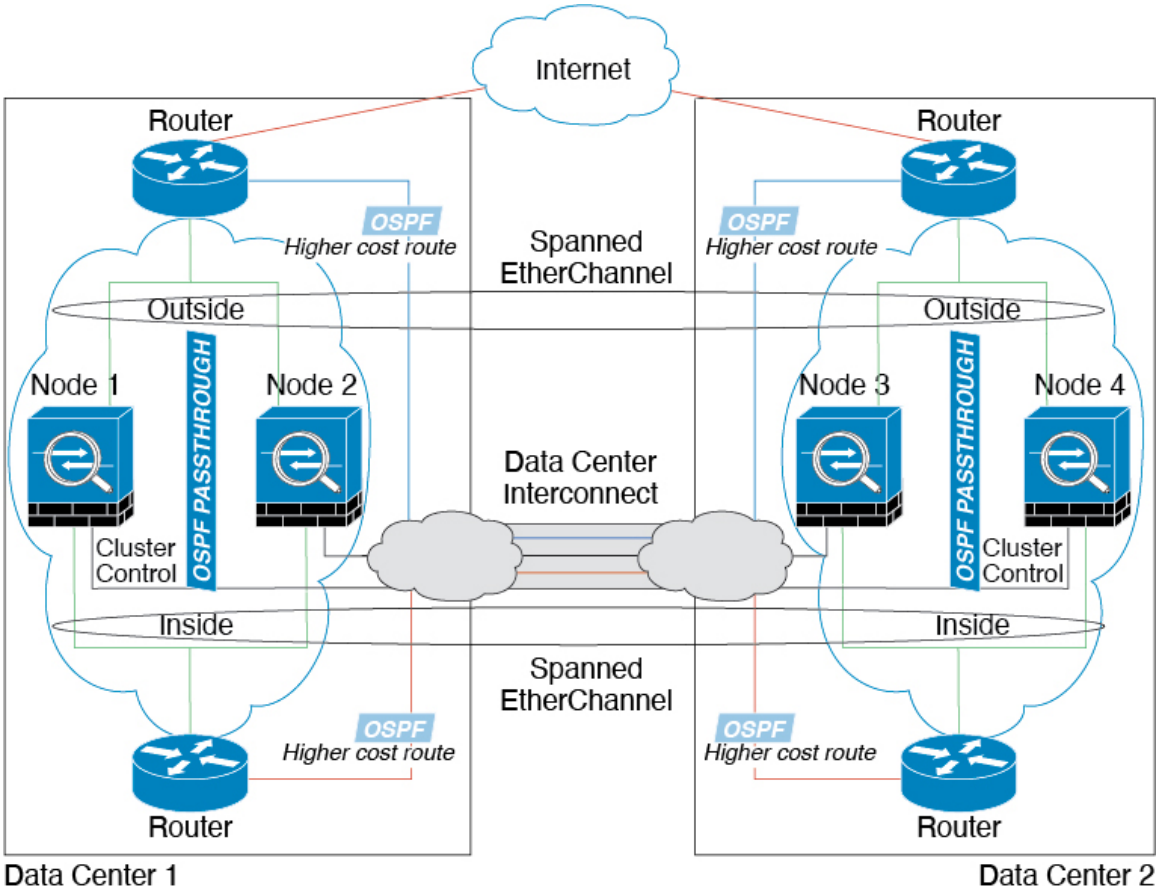
The inside and outside routers at each data center use OSPF, which is passed through the transparent ASAs. Unlike MACs, router IPs are unique on all routers. By assigning a higher cost route across the DCI, traffic stays within each data center unless all cluster members at a given site go down. The lower cost route through the ASAs must traverse the same bridge group at each site for the cluster to maintain asymmetric connections. In the event of a failure of all cluster members at one site, traffic goes from each router over the DCI to the cluster members at the other site.

The implementation of the switches at each site can include:

- Inter-site VSS, vPC, StackWise, or StackWise Virtual—In this scenario, you install one switch at Data Center 1, and the other at Data Center 2. One option is for the cluster nodes at each Data Center to only connect to the local switch, while the redundant switch traffic goes across the DCI. In this case, connections are for the most part kept local to each datacenter. You can optionally connect each node to both switches

across the DCI if the DCI can handle the extra traffic. In this case, traffic is distributed across the data centers, so it is essential for the DCI to be very robust.

