



CHAPTER 9

Completing Interface Configuration (Transparent Mode)

This chapter includes tasks to complete the interface configuration for all models in transparent firewall mode.

This chapter includes the following sections:

- [Information About Completing Interface Configuration in Transparent Mode, page 9-1](#)
- [Licensing Requirements for Completing Interface Configuration in Transparent Mode, page 9-2](#)
- [Guidelines and Limitations, page 9-5](#)
- [Default Settings, page 9-6](#)
- [Completing Interface Configuration in Transparent Mode, page 9-6](#)
- [Monitoring Interfaces, page 9-19](#)
- [Configuration Examples for Interfaces in Transparent Mode, page 9-19](#)
- [Feature History for Interfaces in Transparent Mode, page 9-20](#)



Note

For multiple context mode, complete the tasks in this section in the context execution space. Enter the `changeto context name` command to change to the context you want to configure.

Information About Completing Interface Configuration in Transparent Mode

This section includes the following topics:

- [Bridge Groups in Transparent Mode, page 9-1](#)
- [Security Levels, page 9-2](#)

Bridge Groups in Transparent Mode

If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the ASA, and traffic must exit the ASA before it is routed by an external router back

to another bridge group in the ASA. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration. For complete security policy separation, use security contexts with one bridge group in each context. At least one bridge group is required per context or in single mode.

Each bridge group requires a management IP address. For another method of management, see the [“Management Interface”](#) section.

**Note**

The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported.

Security Levels

Each interface must have a security level from 0 (lowest) to 100 (highest). For example, you should assign your most secure network, such as the inside host network, to level 100. While the outside network connected to the Internet can be level 0. Other networks, such as DMZs can be in between. You can assign interfaces to the same security level. See the [“Allowing Same Security Level Communication”](#) section on page 9-18 for more information.

The level controls the following behavior:

- Network access—By default, there is an implicit permit from a higher security interface to a lower security interface (outbound). Hosts on the higher security interface can access any host on a lower security interface. You can limit access by applying an access list to the interface.

If you enable communication for same security interfaces (see the [“Allowing Same Security Level Communication”](#) section on page 9-18), there is an implicit permit for interfaces to access other interfaces on the same security level or lower.

- Inspection engines—Some application inspection engines are dependent on the security level. For same security interfaces, inspection engines apply to traffic in either direction.
 - NetBIOS inspection engine—Applied only for outbound connections.
 - SQL*Net inspection engine—If a control connection for the SQL*Net (formerly OraServ) port exists between a pair of hosts, then only an inbound data connection is permitted through the ASA.

- Filtering—HTTP(S) and FTP filtering applies only for outbound connections (from a higher level to a lower level).

If you enable communication for same security interfaces, you can filter traffic in either direction.

- **established** command—This command allows return connections from a lower security host to a higher security host if there is already an established connection from the higher level host to the lower level host.

If you enable communication for same security interfaces, you can configure **established** commands for both directions.

Licensing Requirements for Completing Interface Configuration in Transparent Mode

Model	License Requirement
ASA 5505	VLANs: Base License: 3 (2 regular zones and 1 restricted zone that can only communicate with 1 other zone) Security Plus License: 20 VLAN Trunks: Base License: None. Security Plus License: 8. Interfaces of all types ¹ : Base License: 52. Security Plus License: 120.

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, and bridge group interfaces.

