



ASA Cluster for the Firepower 4100/9300 Chassis

Clustering lets you group multiple Firepower 4100/9300 chassis ASAs together as a single logical device. The Firepower 4100/9300 chassis series includes the Firepower 9300 and Firepower 4100 series. 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 47](#).

- [About Clustering on the Firepower 4100/9300 Chassis, on page 1](#)
- [Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis, on page 8](#)
- [Licenses for Clustering on the Firepower 4100/9300 Chassis, on page 9](#)
- [Clustering Guidelines and Limitations, on page 11](#)
- [Configure Clustering on the Firepower 4100/9300 Chassis, on page 16](#)
- [FXOS: Remove a Cluster Unit, on page 39](#)
- [ASA: Manage Cluster Members, on page 40](#)
- [ASA: Monitoring the ASA Cluster on the Firepower 4100/9300 chassis, on page 44](#)
- [Troubleshooting Distributed S2S VPN, on page 46](#)
- [Reference for Clustering, on page 47](#)
- [History for ASA Clustering on the Firepower 4100/9300, on page 61](#)

About Clustering on the Firepower 4100/9300 Chassis

The cluster consists of multiple devices acting as a single logical unit. When you deploy a cluster on the Firepower 4100/9300 chassis, it does the following:

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

For intra-chassis clustering (Firepower 9300 only), this link utilizes the Firepower 9300 backplane for cluster communications.

For inter-chassis clustering, you need to manually assign physical interface(s) to this EtherChannel for communications between chassis.

- Creates the cluster bootstrap configuration within the application.

When you deploy the cluster, the chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings. Some parts of the bootstrap configuration may be user-configurable within the application if you want to customize your clustering environment.

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

For intra-chassis clustering, spanned interfaces are not limited to EtherChannels, like it is for inter-chassis clustering. The Firepower 9300 supervisor uses EtherChannel technology internally to load-balance traffic to multiple modules on a shared interface, so any data interface type works for Spanned mode. For inter-chassis clustering, you must use Spanned EtherChannels for all data interfaces.



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

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

The following sections provide more detail about clustering concepts and implementation. See also [Reference for Clustering, on page 47](#).

Bootstrap Configuration

When you deploy the cluster, the Firepower 4100/9300 chassis supervisor pushes a minimal bootstrap configuration to each unit that includes the cluster name, cluster control link interface, and other cluster settings. Some parts of the bootstrap configuration are user-configurable if you want to customize your clustering environment.

Cluster Members

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

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

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

Some features do not scale in a cluster, and the control unit handles all traffic for those features. See [Centralized Features for Clustering, on page 48](#).

Master and Slave Unit Roles

One member of the cluster is the master unit. The master unit is determined automatically. All other members are slave units.

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

Some features do not scale in a cluster, and the master unit handles all traffic for those features. See [Centralized Features for Clustering, on page 48](#).

Cluster Control Link

The cluster-control link is an EtherChannel (port-channel 48) for unit-to-unit communication. For intra-chassis clustering, this link utilizes the Firepower 9300 backplane for cluster communications. For inter-chassis clustering, you need to manually assign physical interface(s) to this EtherChannel on the Firepower 4100/9300 chassis for communications between chassis.

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

Cluster control link traffic includes both control and data traffic.

Control traffic includes:

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

Data traffic includes:

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

Size the Cluster Control Link

If possible, you should size the cluster control link to match the expected throughput of each chassis so the cluster-control link can handle the worst-case scenarios. For example, if you have the ASA 5585-X with SSP-60, which can pass 14 Gbps per unit maximum in a cluster, then you should also assign interfaces to the cluster control link that can pass at least 14 Gbps. In this case, you could use 2 Ten Gigabit Ethernet interfaces in an EtherChannel for the cluster control link, and use the rest of the interfaces as desired for data links.

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 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 following diagram shows how to use an EtherChannel as a cluster control link in a Virtual Switching System (VSS) or Virtual Port Channel (vPC) environment. All links in the EtherChannel are active. When the switch is part of a VSS or vPC, then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS or vPC. 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.



333222

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 Network

The Firepower 4100/9300 chassis auto-generates the cluster control link interface IP address for each unit based on the chassis ID and slot ID: `127.2.chassis_id.slot_id`. You cannot set this IP address manually, either in FXOS or within the application. The cluster control link network cannot include any routers between units; only Layer 2 switching is allowed. For inter-site traffic, Cisco recommends using Overlay Transport Virtualization (OTV).

Cluster Interfaces

For intra-chassis clustering, you can assign both physical interfaces or EtherChannels (also known as port channels) to the cluster. Interfaces assigned to the cluster are Spanned interfaces that load-balance traffic across all members of the cluster.

For inter-chassis clustering, you can only assign data EtherChannels to the cluster. These Spanned EtherChannels include the same member interfaces on each chassis; on the upstream switch, all of these interfaces are included in a single EtherChannel, so the switch does not know that it is connected to multiple devices.

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

Connecting to a VSS or vPC

We recommend connecting EtherChannels to a VSS or vPC to provide redundancy for your interfaces.

Configuration Replication

All units in the cluster share a single configuration. You can only make configuration changes on the control unit, and changes are automatically synced to all other units 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

You must assign a Management type interface to the cluster. This interface is a special individual interface as opposed to a Spanned interface. The management interface lets you connect directly to each unit.

The Main cluster IP address is a fixed address for the cluster that always belongs to the current control unit. You also configure a range of addresses so that each unit, including the current control unit, can use a Local address from the range. The Main cluster IP address provides consistent 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.

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

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

Control Unit Management Vs. Data Unit Management

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

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

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

RSA Key Replication

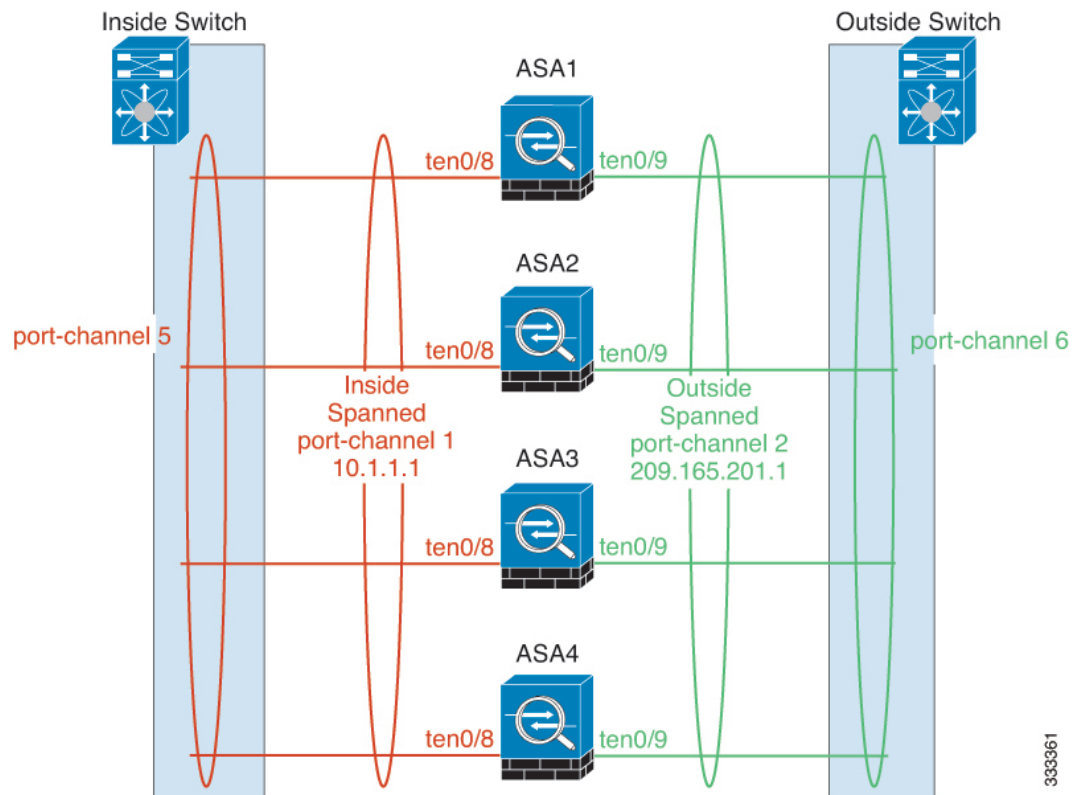
When you create an RSA key on the control unit, the key is replicated to all data units. If you have an SSH session to the Main cluster IP address, you will be disconnected if the control unit fails. The new control unit uses the same key for SSH connections, so that you do not need to update the cached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

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.

Spanned EtherChannels (Recommended)

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.



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 Clustering on the Firepower 4100/9300 Chassis, on page 8](#)
- Inter-Site Guidelines—[Clustering Guidelines and Limitations, on page 11](#)
- Configure Cluster Flow Mobility—[Configure Cluster Flow Mobility, on page 30](#)
- Enable Director Localization—[Configure Basic ASA Cluster Parameters, on page 27](#)

- Enable Site Redundancy—[Configure Basic ASA Cluster Parameters, on page 27](#)

Requirements and Prerequisites for Clustering on the Firepower 4100/9300 Chassis

Maximum Clustering Units Per Model

- Firepower 4100—16 chassis
- Firepower 9300—16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules.

Hardware and Software Requirements for Inter-Chassis Clustering

All chassis in a cluster:

- For the Firepower 4100 series: All chassis must be the same model. For the Firepower 9300: All security modules must be the same type. For example, if you use clustering, all modules in the Firepower 9300 must be SM-40s. You can have different quantities of installed security modules in each chassis, although all modules present in the chassis must belong to the cluster including any empty slots.
- Must run the identical FXOS software except at the time of an image upgrade.
- Must include the same interface configuration for interfaces you assign to the cluster, such as the same Management interface, EtherChannels, active interfaces, speed and duplex, and so on. You can use different network module types on the chassis as long as the capacity matches for the same interface IDs and interfaces can successfully bundle in the same spanned EtherChannel. Note that all data interfaces must be EtherChannels in inter-chassis clustering. If you change the interfaces in FXOS after you enable clustering (by adding or removing interface modules, or configuring EtherChannels, for example), then perform the same changes on each chassis, starting with the data units, and ending with the control unit. Note that if you remove an interface in FXOS, the ASA configuration retains the related commands so that you can make any necessary adjustments; removing an interface from the configuration can have wide effects. You can manually remove the old interface configuration.
- Must use the same NTP server. Do not set the time manually.
- ASA: Each FXOS chassis must be registered with the License Authority or satellite server. There is no extra cost for data units. For permanent license reservation, you must purchase separate licenses for each chassis. For Firepower Threat Defense, all licensing is handled by the Firepower Management Center.

