Basic Interface Configuration (ASA 5512-X and Higher)

This chapter includes tasks for starting your interface configuration for the ASA 5512-X and higher, including configuring Ethernet settings, redundant interfaces, and EtherChannels.

Note For multiple context mode, complete all tasks in this section in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

For ASA cluster interfaces, which have special requirements, see Chapter 9, "ASA Cluster."

This chapter includes the following sections:

Information About Starting ASA 5512-X and Higher Interface Configuration

This section includes the following topics:

Auto-MDI/MDIX Feature

For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Interfaces in Transparent Mode

Interfaces in transparent mode belong to a "bridge group," one bridge group for each network. You can have up to 8 bridge groups of 4 interfaces each per context or in single mode. For more information about bridge groups, see Bridge Groups in Transparent Mode, page 14-1.

Management Interface

Management Interface Overview

You can manage the ASA by connecting to:

- Any through-traffic interface
- A dedicated Management *Slot*/*Port* interface (if available for your model)

You may need to configure management access to the interface according to Chapter 43, "Management Access."

Management *Slot/Port* Interface

Table 10-1 shows the Management interfaces per model.

Table 10-1 Management Interfaces Per Model

Model	Management 0/0[1]	Management 0/1	Management 1/0	Management 1/1	Configurable for Through Traffic[2]	Subinterfaces Allowed
ASA 5505	No	No	No	No	N/A	N/A
ASA 5512-X	**Yes**	No	No	No	No	No
ASA 5515-X	**Yes**	No	No	No	No	No
ASA 5525-X	**Yes**	No	No	No	No	No
ASA 5545-X	**Yes**	No	No	No	No	No
ASA 5555-X	**Yes**	No	No	No	No	No
ASA 5585-X	**Yes**	**Yes**	**Yes**[3]	**Yes**[3]	**Yes**	**Yes**
ASASM	No	No	No	No	N/A	N/A
ASAv	**Yes**	No	No	No	No	No

1. The Management 0/0 interface is configured for ASDM access as part of the default factory configuration. See Factory Default Configurations, page 4-18 for more information.
2. By default, the Management 0/0 interface is configured for management-only traffic (the **management-only** command). For supported models in routed mode, you can remove the limitation and pass through traffic. If your model includes additional Management interfaces, you can use them for through traffic as well. The Management interfaces might not be optimized for through-traffic, however.
3. If you installed an SSP in slot 1, then Management 1/0 and 1/1 provide management access to the SSP in slot 1 only.

Note If you installed a module, then the module management interface(s) provides management access for the module only. For the ASA 5512-X through ASA 5555-X, the software module uses the same physical Management 0/0 interface as the ASA.

Using Any Interface for Management-Only Traffic

You can use any interface as a dedicated management-only interface by configuring it for management traffic, including an EtherChannel interface (see the **management-only** command).

Management Interface for Transparent Mode

In transparent firewall mode, in addition to the maximum allowed through-traffic interfaces, you can also use the Management interface (either the physical interface, a subinterface (if supported for your model), or an EtherChannel interface comprised of Management interfaces (if you have multiple Management interfaces)) as a separate management interface. You cannot use any other interface types as management interfaces.

In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. To provide management per context, you can create subinterfaces of the Management interface and allocate a Management subinterface to each context. Note that the ASA 5512-X through ASA 5555-X do not allow subinterfaces on the Management interface, so for per-context management, you must connect to a data interface.

The management interface is not part of a normal bridge group. Note that for operational purposes, it is part of a non-configurable bridge group.

Note In transparent firewall mode, the management interface updates the MAC address table in the same manner as a data interface; therefore you should not connect both a management and a data interface to the same switch unless you configure one of the switch ports as a routed port (by default Cisco Catalyst switches share a MAC address for all VLAN switch ports). Otherwise, if traffic arrives on the management interface from the physically-connected switch, then the ASA updates the MAC address table to use the *management* interface to access the switch, instead of the data interface. This action causes a temporary traffic interruption; the ASA will not re-update the MAC address table for packets from the switch to the data interface for at least 30 seconds for security reasons.

No Support for Redundant Management Interfaces

Redundant interfaces do not support Management *slot*/*port* interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.

Management 0/0 Interface on the ASA 5512-X through ASA 5555-X

The Management 0/0 interface on the ASA 5512-X through ASA 5555-X has the following characteristics:

- No through traffic support
- No subinterface support
- No priority queue support
- No multicast MAC support
- The software module shares the Management 0/0 interface. Separate MAC addresses and IP addresses are supported for the ASA and module. You must perform configuration of the module IP address within the module operating system. However, physical characteristics (such as enabling the interface) are configured on the ASA.

Redundant Interfaces

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the ASA reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as device-level failover if desired.

Redundant Interface MAC Address

The redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. Alternatively, you can assign a MAC address to the redundant interface, which is used regardless of the member interface MAC addresses (see Configuring the MAC Address, MTU, and TCP MSS, page 13-9 or the Configuring Multiple Contexts, page 7-15). When the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

EtherChannels

An 802.3ad EtherChannel is a logical interface (called a port-channel interface) consisting of a bundle of individual Ethernet links (a channel group) so that you increase the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels.

This section includes the following topics:

Channel Group Interfaces

Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure. For 16 active interfaces, be sure that your switch supports the feature (for example, the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module).

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the ASA EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Cisco Catalyst 6500 switch or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you can connect ASA interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch (see Figure 10-1).

Figure 10-1 Connecting to a VSS/vPC

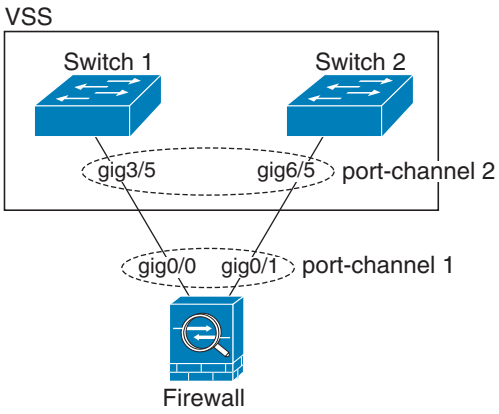

If you use the ASA in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each ASA (see Figure 10-1). On each ASA, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both ASAs (in this case, the EtherChannel will not be established because of the separate ASA system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby ASA.

Figure 10-2 Active/Standby Failover and VSS/vPC

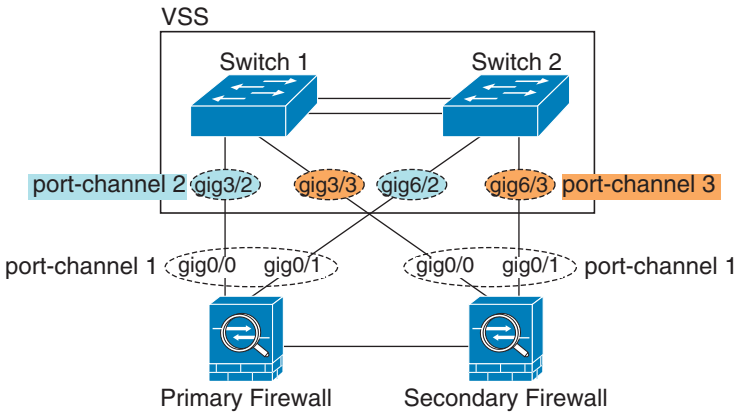

Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel.