- Local VSS, vPC, StackWise, or StackWise Virtual at each site—For better switch redundancy, you can install 2 separate redundant switch pairs at each site. In this case, although the cluster nodes still have a spanned EtherChannel with Data Center 1 chassis connected only to both local switches, and Data Center 2 chassis connected to those local switches, the spanned EtherChannel is essentially "split." Each local redundant switch system sees the spanned EtherChannel as a site-local EtherChannel.

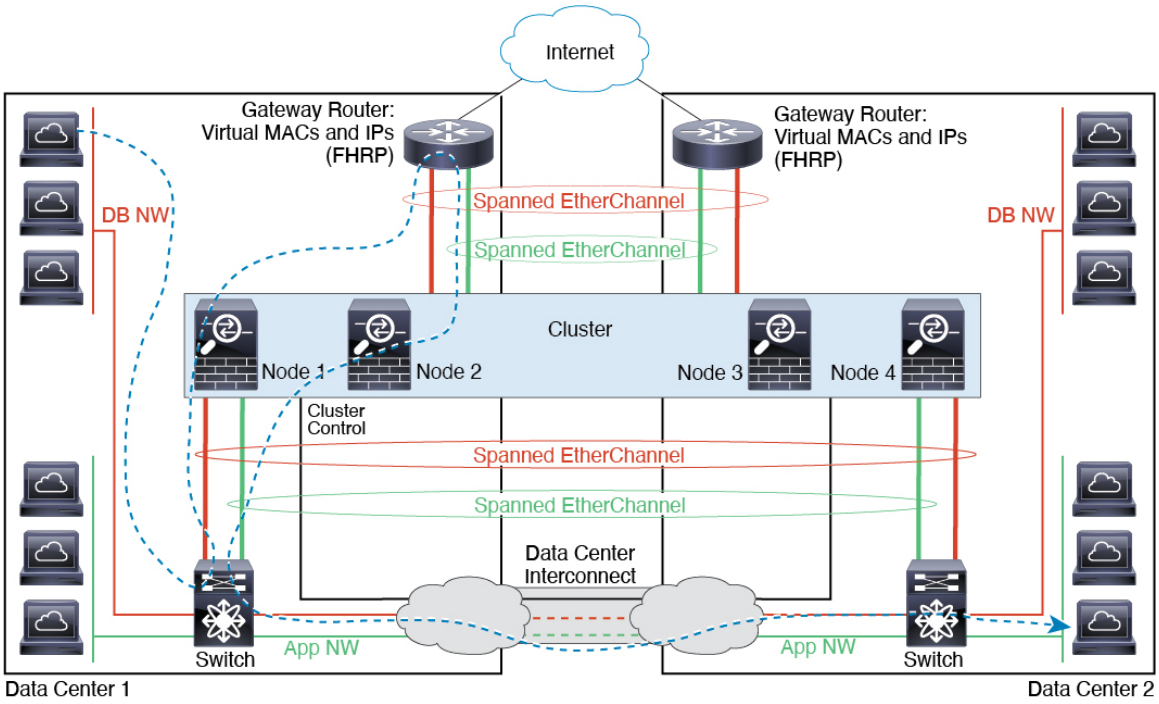


Spanned EtherChannel Transparent Mode East-West Inter-Site Example

The following example shows 2 cluster members at each of 2 data centers placed between the gateway router and two inside networks at each site, the App network and the DB network (East-West insertion). The cluster members are connected by the cluster control link over the DCI. The cluster members at each site connect to the local switches using spanned EtherChannels for both the App and DB networks on the inside and outside. Each EtherChannel is spanned across all chassis in the cluster.

The gateway router at each site uses an FHRP such as HSRP to provide the same destination virtual MAC and IP addresses at each site. A good practice to avoid unintended MAC address flapping is to statically add the gateway routers real MAC addresses to the ASA MAC address table. Without these entries, if the gateway at site 1 communicates with the gateway at site 2, that traffic might pass through the ASA and attempt to reach site 2 from the inside interface and cause problems. The data VLANs are extended between the sites using Overlay Transport Virtualization (OTV) (or something similar). You must add filters to prevent traffic from

traversing the DCI to the other site when the traffic is destined for the gateway router. If the gateway router at one site becomes unreachable, you must remove the filters so traffic can be sent to the other site's gateway router.