Switch Requirements

- Be sure to complete the switch configuration and successfully connect all the EtherChannels from the chassis to the switch(es) before you configure clustering on the Firepower 4100/9300 chassis.
- For supported switch characteristics, see [Cisco FXOS Compatibility](#).

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)).

Licenses for Clustering on the Firepower 4100/9300 Chassis

The clustering feature itself does not require any licenses. To use Strong Encryption and other optional licenses, each Firepower 4100/9300 chassis must be registered with the License Authority or satellite server. There is no extra cost for data units. For permanent license reservation, you must purchase separate licenses for each chassis.

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

In the ASA license configuration, 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:

- **Standard**—Only the control unit requests the Standard license from the server. Because the data units have the Standard license enabled by default, they do not need to register with the server to use it.
- **Context**—Only the control unit requests the Context license from the server. The Standard license includes 10 contexts by default and is present on all cluster members. The value from each unit's Standard 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 Firepower 9300 modules in the cluster. The Standard license includes 10 contexts; for 6 units, these licenses add up to 60 contexts. You configure an additional 20-Context license on the control unit. Therefore, the aggregated cluster license includes 80 contexts. Because the platform limit for one module is 250, the combined license allows a maximum of 250 contexts; the 80 contexts are within the limit. Therefore, you can configure up to 80 contexts on the control unit; each data unit will also have 80 contexts through configuration replication.
 - You have 3 Firepower 4110 units in the cluster. The Standard license includes 10 contexts; for 3 units, these licenses add up to 30 contexts. You configure an additional 250-Context license on the control unit. Therefore, the aggregated cluster license includes 280 contexts. Because the platform limit for one unit is 250, the combined license allows a maximum of 250 contexts; the 280 contexts are over the limit. Therefore, you can only configure up to 250 contexts on the control unit; each data unit will also have 250 contexts through configuration replication. In this case, you should only configure the control unit Context license to be 220 contexts.
- **Carrier**—Required for Distributed S2S VPN. This license is a per-unit entitlement, and each unit requests its own license from the server. This license configuration is replicated to the data units.
- **Strong Encryption (3DES)** (for pre-2.3.0 Cisco Smart Software Manager satellite deployment, or for tracking purposes)—This license is a per-unit entitlement, and each unit requests its own license from the server.

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 12 hours 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.

Licenses for Distributed S2S VPN

A Carrier license is required for Distributed S2S VPN, on each member of the cluster.

Each VPN connection requires two *Other VPN* licensed sessions (the *Other VPN* license is part of the *Base* license), one for the active session and one for the backup session. The maximum VPN session capacity of the cluster can be no more than half of the licensed capacity due to using two licenses for each session.

Clustering Guidelines and Limitations

Switches for Inter-Chassis Clustering

- Make sure connected switches match the MTU for both cluster data interfaces and the cluster control link interface. You should configure the cluster control link interface MTU to be at least 100 bytes higher than the data interface MTU, so make sure to configure the cluster control link connecting switch appropriately. Because the cluster control link traffic includes data packet forwarding, the cluster control link needs to accommodate the entire size of a data packet plus cluster traffic overhead.
- For Cisco IOS XR systems, if you want to set a non-default MTU, set the IOS 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 *IPv4* MTU. This adjustment is not required for Cisco Catalyst and Cisco Nexus switches.
- On the switch(es) for the cluster control link interfaces, you can optionally enable Spanning Tree PortFast on the switch ports connected to the cluster unit to speed up the join process for new units.
- On the switch, we recommend that you use one of the following EtherChannel load-balancing algorithms: **source-dest-ip** or **source-dest-ip-port** (see the Cisco Nexus OS and Cisco IOS **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.
- Some switches do not support dynamic port priority with LACP (active and standby links). You can disable dynamic port priority to provide better compatibility with Spanned EtherChannels.
- 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:

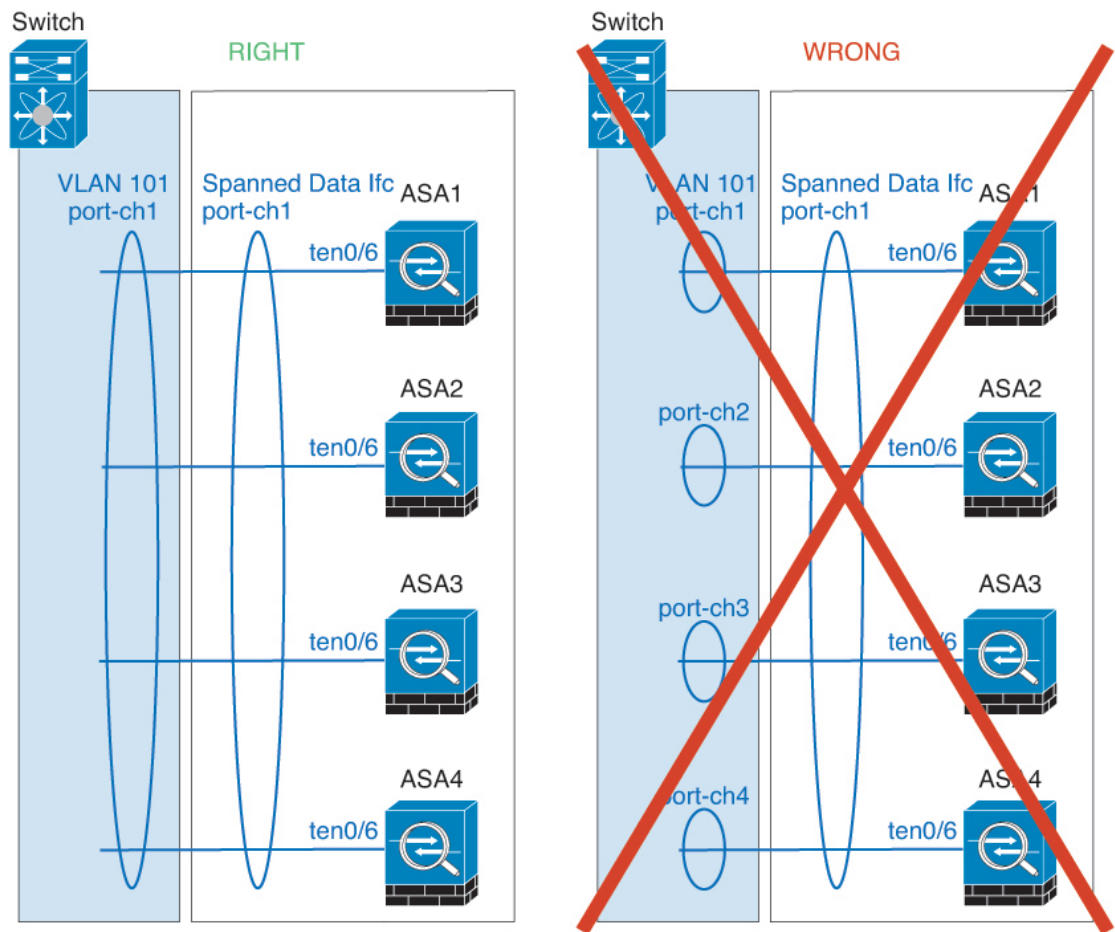
```
router(config)# port-channel id hash-distribution fixed
```

Do not change the algorithm globally; you may want to take advantage of the adaptive algorithm for the VSS peer link.
- Unlike ASA hardware clusters, Firepower 4100/9300 clusters support LACP graceful convergence. So for the Firepower platform, you can leave LACP graceful convergence enabled on connected Cisco Nexus switches.

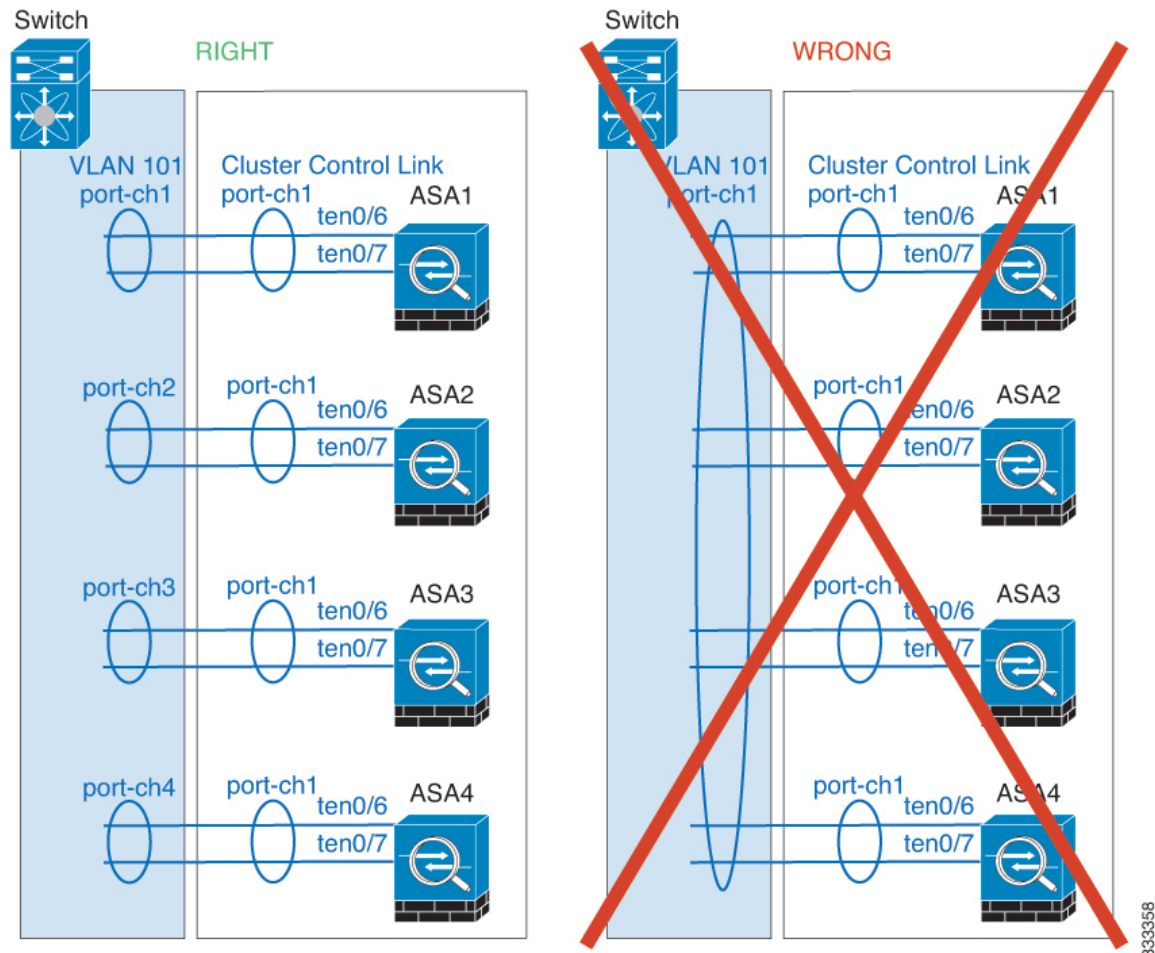
- When you see slow bundling of a Spanned EtherChannel on the switch, you can enable LACP rate fast for an individual interface on the switch. FXOS EtherChannels have the LACP rate set to fast by default. Note that some switches, such as the Nexus series, do not support LACP rate fast when performing in-service software upgrades (ISSUs), so we do not recommend using ISSUs with clustering.