Model	License Requirement
ASA 5510	VLANs: Base License: 50 Security Plus License: 100 Interface Speed: Base License—All interfaces Fast Ethernet. Security Plus License—Ethernet 0/0 and 0/1: Gigabit Ethernet; all others Fast Ethernet. Interfaces of all types ¹ : Base License: 52 Security Plus License: 120
ASA 5520	VLANs: Base License: 150. Interfaces of all types ¹ : Base License: 640
ASA 5540	VLANs: Base License: 200 Interfaces of all types ¹ : Base License: 840
ASA 5550	VLANs: Base License: 400 Interfaces of all types ¹ : Base License: 1640

Model	License Requirement
ASA 5580	VLANs: Base License: 1024 Interfaces of all types ¹ : Base License: 4176
ASA 5512-X	VLANs: Base License: 50 Interfaces of all types ¹ : Base License: 328
ASA 5515-X	VLANs: Base License: 100 Interfaces of all types ¹ : Base License: 528
ASA 5525-X	VLANs: Base License: 200 Interfaces of all types ¹ : Base License: 928
ASA 5545-X	VLANs: Base License: 300 Interfaces of all types ¹ : Base License: 1328
ASA 5555-X	VLANs: Base License: 500 Interfaces of all types ¹ : Base License: 2128
ASA 5585-X	VLANs: Base License: 1024 Interface Speed for SSP-10 and SSP-20: Base License—1-Gigabit Ethernet for fiber interfaces 10 GE I/O License—10-Gigabit Ethernet for fiber interfaces (SSP-40 and SSP-60 support 10-Gigabit Ethernet by default.) Interfaces of all types ¹ : Base License: 4176

1. The maximum number of combined interfaces; for example, VLANs, physical, redundant, bridge group, and EtherChannel interfaces.

Guidelines and Limitations

This section includes the guidelines and limitations for this feature.

Context Mode Guidelines

- For the ASA 5510 and higher in multiple context mode, configure the physical interfaces in the system execution space according to [Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\)”](#). Then, configure the logical interface parameters in the context execution space according to this chapter.

The ASA 5505 does not support multiple context mode.

- You can only configure context interfaces that you already assigned to the context in the system configuration using the **allocate-interface** command.

Firewall Mode Guidelines

- You can configure up to 8 bridge groups in single mode or per context in multiple mode. Note that you must use at least 1 bridge group; data interfaces must belong to a bridge group.



Note Although you can configure multiple bridge groups on the ASA 5505, the restriction of 2 data interfaces in transparent mode on the ASA 5505 means you can only effectively use 1 bridge group.

- Each bridge group can include up to 4 interfaces.
- For IPv4, a management IP address is required for each bridge group for both management traffic and for traffic to pass through the ASA.

Unlike routed mode, which requires an IP address for each interface, a transparent firewall has an IP address assigned to the entire bridge group. The ASA uses this IP address as the source address for packets originating on the ASA, such as system messages or AAA communications. In addition to the bridge group management address, you can optionally configure a management interface for some models; see the [“Management Interface” section on page 6-2](#) for more information.

The management IP address must be on the same subnet as the connected network. You cannot set the subnet to a host subnet (255.255.255.255). The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported. See the [“Configuring Bridge Groups” section on page 9-7](#) for more information about management IP subnets.

- For IPv6, at a minimum you need to configure link-local addresses for each interface for through traffic. For full functionality, including the ability to manage the ASA, you need to configure a global IPv6 address for each bridge group.
- For multiple context mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode, each context typically uses a different subnet. You can use overlapping subnets, but your network topology requires router and NAT configuration to make it possible from a routing standpoint.

Failover Guidelines

Do not finish configuring failover interfaces with the procedures in this chapter. See the “[Configuring Active/Standby Failover](#)” section on page 62-7 or the “[Configuring Active/Active Failover](#)” section on page 63-8 to configure the failover and state links. In multiple context mode, failover interfaces are configured in the system configuration.

IPv6 Guidelines

- Supports IPv6.
- No support for IPv6 anycast addresses in transparent mode.

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 the “[Factory Default Configurations](#)” section on page 2-10.

Default Security Level

The default security level is 0. If you name an interface “inside” and you do not set the security level explicitly, then the ASA sets the security level to 100.

**Note**

If you change the security level of an interface, and you do not want to wait for existing connections to time out before the new security information is used, you can clear the connections using the **clear local-host** command.