- On—The EtherChannel is always on, and LACP is not used. An “on” EtherChannel can only establish a connection with another “on” EtherChannel.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. “On” mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load Balancing

The ASA distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable; see Customizing the EtherChannel, page 10-21). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value* **mod** *active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

For a spanned EtherChannel in clustering, load balancing occurs on a per ASA basis. For example, if you have 32 active interfaces in the spanned EtherChannel across 8 ASAs, with 4 interfaces per ASA in the EtherChannel, then load balancing only occurs across the 4 interfaces on the ASA.

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC Address

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. In multiple context mode, you can automatically assign unique MAC addresses to interfaces, including an EtherChannel port interface. We recommend manually, or in multiple context mode, automatically configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size

MTU Overview

The maximum transmission unit (MTU) specifies the maximum frame payload size that the ASA can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, FCS, or VLAN tagging. The Ethernet header is 14 bytes and the FCS is 4 bytes. When you set the MTU to 1500, the expected frame size is 1518 bytes including the headers. If you are using VLAN tagging (which adds an additional 4 bytes), then when you set the MTU to 1500, the expected frame size is 1522. Do not set the MTU value higher to accommodate these headers. For information about accommodating TCP headers for encapsulation, do not alter the MTU setting; instead change the TCP Maximum Segment Size (the TCP Maximum Segment Size Overview, page 10-9).

If an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can cause performance degradation. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

Note The ASA can receive frames larger than the configured MTU as long as there is room in memory. See Enabling Jumbo Frame Support, page 10-24 to increase memory for larger frames.

Default MTU

The default MTU on the ASA is 1500 bytes. This value does not include the 18 or more bytes for the Ethernet header, CRC, VLAN tagging, and so on.

Path MTU Discovery

The ASA supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

Setting the MTU and Jumbo Frames

See Configuring the MAC Address, MTU, and TCP MSS, page 13-9. For multiple context mode, set the MTU within each context.

See Enabling Jumbo Frame Support, page 10-24. For multiple context mode, set the jumbo frame support in the system execution space.

See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all ASA interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—If you enable jumbo frames, you can set the MTU up to 9198 bytes.

TCP Maximum Segment Size Overview

The TCP maximum segment size (TCP MSS) is the size of the TCP payload *before* any TCP headers are added. UDP packets are not affected. The client and the server exchange TCP MSS values during the three-way handshake when establishing the connection.

You can set the TCP MSS on the ASA. If either endpoint of a connection requests a TCP MSS that is larger than the value set on the ASA, the ASA overwrites the TCP MSS in the request packet with the ASA maximum. If the host or server does not request a TCP MSS, then the ASA assumes the RFC 793-default value of 536 bytes, but does not modify the packet. You can also configure the minimum TCP MSS; if a host or server requests a very small TCP MSS, the ASA can adjust the value up. By default, the minimum TCP MSS is not enabled.

For example, you configure the default MTU of 1500 bytes. A host requests an MSS of 1700. If the ASA maximum TCP MSS is 1380, then the ASA changes the MSS value in the TCP request packet to 1380. The server then sends 1380-byte packets.

Default TCP MSS

By default, the maximum TCP MSS on the ASA is 1380 bytes. This default accommodates VPN connections where the headers can add up to 120 bytes; this value fits within the default MTU of 1500 bytes.

Setting the TCP MSS for VPN and Non-VPN Traffic

See Configuring the MAC Address, MTU, and TCP MSS, page 13-9. For multiple context mode, set the TCP MSS within each context.

See the following guidelines:

- Non-VPN traffic—If you do not use VPN and do not need extra space for headers, then you should disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-VPN packets usually fit this TCP MSS.
- VPN traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to a higher value, then you need to set the TCP MSS to accommodate the new MTU.

Examples

The following example enables jumbo frames, increases the MTU on all interfaces, and disables the TCP MSS for non-VPN traffic (by setting the TCP MSS to 0, which means there is no limit):

```
jumbo frame-reservation
mtu inside 9198
mtu outside 9198
sysopt connection tcpmss 0
```

The following example enables jumbo frames, increases the MTU on all interfaces, and changes the TCP MSS for VPN traffic to 9078 (the MTU minus 120):

```
jumbo frame-reservation
mtu inside 9198
mtu outside 9198
sysopt connection tcpmss 9078
```

Licensing Requirements for ASA 5512-X and Higher Interfaces

Model	License Requirement
ASA 5512-X	VLANs[1]: Base License: 50 Security Plus License: 100 Interfaces of all types[2]: Base License: 716 Security Plus License: 916
ASA 5515-X	VLANs[1]: Base License: 100 Interfaces of all types[2]: Base License: 916
ASA 5525-X	VLANs[1]: Base License: 200 Interfaces of all types[2]: Base License: 1316
ASA 5545-X	VLANs[1]: Base License: 300 Interfaces of all types[2]: Base License: 1716
ASA 5555-X	VLANs[1]: Base License: 500 Interfaces of all types[2]: Base License: 2516
ASA 5585-X	VLANs[1]: Base and Security Plus License: 1024 Interface Speed for SSP-10 and SSP-20: Base License—1-Gigabit Ethernet for fiber interfaces 10 GE I/O License (Security Plus)—10-Gigabit Ethernet for fiber interfaces (SSP-40 and SSP-60 support 10-Gigabit Ethernet by default.) Interfaces of all types[2]: Base and Security Plus License: 4612

1. For an interface to count against the VLAN limit, you must assign a VLAN to it. For example:
   interface gigabitethernet 0/0.100
   vlan 100

2. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces. Every **interface** command defined in the configuration counts against this limit. For example, both of the following interfaces count even if the GigabitEthernet 0/0 interface is defined as part of port-channel 1:
interface gigabitethernet 0/0
and
interface port-channel 1

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

In multiple context mode, configure the physical interfaces in the system execution space according to the Starting Interface Configuration (ASA 5512-X and Higher), page 10-13. Then, configure the logical interface parameters in the context execution space according to Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Firewall Mode Guidelines

- For transparent mode, you can configure up to 8 bridge groups per context or for a single mode device.
- Each bridge group can include up to 4 interfaces.
- For multiple context, transparent mode, each context must use different interfaces; you cannot share an interface across contexts.