EtherChannels for Inter-Chassis Clustering

- For connecting switches, set the EtherChannel mode to Active; On mode is not supported on the Firepower 4100/9300 chassis, even for the cluster control link.
- 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 Clustering

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 unit 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 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 units. 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 Firepower 4100/9300 chassis or the switch, adding an additional switch to form a VSS or vPC) you should disable the health check feature, and also disable interface monitoring for the disabled interfaces. When the topology change is complete, and the configuration change is synced to all units, you can re-enable the 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 connections; for example, dropping a FIN/ACK packet for an FTP connection will make the FTP client hang. In this case, you need to reestablish the FTP connection.
- If you use a Windows 2003 server connected to a Spanned EtherChannel interface, when the syslog server port is down, and the server does not throttle ICMP error messages, then large numbers of ICMP

messages are sent back to the cluster. These messages can result in some units of the cluster experiencing high CPU, which can affect performance. We recommend that you throttle ICMP error messages.

- We recommend connecting EtherChannels to a VSS or vPC for redundancy.
- Within a chassis, you cannot cluster some security modules and run other security modules in standalone mode; you must include all security modules in the cluster.

Defaults

- The cluster health check feature is enabled by default with the holdtime of 3 seconds. Interface health monitoring is enabled on all interfaces by default.
- Connection rebalancing is disabled by default. If you enable connection rebalancing, the default time between load information exchanges is 5 seconds.
- The cluster auto-rejoin feature for a failed cluster control link is set to unlimited attempts every 5 minutes.
- The cluster auto-rejoin feature for a failed data interface is set to 3 attempts every 5 minutes, with the increasing interval set to 2.
- Connection replication delay of 5 seconds is enabled by default for HTTP traffic.

Configure Clustering on the Firepower 4100/9300 Chassis

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit. This section describes the default bootstrap configuration and optional customization you can perform on the ASA. This section also describes how to manage cluster members from within the ASA. You can also manage cluster membership from the Firepower 4100/9300 chassis. See the Firepower 4100/9300 chassis documentation for more information.

Procedure

-
- Step 1** [FXOS: Add an ASA Cluster, on page 16](#)
 - Step 2** [ASA: Change the Firewall Mode and Context Mode, on page 24](#)
 - Step 3** [ASA: Configure Data Interfaces, on page 25](#)
 - Step 4** [ASA: Customize the Cluster Configuration, on page 27](#)
 - Step 5** [ASA: Manage Cluster Members, on page 40](#)
-

FXOS: Add an ASA Cluster

You can add a single Firepower 9300 chassis as an intra-chassis cluster, or add multiple chassis for inter-chassis clustering. For inter-chassis clustering, you must configure each chassis separately. Add the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

Create an ASA Cluster

Set the scope to the image version.

You can easily deploy the cluster from the Firepower 4100/9300 chassis supervisor. All initial configuration is automatically generated for each unit.

For inter-chassis clustering, you must configure each chassis separately. Deploy the cluster on one chassis; you can then copy the bootstrap configuration from the first chassis to the next chassis for ease of deployment.

In a Firepower 9300 chassis, you must enable clustering for all 3 module slots, even if you do not have a module installed. If you do not configure all 3 modules, the cluster will not come up.

For multiple context mode, you must first deploy the logical device, and then enable multiple context mode in the ASA application.

To change the ASA to transparent firewall mode, complete the initial deployment, and then change the firewall mode within the ASA CLI.

When you deploy a cluster, the Firepower 4100/9300 chassis supervisor configures each ASA application with the following bootstrap configuration. You can later modify parts of the bootstrap configuration from the ASA, if desired (shown in **Bold** text).

```
interface Port-channel48
  description Clustering Interface
  cluster group <service_type_name>
  key <secret>
  local-unit unit-<chassis#-module#>
  site-id <number>
  cluster-interface port-channel48 ip 127.2.<chassis#>.<module#> 255.255.255.0
  priority <auto>
  health-check holdtime 3
  health-check data-interface auto-rejoin 3 5 2
  health-check cluster-interface auto-rejoin unlimited 5 1
  enable

ip local pool cluster_ipv4_pool <ip_address>-<ip_address> mask <mask>

interface <management_ifc>
  management-only individual
  nameif management
  security-level 0
  ip address <ip_address> <mask> cluster-pool cluster_ipv4_pool
  no shutdown

http server enable
http 0.0.0.0 0.0.0.0 management
route management <management_host_ip> <mask> <gateway_ip> 1
```



Note The **local-unit** name can only be changed if you disable clustering.

Before you begin

- Download the application image you want to use for the logical device from Cisco.com, and then upload that image to the Firepower 4100/9300 chassis.
- Gather the following information:

- Management interface ID, IP address, and network mask
- Gateway IP address

Procedure

Step 1

Configure interfaces.

- a) Add at least one Data type interface or EtherChannel (also known as a port-channel) before you deploy the cluster. See [Add an EtherChannel \(Port Channel\)](#) or [Configure a Physical Interface](#).

For inter-chassis clustering, all data interfaces must be Spanned EtherChannels with at least one member interface. Add the same EtherChannels on each chassis. Combine the member interfaces from all cluster units into a single EtherChannel on the switch. See [Clustering Guidelines and Limitations, on page 11](#) for more information about EtherChannels for inter-chassis clustering.

- b) Add a Management type interface or EtherChannel. See [Add an EtherChannel \(Port Channel\)](#) or [Configure a Physical Interface](#).

The management interface is required. Note that this management interface is not the same as the chassis management interface that is used only for chassis management (in FXOS, you might see the chassis management interface displayed as MGMT, management0, or other similar names).

For inter-chassis clustering, add the same Management interface on each chassis.

- c) For inter-chassis clustering, add a member interface to the cluster control link EtherChannel (by default, port-channel 48). See [Add an EtherChannel \(Port Channel\)](#).

Do not add a member interface for intra-chassis clustering. If you add a member, the chassis assumes this cluster will be inter-chassis, and will only allow you to use Spanned EtherChannels, for example.

On the **Interfaces** tab, the port-channel 48 cluster type interface shows the **Operation State** as **failed** if it does not include any member interfaces. For intra-chassis clustering, this EtherChannel does not require any member interfaces, and you can ignore this Operational State.

Add the same member interfaces on each chassis. The cluster control link is a device-local EtherChannel on each chassis. Use separate EtherChannels on the switch per device. See [Clustering Guidelines and Limitations, on page 11](#) for more information about EtherChannels for inter-chassis clustering.

Step 2

Choose **Logical Devices**.

Step 3

Click **Add Device**, and set the following parameters:

Add Device

Device Name: asa_cluster

Template: Cisco: Adaptive Security Appliance

Image Version: 9.12.1

Instance Type: Native

Usage: Standalone Cluster

Do you want to: Create New Cluster Join Existing Cluster

OK Cancel

- a) Provide a **Device Name**.

This name is used internally by the chassis supervisor to configure management settings and to assign interfaces; it is not the device name used in the application configuration.

- b) For the **Template**, choose **Cisco Adaptive Security Appliance**.
- c) Choose the **Image Version**.
- d) For the **Instance Type**, only the **Native** type is supported.
- e) For the **Usage**, click the **Cluster** radio button.
- f) Click the **Create New Cluster** radio button.
- g) Click **OK**.

You see the Provisioning - *device name* window.

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

All valid interfaces are assigned by default.

Step 5 Click the device icon in the center of the screen.

A dialog box appears where you can configure initial bootstrap settings. These settings are meant for initial deployment only, or for disaster recovery. For normal operation, you can later change most values in the application CLI configuration.

Step 6 On the **Cluster Information** page, complete the following.

Cisco: Adaptive Security Appliance - Bootstrap Configuration ? X

Cluster Information Settings

Security Module

Security Module-1,Security Module-2,Security Module-3

Interface Information

Chassis ID:

Site ID:

Cluster Key:

Confirm Cluster Key:

Cluster Group Name:

Management Interface:

CCL Subnet IP:

DEFAULT

Address Type:

IPv4

Management IP Pool: -

Virtual IPv4 Address:

Network Mask:

Network Gateway:

OK Cancel

- a) For inter-chassis clustering, in the **Chassis ID** field, enter a chassis ID. Each chassis in the cluster must use a unique ID.

This field only appears if you added a member interface to cluster control link Port-Channel 48.

- b) For inter-site clustering, in the **Site ID** field, enter the site ID for this chassis between 1 and 8.
- c) In the **Cluster Key** field, configure an authentication key for control traffic on the cluster control link.
- The shared secret is an ASCII string from 1 to 63 characters. The shared secret is used to generate the key. This option does not affect datapath traffic, including connection state update and forwarded packets, which are always sent in the clear.
- d) Set the **Cluster Group Name**, which is the cluster group name in the logical device configuration.
- The name must be an ASCII string from 1 to 38 characters.
- e) Choose the **Management Interface**.

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

- f) Choose the **Address Type** for the management interface.

This information is used to configure a management interface in the ASA configuration. Set the following information:

- **Management IP Pool**—Configure a pool of Local IP addresses, one of which will be assigned to each cluster unit for the interface, by entering the starting and ending addresses separated by a hyphen.

Include at least as many addresses as there are units in the cluster. Note that for the Firepower 9300, you must include 3 addresses per chassis, even if you do not have all module slots filled. If you plan to expand the cluster, include additional addresses. The Virtual IP address (known as the Main cluster IP address) that belongs to the current control unit is *not* a part of this pool; be sure to reserve an IP address on the same network for the Main cluster IP address. You can use IPv4 and/or IPv6 addresses.