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
- IGMP multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- PIM multicast control plane protocol processing (data plane forwarding is distributed across the cluster)
- Dynamic routing

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.

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

Multicast Routing in Spanned EtherChannel Mode

In Spanned EtherChannel mode: The control unit handles all multicast routing packets and data packets until fast-path forwarding is established. After the connection is established, each data unit can forward multicast data packets.

NAT and clustering

NAT can impact the overall throughput of the cluster. Inbound and outbound NAT packets can be sent to different devices 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 device 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.

Dynamic Routing and Clustering

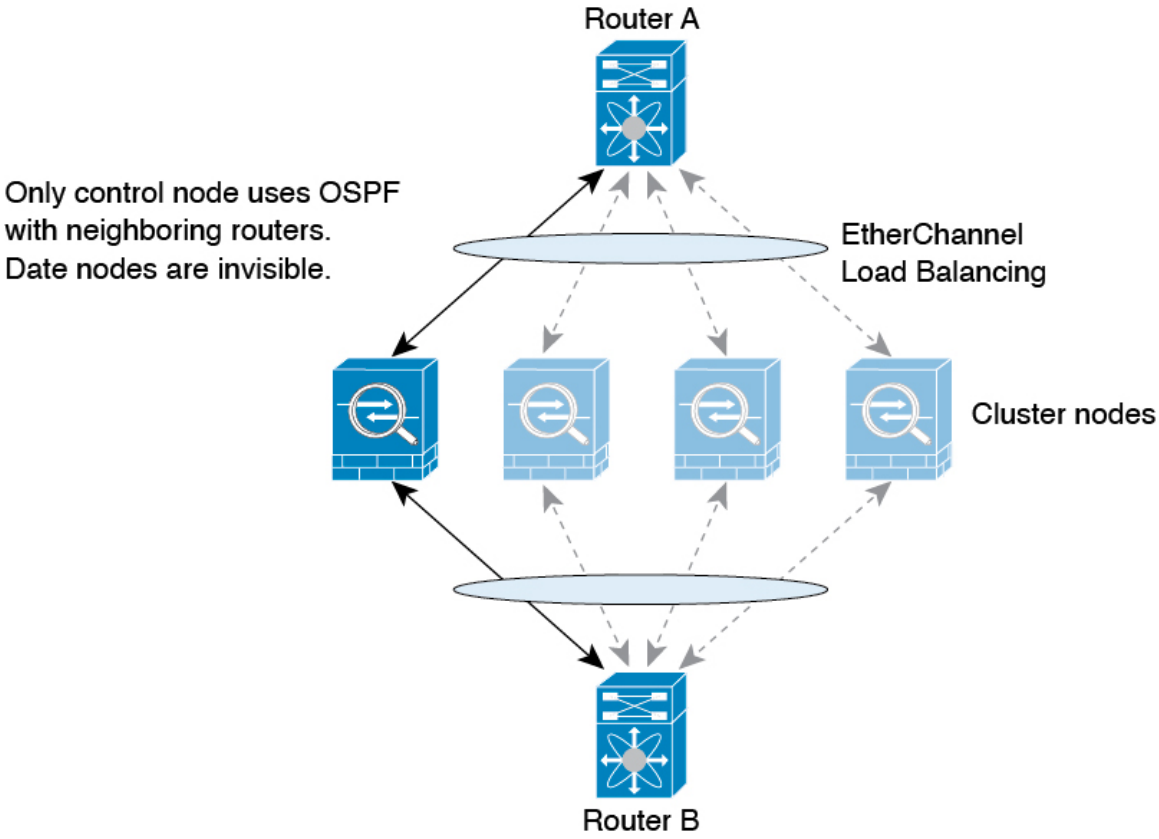
This section describes how to use dynamic routing with clustering.

Dynamic routing in spanned EtherChannel mode

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

Note IS-IS is not supported in Spanned EtherChannel mode.