Failover Guidelines

- When you use a redundant or EtherChannel interface as a failover link, it must be pre-configured on both units in the failover pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the failover link itself is required for replication*.
- If you use a redundant or EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal.
- You can monitor redundant or EtherChannel interfaces for failover using the **monitor-interface** command; be sure to reference the logical redundant interface name. When an active member interface fails over to a standby interface, this activity does not cause the redundant or EtherChannel interface to appear to be failed when being monitored for device-level failover. Only when all physical interfaces fail does the redundant or EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a failover or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a failover link. To alter the configuration, you need to either shut down the EtherChannel while you make changes, or temporarily disable failover; either action prevents failover from occurring for the duration.
- You cannot share a failover or state interface with a data interface.

Clustering Guidelines

- When you use a redundant or EtherChannel interface as the cluster control link, it must be pre-configured on all units in the cluster; you cannot configure it on the primary unit and expect it to replicate to member units because *the cluster control link itself is required for replication*.
- To configure a spanned EtherChannel, see Configuring Spanned EtherChannels, page 9-42.
- To configure an individual cluster interface, see Configuring Individual Interfaces (Recommended for the Management Interface), page 9-40.

Redundant Interface Guidelines

- You can configure up to 8 redundant interface pairs.
- All ASA configuration refers to the logical redundant interface instead of the member physical interfaces.
- You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the ASA if they do not use the same physical interfaces.
- If you shut down the active interface, then the standby interface becomes active.
- Redundant interfaces do not support Management *slot*/*port* interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.
- For failover guidelines, see Failover Guidelines, page 10-11.
- For clustering guidelines, see Clustering Guidelines, page 10-12.

EtherChannel Guidelines

- You can configure up to 48 EtherChannels.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.
- The device to which you connect the ASA EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Cisco Catalyst 6500 switch or Cisco Nexus 7000 switch.
- The ASA does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the ASA will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch. In multiple context mode, these messages are not included in a packet capture, so that you cannot diagnose the issue easily.
- The ASA does not support connecting an EtherChannel to a switch stack. If the ASA EtherChannel is connected cross stack, and if the Master switch is powered down, then the EtherChannel connected to the remaining switch will not come up.
- All ASA configuration refers to the logical EtherChannel interface instead of the member physical interfaces.
- You cannot use a redundant interface as part of an EtherChannel, nor can you use an EtherChannel as part of a redundant interface. You cannot use the same physical interfaces in a redundant interface and an EtherChannel interface. You can, however, configure both types on the ASA if they do not use the same physical interfaces.

- For failover guidelines, see Failover Guidelines, page 10-11.
- For clustering guidelines, see Clustering Guidelines, page 10-12.

Default Settings

This section lists default settings for interfaces if you do not have a factory default configuration. For information about the factory default configurations, see Factory Default Configurations, page 4-18.

Default State of Interfaces

The default state of an interface depends on the type and the context mode.

In multiple context mode, all allocated interfaces are enabled by default, no matter what the state of the interface is in the system execution space. However, for traffic to pass through the interface, the interface also has to be enabled in the system execution space. If you shut down an interface in the system execution space, then that interface is down in all contexts that share it.

In single mode or in the system execution space, interfaces have the following default states:

- Physical interfaces—Disabled.
- Redundant Interfaces—Enabled. However, for traffic to pass through the redundant interface, the member physical interfaces must also be enabled.
- Subinterfaces—Enabled. However, for traffic to pass through the subinterface, the physical interface must also be enabled.
- EtherChannel port-channel interfaces—Enabled. However, for traffic to pass through the EtherChannel, the channel group physical interfaces must also be enabled.

Default Speed and Duplex

- By default, the speed and duplex for copper (RJ-45) interfaces are set to auto-negotiate.
- For fiber interfaces for the 5585-X, the speed is set for automatic link negotiation.

Default Connector Type

Some models include two connector types: copper RJ-45 and fiber SFP. RJ-45 is the default. You can configure the ASA to use the fiber SFP connectors.

Default MAC Addresses

By default, the physical interface uses the burned-in MAC address, and all subinterfaces of a physical interface use the same burned-in MAC address.

Starting Interface Configuration (ASA 5512-X and Higher)

This section includes the following topics:

- Enabling Jumbo Frame Support, page 10-24
- Converting In-Use Interfaces to a Redundant or EtherChannel Interface, page 10-25

Task Flow for Starting Interface Configuration

Note If you have an existing configuration, and want to convert interfaces that are in use to a redundant or EtherChannel interface, perform your configuration offline to minimize disruption. See Converting In-Use Interfaces to a Redundant or EtherChannel Interface, page 10-25.

To start configuring interfaces, perform the following steps:

Step 1 (Multiple context mode) Complete all tasks in this section in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Step 2 Enable the physical interface, and optionally change Ethernet parameters. See Enabling the Physical Interface and Configuring Ethernet Parameters, page 10-14.

Physical interfaces are disabled by default.

Step 3 (Optional) Configure redundant interface pairs. See Configuring a Redundant Interface, page 10-17.

A logical redundant interface pairs an active and a standby physical interface. When the active interface fails, the standby interface becomes active and starts passing traffic.

Step 4 (Optional) Configure an EtherChannel. See Configuring an EtherChannel, page 10-19.

An EtherChannel groups multiple Ethernet interfaces into a single logical interface.

Step 5 (Optional) Configure VLAN subinterfaces. See Configuring VLAN Subinterfaces and 802.1Q Trunking, page 10-22.

Step 6 (Optional) Enable jumbo frame support according to the Enabling Jumbo Frame Support, page 10-24.

Step 7 (Multiple context mode only) To complete the configuration of interfaces in the system execution space, perform the following tasks that are documented in Chapter 7, "Multiple Context Mode":

- To assign interfaces to contexts, see Configuring a Security Context, page 7-19.
- (Optional) To automatically assign unique MAC addresses to context interfaces, see Automatically Assigning MAC Addresses to Context Interfaces, page 7-24.

The MAC address is used to classify packets within a context. If you share an interface, but do not have unique MAC addresses for the interface in each context, then the destination IP address is used to classify packets. Alternatively, you can manually assign MAC addresses within the context according to the Configuring the MAC Address, MTU, and TCP MSS, page 13-9.

Step 8 Complete the interface configuration according to Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Enabling the Physical Interface and Configuring Ethernet Parameters

This section describes how to:

- Enable the physical interface

- Set a specific speed and duplex (if available)
- Enable pause frames for flow control

Prerequisites

For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	**interface** *physical_interface* **Example:** `ciscoasa(config)# interface gigabitethernet 0/0`	Specifies the interface you want to configure. where the *physical_interface* ID includes the type, slot, and port number as *type*[*slot*/]*port*. The physical interface types include the following: • **gigabitethernet** • **tengigabitethernet** • **management** Enter the type followed by *slot*/*port*, for example, **gigabitethernet0/1**. A space is optional between the type and the slot/port.
Step 2	(Optional) **media-type sfp** **Example:** `ciscoasa(config-if)# media-type sfp`	Sets the media type to SFP, if available for your model. To restore the default RJ-45, enter the **media-type rj45** command.
Step 3	(Optional) **speed** {**auto** | **10** | **100** | **1000** | **nonegotiate**} **Example:** `ciscoasa(config-if)# speed 100`	Sets the speed. For RJ-45 interfaces, the default setting is **auto**. For SFP interfaces, the default setting is **no speed nonegotiate**, which sets the speed to the maximum speed and enables link negotiation for flow-control parameters and remote fault information. The **nonegotiate** keyword is the only keyword available for SFP interfaces. The **speed nonegotiate** command disables link negotiation.
Step 4	(Optional) **duplex** {**auto** | **full** | **half**} **Example:** `ciscoasa(config-if)# duplex full`	Sets the duplex for RJ-45 interfaces. The **auto** setting is the default. **Note** The duplex setting for an EtherChannel interface must be Full or Auto.