- **Network Mask or Prefix Length**

- **Network Gateway**

- **Virtual IP address**—Set the management IP address of the current control unit. This IP address must be on the same network as the cluster pool addresses, but not be part of the pool.

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


- a) Enter and confirm a **Password** for the admin user and for the enable password.

The pre-configured ASA admin user is useful for password recovery; if you have FXOS access, you can reset the admin user password if you forget it.

Step 8 Click **OK** to close the configuration dialog box.

Step 9 Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for the status of the new logical device. When the logical device shows its **Status** as **online**, you can add the remaining cluster chassis, or for intra-chassis clustering start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.

**Step 10**


For inter-chassis clustering, add the next chassis to the cluster:

- a) On the first chassis Firepower Chassis Manager, click the **Show Configuration icon** at the top right; copy the displayed cluster configuration.
- b) Connect to the Firepower Chassis Manager on the next chassis, and add a logical device according to this procedure.
- c) Choose **Join an Existing Cluster**.
- d) Click the **Copy config** check box, and click **OK**. If you uncheck this check box, you must manually enter the settings to match the first chassis configuration.
- e) In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.
- f) Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:
 - **Chassis ID**—Enter a unique chassis ID.
 - **Site ID**—Enter the correct site ID.
 - **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

- g) Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



Step 11 Connect to the control unit ASA to customize your clustering configuration.

Add More Cluster Members

Add or replace an ASA cluster member.



Note This procedure only applies to adding or replacing a *chassis*; if you are adding or replacing a module to a Firepower 9300 where clustering is already enabled, the module will be added automatically.

Before you begin

- Make sure your existing cluster has enough IP addresses in the management IP address pool for this new member. If not, you need to edit the existing cluster bootstrap configuration on each chassis before you add this new member. This change causes a restart of the logical device.
- The interface configuration must be the same on the new chassis. You can export and import FXOS chassis configuration to make this process easier.
- For multiple context mode, enable multiple context mode in the ASA application on the first cluster member; additional cluster members will inherit the multiple context mode configuration automatically.

Procedure


- Step 1** On an existing cluster chassis Firepower Chassis Manager, choose **Logical Devices** to open the **Logical Devices** page.
- Step 2** Click the Show Configuration icon () at the top right; copy the displayed cluster configuration.
- Step 3** Connect to the Firepower Chassis Manager on the new chassis, and click **Add Device**.
- Step 4** For the **Device Name**, provide a name for the logical device.
- Step 5** For the **Template**, choose **Cisco Adaptive Security Appliance**.

- Step 6** For the **Image Version**, choose the ASA software version.
- Step 7** For the **Device Mode**, click the **Cluster** radio button.
- Step 8** Choose **Join an Existing Cluster**.
- Step 9** Click the **Copy config** check box, and click **OK**. If you uncheck this check box, you must manually enter the settings to match the first chassis configuration.
- Step 10** In the **Copy Cluster Details** box, paste in the cluster configuration from the first chassis, and click **OK**.
- Step 11** Click the device icon in the center of the screen. The cluster information is mostly pre-filled, but you must change the following settings:
- **Chassis ID**—Enter a unique chassis ID.
 - **Site ID**—Enter the correct site ID.
 - **Cluster Key**—(Not prefilled) Enter the same cluster key.

Click **OK**.

- Step 12** Click **Save**.

The chassis deploys the logical device by downloading the specified software version and pushing the bootstrap configuration and management interface settings to the application instance. Check the **Logical Devices** page for each cluster member for the status of the new logical device. When the logical device for each cluster member shows its **Status** as **online**, you can start configuring the cluster in the application. You may see the "Security module not responding" status as part of the process; this status is normal and is temporary.



ASA: Change the Firewall Mode and Context Mode

By default, the FXOS chassis deploys a cluster in routed or transparent firewall mode, and single context mode.

- Change the firewall mode— To change the mode after you deploy, change the mode on the control unit; the mode is automatically changed on all data units to match. See [Set the Firewall Mode \(Single Mode\)](#). In multiple context mode, you set the firewall mode per context. See [Configure a Security Context](#).

- Change to multiple context mode—To change to multiple context mode after you deploy, change the mode on the control unit; the mode is automatically changed on all data units to match. See [Enable Multiple Context Mode](#).

ASA: Configure Data Interfaces

This procedure configures basic parameters for each data interface that you assigned to the cluster when you deployed it in FXOS. For inter-chassis clustering, data interfaces are always Spanned EtherChannel interfaces.



Note The management interface was pre-configured when you deployed the cluster. You can also change the management interface parameters in ASA, but this procedure focuses on data interfaces. The management interface is an individual interface, as opposed to a Spanned interface. See [Management Interface, on page 5](#) for more information.

Before you begin

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

Procedure

- Step 1** 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** Select the interface, and click **Edit**.
The **Edit Interface** dialog box appears.
- Step 3** Set the following:
- (For EtherChannels) **MIO Port-channel ID**—Enter the same ID used in FXOS.
 - **Enable Interface** (checked by default)
- 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.

- In the **ASA Cluster** area, for inter-site clustering set **Site specific MAC Addresses**, and for routed mode, the 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.

- Click **OK** to accept your changes.
- Allocate interfaces.
- 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.
- 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.

- Check the **Enable IPv6** check box.
- 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.

- 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**.

ASA: Customize the Cluster Configuration

If you want to change bootstrap settings after you deploy the cluster or configure additional options, such as clustering health monitoring, TCP connection replication delay, flow mobility, and other optimizations, you can do so on the control unit.

Configure Basic ASA Cluster Parameters

You can customize cluster settings on the control unit.

Before you begin

- For multiple context mode, complete this procedure in the system execution space on the control unit. 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 local-unit **Member Name** and several other options can only be set on the FXOS chassis, or they can only be changed on the ASA if you disable clustering, so they are not included in the following procedure.

Procedure

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

Step 2 (Optional) Configure the following optional parameters:

- **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 unit to the data units.
- **Enable health monitoring of this device within the cluster**—Enables the cluster unit health check feature, and determines the amount of time between unit heartbeat status messages, between .3 and 45 seconds; The default is 3 seconds. **Note:** When you are adding new units 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 unit health, the ASA cluster units send heartbeat messages on the cluster control link to other units. If a unit does not receive any heartbeat messages from a peer unit within the holdtime period, the peer unit is considered unresponsive or dead.

- **Debounce Time**—Configures the debounce time before the ASA considers an interface to be failed and the unit 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 unit is removed from the cluster. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.
- **Replicate console output**—Enables console replication from data units to the control unit. 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 units send the console messages to the control unit 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 unit to the data units.
- **Enable Clustering Flow Mobility**. See [Configure LISP Inspection, on page 32](#).
- **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.
- **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.

Step 3 In the **Cluster Control Link** area, you can configure the cluster control link MTU. Other options in this area cannot be configured on the ASA.

- **MTU**—Specify the maximum transmission unit for the cluster control link interface to be at least 100 bytes higher than the highest MTU of the data interfaces. We suggest setting the MTU to the maximum of 9184; the minimum value is 1400 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.

For example, because the maximum MTU is 9184, then the highest data interface MTU can be 9084, while the cluster control link can be set to 9184.

Step 4 (Optional) In the **Cluster LACP** area, you can enable static port priority. 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. Other options in this area cannot be configured on the ASA without disabling clustering.

- **Enable static port priority**—Disables dynamic port priority in LACP. Some switches do not support dynamic port priority, so this parameter improves switch compatibility. Moreover, it enables support of more than 8 active spanned EtherChannel members, up to 32 members. Without this parameter, only 8 active members and 8 standby members are supported. If you enable this parameter, then you cannot use any standby members; all members are active. This parameter is not part of the bootstrap configuration, and is replicated from the control unit to the data units.

Step 5 Click **Apply**.

Configure Interface Health Monitoring and Auto-Rejoin Settings

You might want to disable health monitoring of non-essential interfaces, for example, the management interface. You can monitor any port-channel 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 Select an interface in the **Monitored Interfaces** box, 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 unit, but there are active ports under the same logical interface on other units, then the unit is removed from the cluster. If a unit 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 unit 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 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, Firepower 4100/9300 chassis, or the switch, or adding an additional switch to form a VSS or vPC) 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 units, 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 unit 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. This setting is enabled by default for 5 seconds.
- **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 units 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 unit 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

- Set the site ID for the chassis on Firepower 4100/9300 chassis supervisor.
- 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**.

Configure Distributed Site-to-Site VPN

By default, the ASA cluster uses Centralized Site-to-Site VPN mode. To take advantage of the scalability of clustering, you can enable Distributed Site-to-Site VPN mode. In this mode, S2S IPsec IKEv2 VPN connections are distributed across members of an ASA cluster. Distributing VPN connections across the members of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond Centralized VPN capabilities.

About Distributed Site-to-Site VPN

Distributed VPN Connection Roles

When running in Distributed VPN mode the following roles are assigned to the cluster members:

- **Active Session Owner**—The unit that initially receives the connection, or that has transitioned a backup session to an active session. The owner maintains state and processes packets for the complete session, including the IKE and IPsec tunnels and all traffic associated with them.

- **Backup Session Owner**—The unit that is handling the backup session for an existing active session. Depending on the backup strategy chosen, this may be a unit in the same chassis as the active session owner, or a unit in another chassis. If the active session owner fails, the backup session owner becomes the active session owner, and a new backup session is established on a different unit.
- **Forwarder**—If traffic associated with a VPN session is sent to a unit that does not own the VPN session, that unit will use the Cluster Control Link (CCL) to forward the traffic to the member which owns the VPN session
- **Orchestrator**—The orchestrator (always the control unit of the cluster) is responsible for calculating which sessions will move and where to when executing an Active Session Redistribution (ASR). It sends a request to the owner member X to move N sessions to member Y. Member X will respond back to the orchestrator when complete, specifying how many sessions it was able to move.

Distributed VPN Session Characteristics

Distributed S2S VPN Sessions have the following characteristics. Otherwise, VPN connections behave as they normally do if not on an ASA cluster.

