# Routed and Transparent Mode Interfaces

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

> **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.

## About Routed and Transparent Mode Interfaces

The ASA supports two types of interfaces: routed and bridged.

Each Layer 3 routed interface requires an IP address on a unique subnet.

Bridged interfaces belong to a bridge group, and all interfaces are on the same network. The bridge group is represented by a Bridge Virtual Interface (BVI) that has an IP address on the bridge network. Routed mode only supports routed interfaces. Transparent firewall mode only supports bridge group and BVI interfaces.

### Security Levels

Each interface must have a security level from 0 (lowest) to 100 (highest), including bridge group member interfaces. 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.

In transparent mode, the BVI interface does not have a security level because it does not participate in routing between interfaces.

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 ACL to the interface.

  If you enable communication for same-security interfaces, 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.

## Dual IP Stack (IPv4 and IPv6)

The ASA supports both IPv6 and IPv4 addresses on an interface. Make sure you configure a default route for both IPv4 and IPv6.

# Guidelines and Limitations for Routed and Transparent Mode Interfaces

### Context Mode

- In multiple context mode, you can only configure context interfaces that you already assigned to the context in the system configuration according to Configure Multiple Contexts.
- PPPoE is not supported in multiple context mode.
- For multiple context mode in transparent mode, each context must use different interfaces; you cannot share an interface across contexts.
- For multiple context mode in transparent 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

- Do not configure failover links with the procedures in this chapter. See the Failover chapter for more information.
- When you use Failover, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported.

### IPv6

- IPv6 is supported on all interfaces.

- You can only configure IPv6 addresses manually in transparent mode.
- The ASA does not support IPv6 anycast addresses.

### VLAN IDs for the ASASM

You can add any VLAN ID to the configuration, but only VLANs that are assigned to the ASA by the switch can pass traffic. To view all VLANs assigned to the ASA, use the **show vlan** command.

If you add an interface for a VLAN that is not yet assigned to the ASA by the switch, the interface will be in the down state. When you assign the VLAN to the ASA, the interface changes to an up state. See the **show interface** command for more information about interface states.

### Transparent Mode and Bridge Group Guidelines

- You can create up to 250 bridge groups, with 4 interfaces per bridge group.
- Each directly-connected network must be on the same subnet.
- The ASA does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.
- An IP address for the BVI is required for each bridge group for to-the-device and from-the-device management traffic, as well as for data traffic to pass through the ASA. For IPv4 traffic, specify an IPv4 address. For IPv6 traffic, specify an IPv6 address.
- You can only configure IPv6 addresses manually.
- The BVI 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).
- Management interfaces are not supported as bridge group members.
- In transparent mode, you must use at least 1 bridge group; data interfaces must belong to a bridge group.
- In transparent mode, do not specify the BVI IP address as the default gateway for connected devices; devices need to specify the router on the other side of the ASA as the default gateway.
- In transparent mode, the *default* route, which is required to provide a return path for management traffic, is only applied to management traffic from one bridge group network. This is because the default route specifies an interface in the bridge group as well as the router IP address on the bridge group network, and you can only define one default route. If you have management traffic from more than one bridge group network, you need to specify a regular static route that identifies the network from which you expect management traffic.
- In transparent mode, PPPoE is not supported for the Management interface.
- Bidirectional Forwarding Detection (BFD) echo packets are not allowed through the ASA when using bridge group members. If there are two neighbors on either side of the ASA running BFD, then the ASA will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

### 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 conn** command.

### Additional Guidelines and Requirements

- The ASA supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support).

# Configure Routed Mode Interfaces

To configure routed mode interfaces, perform the following steps.

## Configure General Routed Mode Interface Parameters

This procedure describes how to set the name, security level, IPv4 address, and other options.

### Before you begin

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.

### Procedure

**Step 1** Enter interface configuration mode:

**interface** *id*

**Example:**

```
ciscoasa(config)# interface gigabithethernet 0/0
```

The interface ID can be:

- **redundant**
- **port-channel**
- *physical*—For example, **ethernet**, **gigabitethernet**, **tengigabitethernet**, **management**. Refer to the hardware installation guide for your model for interface names.
- *physical***.***subinterface*—For example, **gigabitethernet0/0.100**.
- **vni**
- **vlan**
- *mapped_name*—For multiple context mode.

**Step 2** Name the interface:

**nameif** *name*

**Example:**

```
ciscoasa(config-if)# nameif inside
```

The *name* 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.

**Step 3** Set the IP address using one of the following methods.

For failover and clustering, you must set the IP address manually; DHCP and PPPoE are not supported.

- Set the IP address manually:

  **ip address** *ip_address* [*mask*] [**standby** *ip_address*]

  Example:

  ```
  ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0 standby 10.1.1.2
  ```

  The standby *ip_address* argument is used for failover. If you do not set the standby IP address, the active unit cannot monitor the standby interface using network tests; it can only track the link state.

  The *ip_address* and *mask* arguments set the interface IP address and subnet mask.

- Obtain an IP address from a DHCP server:

  **ip address dhcp** [**setroute**]

  Example:

  ```
  ciscoasa(config-if)# ip address dhcp
  ```

  The **setroute** keyword lets the ASA use the default route supplied by the DHCP server.