	Command	Purpose
Step 5	(Optional) **flowcontrol send on** [*low_water high_water pause_time*] [**noconfirm**] **Example:** `ciscoasa(config-if)# flowcontrol send on 95 200 10000`	Enables pause (XOFF) frames for flow control on GigabitEthernet and TenGigabitEthernet interfaces. If you have a traffic burst, dropped packets can occur if the burst exceeds the buffering capacity of the FIFO buffer on the NIC and the receive ring buffers. Enabling pause frames for flow control can alleviate this issue. Pause (XOFF) and XON frames are generated automatically by the NIC hardware based on the FIFO buffer usage. A pause frame is sent when the buffer usage exceeds the high-water mark. The default *high_water* value is 128 KB (10 GigabitEthernet) and 24 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and 47 KB (1 GigabitEthernet). After a pause is sent, an XON frame can be sent when the buffer usage is reduced below the low-water mark. By default, the *low_water* value is 64 KB (10 GigabitEthernet) and 16 KB (1 GigabitEthernet); you can set it between 0 and 511 (10 GigabitEthernet) or 0 and 47 KB (1 GigabitEthernet). The link partner can resume traffic after receiving an XON, or after the XOFF expires, as controlled by the timer value in the pause frame. The default *pause_time* value is 26624; you can set it between 0 and 65535. If the buffer usage is consistently above the high-water mark, pause frames are sent repeatedly, controlled by the pause refresh threshold value. When you use this command, you see the following warning: `Changing flow-control parameters will reset the interface. Packets may be lost during the reset. Proceed with flow-control changes?` To change the parameters without being prompted, use the **noconfirm** keyword. **Note** Only flow control frames defined in 802.3x are supported. Priority-based flow control is not supported.
Step 6	**no shutdown** **Example:** `ciscoasa(config-if)# no shutdown`	Enables the interface. To disable the interface, enter the **shutdown** command. If you enter the **shutdown** command, you also shut down all subinterfaces. If you shut down an interface in the system execution space, then that interface is shut down in all contexts that share it.

What to Do Next

Optional Tasks:

- Configure redundant interface pairs. See Configuring a Redundant Interface, page 10-17.
- Configure an EtherChannel. See Configuring an EtherChannel, page 10-19.
- Configure VLAN subinterfaces. See Configuring VLAN Subinterfaces and 802.1Q Trunking, page 10-22.
- Configure jumbo frame support. See Enabling Jumbo Frame Support, page 10-24.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Configuring a Redundant Interface

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the ASA reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as failover if desired.

This section describes how to configure redundant interfaces and includes the following topics:

Configuring a Redundant Interface

This section describes how to create a redundant interface. By default, redundant interfaces are enabled.

Guidelines and Limitations

- You can configure up to 8 redundant interface pairs.
- Redundant interface delay values are configurable, but by default the ASA inherits the default delay values based on the physical type of its member interfaces.
- See also the Redundant Interface Guidelines, page 10-12.

Prerequisites

- Both member interfaces must be of the same physical type. For example, both must be GigabitEthernet.
- You cannot add a physical interface to the redundant interface if you configured a name for it. You must first remove the name using the **no nameif** command.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Detailed Steps

	Command	Purpose
Step 1	**interface redundant** *number* **Example:** `ciscoasa(config)# interface redundant 1`	Adds the logical redundant interface, where the *number* argument is an integer between 1 and 8. **Note** You need to add at least one member interface to the redundant interface before you can configure logical parameters for it such as a name.
Step 2	**member-interface** *physical_interface* **Example:** `ciscoasa(config-if)# member-interface gigabitethernet 0/0`	Adds the first member interface to the redundant interface. See Enabling the Physical Interface and Configuring Ethernet Parameters, page 10-14 for a description of the physical interface ID. Redundant interfaces do not support Management *slot/port* interfaces as members. After you add the interface, any configuration for it (such as an IP address) is removed.
Step 3	**member-interface** *physical_interface* **Example:** `ciscoasa(config-if)# member-interface gigabitethernet 0/1`	Adds the second member interface to the redundant interface. Make sure the second interface is the same physical type as the first interface. To remove a member interface, enter the **no member-interface** *physical_interface* command. You cannot remove both member interfaces from the redundant interface; the redundant interface requires at least one member interface.

Examples

The following example creates two redundant interfaces:

```
ciscoasa(config)# interface redundant 1
ciscoasa(config-if)# member-interface gigabitethernet 0/0
ciscoasa(config-if)# member-interface gigabitethernet 0/1
ciscoasa(config-if)# interface redundant 2
ciscoasa(config-if)# member-interface gigabitethernet 0/2
ciscoasa(config-if)# member-interface gigabitethernet 0/3
```

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See Configuring VLAN Subinterfaces and 802.1Q Trunking, page 10-22.
- Configure jumbo frame support. See Enabling Jumbo Frame Support, page 10-24.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, “Routed Mode Interfaces,” or Chapter 14, “Transparent Mode Interfaces.”

Changing the Active Interface

By default, the active interface is the first interface listed in the configuration, if it is available. To view which interface is active, enter the following command:

```
ciscoasa# show interface redundantnumber detail | grep Member
```

For example:

```
ciscoasa# show interface redundant1 detail | grep Member
      Members GigabitEthernet0/3(Active), GigabitEthernet0/2
```

To change the active interface, enter the following command:

```
ciscoasa# redundant-interface redundantnumber active-member physical_interface
```

where the **redundant***number* argument is the redundant interface ID, such as **redundant1**.

The *physical_interface* is the member interface ID that you want to be active.

Configuring an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

This section includes the following topics:

Adding Interfaces to the EtherChannel

This section describes how to create an EtherChannel port-channel interface and assign interfaces to the EtherChannel. By default, port-channel interfaces are enabled.

Guidelines and Limitations

- You can configure up to 48 EtherChannels.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only eight interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- To configure a spanned EtherChannel for clustering, see Configuring Spanned EtherChannels, page 9-42 instead of this procedure.
- See also the EtherChannel Guidelines, page 10-12.