- VPN sessions are distributed across the cluster at the session level. Meaning the same cluster member handles the IKE and IPsec tunnels, and all their traffic, for a VPN connection. If VPN session traffic is sent to a cluster member that does not own that VPN session, traffic is forwarded to the cluster member that owns the VPN session.
- VPN sessions have a Session ID that is unique across the cluster. Using the session ID, traffic is validated, forwarding decisions are made, and IKE negotiation is completed.
- In an S2S VPN hub and spoke configuration, when clients connect through the ASA cluster (called hair-pinning), the session traffic flowing in and the session traffic flowing out may be on different cluster members.
- You can require that the backup session to be allocated on a security module in another chassis; this provides protection against chassis failure. Or, you can choose to allocate backup sessions on any node in the cluster; this provides protection against node failure only. When there are two chassis in the cluster, remote-chassis backup is strongly recommended.
- Only IKEv2 IPsec S2S VPN is supported in Distributed S2S VPN mode, IKEv1 is not. IKEv1 S2S is supported in centralized VPN mode.
- Each security module supports up to 6K VPN sessions for a maximum of approximately 36K sessions across 6 members. The actual number of sessions supported on a cluster member is determined by platform capacity, allocated licenses, and per context resource allocation. When utilization is close to the limit, there may be cases where session creation fails, even though the maximum capacity has not been reached on each cluster unit. This is because active session allocation is determined by external switching, and backup session allocation is determined by an internal cluster algorithm. Customers are encouraged to size their utilization accordingly and allow room for uneven distribution.

Distributed VPN Handling of Cluster Events

Table 1:

Event	Distributed VPN
Member failure	For all active sessions on this failed member, the backup sessions (on another member) become active and backup sessions are reallocated on another unit according to the backup strategy.
Chassis failure	<p>When a remote-chassis backup strategy is being used, for all active sessions on the failed chassis, the backup sessions (on a member in the other chassis) become active. When the units are replaced, backup sessions for these now active sessions will be reallocated on members in the replaced chassis.</p> <p>When a flat backup strategy is being used, if both the active and backup sessions are on the failed chassis, the connection will drop. All active sessions with backup sessions on a member in the other chassis, fallback to these sessions. New backup sessions will be allocated on another member in the surviving chassis.</p>
Inactivate a cluster member	For all active sessions on the cluster member being inactivated, backup sessions (on another member) become active and reallocate backup sessions on another unit according to the backup strategy.
Cluster member join	<p>If the VPN cluster mode is not set to distributed, the control unit will request a mode change.</p> <p>If, or once the VPN mode is compatible, the cluster member will be assigned active and backup sessions in the flow of normal operations.</p>

Unsupported Inspections

The following types of inspections are not supported or are disabled in Distributed S2S VPN mode:

- CTIQBE
- DCERPC
- H323, H225, and RAS
- IPsec pass-through
- MGCP
- MMP
- NetBIOS
- PPTP
- RADIUS
- RSH
- RTSP
- SCCP (Skinny)

- SUNRPC
- TFTP
- WAAS
- WCCP
- XDMCP

IPsec IKEv2 Modifications

IKEv2 is modified while in Distributed S2S VPN mode in the following ways:

- An identity is used in place of IP/port tuples. This allows for proper forwarding decisions on the packets, and cleanup of previous connections that may be on other cluster members.
- The (SPI) identifiers that identify a single IKEv2 session are locally generated, random 8-byte values that are unique across the cluster. An SPI embeds a time stamp and a cluster member ID. Upon receipt of an IKE negotiation packet, if the time stamp or cluster member ID check fails, the packet is dropped and a message is logged indicating the reason.
- IKEv2 processing has been modified to prevent NAT-T negotiations from failing by being split across cluster members. A new ASP classify domain, *cluster_isakmp_redirect*, and rules are added when IKEv2 is enabled on an interface.

Model Support

The only device supported for Distributed VPN is the Firepower 9300. Distributed VPN supports a maximum of 6 modules on up to 2 chassis. You can have different quantities of installed security modules in each chassis, although we recommend an equal distribution.

Inter-site clustering is not supported.

Firewall Mode

Distributed S2S VPN is supported in routed mode only.

Context Mode

Distributed S2S VPN operates in both single and multiple context modes. However, in multiple context mode, active session redistribution is done at the system level, not at the context level. This prevents an active session associated with a context from moving to a cluster member that contains active sessions associated with a different context, unknowingly creating an unsupportable load.

High Availability

The following capabilities provide resiliency against single failure of a security module or chassis:

- VPN Sessions that are backed up on another security module in the cluster, on any chassis, withstand security module failures.
- VPN Sessions that are backed up on another chassis withstand chassis failures.
- The control unit can change without losing VPN S2S sessions.

If an additional failure occurs before the cluster has stabilized, connections may be lost if the both active and backup sessions are on the failed units.

All attempts are made to ensure no sessions are lost when a member leaves the cluster in a graceful manner such as disabling the VPN cluster mode, reloading a cluster member, and other anticipated chassis changes. During these types of operations, sessions will not be lost as long as the cluster is given time to re-establish session backups between operations. If a graceful exit is triggered on the last cluster member, it will gracefully tear down existing sessions.

Dynamic PAT

Is not available while in Distributed VPN mode.

CMPv2

The CMPv2 ID certificate and key pairs are synchronized across the cluster members. However, only the control unit in the cluster automatically renews and rekeys the CMPv2 certificate. The control unit synchronizes these new ID certificates and keys to all cluster members on a renewal. In this way, all members in the cluster utilize the CMPv2 certificates for authentication, and also any member is capable of taking over as the control unit.

Enable Distributed S2S VPN

Enable Distributed Site-to-Site VPN to take advantage of the scalability of clustering for VPN sessions.



Note Changing the VPN mode between centralized and distributed causes all existing sessions to be torn down. Changing the backup mode is dynamic and will not terminate sessions.

Before you begin

- You must have a Carrier License configured on all members of the cluster.
- Your S2S VPN configuration must be set.

Procedure

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

Step 2 In the **VPN Cluster Mode** area, choose the **VPN Mode** for the cluster, **Centralized** or **Distributed**.

Step 3 Choose the **Backup Distribution Mode**, **Flat** or **Remote-chassis**.

In flat backup mode, standby sessions are established on any other cluster member. This will protect users from blade failures, however, chassis failure protection is not guaranteed.

In remote-chassis backup mode standby sessions are established on a member of another chassis in the cluster. This will protect users from both blade failures and chassis failures.

If remote-chassis is configured in a single chassis environment (intentionally configured or the result of a failure), no backups will be created until another chassis joins.

Redistribute Distributed S2S VPN Sessions

Active Session Redistribution (ASR) redistributes the active VPN session load across the cluster members. Due to the dynamic nature of beginning and ending sessions, ASR is a best effort balancing of the sessions across all cluster members. Repeated redistribution actions will optimize the balance.

Redistribution can be run at any time, should be run after any topology change in the cluster, and is recommended after a new member joins the cluster. The goal of redistribution is to create a stable VPN cluster. A stable VPN cluster has an almost equal number of active and backup sessions across the nodes.

To move a session, the backup session becomes the active one and another node is selected to host a new backup session. Moving sessions is dependent on the location of the active session's backup and the number of active sessions already on that particular backup node. If the backup session node is unable to host the active session for some reason, the original node remains owner of the session.

In multiple-context mode, active session redistribution is done at the system level, not the individual context level. It is not done at the context level because an active session in one context could be moved a member that contains many more active sessions in a different context, creating more load on that cluster member.

Before you begin

- Enable system logs if you would like to monitor redistribution activity.
- This procedure must be carried out on the control unit of the cluster.

Procedure

Step 1 Choose **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to view how active and backup sessions are distributed across the cluster.

Depending on the number of sessions to redistribute and the load on the cluster, this may take some time. Syslogs containing the following phrases (and other system details not shown here) are provided as redistribution activity occurs:

Syslog Phrase	Notes
VPN session redistribution started	Control unit only
Sent request to move <i>number</i> sessions from <i>orig-member-name</i> to <i>dest-member-name</i>	Control unit only
Failed to send session redistribution message to <i>member-name</i>	Control unit only
Received request to move <i>number</i> sessions from <i>orig-member-name</i> to <i>dest-member-name</i>	Data unit only
Moved <i>number</i> sessions to <i>member-name</i>	The number of active sessions moved to the named cluster.
Failed to receive session move response from <i>dest-member-name</i>	Control unit only
VPN session completed	Control unit only
Cluster topology change detected. VPN session redistribution aborted.	

Step 2 Click **Re-Distribute**.

Step 3 Refresh the **Monitoring > ASA Cluster > ASA Cluster > Cluster Summary > VPN Cluster Summary** to see the results of the redistribution activity.

If your redistribution was successful, and there has been no substantial system or session activity, your system will be balanced and this action is complete.

Otherwise, repeat the redistribution process to obtain a balanced and stable system.

FXOS: Remove a Cluster Unit

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

Temporary Removal

A cluster unit will be automatically removed from the cluster due to a hardware or network failure, for example. This removal is temporary until the conditions are rectified, and it can rejoin the cluster. You can also manually disable clustering.

To check whether a device is currently in the cluster, check the cluster status on the Firepower Chassis Manager **Logical Devices** page:

The screenshot shows the 'Logical Devices' page in the Firepower Chassis Manager. It displays a table with columns for 'Management Port' and 'Status'. The 'Management Port' is 'Ethernet1/4' and the 'Status' is 'online'. To the right of the status is a slider control that is currently 'enabled' (checked). Below the table, there is an 'Attributes' section listing various system parameters:

Management Port	Status
Ethernet1/4	online

Attributes

- Cluster Operational Status : not-in-cluster
- FIREPOWER-MGMT-IP : 10.89.5.20
- CLUSTER-ROLE : none
- CLUSTER-IP : 127.2.1.1
- MGMT-URL : https://10.89.5.35/
- UUID : 8e459170-451d-11e9-8475-f22f06c32630

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

When a device becomes inactive, all data interfaces are shut down; only the Management interface can send and receive traffic. To resume traffic flow, re-enable clustering. The Management interface remains up using the IP address the unit received from the bootstrap configuration. However if you reload, and the unit is still inactive in the cluster (for example, you saved the configuration with clustering disabled), the Management interface is disabled.

To reenabling clustering, on the ASA enter **cluster group *name*** and then **enable**.

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

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

Permanent Removal

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

- Delete the logical device—In Firepower Chassis Manager on the **Logical Devices** page, click the **Delete** (🗑️). You can then deploy a standalone logical device, a new cluster, or even add a new logical device to the same cluster.
- Remove the chassis or security module from service—If you remove a device from service, you can add replacement hardware as a new member of the cluster.

ASA: Manage Cluster Members

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

Become an Inactive Member

To become an inactive member of the cluster, disable clustering on the unit 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 unit altogether from the cluster. 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 (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 Unit from the Control Unit

To deactivate a data unit, 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 unit received from the cluster IP pool. However if you reload, and the unit 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**.