Completing Interface Configuration in Transparent Mode

This section includes the following topics:

- [Task Flow for Completing Interface Configuration](#), page 9-6
- [Configuring Bridge Groups](#), page 9-7
- [Configuring General Interface Parameters](#), page 9-8
- [Configuring a Management Interface \(ASA 5510 and Higher\)](#), page 9-11
- [Configuring the MAC Address and MTU](#), page 9-12
- [Configuring IPv6 Addressing](#), page 9-15
- [Allowing Same Security Level Communication](#), page 9-18

Task Flow for Completing Interface Configuration

-
- Step 1** Set up your interfaces depending on your model:
- ASA 5510 and higher—[Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 7, “Starting Interface Configuration \(ASA 5505\).”](#)

- Step 2** (Multiple context mode) Allocate interfaces to the context according to the “[Configuring Multiple Contexts](#)” section on page 5-14.
 - Step 3** (Multiple context mode) Enter the **changeto context** *name* command to change to the context you want to configure. Configure one or more bridge groups, including the IPv4 address. See the “[Configuring Bridge Groups](#)” section on page 9-7.
 - Step 4** Configure general interface parameters, including the interface name and security level. See the “[Configuring General Interface Parameters](#)” section on page 9-8.
 - Step 5** (Optional; not supported for the ASA 5505) Configure a management interface. See the “[Configuring a Management Interface \(ASA 5510 and Higher\)](#)” section on page 9-11.
 - Step 6** (Optional) Configure the MAC address and the MTU. See the “[Configuring the MAC Address and MTU](#)” section on page 9-12.
 - Step 7** (Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 9-15.
 - Step 8** (Optional) Allow same security level communication, either by allowing communication between two interfaces or by allowing traffic to enter and exit the same interface. See the “[Allowing Same Security Level Communication](#)” section on page 9-18.
-

Configuring Bridge Groups

Each bridge group requires a management IP address. The ASA uses this IP address as the source address for packets originating from the bridge group. The management IP address must be on the same subnet as the connected network. For IPv4 traffic, the management IP address is required to pass any traffic. For IPv6 traffic, you must, at a minimum, configure the link-local addresses to pass traffic, but a global management address is recommended for full functionality, including remote management and other management operations.

Guidelines and Limitations

You can configure up to 8 bridge groups in single mode or per context in multiple mode. Note that you must use at least one bridge group; data interfaces must belong to a bridge group.

**Note**

For a separate management interface (for supported models), a non-configurable bridge group (ID 101) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

Detailed Steps

	Command	Purpose
Step 1	<code>interface bvi <i>bridge_group_number</i></code> Example: hostname(config)# interface bvi 1	Creates a bridge group, where <i>bridge_group_number</i> is an integer between 1 and 100.
Step 2	<code>ip address <i>ip_address</i> [<i>mask</i>] [<i>standby ip_address</i>]</code> Example: hostname(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2	Specifies the management IP address for the bridge group. Do not assign a host address (/32 or 255.255.255.255) to the bridge group. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and transparent firewall) such as a /30 subnet (255.255.255.252). The ASA drops all ARP packets to or from the first and last addresses in a subnet. Therefore, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the ASA drops the ARP request from the downstream router to the upstream router. The ASA does not support traffic on secondary networks; only traffic on the same network as the management IP address is supported. The standby keyword and address is used for failover.

Examples

The following example sets the management address and standby address of bridge group 1:

```
hostname(config)# interface bvi 1
hostname(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2
```

What to Do Next

Configure general interface parameters. See the [“Configuring General Interface Parameters”](#) section on page 9-8.

Configuring General Interface Parameters

This procedure describes how to set the name, security level, and bridge group for each transparent interface.

To configure a separate management interface, see the [“Configuring a Management Interface \(ASA 5510 and Higher\)”](#) section on page 9-11.