Prerequisites

- All interfaces in the channel group must be the same type, speed, and duplex. Half duplex is not supported.
- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name using the **no nameif** command.
- For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Detailed Steps

	Command	Purpose
Step 1	**interface** *physical_interface* **Example:** `ciscoasa(config)# interface gigabitethernet 0/0`	Specifies the interface you want to add to the channel group, where the *physical_interface* ID includes the type, slot, and port number as *type*[*slot*/]*port*. This first interface in the channel group determines the type and speed for all other interfaces in the group. In transparent mode, if you create a channel group with multiple Management interfaces, then you can use this EtherChannel as the management-only interface.
Step 2	**channel-group** *channel_id* **mode** {**active** | **passive** | **on**} **Example:** `ciscoasa(config-if)# channel-group 1 mode active`	Assigns this physical interface to an EtherChannel with the *channel_id* between 1 and 48. If the port-channel interface for this channel ID does not yet exist in the configuration, one will be added: **interface port-channel** *channel_id* We recommend using **active** mode. For information about active, passive, and on modes, see Link Aggregation Control Protocol, page 10-6.
Step 3	(Optional) **lacp port-priority** *number* **Example:** `ciscoasa(config-if)# lacp port-priority 12345`	Sets the priority for a physical interface in the channel group between 1 and 65535. The default is 32768. The higher the number, the lower the priority. The ASA uses this setting to decide which interfaces are active and which are standby if you assign more interfaces than can be used. If the port priority setting is the same for all interfaces, then the priority is determined by the interface ID (slot/port). The lowest interface ID is the highest priority. For example, GigabitEthernet 0/0 is a higher priority than GigabitEthernet 0/1. If you want to prioritize an interface to be active even though it has a higher interface ID, then set this command to have a lower value. For example, to make GigabitEthernet 1/3 active before GigabitEthernet 0/7, then make the **lacp port-priority** value be 12345 on the 1/3 interface vs. the default 32768 on the 0/7 interface. If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. See the **lacp system-priority** command in the Customizing the EtherChannel, page 10-21.
Step 4	Repeat steps 1 through 3 for each interface you want to add to the channel group.	Each interface in the channel group must be the same type and speed. Half duplex is not supported. If you add an interface that does not match, it will be placed in a suspended state.

What to Do Next

Optional Tasks:

- Customize the EtherChannel interface. See Customizing the EtherChannel, page 10-21.
- Configure VLAN subinterfaces. See Configuring VLAN Subinterfaces and 802.1Q Trunking, page 10-22.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Customizing the EtherChannel

This section describes how to set the maximum number of interfaces in the EtherChannel, the minimum number of operating interfaces for the EtherChannel to be active, the load balancing algorithm, and other optional parameters.

Detailed Steps

	Command	Purpose
Step 1	**interface port-channel** *channel_id* **Example:** `ciscoasa(config)# interface port-channel 1`	Specifies the port-channel interface. This interface was created automatically when you added an interface to the channel group. If you have not yet added an interface, then this command creates the port-channel interface. **Note** You need to add at least one member interface to the port-channel interface before you can configure logical parameters for it such as a name.
Step 2	**lacp max-bundle** *number* **Example:** `ciscoasa(config-if)# lacp max-bundle 6`	Specifies the maximum number of active interfaces allowed in the channel group, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
Step 3	**port-channel min-bundle** *number* **Example:** `ciscoasa(config-if)# port-channel min-bundle 2`	Specifies the minimum number of active interfaces required for the port-channel interface to become active, between 1 and 16. The default is 1. If the active interfaces in the channel group falls below this value, then the port-channel interface goes down, and could trigger a device-level failover.

	Command	Purpose
Step 4	**port-channel load-balance** {**dst-ip** | **dst-ip-port** | **dst-mac** | **dst-port** | **src-dst-ip** | **src-dst-ip-port** | **src-dst-mac** | **src-dst-port** | **src-ip** | **src-ip-port** | **src-mac** | **src-port** | **vlan-dst-ip** | **vlan-dst-ip-port** | **vlan-only** | **vlan-src-dst-ip** | **vlan-src-dst-ip-port** | **vlan-src-ip** | **vlan-src-ip-port**} **Example:** `ciscoasa(config-if)# port-channel load-balance src-dst-mac`	Configures the load-balancing algorithm. By default, the ASA balances the packet load on interfaces according to the source and destination IP address (**src-dst-ip**) of the packet. If you want to change the properties on which the packet is categorized, use this command. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see Load Balancing, page 10-7.
Step 5	**lacp system-priority** *number* **Example:** `ciscoasa(config)# lacp system-priority 12345`	Sets the LACP system priority, from 1 to 65535. The default is 32768. The higher the number, the lower the priority. This command is global for the ASA. If the device at the other end of the EtherChannel has conflicting port priorities, the system priority is used to determine which port priorities to use. For interface priorities within an EtherChannel, see the **lacp port-priority** command in the Adding Interfaces to the EtherChannel, page 10-19.
Step 6	(Optional) You can set the Ethernet properties for the port-channel interface to override the properties set on the individual interfaces.	See Enabling the Physical Interface and Configuring Ethernet Parameters, page 10-14 for Ethernet commands. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

What to Do Next

Optional Task:

- Configure VLAN subinterfaces. See Configuring VLAN Subinterfaces and 802.1Q Trunking, page 10-22.
- Configure jumbo frame support. See Enabling Jumbo Frame Support, page 10-24.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, “Routed Mode Interfaces,” or Chapter 14, “Transparent Mode Interfaces.”

Configuring VLAN Subinterfaces and 802.1Q Trunking

Subinterfaces let you divide a physical, redundant, or EtherChannel interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or ASAs. This feature is particularly useful in multiple context mode so that you can assign unique interfaces to each context.

Guidelines and Limitations

- Maximum subinterfaces—To determine how many VLAN subinterfaces are allowed for your model, see Licensing Requirements for ASA 5512-X and Higher Interfaces, page 10-10.
- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. This property is also true for the active physical interface in a redundant interface pair and for EtherChannel links. Because the physical, redundant, or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical, redundant, or EtherChannel interface does not pass traffic by leaving out the **nameif** command. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the **nameif** command as usual. See Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces," for more information about completing the interface configuration.
- (ASA 5512-X through ASA 5555-X) You cannot configure subinterfaces on the Management 0/0 interface.

Prerequisites

For multiple context mode, complete this procedure in the system execution space. To change from the context to the system execution space, enter the **changeto system** command.

Detailed Steps

	Command	Purpose
Step 1	**interface** {*physical_interface* | **redundant** *number* | **port-channel** *number*}.*subinterface* **Example:** `ciscoasa(config)# interface gigabitethernet 0/1.100`	Specifies the new subinterface. See Enabling the Physical Interface and Configuring Ethernet Parameters, page 10-14 section for a description of the physical interface ID. The **redundant** *number* argument is the redundant interface ID, such as **redundant 1**. The **port-channel** *number* argument is the EtherChannel interface ID, such as **port-channel 1**. The *subinterface* ID is an integer between 1 and 4294967293.
Step 2	**vlan** *vlan_id* **Example:** `ciscoasa(config-subif)# vlan 101`	Specifies the VLAN for the subinterface. The *vlan_id* is an integer between 1 and 4094. Some VLAN IDs might be reserved on connected switches, so check the switch documentation for more information. You can only assign a single VLAN to a subinterface, and you cannot assign the same VLAN to multiple subinterfaces. You cannot assign a VLAN to the physical interface. Each subinterface must have a VLAN ID before it can pass traffic. To change a VLAN ID, you do not need to remove the old VLAN ID with the **no** option; you can enter the **vlan** command with a different VLAN ID, and the ASA changes the old ID.