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

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

Step 3 Click **Apply**.

Rejoin the Cluster

If a unit 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 reenabling clustering. Other interfaces are shut down. The exception is if you manually disabled clustering in ASDM, then you can reenabling 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 reenabling 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 reenabling clustering in ASDM by connecting ASDM to the unit you want to reenabling.
- You cannot reenabling clustering for a data unit from the control unit unless you add it as a new member.
- Choose **Configuration > Device Management > High Availability and Scalability > ASA Cluster**.
 - Check the **Participate in ASA cluster** check box.
 - 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
```

Change the Control Unit



Caution

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

To change the control unit, 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 unit to become master, and click the button to make it the control unit.
 - Step 3** You are prompted to confirm the control unit 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 members in the cluster, or to a specific member, perform the following steps. Sending a **show** command to all members collects all output and displays it on the console of the current unit. (Note that alternatively there are show commands that you can enter on the control unit to view cluster-wide statistics.) 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 members, or if you specify the unit name, a specific member:

```
cluster exec [unit unit_name] command
```

Example:

```
cluster exec show xlate
```

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

Examples

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

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

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

The following sample output for the **cluster exec show memory** command shows memory information for each member in the cluster:

```

cluster exec show memory
unit-1-1 (LOCAL):*****
Free memory:      108724634538 bytes (92%)
Used memory:      9410087158 bytes ( 8%)
-----
Total memory:     118111600640 bytes (100%)

unit-1-3:*****
Free memory:      108749922170 bytes (92%)
Used memory:      9371097334 bytes ( 8%)
-----
Total memory:     118111600640 bytes (100%)

unit-1-2:*****
Free memory:      108426753537 bytes (92%)
Used memory:      9697869087 bytes ( 8%)
-----
Total memory:     118111600640 bytes (100%)

```

ASA: Monitoring the ASA Cluster on the Firepower 4100/9300 chassis

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 unit to which you are connected, as well as other units in the cluster. You can also change the primary unit from this pane.

- **Cluster Dashboard**

On the home page on the primary unit, 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 unit, which is then automatically enabled on all of the data units 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 members.

- **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 members.

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 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.

Monitoring Distributed S2S VPN

See the following screens for monitoring VPN cluster status:

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

Displays the distribution of the session across the cluster and provides you with the ability to re-distribute the sessions.

- **Monitoring > VPN > VPN Statistics > Sessions**

Both control and data cluster units are listed. Click any member for details.

Configuring Logging for Clustering

See the followingscreen for configuring logging for clustering:

Configuration > Device Management > Logging > Syslog Setup

Each unit 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 units in the cluster.

Troubleshooting Distributed S2S VPN

Distributed VPN Notifications

You will be notified with messages containing the identified phrases when the following error situations occur on a cluster running distributed VPN:

Situation	Notification
If an existing or joining cluster data unit is not in distributed VPN mode when attempting to join the cluster:	New cluster member (<i>member-name</i>) rejected due to vpn mode mismatch. and Master (<i>control-name</i>) rejects enrollment request from unit (<i>unit-name</i>) for the reason: the vpn mode capabilities are not compatible with the master configuration
If licensing is not properly configured on a cluster member for Distributed VPN:	ERROR: Master requested cluster vpn-mode change to distributed. Unable to change mode due to missing Carrier License.
If the time stamp or member ID is invalid in the SPI of a received IKEv2 packet:	Expired SPI received or Corrupted SPI detected
If the cluster is unable to create a backup session:	Failed to create the backup for an IKEv2 session.
IKEv2 Initial Contact (IC) processing error:	IKEv2 Negotiation aborted due to ERROR: Stale backup session found on backup
Redistribution problems:	Failed to send session redistribution message to <i>member-name</i> Failed to receive session move response from <i>member-name</i> (master only)
If the topology changes during redistribution of the sessions:	Cluster topology change detected. VPN session redistribution aborted.

You may be encountering one of the following situations:

- L2L VPN sessions are being distributed to only one of the chassis in a cluster when the N7K Switch is configured with L4port as a load balancing algorithm using the **port-channel load-balance src-dst l4port** command. . An example of the cluster session allocation looks like below:

```
SSP-Cluster/slave(cfg-cluster)# show cluster vpn-sessiondb distribution
Member 0 (unit-1-3): active: 0
Member 1 (unit-2-2): active: 13295; backups at: 0 (2536), 2 (2769), 3 (2495), 4 (2835),
5 (2660)
Member 2 (unit-2-3): active: 12174; backups at: 0 (2074), 1 (2687), 3 (2207), 4 (3084),
5 (2122)
Member 3 (unit-2-1): active: 13416; backups at: 0 (2419), 1 (3013), 2 (2712), 4 (2771),
5 (2501)
Member 4 (unit-1-1): active: 0
Member 5 (unit-1-2): active: 0
```

Since L2L IKEv2 VPN uses port 500 for both source and destination ports, IKE packets are only sent to one of the links in the port channel connected between the N7K and the chassis.

Change the N7K load balancing algorithm to IP and L4 port using the **port-channel load-balance src-dst ip-l4port**. Then the IKE packets are sent to all the links and thus both Firepower9300 chassis.

For a more immediate adjustment, on the control unit of the ASA cluster execute: **cluster redistribute vpn-sessiondb** to redistribute active VPN sessions to the cluster members of the other chassis.

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 unit. Other features might have caveats for proper usage.

Unsupported Features with Clustering

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

- Unified Communication features that rely on TLS Proxy
- Remote access VPN (SSL VPN and IPsec VPN)
- IS-IS routing
- 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
- Integrated Routing and Bridging
- Dead Connection Detection (DCD)
- FIPS mode

Centralized Features for Clustering

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



Note

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

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

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

- The following application inspections:
 - DCERPC
 - NetBIOS
 - PPTP
 - RADIUS
 - RSH
 - SUNRPC
 - TFTP
 - XDMCP

- Dynamic routing

- Static route tracking
- 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)
- Authentication and Authorization for network access. Accounting is decentralized.
- Filtering Services
- Site-to-site IKEv1/IKEv2 VPN

In centralized mode, VPN connections are established with the control unit of the cluster only. This is the default mode for VPN clustering. Site-to-site VPN can also be deployed in Distributed VPN Mode, where S2S IKEv2 VPN connections are distributed across members.

Features Applied to Individual Units

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

- QoS—The QoS policy is synced across the cluster as part of configuration replication. However, the policy is enforced on each unit 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 units and with traffic evenly distributed, the conform rate actually becomes 3 times the rate for the cluster.
- Threat detection—Threat detection works on each unit independently; for example, the top statistics is unit-specific. Port scanning detection, for example, does not work because scanning traffic will be load-balanced between all units, and one unit will not see all traffic.
- Resource management—Resource management in multiple context mode is enforced separately on each unit based on local usage.
- LISP traffic—LISP traffic on UDP port 4342 is inspected by each receiving unit, but is not assigned a director. Each unit 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 unit with replication of the data structures to the cluster data units. If a control unit is elected, the new control unit 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 unit change occurs.

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

Connection Settings

Connection limits are enforced cluster-wide (see **Configuration > Firewall > Service Policy** page). Each unit 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 unit 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 unit.

Identity Firewall and Clustering

Only the control unit retrieves the user-group from the AD and the user-ip mapping from the AD agent. The control unit then populates the user information to data units, and data units can make a match decision for user identity based on the security policy.

Multicast Routing and Clustering

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

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

- PAT with Port Block Allocation—See the following guidelines for this feature:
 - Maximum-per-host limit is not a cluster-wide limit, and is enforced on each unit individually. Thus, in a 3-node cluster with the maximum-per-host limit configured as 1, if the traffic from a host is load-balanced across all 3 units, then it can get allocated 3 blocks with 1 in each unit.
 - Port blocks created on the backup unit from the backup pools are not accounted for when enforcing the maximum-per-host limit.
 - When a PAT IP address owner goes down, the backup unit will own the PAT IP address, corresponding port blocks, and xlates. If it runs out of ports on its normal PAT address, it can use the address that it took over to service new requests. As the connections eventually time out, the blocks get freed.
 - On-the-fly PAT rule modifications, where the PAT pool is modified with a completely new range of IP addresses, will result in xlate backup creation failures for the xlate backup requests that were still in transit while the new pool became effective. This behavior is not specific to the port block

allocation feature, and is a transient PAT pool issue seen only in cluster deployments where the pool is distributed and traffic is load-balanced across the cluster units.

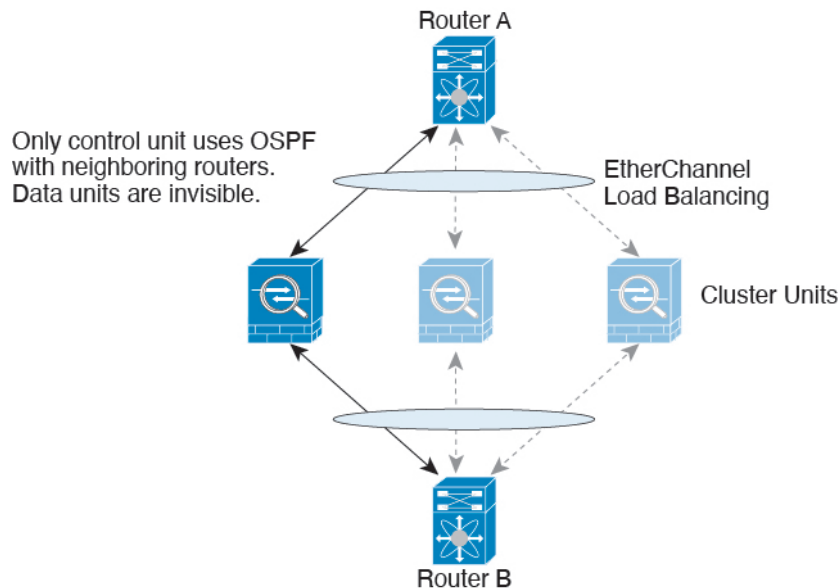
- NAT pool address distribution for dynamic PAT—The control unit evenly pre-distributes addresses across the cluster. If a member receives a connection and they have no addresses assigned, then the connection is forwarded to the control unit for PAT. If a cluster member leaves the cluster (due to failure), a backup member will get the PAT IP address, and if the backup exhausts its normal PAT IP address, it can make use of the new address. Make sure to include at least as many NAT addresses as there are units in the cluster, plus at least one extra address, to ensure that each unit receives an address, and that a failed unit can get a new address if its old address is in use by the member that took over the address. Use the **show nat pool cluster** command to see the address allocations.
- Reusing a PAT pool in multiple rules—To use the same PAT pool in multiple rules, you must be careful about the interface selection in the rules. You must either use specific interfaces in all rules, or "any" in all rules. You cannot mix specific interfaces and "any" across the rules, or the system might not be able to match return traffic to the right node in the cluster. Using unique PAT pools per rule is the most reliable option.
- No round-robin—Round-robin for a PAT pool is not supported with clustering.
- Dynamic NAT xlates managed by the control unit—The control unit maintains and replicates the xlate table to data units. When a data unit receives a connection that requires dynamic NAT, and the xlate is not in the table, it requests the xlate from the control unit. The data unit owns the connection.
- Stale xlates—The xlate idle time on the connection owner does not get updated. Thus, the idle time might exceed the idle timeout. An idle timer value higher than the configured timeout with a refcnt of 0 is an indication of a stale xlate.
- Per-session PAT feature—Although not exclusive to clustering, the per-session PAT feature improves the scalability of PAT and, for clustering, allows each data unit to own PAT connections; by contrast, multi-session PAT connections have to be forwarded to and owned by the control unit. By default, all TCP traffic and UDP DNS traffic use a per-session PAT xlate, whereas ICMP and all other UDP traffic uses multi-session. You can configure per-session NAT rules to change these defaults for TCP and UDP, but you cannot configure per-session PAT for ICMP. For traffic that benefits from multi-session PAT, such as H.323, SIP, or Skinny, you can disable per-session PAT for the associated TCP ports (the UDP ports for those H.323 and SIP are already multi-session by default). For more information about per-session PAT, see the firewall configuration guide.
- No static PAT for the following inspections—
 - FTP
 - PPTP
 - RSH
 - SQLNET
 - TFTP
 - XDMCP
 - SIP

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

Dynamic Routing and Clustering

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

Figure 1: Dynamic Routing



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

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

SCTP and Clustering

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

SIP Inspection and Clustering

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

TLS Proxy configuration is not supported.

SNMP and Clustering

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

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

When using SNMPv3 with clustering, if you add a new cluster unit after the initial cluster formation, then SNMPv3 users are not replicated to the new unit. You must re-add them on the control unit to force the users to replicate to the new unit, or directly on the data unit.

STUN and Clustering

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

Syslog and NetFlow and Clustering

- **Syslog**—Each unit in the cluster generates its own syslog messages. You can configure logging so that each unit 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 units in the cluster. If you configure logging to use the hostname as the device ID, syslog messages generated by all units look as if they come from a single unit. If you configure logging to use the local-unit name that is assigned in the cluster bootstrap configuration as the device ID, syslog messages look as if they come from different units.
- **NetFlow**—Each unit 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 unit learns security group tag (SGT) information. The control unit then populates the SGT to data units, and data units can make a match decision for SGT based on the security policy.

VPN and Clustering on the FXOS Chassis

An ASA FXOS Cluster supports one of two mutually exclusive modes for S2S VPN, centralized or distributed:

- **Centralized VPN Mode.** The default mode. In centralized mode, VPN connections are established with the control unit of the cluster only.

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

When you connect a VPN tunnel to a Spanned interface address, connections are automatically forwarded to the control unit. VPN-related keys and certificates are replicated to all units.
- **Distributed VPN Mode.** In this mode, S2S IPsec IKEv2 VPN connections are distributed across members of an ASA cluster providing scalability. Distributing VPN connections across the members of a cluster allows both the capacity and throughput of the cluster to be fully utilized, significantly scaling VPN support beyond Centralized VPN capabilities.



-
- Note** Centralized VPN clustering mode supports S2S IKEv1 and S2S IKEv2.
 Distributed VPN clustering mode supports S2S IKEv2 only.
 Distributed VPN clustering mode is supported on the Firepower 9300 only.
 Remote access VPN is not supported in centralized or distributed VPN clustering mode.
-

Performance Scaling Factor

When you combine multiple units into a cluster, you can expect the total cluster performance to be approximately:

- 80% of the combined TCP or CPS throughput
- 90% of the combined UDP throughput
- 60% of the combined Ethernet MIX (EMIX) throughput, depending on the traffic mix.

For example, for TCP throughput, the Firepower 9300 with 3 SM-44 modules can handle approximately 135 Gbps of real world firewall traffic when running alone. For 2 chassis, the maximum combined throughput will be approximately 80% of 270 Gbps (2 chassis x 135 Gbps): 216 Gbps.

Control Unit Election

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

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



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

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



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

High Availability Within the Cluster

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

Chassis-Application Monitoring

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

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

Unit Health Monitoring

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

If units cannot reach each other over the cluster control link because of a network failure and not because a unit has actually failed, then the cluster may go into a "split brain" scenario where isolated data units will elect their own control units. For example, if a router fails between two cluster locations, then the original control unit at location 1 will remove the location 2 data units from the cluster. Meanwhile, the units at location 2 will elect their own control unit and form their own cluster. Note that asymmetric traffic may fail in this scenario. After the cluster control link is restored, then the control unit that has the higher priority will keep the control unit's role. See [Control Unit Election, on page 54](#) for more information.

Interface Monitoring

Each unit monitors the link status of all hardware interfaces in use, and reports status changes to the control unit. For inter-chassis clustering, Spanned EtherChannels use the cluster Link Aggregation Control Protocol (cLACP). Each chassis monitors the link status and the cLACP protocol messages to determine if the port is still active in the EtherChannel, and informs the ASA application if the interface is down. When you enable health monitoring, all physical interfaces are monitored by default (including the main EtherChannel for EtherChannel interfaces). Only named interfaces that are in an Up state can be monitored. For example, all member ports of an EtherChannel must fail before a *named* EtherChannel is removed from the cluster (depending on your minimum port bundling setting). You can optionally disable monitoring per interface.

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

For inter-chassis clustering, if you add or delete an EtherChannel from the cluster, interface health-monitoring is suspended for 95 seconds to ensure that you have time to make the changes on each chassis.

Decorator Application Monitoring

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

Status After Failure

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

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

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



Note

When the ASA becomes inactive and fails to automatically rejoin the cluster, all data interfaces are shut down; only the management-only interface can send and receive traffic. The management interface remains up using the IP address the unit received from the cluster IP pool. However if you reload, and the unit 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 member is removed from the cluster, how it can rejoin the cluster depends on why it was removed:

- Failed cluster control link when initially joining—After you resolve the problem with the cluster control link, you must manually rejoin the cluster by re-enabling clustering.