For the ASA 5510 and higher, you must configure interface parameters for the following interface types:

- Physical interfaces
- VLAN subinterfaces
- Redundant interfaces
- EtherChannel interfaces

For the ASA 5505, you must configure interface parameters for the following interface types:

- VLAN interfaces

Guidelines and Limitations

- You can configure up to four interfaces per bridge group.
- For the ASA 5550, for maximum throughput, be sure to balance your traffic over the two interface slots; for example, assign the inside interface to slot 1 and the outside interface to slot 0.
- For information about security levels, see the [“Security Levels” section on page 9-2](#).
- If you are using failover, do not use this procedure to name interfaces that you are reserving for failover and Stateful Failover communications. See the [“Configuring Active/Standby Failover” section on page 62-7](#) or the [“Configuring Active/Active Failover” section on page 63-8](#) to configure the failover and state links.

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—[Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 7, “Starting Interface Configuration \(ASA 5505\).”](#)
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the [“Configuring Multiple Contexts” section on page 5-14](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505:</p> <pre>hostname(config)# interface vlan number</pre> <p>Example:</p> <pre>hostname(config)# interface vlan 100</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID. Do not use this procedure for Management interfaces; see the “Configuring a Management Interface (ASA 5510 and Higher)” section on page 9-11 to configure the Management interface.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>bridge-group number</pre> <p>Example:</p> <pre>hostname(config-if)# bridge-group 1</pre>	<p>Assigns the interface to a bridge group, where <i>number</i> is an integer between 1 and 100. You can assign up to four interfaces to a bridge group. You cannot assign the same interface to more than one bridge group.</p>
Step 3	<pre>nameif name</pre> <p>Example:</p> <pre>hostname(config-if)# nameif inside</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value. Do not enter the no form, because that command causes all commands that refer to that name to be deleted.</p>
Step 4	<pre>security-level number</pre> <p>Example:</p> <pre>hostname(config-if)# security-level 50</pre>	<p>Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest).</p>

What to Do Next

- (Optional) Configure a management interface. See the “[Configuring a Management Interface \(ASA 5510 and Higher\)](#)” section on page 9-11.
- (Optional) Configure the MAC address and the MTU. See the “[Configuring the MAC Address and MTU](#)” section on page 9-12.
- (Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 9-15.

Configuring a Management Interface (ASA 5510 and Higher)

You can configure one management interface separate from the bridge group interfaces in single mode or per context. For more information, see the [“Management Interface” section on page 6-2](#).

Restrictions

- See the [“Management Interface” section on page 6-2](#).
- Do not assign this interface to a bridge group; a non-configurable bridge group (ID 101) is automatically added to your configuration. This bridge group is not included in the bridge group limit.
- If your model does not include a Management interface, you must manage the transparent firewall from a data interface; skip this procedure. (For example, on the ASA 5505.)
- 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.

Prerequisites

- Complete the procedures in [Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the [“Configuring Multiple Contexts” section on page 5-14](#).
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<pre>interface {{port-channel number management slot/port}[.subinterface] mapped_name} Example: hostname(config)# interface management 0/0.1</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode for the management interface.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1. The EtherChannel interface must have only Management member interfaces.</p> <p>Redundant interfaces do not support Management <i>slot/port</i> interfaces as members. You also cannot set a redundant interface comprised of non-Management interfaces as management-only.</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>nameif name Example: hostname(config-if)# nameif management</pre>	<p>Names the interface.</p> <p>The <i>name</i> is a text string up to 48 characters, and is not case-sensitive. You can change the name by reentering this command with a new value. Do not enter the no form, because that command causes all commands that refer to that name to be deleted.</p>