Figure 5: Dynamic routing in spanned EtherChannel mode



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

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

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

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 and NetFlow and clustering is a network management configuration that enables logging and flow monitoring functionality within clustered network environments.

- 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.

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

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

Performance scaling factor

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

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

Summary

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

Workflow

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 clusters

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.

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

When you enable health monitoring, all physical interfaces (including the main EtherChannel) are monitored by default; you can optionally disable monitoring per interface. Only named interfaces can be monitored. For example, the named EtherChannel must fail to be considered failed, which means all member ports of an EtherChannel must fail to trigger cluster removal (depending on your minimum port bundling setting).

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. If the interface is down on an established member, then the ASA removes the member after 9 seconds. 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. For non-EtherChannels, the node is removed after 500 ms, regardless of the member state.

Status after failure

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

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

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

Note When the device 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.

See Configure Basic ASA Cluster Parameters, on page 30.

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)	Yes	Backup session becomes the active session, then a new backup session is created.

How the cluster manages connections

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

Connection roles

These roles are 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

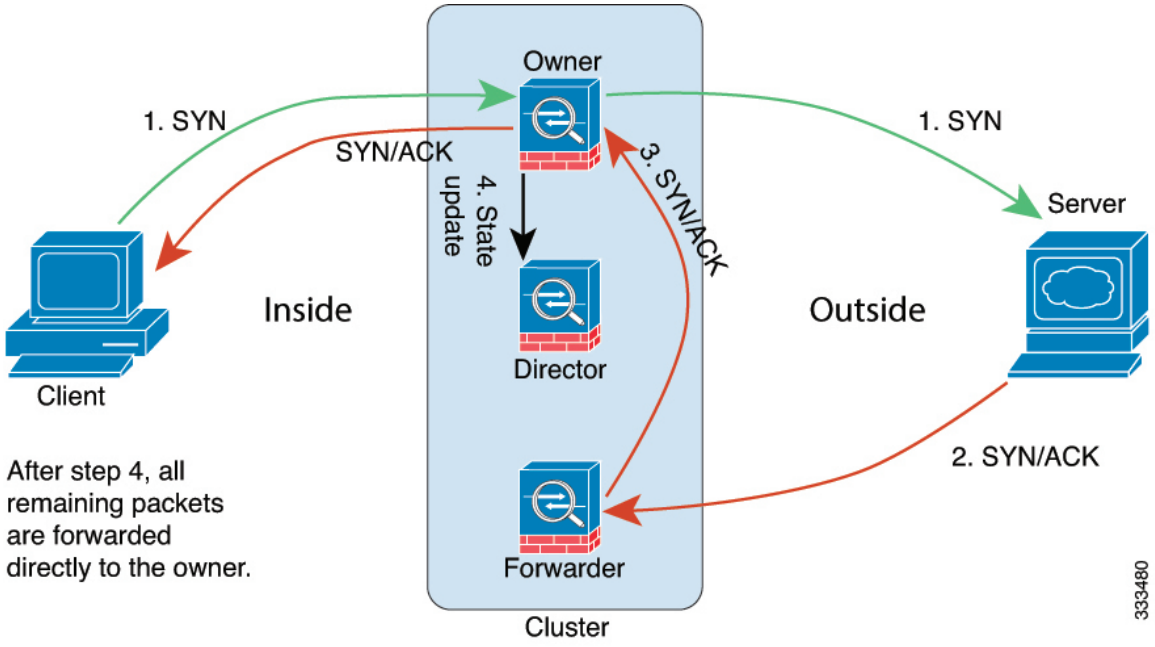
When a new connection is directed to a node of the cluster via load balancing, that node owns both directions of the connection. If any connection packets arrive at a different node, they are forwarded to the owner node over the cluster control link. For best performance, proper external load balancing 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

Summary

This example shows the establishment of a new connection.