What to Do Next

Optional Task:

- Configure jumbo frame support. See Enabling Jumbo Frame Support, page 10-24.

Required Tasks:

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Enabling Jumbo Frame Support

A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as ACLs. See Controlling Fragmentation with the Maximum Transmission Unit and TCP Maximum Segment Size, page 10-7 for more information.

Prerequisites

- In multiple context mode, set this option in the system execution space.
- Changes in this setting require you to reload the ASA.
- Be sure to set the MTU for each interface that needs to transmit jumbo frames to a higher value than the default 1500; for example, set the value to 9198 using the **mtu** command. See Configuring the MAC Address, MTU, and TCP MSS, page 13-9. In multiple context mode, set the MTU within each context.
- Be sure to adjust the TCP MSS, either to disable it for non-VPN traffic (**sysopt connection tcpmss 0**), or to increase it in accord with the MTU according to the Configuring the MAC Address, MTU, and TCP MSS, page 13-9.

Detailed Steps

Command	Purpose
jumbo-frame reservation **Example:** `ciscoasa(config)# jumbo-frame reservation`	Enables jumbo frame support. To disable jumbo frames, use the **no** form of this command.

Examples

The following example enables jumbo frame reservation, saves the configuration, and reloads the ASA:

```
ciscoasa(config)# jumbo-frame reservation
WARNING: this command will take effect after the running-config is saved
and the system has been rebooted. Command accepted.

ciscoasa(config)# write memory
Building configuration...
Cryptochecksum: 718e3706 4edb11ea 69af58d0 0a6b7cb5

70291 bytes copied in 3.710 secs (23430 bytes/sec)
[OK]
ciscoasa(config)# reload
Proceed with reload? [confirm] Y
```

What to Do Next

- For multiple context mode, assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Configuring Multiple Contexts, page 7-15.
- For single context mode, complete the interface configuration. See Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Converting In-Use Interfaces to a Redundant or EtherChannel Interface

If you have an existing configuration and want to take advantage of the redundant or EtherChannel interface feature for interfaces that are currently in use, you will have some amount of downtime when you convert to the logical interfaces.

This section provides an overview of how to convert your existing interfaces to a redundant or EtherChannel interface with minimal downtime. See Configuring a Redundant Interface, page 10-17 and the Configuring an EtherChannel, page 10-19 for more information.

Detailed Steps (Single Mode)

We recommend that you update your configuration offline as a text file, and reimport the whole configuration for the following reasons:

- Because you cannot add a named interface as a member of a redundant or EtherChannel interface, you must remove the name from the interface. When you remove the name from the interface, any command that referred to that name is deleted. Because commands that refer to interface names are widespread throughout the configuration and affect multiple features, removing a name from an in-use interface at the CLI or in ASDM would cause significant damage to your configuration, not to mention significant downtime while you reconfigure all your features around a new interface name.
- Changing your configuration offline lets you use the same interface names for your new logical interfaces, so that you do not need to touch the feature configurations that refer to interface names. You only need to change the interface configuration.
- Clearing the running configuration and immediately applying a new configuration will minimize the downtime of your interfaces. You will not be waiting to configure the interfaces in real time.

Step 1 Connect to the ASA; if you are using failover, connect to the active ASA.

Step 2 If you are using failover, disable failover by entering the **no failover** command.

Step 3 Copy the running configuration by entering the **more system:running-config** command and copying the display output to a text editor.

Be sure to save an extra copy of the old configuration in case you make an error when you edit it.

Step 4 For each in-use interface that you want to add to a redundant or EtherChannel interface, cut and paste all commands under the **interface** command to the end of the interface configuration section for use in creating your new logical interfaces. The only exceptions are the following commands, which should stay with the physical interface configuration:

- **media-type**
- **speed**
- **duplex**
- **flowcontrol**

Note You can only add *physical* interfaces to an EtherChannel or redundant interface; you cannot have VLANs configured for the physical interfaces.

Be sure to match the above values for all interfaces in a given EtherChannel or redundant interface. Note that the duplex setting for an EtherChannel interface must be Full or Auto.

For example, you have the following interface configuration. The bolded commands are the ones we want to use with three new EtherChannel interfaces, and that you should cut and paste to the end of the interface section.

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/1
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 no shutdown
!
interface Management0/1
 shutdown
```

```
 no nameif
 no security-level
 no ip address
```

Step 5 Above each pasted command section, create your new logical interfaces by entering one of the following commands:

- **interface redundant** *number* [1-8]
- **interface port-channel** *channel_id* [1-48]

For example:

```
...

interface port-channel 1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
 no shutdown
!
interface port-channel 2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface port-channel 3
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 no shutdown
```

Step 6 Assign the physical interfaces to the new logical interfaces:

- Redundant interface—Enter the following commands under the new **interface redundant** command:

  member-interface *physical_interface1*
  member-interface *physical_interface2*

  Where the physical interfaces are any two interfaces of the same type (either formerly in use or unused). You cannot assign a Management interface to a redundant interface.

  For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside redundant interfaces:

  ```
  interface redundant 1
   nameif outside
   security-level 0
   ip address 10.86.194.225 255.255.255.0
   member-interface GigabitEthernet0/0
   member-interface GigabitEthernet0/2

  interface redundant 2
   nameif inside
   security-level 100
   ip address 192.168.1.3 255.255.255.0
   member-interface GigabitEthernet0/1
   member-interface GigabitEthernet0/3
  ```

- EtherChannel interface—Enter the following command under each interface you want to add to the EtherChannel (either formerly in use or unused). You can assign up to 16 interfaces per EtherChannel, although only eight can be active; the others are in a standby state in case of failure.

channel-group *channel_id* **mode active**

For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside EtherChannel interfaces:

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 shutdown
 no nameif
 no security-level
 no ip address

...
```

Step 7 Enable each formerly unused interface that is now part of a logical interface by adding **no** in front of the **shutdown** command.