	Command	Purpose
Step 3	Do one of the following:	
	<p>ip address <i>ip_address</i> [<i>mask</i>] [standby <i>ip_address</i>]</p> <p>Example: <pre>hostname(config-if)# ip address 10.1.1.1 255.255.255.0 standby 10.1.1.2</pre></p>	<p>Sets the IP address manually.</p> <p>Note For use with failover, you must set the IP address and standby address manually; DHCP is not supported.</p> <p>The <i>ip_address</i> and <i>mask</i> arguments set the interface IP address and subnet mask.</p> <p>The standby <i>ip_address</i> argument is used for failover. See the “Configuring Active/Standby Failover” section on page 62-7 or the “Configuring Active/Active Failover” section on page 63-8 for more information.</p>
	<p>ip address dhcp [setroute]</p> <p>Example: <pre>hostname(config-if)# ip address dhcp</pre></p>	<p>Obtains an IP address from a DHCP server.</p> <p>The setroute keyword lets the ASA use the default route supplied by the DHCP server.</p> <p>Reenter this command to reset the DHCP lease and request a new lease.</p> <p>If you do not enable the interface using the no shutdown command before you enter the ip address dhcp command, some DHCP requests might not be sent.</p>
Step 4	<p>security-level <i>number</i></p> <p>Example: <pre>hostname(config-if)# security-level 50</pre></p>	<p>Sets the security level, where <i>number</i> is an integer between 0 (lowest) and 100 (highest).</p>

What to Do Next

- (Optional) Configure the MAC address and the MTU. See the “[Configuring the MAC Address and MTU](#)” section on page 9-12.
- (Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 9-15.

Configuring the MAC Address and MTU

This section describes how to configure MAC addresses for interfaces and how to set the MTU.

Information About 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.

A 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. If you assign a MAC address to the redundant interface using this command, then it is used regardless of the member interface MAC addresses.

For an EtherChannel, 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.

In multiple context mode, if you share an interface between contexts, you can assign a unique MAC address to the interface in each context. This feature lets the ASA easily classify packets into the appropriate context. Using a shared interface without unique MAC addresses is possible, but has some limitations. See the [“How the ASA Classifies Packets”](#) section on page 5-3 for more information. You can assign each MAC address manually, or you can automatically generate MAC addresses for shared interfaces in contexts. See the [“Automatically Assigning MAC Addresses to Context Interfaces”](#) section on page 5-22 to automatically generate MAC addresses. If you automatically generate MAC addresses, you can use this procedure to override the generated address.

For single context mode, or for interfaces that are not shared in multiple context mode, you might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address.

Information About the MTU

The MTU is the maximum datagram size that is sent on a connection. Data that is larger than the MTU value is fragmented before being sent.

The ASA supports IP path MTU discovery (as defined in RFC 1191), which allows a host to dynamically discover and cope with the differences in the maximum allowable MTU size of the various links along the path. Sometimes, the ASA cannot forward a datagram because the packet is larger than the MTU that you set for the interface, but the “don't fragment” (DF) bit is set. The network software sends a message to the sending host, alerting it to the problem. The host has to fragment packets for the destination so that they fit the smallest packet size of all the links along the path.

The default MTU is 1500 bytes in a block for Ethernet interfaces. This value is sufficient for most applications, but you can pick a lower number if network conditions require it.

To enable jumbo frames, see the [“Enabling Jumbo Frame Support \(Supported Models\)”](#) section on page 6-32. 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. Jumbo frames require extra memory to process, and assigning more memory for jumbo frames might limit the maximum use of other features, such as access lists. To use jumbo frames, set the value higher, for example, to 9000 bytes.

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—[Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 7, “Starting Interface Configuration \(ASA 5505\).”](#)
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the [“Configuring Multiple Contexts”](#) section on page 5-14.