- Failed cluster control link after joining the cluster—The 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 unit—If the unit was removed from the cluster because of a unit health check failure, then rejoining the cluster depends on the source of the failure. For example, a temporary power failure means the unit will rejoin the cluster when it starts up again as long as the cluster control link is up. The unit attempts to rejoin the cluster every 5 seconds.
- Failed Chassis-Application Communication—When the ASA detects that the chassis-application health has recovered, the ASA tries to rejoin the cluster automatically.
- Failed decorator application—The ASA rejoins the cluster when it senses that the decorator application is back up.

- Internal error—Internal failures include: application sync timeout; inconsistent application statuses; and so on. A unit will attempt to rejoin the cluster automatically at the following intervals: 5 minutes, 10 minutes, and then 20 minutes. This behavior is configurable.

Data Path Connection State Replication

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

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

Table 2: Features Replicated Across the Cluster

Traffic	State Support	Notes
Up time	Yes	Keeps track of the system up time.
ARP Table	Yes	—
MAC address table	Yes	—
User Identity	Yes	Includes AAA rules (uauth) and identity firewall.
IPv6 Neighbor database	Yes	—
Dynamic routing	Yes	—
SNMP Engine ID	No	—
Centralized VPN (Site-to-Site)	No	VPN sessions will be disconnected if the control unit fails.
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 members of the cluster. Connection roles determine how connections are handled in both normal operation and in a high availability situation.

Connection Roles

See the following roles defined for each connection:

- Owner—Usually, the unit 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 units receive packets from the connection, the director chooses a new owner from those units.

- **Backup owner**—The unit 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 unit 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 unit 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 inter-chassis clustering on the Firepower 9300, which can include up to 3 cluster units 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 unit 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 unit 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, and sends a message to the director to register the new connection. If packets arrive at any unit other than the owner, the unit queries the director about which unit 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 unit 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, then there are two director roles: the local director and the global director. The owner always chooses a local director at the same site as itself (based on site ID). The global director can be at any site, and might even be the same unit as the local director. If the original owner fails, then the local director chooses a new connection owner at the same site.

- **Forwarder**—A unit 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, then the forwarder always queries the local director. The forwarder only queries the global director if the local director does not know the owner, for example, if a cluster member receives packets for a connection that is owned on a different site. 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 units 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 units, 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 units. 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.

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 unit 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 unit. If a multi-session PAT connection is initially received by a data unit, then the data unit forwards the connection to the control unit.
By default, UDP (except for DNS UDP) and ICMP traffic use multi-session PAT, so these connections are always owned by the control unit.

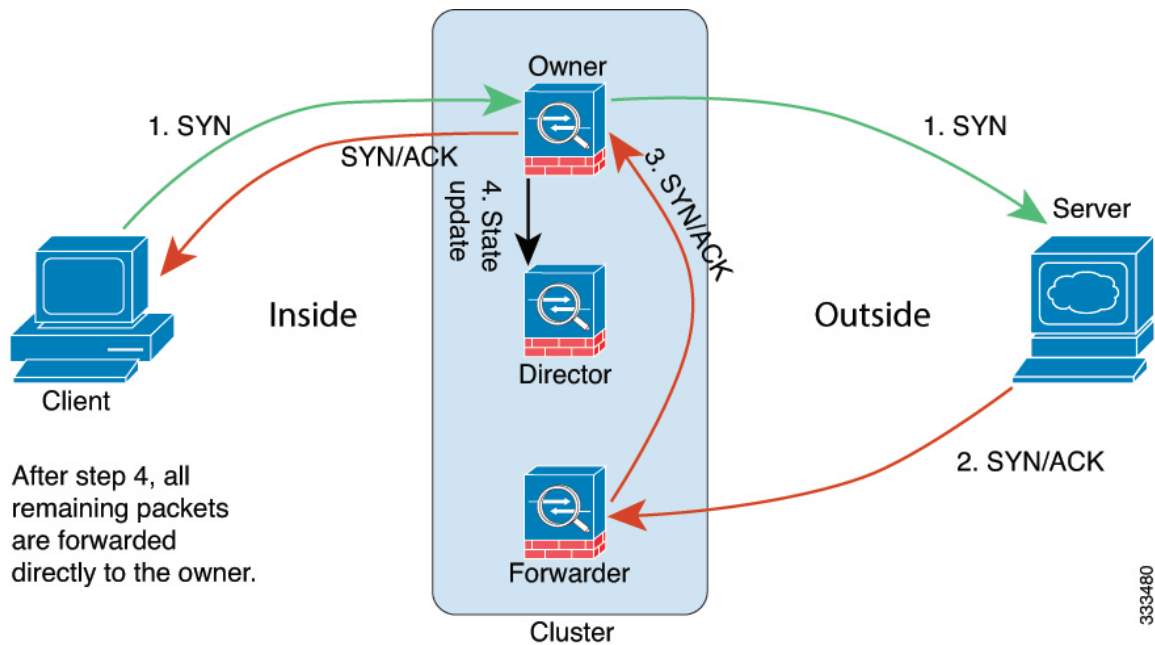
You can change the per-session PAT defaults for TCP and UDP so connections for these protocols are handled per-session or multi-session depending on the configuration. For ICMP, you cannot change from the default multi-session PAT. For more information about per-session PAT, see the firewall configuration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

New Connection Ownership

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

Sample Data Flow

The following example shows the establishment of a new connection.



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

History for ASA Clustering on the Firepower 4100/9300

Feature Name	Version	Feature Information
Automatically rejoin the cluster after an internal failure	9.9(2)	<p>Formerly, many error conditions caused a cluster unit to be removed from the cluster, and you were required to manually rejoin the cluster after resolving the issue. Now, a unit will attempt to rejoin the cluster automatically at the following intervals by default: 5 minutes, 10 minutes, and then 20 minutes. These values are configurable. Internal failures include: application sync timeout; inconsistent application statuses; and so on.</p> <p>New or modified screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Auto Rejoin</p>
Show transport related statistics for cluster reliable transport protocol messages	9.9(2)	<p>You can now view per-unit cluster reliable transport buffer usage so you can identify packet drop issues when the buffer is full in the control plane.</p> <p>New or modified command: show cluster info transport cp detail</p>
cluster remove unit command behavior matches no enable behavior	9.9(1)	<p>The cluster remove unit command now removes a unit from the cluster until you manually reenables clustering or reload, similar to the no enable command. Previously, if you redeployed the bootstrap configuration from FXOS, clustering would be reenables. Now, the disabled status persists even in the case of a bootstrap configuration redeployment. Reloading the ASA, however, will reenables clustering.</p> <p>New/Modified screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster</p>
Improved chassis health check failure detection for the Firepower chassis	9.9(1)	<p>You can now configure a lower holdtime for the chassis health check: 100 ms. The previous minimum was 300 ms. Note that the minimum combined time (<i>interval x retry-count</i>) cannot be less than 600 ms.</p> <p>New or modified command: app-agent heartbeat interval</p> <p>No ASDM support.</p>
Inter-site redundancy for clustering	9.9(1)	<p>Inter-site redundancy ensures that a backup owner for a traffic flow will always be at the other site from the owner. This feature guards against site failure.</p> <p>New or modified screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster</p>

Feature Name	Version	Feature Information
Distributed Site-to-Site VPN with clustering on the Firepower 9300	9.9(1)	<p>An ASA cluster on the Firepower 9300 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 unit (as in centralized mode). This significantly scales VPN support beyond Centralized VPN capabilities and provides high availability. Distributed S2S VPN runs on a cluster of up to two chassis, each containing up to three modules (six total cluster members), each module supporting up to 6K active sessions (12K total), for a maximum of approximately 36K active sessions (72K total).</p> <p>New or modified screens:</p> <p>Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary</p> <p>Monitoring > VPN > VPN Statistics > Sessions</p> <p>Configuration > Device Management > High Availability and Scalability > ASA Cluster Wizards > Site-to-Site</p> <p>Monitoring > VPN > VPN Statistics > Sessions</p> <p>Monitoring > ASA Cluster > ASA Cluster > VPN Cluster Summary</p> <p>Monitoring > ASA Cluster > ASA Cluster > System Resource Graphs > CPU/Memory</p> <p>Monitoring > Logging > Real-Time Log Viewer</p>
Improved cluster unit health-check failure detection	9.8(1)	<p>You can now configure a lower holdtime for the unit health check: .3 seconds minimum. The previous minimum was .8 seconds. This feature changes the unit health check messaging scheme to <i>heartbeats</i> in the data plane from <i>keepalives</i> in the control plane. Using heartbeats improves the reliability and the responsiveness of clustering by not being susceptible to control plane CPU hogging and scheduling delays. Note that configuring a lower holdtime increases cluster control link messaging activity. We suggest that you analyze your network before you configure a low holdtime; for example, make sure a ping from one unit to another over the cluster control link returns within the <i>holdtime/3</i>, because there will be three heartbeat messages during one holdtime interval. If you downgrade your ASA software after setting the hold time to .3 - .7, this setting will revert to the default of 3 seconds because the new setting is unsupported.</p> <p>We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster</p>
Configurable debounce time to mark an interface as failed for the Firepower 4100/9300 chassis	9.8(1)	<p>You can now configure the debounce time before the ASA considers an interface to be failed, and the unit 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 unit is removed from the cluster. The default debounce time is 500 ms, with a range of 300 ms to 9 seconds.</p> <p>New or modified screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster</p>

Feature Name	Version	Feature Information
Inter-site clustering improvement for the ASA on the Firepower 4100/9300 chassis	9.7(1)	<p>You can now configure the site ID for each Firepower 4100/9300 chassis when you deploy the ASA cluster. Previously, you had to configure the site ID within the ASA application; this new feature eases initial deployment. Note that you can no longer set the site ID within the ASA configuration. Also, for best compatibility with inter-site clustering, we recommend that you upgrade to ASA 9.7(1) and FXOS 2.1.1, which includes several improvements to stability and performance.</p> <p>We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Configuration</p>
Director localization: inter-site clustering improvement for data centers	9.7(1)	<p>To improve performance and keep traffic within a site 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 <i>any</i> 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 chooses 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.</p> <p>We modified the following screen: Configuration > Device Management > High Availability and Scalability > Cluster Configuration</p>
Support for 16 chassis for the Firepower 4100 series	9.6(2)	<p>You can now add up to 16 chassis to the cluster for the Firepower 4100 series.</p> <p>We did not modify any screens.</p>
Support for the Firepower 4100 series	9.6(1)	<p>With FXOS 1.1.4, the ASA supports inter-chassis clustering on the Firepower 4100 series for up to 6 chassis.</p> <p>We did not modify any screens.</p>
Support for site-specific IP addresses in Routed, Spanned EtherChannel mode	9.6(1)	<p>For inter-site clustering in routed mode with Spanned EtherChannels, you can now configure site-specific IP addresses in addition to site-specific MAC addresses. The addition of site IP addresses allows you to use ARP inspection on the Overlay Transport Virtualization (OTV) devices to prevent ARP responses from the global MAC address from traveling over the Data Center Interconnect (DCI), which can cause routing problems. ARP inspection is required for some switches that cannot use VACLs to filter MAC addresses.</p> <p>We modified the following screen: Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit EtherChannel Interface > Advanced</p>
Inter-chassis clustering for 16 modules, and inter-site clustering for the Firepower 9300 ASA application	9.5(2.1)	<p>With FXOS 1.1.3, you can now enable inter-chassis, and by extension inter-site clustering. You can include up to 16 modules. For example, you can use 1 module in 16 chassis, or 2 modules in 8 chassis, or any combination that provides a maximum of 16 modules.</p> <p>We did not modify any screens.</p>
Site-specific MAC addresses for inter-site clustering support for Spanned EtherChannel in Routed firewall mode	9.5(2)	<p>You can now use inter-site clustering for Spanned EtherChannels in routed mode. To avoid MAC address flapping, configure a site ID for each cluster member so that a site-specific MAC address for each interface can be shared among a site's units.</p> <p>We modified the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Configuration</p>

Feature Name	Version	Feature Information
ASA cluster customization of the auto-rejoin behavior when an interface or the cluster control link fails	9.5(2)	<p>You can now customize the auto-rejoin behavior when an interface or the cluster control link fails.</p> <p>We introduced the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster > Auto Rejoin</p>
The ASA cluster supports GTPv1 and GTPv2	9.5(2)	<p>The ASA cluster now supports GTPv1 and GTPv2 inspection.</p> <p>We did not modify any screens.</p>
Cluster replication delay for TCP connections	9.5(2)	<p>This feature helps eliminate the “unnecessary work” related to short-lived flows by delaying the director/backup flow creation.</p> <p>We introduced the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster Replication</p>
LISP Inspection for Inter-Site Flow Mobility	9.5(2)	<p>Cisco Locator/ID Separation Protocol (LISP) architecture separates the device identity from its location into two different numbering spaces, making server migration transparent to clients. The ASA can inspect LISP traffic for location changes and then use this information for seamless clustering operation; the ASA cluster members inspect LISP traffic passing between the first hop router and the egress tunnel router (ETR) or ingress tunnel router (ITR), and then change the flow owner to be at the new site.</p> <p>We introduced or modified the following screens:</p> <p>Configuration > Device Management > High Availability and Scalability > ASA Cluster > Cluster Configuration</p> <p>Configuration > Firewall > Objects > Inspect Maps > LISP</p> <p>Configuration > Firewall > Service Policy Rules > Protocol Inspection</p> <p>Configuration > Firewall > Service Policy Rules > Cluster</p> <p>Monitoring > Routing > LISP-EID Table</p>
Carrier Grade NAT enhancements now supported in failover and ASA clustering	9.5(2)	<p>For carrier-grade or large-scale PAT, you can allocate a block of ports for each host, rather than have NAT allocate one port translation at a time (see RFC 6888). This feature is now supported in failover and ASA cluster deployments.</p> <p>We did not modify any screens.</p>
Configurable level for clustering trace entries	9.5(2)	<p>By default, all levels of clustering events are included in the trace buffer, including many low level events. To limit the trace to higher level events, you can set the minimum trace level for the cluster.</p> <p>We did not modify any screens.</p>
Intra-chassis ASA Clustering for the Firepower 9300	9.4(1.150)	<p>You can cluster up to 3 security modules within the Firepower 9300 chassis. All modules in the chassis must belong to the cluster.</p> <p>We introduced the following screen: Configuration > Device Management > High Availability and Scalability > ASA Cluster Replication</p>