For example, your final EtherChannel configuration is:

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
```

```
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 no shutdown
 no nameif
 no security-level
 no ip address
!
interface port-channel 1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
!
interface port-channel 2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
!
interface port-channel 3
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
```

Note Other optional EtherChannel parameters can be configured after you import the new configuration. See Configuring an EtherChannel, page 10-19.

Step 8 At the ASA CLI prompt, perform the following steps depending on your connection (console or remote).

- Console connection:
  a. Copy the entire new configuration to the clipboard, including the altered interface section.
  b. Clear the running configuration by entering:

     ```
     ciscoasa(config)# clear configure all
     ```

     Traffic through the ASA stops at this point.
  c. Paste in the new configuration at the prompt.

     Traffic through the ASA resumes.
- Remote connection:
  a. Save the new configuration to a TFTP or FTP server, so that you can copy it to the startup configuration on the ASA. For example, you can run a TFTP or FTP server on your PC.
  b. Clear the startup configuration by entering:

     ```
     ciscoasa(config)# write erase
     ```
  c. Copy the new configuration to the startup configuration by entering:

     ```
     ciscoasa(config)# copy url startup-config
     ```

     See Copying a File to the ASA, page 44-17.
  d. Reload the ASA using the **reload** command. Do not save the running configuration.

Step 9 Reenable failover by entering the **failover** command.

Detailed Steps (Multiple Mode)

We recommend that you update your system and context configurations offline as text files, and reimport them for the following reasons:

- Because you cannot add an allocated interface as a member of a redundant or EtherChannel interface, you must deallocate the interface from any contexts. When you deallocate the interface, any context command that referred to that interface is deleted. Because commands that refer to interfaces are widespread throughout the configuration and affect multiple features, removing an allocation from an in-use interface at the CLI or in ASDM would cause significant damage to your configuration, not to mention significant downtime while you reconfigure all your features around a new interface.
- Changing your configuration offline lets you use the same interface names for your new logical interfaces, so that you do not need to touch the feature configurations that refer to interface names. You only need to change the interface configuration.
- Clearing the running system configuration and immediately applying a new configuration will minimize the downtime of your interfaces. You will not be waiting to configure the interfaces in real time.

Step 1 Connect to the ASA, and change to the system; if you are using failover, connect to the active ASA.

Step 2 If you are using failover, disable failover by entering the **no failover** command.

Step 3 In the system, copy the running configuration by entering the **more system:running-config** command and copying the display output to a text editor.

Be sure to save an extra copy of the old configuration in case you make an error when you edit it.

For example, you have the following interface configuration and allocation in the system configuration, with shared interfaces between two contexts.

System

```
interface GigabitEthernet0/0
    no shutdown
interface GigabitEthernet0/1
    no shutdown
interface GigabitEthernet0/2
    shutdown
interface GigabitEthernet0/3
    shutdown
interface GigabitEthernet0/4
    shutdown
interface GigabitEthernet0/5
    shutdown
interface Management0/0
    no shutdown
interface Management1/0
    shutdown
!
context customerA
    allocate-interface gigabitethernet0/0 int1
    allocate-interface gigabitethernet0/1 int2
    allocate-interface management0/0 mgmt
context customerB
    allocate-interface gigabitethernet0/0
    allocate-interface gigabitethernet0/1
    allocate-interface management0/0
```

Step 4 Get copies of *all* context configurations that will use the new EtherChannel or redundant interface. See Backing Up Configurations or Other Files, page 44-25.

For example, you download the following context configurations (interface configuration shown):

CustomerA Context

```
interface int1
 nameif outside
 security-level 0
 ip address 10.86.194.225 255.255.255.0
!
interface int2
 nameif inside
 security-level 100
 ip address 192.168.1.3 255.255.255.0
 no shutdown
!
interface mgmt
 nameif mgmt
 security-level 100
 ip address 10.1.1.5 255.255.255.0
 management-only
```

CustomerB Context

```
interface GigabitEthernet0/0
 nameif outside
 security-level 0
 ip address 10.20.15.5 255.255.255.0
!
interface GigabitEthernet0/1
```

```
 nameif inside
 security-level 100
 ip address 192.168.6.78 255.255.255.0
!
interface Management0/0
 nameif mgmt
 security-level 100
 ip address 10.8.1.8 255.255.255.0
 management-only
```

Step 5 In the system configuration, create the new logical interfaces according to the Configuring a Redundant Interface, page 10-17 or the Configuring an EtherChannel, page 10-19. Be sure to enter the **no shutdown** command on any additional physical interfaces you want to use as part of the logical interface.

Note You can only add *physical* interfaces to an EtherChannel or redundant interface; you cannot have VLANs configured for the physical interfaces.

Be sure to match physical interface parameters such as speed and duplex for all interfaces in a given EtherChannel or redundant interface. Note that the duplex setting for an EtherChannel interface must be Full or Auto.

For example, the new configuration is:

System

```
interface GigabitEthernet0/0
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/1
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/2
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/3
 channel-group 1 mode active
 no shutdown
!
interface GigabitEthernet0/4
 channel-group 2 mode active
 no shutdown
!
interface GigabitEthernet0/5
 channel-group 2 mode active
 no shutdown
!
interface Management0/0
 channel-group 3 mode active
 no shutdown
!
interface Management0/1
 channel-group 3 mode active
 no shutdown
!
interface port-channel 1
interface port-channel 2
interface port-channel 3
```

Step 6 Change the interface allocation per context to use the new EtherChannel or redundant interfaces. See Configuring a Security Context, page 7-19.

For example, to take advantage of existing cabling, you would continue to use the formerly in-use interfaces in their old roles as part of the inside and outside redundant interfaces:

```
context customerA
    allocate-interface port-channel1 int1
    allocate-interface port-channel2 int2
    allocate-interface port-channel3 mgmt
context customerB
    allocate-interface port-channel1
    allocate-interface port-channel2
    allocate-interface port-channel3
```

Note You might want to take this opportunity to assign mapped names to interfaces if you have not done so already. For example, the configuration for customerA does not need to be altered at all; it just needs to be reapplied on the ASA. The customerB configuration, however, needs to have all of the interface IDs changed; if you assign mapped names for customerB, you still have to change the interface IDs in the context configuration, but mapped names might help future interface changes.

Step 7 For contexts that do not use mapped names, change the context configuration to use the new EtherChannel or redundant interface ID. (Contexts that use mapped interface names do not require any alteration.)

For example:

CustomerB Context

```
interface port-channel1
 nameif outside
 security-level 0
 ip address 10.20.15.5 255.255.255.0
!
interface port-channel2
 nameif inside
 security-level 100
 ip address 192.168.6.78 255.255.255.0
!
interface port-channel3
 nameif mgmt
 security-level 100
 ip address 10.8.1.8 255.255.255.0
 management-only
```

Step 8 Copy the new context configuration files over the old ones. For example, if your contexts are on an FTP server, copy over the existing files (making backups as desired) using FTP. If your contexts are in flash memory, you can use the **copy** command and run a TFTP or FTP server on your PC, or use secure copy. See Copying a File to the ASA, page 44-17. This change only affects the startup configuration; the running configuration is still using the old context configuration.

Step 9 At the ASA system CLI prompt, perform the following steps depending on your connection (console or remote).

- Console connection:
  a. Copy the entire new system configuration to the clipboard, including the altered interface section.

b. Clear the running configuration (both system and contexts) by entering:

```
ciscoasa(config)# clear configure all
```

Traffic through the ASA stops at this point.

c. Paste in the new system configuration at the prompt.

All of the new context configurations now reload. When they are finished reloading, traffic through the ASA resumes.

- Remote connection:

a. Save the new system configuration to a TFTP or FTP server, so that you can copy it to the startup configuration on the ASA. For example, you can run a TFTP or FTP server on your PC.

b. Clear the startup configuration by entering:

```
ciscoasa(config)# write erase
```

c. Copy the new system configuration to the startup configuration by entering:

```
ciscoasa(config)# copy url startup-config
```

See Copying a File to the ASA, page 44-17.

d. Reload the ASA using the **reload** command. Do not save the running configuration.

Step 10 Reenable failover by entering the **failover** command.

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
show interface	Displays interface statistics.
show interface ip brief	Displays interface IP addresses and status.
show lacp {[*channel_group_number*] {**counters** | **internal** | **neighbor**} | **sys-id**}	For EtherChannel, displays LACP information such as traffic statistics, system identifier and neighbor details.
show port-channel [*channel_group_number*] [**brief** | **detail** | **port** | **protocol** | **summary**]	For EtherChannel, displays EtherChannel information in a detailed and one-line summary form. This command also displays the port and port-channel information.
show port-channel *channel_group_number* **load-balance** [**hash-result** {**ip** | **ipv6** | **l4port** | **mac** | **mixed** | **vlan-only**} *parameters*]	For EtherChannel, displays port-channel load-balance information along with the hash result and member interface selected for a given set of parameters.

Configuration Examples for ASA 5512-X and Higher Interfaces

This section includes the following topics:

Physical Interface Parameters Example

The following example configures parameters for the physical interface in single mode:

```
interface gigabitethernet 0/1
    speed 1000
    duplex full
    no shutdown
```

Subinterface Parameters Example

The following example configures parameters for a subinterface in single mode:

```
interface gigabitethernet 0/1.1
    vlan 101
    no shutdown
```

Multiple Context Mode Example

The following example configures interface parameters in multiple context mode for the system configuration, and allocates the gigabitethernet 0/1.1 subinterface to contextA:

```
interface gigabitethernet 0/1
    speed 1000
    duplex full
    no shutdown
interface gigabitethernet 0/1.1
    vlan 101
context contextA
    allocate-interface gigabitethernet 0/1.1
```

EtherChannel Example

The following example configures three interfaces as part of an EtherChannel. It also sets the system priority to be a higher priority, and GigabitEthernet 0/2 to be a higher priority than the other interfaces in case more than eight interfaces are assigned to the EtherChannel.

```
lacp system-priority 1234
interface GigabitEthernet0/0
    channel-group 1 mode active
interface GigabitEthernet0/1
```

```
    channel-group 1 mode active
interface GigabitEthernet0/2
    lacp port-priority 1234
    channel-group 1 mode passive
interface Port-channel1
    lacp max-bundle 4
    port-channel min-bundle 2
    port-channel load-balance dst-ip
```

Where to Go Next

- For multiple context mode:
  a. Assign interfaces to contexts and automatically assign unique MAC addresses to context interfaces. See Chapter 7, "Multiple Context Mode."
  b. Complete the interface configuration according to Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."
- For single context mode, complete the interface configuration according to Chapter 13, "Routed Mode Interfaces," or Chapter 14, "Transparent Mode Interfaces."

Feature History for ASA 5512-X and Higher Interfaces

Table 10-2 lists the release history for this feature.

Table 10-2 Feature History for Interfaces

Feature Name	Releases	Feature Information
Increased VLANs	7.0(5)	Increased the following limits: • ASA5510 Base license VLANs from 0 to 10. • ASA5510 Security Plus license VLANs from 10 to 25. • ASA5520 VLANs from 25 to 100. • ASA5540 VLANs from 100 to 200.
Increased interfaces for the Base license on the ASA 5510	7.2(2)	For the Base license on the ASA 5510, the maximum number of interfaces was increased from 3 plus a management interface to unlimited interfaces.
Increased VLANs	7.2(2)	VLAN limits were increased for the ASA 5510 (from 10 to 50 for the Base license, and from 25 to 100 for the Security Plus license), the ASA 5520 (from 100 to 150), the ASA 5550 (from 200 to 250).

Table 10-2 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	The ASA 5510 ASA now supports GE (Gigabit Ethernet) for port 0 and 1 with the Security Plus license. If you upgrade the license from Base to Security Plus, the capacity of the external Ethernet0/0 and Ethernet0/1 ports increases from the original FE (Fast Ethernet) (100 Mbps) to GE (1000 Mbps). The interface names will remain Ethernet 0/0 and Ethernet 0/1. Use the **speed** command to change the speed on the interface and use the **show interface** command to see what speed is currently configured for each interface.
Redundant interfaces	8.0(2)	A logical redundant interface pairs an active and a standby physical interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the ASA reliability. This feature is separate from device-level failover, but you can configure redundant interfaces as well as failover if desired. You can configure up to eight redundant interface pairs.
Jumbo packet support for the ASA 5580	8.1(1)	The Cisco ASA 5580 supports jumbo frames. A jumbo frame is an Ethernet packet larger than the standard maximum of 1518 bytes (including Layer 2 header and FCS), up to 9216 bytes. You can enable support for jumbo frames for all interfaces by increasing the amount of memory to process Ethernet frames. Assigning more memory for jumbo frames might limit the maximum use of other features, such as ACLs. This feature is also supported on the ASA 5585-X. We introduced the following command: **jumbo-frame reservation**.
Increased VLANs for the ASA 5580	8.1(2)	The number of VLANs supported on the ASA 5580 are increased from 100 to 250.
Support for Pause Frames for Flow Control on the ASA 5580 Ten Gigabit Ethernet Interfaces	8.2(2)	You can now enable pause (XOFF) frames for flow control. This feature is also supported on the ASA 5585-X. We introduced the following command: **flowcontrol**.
Support for Pause Frames for Flow Control on Gigabit Ethernet Interfaces	8.2(5)/8.4(2)	You can now enable pause (XOFF) frames for flow control for Gigabit Ethernet interfaces on all models. We modified the following command: **flowcontrol**.

Table 10-2 Feature History for Interfaces (continued)

Feature Name	Releases	Feature Information
EtherChannel support	8.4(1)	You can configure up to 48 802.3ad EtherChannels of eight active interfaces each. We introduced the following commands: **channel-group**, **lacp port-priority**, **interface port-channel**, **lacp max-bundle**, **port-channel min-bundle**, **port-channel load-balance**, **lacp system-priority**, **clear lacp counters**, **show lacp**, **show port-channel**. **Note** EtherChannel is not supported on the ASA 5505.
Support for 16 active links in an EtherChannel	9.2(1)	You can now configure up to 16 active links in an EtherChannel. Previously, you could have 8 active links and 8 standby links. Be sure that your switch can support 16 active links (for example the Cisco Nexus 7000 with F2-Series 10 Gigabit Ethernet Module). **Note** If you upgrade from an earlier ASA version, the maximum active interfaces is set to 8 for compatibility purposes (the **lacp max-bundle** command). We modified the following commands: **lacp max-bundle** and **port-channel min-bundle**.