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	<p>For the ASA 5510 and higher:</p> <pre>interface {{redundant number port-channel number physical_interface} [.subinterface] mapped_name}</pre> <p>For the ASA 5505:</p> <pre>hostname(config)# interface vlan number</pre> <p>Example:</p> <pre>hostname(config)# interface vlan 100</pre>	<p>If you are not already in interface configuration mode, enters interface configuration mode.</p> <p>The redundant number argument is the redundant interface ID, such as redundant 1.</p> <p>The port-channel number argument is the EtherChannel interface ID, such as port-channel 1.</p> <p>See the “Enabling the Physical Interface and Configuring Ethernet Parameters” section for a description of the physical interface ID.</p> <p>Append the <i>subinterface</i> ID to the physical or redundant interface ID separated by a period (.).</p> <p>In multiple context mode, enter the <i>mapped_name</i> if one was assigned using the allocate-interface command.</p>
Step 2	<pre>mac-address mac_address [standby mac_address]</pre> <p>Example:</p> <pre>hostname(config-if)# mac-address 000C.F142.4CDE</pre>	<p>Assigns a private MAC address to this interface. The <i>mac_address</i> is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE.</p> <p>The first two bytes of a manual MAC address cannot be A2 if you also want to use auto-generated MAC addresses.</p> <p>For use with failover, set the standby MAC address. If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.</p>
Step 3	<pre>mtu interface_name bytes</pre> <p>Example:</p> <pre>hostname(config)# mtu inside 9200</pre>	<p>Sets the MTU between 300 and 65,535 bytes. The default is 1500 bytes.</p> <p>Note When you set the MTU for a redundant or port-channel interface, the ASA applies the setting to all member interfaces.</p> <p>For models that support jumbo frames, if you enter a value for any interface that is greater than 1500, then you need to enable jumbo frame support. See the “Enabling Jumbo Frame Support (Supported Models)” section on page 6-32.</p>

What to Do Next

(Optional) Configure IPv6 addressing. See the “[Configuring IPv6 Addressing](#)” section on page 9-15.

Configuring IPv6 Addressing

This section describes how to configure IPv6 addressing. For more information about IPv6, see the “[Information About IPv6 Support](#)” section on page 21-9 and the “[IPv6 Addresses](#)” section on page B-5.

This section includes the following topics:

- [Information About IPv6, page 9-15](#)
- [Configuring a Global IPv6 Address and Other Options, page 9-17](#)

Information About IPv6

This section includes information about how to configure IPv6, and includes the following topics:

- [IPv6 Addressing, page 9-15](#)
- [Duplicate Address Detection, page 9-15](#)
- [Modified EUI-64 Interface IDs, page 9-16](#)
- [Unsupported Commands, page 9-16](#)

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- **Global**—The global address is a public address that you can use on the public network. This address needs to be configured for each bridge group, and not per-interface. You can also configure a global IPv6 address for the management interface.
- **Link-local**—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the ND functions such as address resolution and neighbor discovery. Because the link-local address is only available on a segment, and is tied to the interface MAC address, you need to configure the link-local address per interface.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on each interface, so you do not also need to specifically configure a link-local address. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

**Note**

If you want to only configure the link-local addresses, see the **ipv6 enable** (to auto-configure) or **ipv6 address link-local** (to manually configure) command in the command reference.

Duplicate Address Detection

During the stateless autoconfiguration process, duplicate address detection (DAD) verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces (the new addresses remain in a tentative state while duplicate address detection is performed). Duplicate address detection is performed first on the new link-local address. When the link local address is verified as unique, then duplicate address detection is performed all the other IPv6 unicast addresses on the interface.

Duplicate address detection is suspended on interfaces that are administratively down. While an interface is administratively down, the unicast IPv6 addresses assigned to the interface are set to a pending state. An interface returning to an administratively up state restarts duplicate address detection for all of the unicast IPv6 addresses on the interface.