Workflow



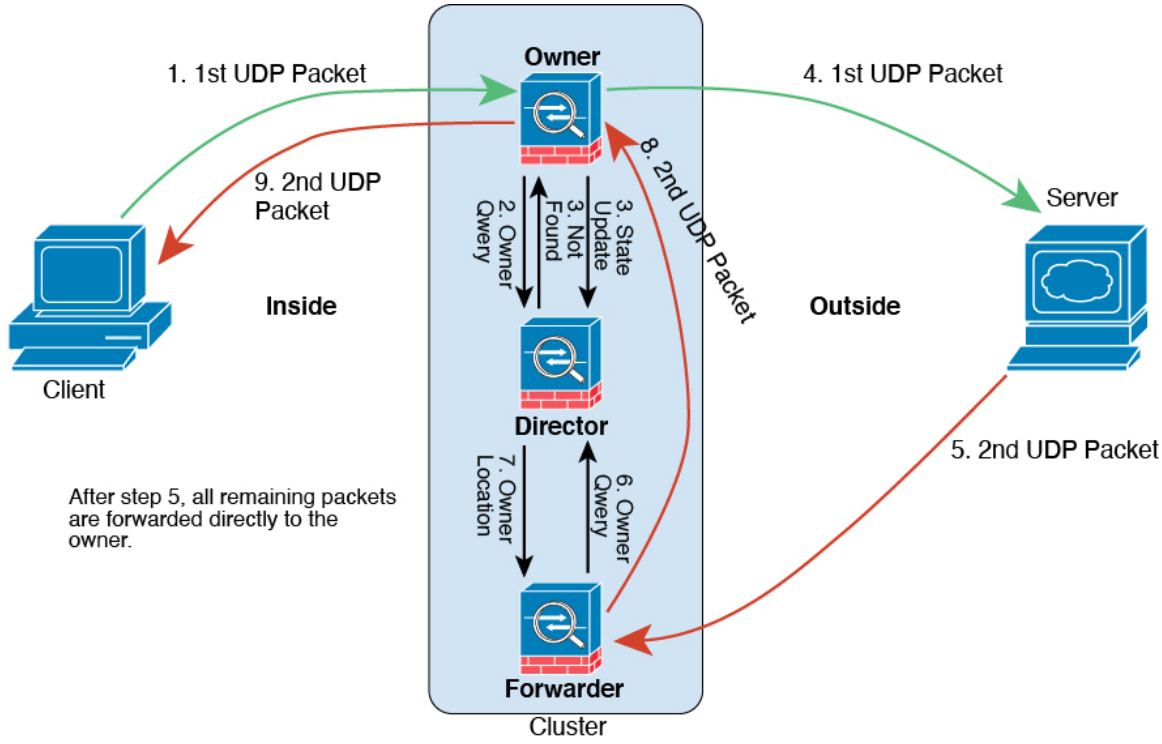
1. The SYN packet originates from the client and is delivered to one device (based on the load balancing method), which becomes the owner. The owner creates a flow, encodes owner information into a SYN cookie, and forwards the packet to the server.
2. The SYN-ACK packet originates from the server and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

Summary

This example shows the establishment of a new connection.

Workflow



1. The first UDP packet originates from the client and is delivered to one device (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 Clustering for the Secure Firewall 3100/4200

Feature Name	Version	Feature Information
Distributed site-to-site VPN with clustering on the Secure Firewall 4200	9.23(1)	An ASA cluster on the Secure Firewall 4200 supports site-to-site VPN in distributed mode. Distributed mode provides the ability to have many site-to-site IPsec IKEv2 VPN connections distributed across members of an ASA cluster, not just on the control node (as in centralized mode). This significantly scales VPN support beyond centralized VPN capabilities and provides high availability. Added/modified screens: **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > VPN > VPN Statistics > Sessions** **Configuration > Device Management > High Availablility and Scalability > ASA Cluster** **Wizards > Site-to-Site** **Monitoring > VPN > VPN Statistics > Sessions** **Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary** **Monitoring > ASA Cluster > ASA Cluster > System Resource Graphs > CPU/Memory** **Monitoring > Logging > Real-Time Log Viewer**
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**
Support for clustering on the Secure Firewall 4200 was introduced	9.20(1)	You can cluster up to 8 Secure Firewall 4200 nodes in Spanned EtherChannel mode.
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**
Support for clustering on the Secure Firewall 3100 was introduced	9.17(1)	You can cluster up to 8 Secure Firewall 3100 nodes in Spanned EtherChannel mode.