When a duplicate address is identified, the state of the address is set to DUPLICATE, the address is not used, and the following error message is generated:

```
%ASA-4-325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used. However, all configuration commands associated with the duplicate address remain as configured while the state of the address is set to DUPLICATE.

If the link-local address for an interface changes, duplicate address detection is performed on the new link-local address and all of the other IPv6 address associated with the interface are regenerated (duplicate address detection is performed only on the new link-local address).

The ASA uses neighbor solicitation messages to perform duplicate address detection. By default, the number of times an interface performs duplicate address detection is 1.

Modified EUI-64 Interface IDs

RFC 3513: Internet Protocol Version 6 (IPv6) Addressing Architecture requires that the interface identifier portion of all unicast IPv6 addresses, except those that start with binary value 000, be 64 bits long and be constructed in Modified EUI-64 format. The ASA can enforce this requirement for hosts attached to the local link.

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64 format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
%ASA-3-325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link. Packets received from hosts behind a router will fail the address format verification, and be dropped, because their source MAC address will be the router MAC address and not the host MAC address.

Unsupported Commands

The following IPv6 commands are not supported in transparent firewall mode, because they require router capabilities:

- **ipv6 address autoconfig**
- **ipv6 nd prefix**
- **ipv6 nd ra-interval**
- **ipv6 nd ra-lifetime**
- **ipv6 nd suppress-ra**

The **ipv6 local pool VPN** command is not supported, because transparent mode does not support VPN.

Configuring a Global IPv6 Address and Other Options

To configure a global IPv6 address and other options for a bridge group or management interface, perform the following steps.



Note

Configuring the global address automatically configures the link-local address, so you do not need to configure it separately.

Restrictions

The ASA does not support IPv6 anycast addresses.

Prerequisites

- Set up your interfaces depending on your model:
 - ASA 5510 and higher—[Chapter 6, “Starting Interface Configuration \(ASA 5510 and Higher\).”](#)
 - ASA 5505—[Chapter 7, “Starting Interface Configuration \(ASA 5505\).”](#)
- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to the [“Configuring Multiple Contexts” section on page 5-14.](#)
- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, enter the **changeto context name** command.

Detailed Steps

	Command	Purpose
Step 1	For the bridge group: <pre>interface bvi bridge_group_id</pre> For the management interface: <pre>interface management_interface_id</pre> Example: <pre>hostname(config)# interface bvi 1</pre>	If you are not already in interface configuration mode, enters interface configuration mode.
Step 2	<pre>ipv6 address ipv6-address/prefix-length [standby ipv6-address]</pre> Example: <pre>hostname(config-if)# ipv6 address 2001:0DB8::BA98:0:3210/48</pre>	Assigns a global address to the interface. When you assign a global address, the link-local address is automatically created for the interface (for a bridge group, for each member interface). standby specifies the interface address used by the secondary unit or failover group in a failover pair. Note The eui-64 keyword to use the Modified EUI-64 interface ID for the interface ID is not supported in transparent mode. See the “IPv6 Addresses” section on page B-5 for more information about IPv6 addressing.

	Command	Purpose
Step 3	(Optional) <code>ipv6 nd suppress-ra</code> Example: hostname(config-if)# ipv6 nd suppress-ra	Suppresses Router Advertisement messages on an interface. By default, Router Advertisement messages are automatically sent in response to router solicitation messages. You may want to disable these messages on any interface for which you do not want the ASA to supply the IPv6 prefix (for example, the outside interface).
Step 4	(Optional) <code>ipv6 nd dad attempts value</code> Example: hostname(config-if)# ipv6 nd dad attempts 3	Changes the number of duplicate address detection attempts. The <i>value</i> argument can be any value from 0 to 600. Setting the <i>value</i> argument to 0 disables duplicate address detection on the interface. By default, the number of times an interface performs duplicate address detection is 1. See the “Duplicate Address Detection” section on page 9-15 for more information.
Step 5	(Optional) <code>ipv6 nd ns-interval value</code> Example: hostname(config-if)# ipv6 nd ns-interval 2000	Changes the neighbor solicitation message interval. When you configure an interface to send out more than one duplicate address detection attempt with the <code>ipv6 nd dad attempts</code> command, this command configures the interval at which the neighbor solicitation messages are sent out. By default, they are sent out once every 1000 milliseconds. The <i>value</i> argument can be from 1000 to 3600000 milliseconds. Note Changing this value changes it for all neighbor solicitation messages sent out on the interface, not just those used for duplicate address detection.
Step 6	(Optional) <code>ipv6 enforce-eui64 if_name</code> Example: hostname(config)# ipv6 enforce-eui64 inside	Enforces the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link. The <i>if_name</i> argument is the name of the interface, as specified by the <code>nameif</code> command, on which you are enabling the address format enforcement. See the “Modified EUI-64 Interface IDs” section on page 9-16 for more information.

Allowing Same Security Level Communication

By default, interfaces on the same security level cannot communicate with each other, and packets cannot enter and exit the same interface. This section describes how to enable inter-interface communication when interfaces are on the same security level.

Information About Inter-Interface Communication

Allowing interfaces on the same security level to communicate with each other is useful if you want traffic to flow freely between all same security interfaces without access lists.

If you enable same security interface communication, you can still configure interfaces at different security levels as usual.

Detailed Steps

Command	Purpose
<code>same-security-traffic permit inter-interface</code>	Enables interfaces on the same security level so that they can communicate with each other.

Monitoring Interfaces

To monitor interfaces, enter one of the following commands:

Command	Purpose
<code>show interface</code>	Displays interface statistics.
<code>show interface ip brief</code>	Displays interface IP addresses and status.
<code>show bridge-group</code>	Shows bridge group information.

Configuration Examples for Interfaces in Transparent Mode

The following example includes two bridge groups of three interfaces each, plus a management-only interface:

```
interface gigabitethernet 0/0
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 0/1
  nameif outside1
  security-level 0
  bridge-group 1
  no shutdown
interface gigabitethernet 0/2
  nameif dmz1
  security-level 50
  bridge-group 1
  no shutdown
interface bvi 1
  ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2

interface gigabitethernet 1/0
  nameif inside2
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/1
  nameif outside2
  security-level 0
  bridge-group 2
  no shutdown
interface gigabitethernet 1/2
  nameif dmz2
  security-level 50
  bridge-group 2
```

```

no shutdown
interface bvi 2
  ip address 10.3.5.8 255.255.255.0 standby 10.3.5.9

interface management 0/0
  nameif mgmt
  security-level 100
  ip address 10.2.1.1 255.255.255.0 standby 10.2.1.2
no shutdown

```

Feature History for Interfaces in Transparent Mode

Table 9-1 lists each feature change and the platform release in which it was implemented.

Table 9-1 Feature History for Interfaces in Transparent Mode

Feature Name	Platform Releases	Feature Information
Increased VLANs	7.0(5)	<p>Increased the following limits:</p> <ul style="list-style-type: none"> • 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 VLANs	7.2(2)	<p>The maximum number of VLANs for the Security Plus license on the ASA 5505 was increased from 5 (3 fully functional; 1 failover; one restricted to a backup interface) to 20 fully functional interfaces. In addition, the number of trunk ports was increased from 1 to 8. Now there are 20 fully functional interfaces, you do not need to use the backup interface command to cripple a backup ISP interface; you can use a fully-functional interface for it. The backup interface command is still useful for an Easy VPN configuration.</p> <p>VLAN limits were also 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).</p>
Gigabit Ethernet Support for the ASA 5510 Security Plus License	7.2(3)	<p>The ASA 5510 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.</p>

Table 9-1 Feature History for Interfaces in Transparent Mode (continued)

Feature Name	Platform Releases	Feature Information
Native VLAN support for the ASA 5505	7.2(4)/8.0(4)	You can now include the native VLAN in an ASA 5505 trunk port. We introduced the following command: switchport trunk native vlan .
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 access lists. 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.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
Support for Pause Frames for Flow Control on the ASA 5580 10-Gigabit Ethernet Interfaces	8.2(2)	You can now enable pause (XOFF) frames for flow control. We introduced the following command: flowcontrol .
Bridge groups for transparent mode	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We introduced the following commands: interface bvi , show bridge-group .