  Reenter this command to reset the DHCP lease and request a new lease.

  **Note** 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.

- Obtain an IP address from a PPPoE server:

  **ip address pppoe** [**setroute**]

  Example:

  ```
  ciscoasa(config-if)# ip address pppoe setroute
  ```

  You can alternatively enable PPPoE by manually entering the IP address:

  **ip address** *ip_address mask* **pppoe**

  Example:

```
ciscoasa(config-if)# ip address 10.1.1.78 255.255.255.0 pppoe
```

The **setroute** option sets the default routes when the PPPoE client has not yet established a connection. When using the **setroute** option, you cannot have a statically defined route in the configuration.

**Note** If PPPoE is enabled on two interfaces (such as a primary and backup interface), and you do not configure dual ISP support, then the ASA can only send traffic through the first interface to acquire an IP address.

**Step 4** Set the security level:

**security-level** *number*

**Example:**

```
ciscoasa(config-if)# security-level 50
```

The *number* is an integer between 0 (lowest) and 100 (highest)..

**Step 5** (Optional) Set an interface to management-only mode so that it does not pass through traffic:

**management-only**

By default, Management interfaces are configured as management-only.

### Examples

The following example configures parameters for VLAN 101:

```
ciscoasa(config)# interface vlan 101
ciscoasa(config-if)# nameif inside
ciscoasa(config-if)# security-level 100
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0
```

The following example configures parameters in multiple context mode for the context configuration. The interface ID is a mapped name.

```
ciscoasa/contextA(config)# interface int1
ciscoasa/contextA(config-if)# nameif outside
ciscoasa/contextA(config-if)# security-level 100
ciscoasa/contextA(config-if)# ip address 10.1.2.1 255.255.255.0
```

### Related Topics

Configure IPv6 Addressing, on page 12
Enable the Physical Interface and Configure Ethernet Parameters
Configure PPPoE, on page 7

# Configure PPPoE

If the interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, configure the following parameters.

### Procedure

**Step 1** Define the Virtual Private Dialup Network (VPDN) group name of your choice to represent this connection:

**vpdn group** *group_name* **request dialout pppoe**

**Example:**

```
ciscoasa(config)# vpdn group pppoe-sbc request dialout pppoe
```

**Step 2** If your ISP requires authentication, select an authentication protocol:

**vpdn group** *group_name* **ppp authentication** {**chap** | **mschap** | **pap**}

**Example:**

```
ciscoasa(config)# vpdn group pppoe-sbc ppp authentication chap
```

Enter the appropriate keyword for the type of authentication used by your ISP.

When using CHAP or MS-CHAP, the username may be referred to as the remote system name, while the password may be referred to as the CHAP secret.

**Step 3** Associate the username assigned by your ISP to the VPDN group:

**vpdn group** *group_name* **localname** *username*

**Example:**

```
ciscoasa(config)# vpdn group pppoe-sbc localname johncrichton
```

**Step 4** Create a username and password pair for the PPPoE connection:

**vpdn username** *username* **password** *password* [**store-local**]

**Example:**

```
ciscoasa(config)# vpdn username johncrichton password moya
```

The **store-local** option stores the username and password in a special location of NVRAM on the ASA. If an Auto Update Server sends a **clear config** command to the ASA and the connection is then interrupted, the ASA can read the username and password from NVRAM and re-authenticate to the Access Concentrator.

# Configure Transparent Mode Bridge Group Interfaces

A bridge group is a group of interfaces that the ASA bridges instead of routes. Bridge groups are only supported in Transparent Firewall Mode. For more information about bridge groups, see About Bridge Groups.

To configure bridge groups and associated interfaces, perform these steps.

## Configure the Bridge Virtual Interface (BVI)

Each bridge group requires a BVI for which you configure an IP address. The ASA uses this IP address as the source address for packets originating from the bridge group. The BVI IP address must be on the same subnet as the connected network. For IPv4 traffic, the BVI 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.

Some models include a bridge group and BVI in the default configuration. You can create additional bridge groups and BVIs and reassign member interfaces between the groups.

**Note** For a separate management interface in transparent mode (for supported models), a non-configurable bridge group (ID 301) is automatically added to your configuration. This bridge group is not included in the bridge group limit.

### Procedure

**Step 1** Create a BVI:

**interface bvi** *bridge_group_number*

**Example:**

```
ciscoasa(config)# interface bvi 2
```

The *bridge_group_number* is an integer between 1 and 250. You will later assign physical interfaces to this bridge group number.

**Step 2** Specify the IP address for the BVI:

**ip address** *ip_address* [*mask*] [**standby** *ip_address*]

**Example:**

```
ciscoasa(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2
```

Do not assign a host address (/32 or 255.255.255.255) to the BVI. Also, do not use other subnets that contain fewer than 3 host addresses (one each for the upstream router, downstream router, and BVI) 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 **standby** keyword and address is used for failover.

### Example

The following example sets the BVI 2 address and standby address:

```
ciscoasa(config)# interface bvi 2
ciscoasa(config-if)# ip address 10.1.3.1 255.255.255.0 standby 10.1.3.2
```

## Configure General Bridge Group Member Interface Parameters

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

### Before you begin

- The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, VNI interfaces, EtherChannels, and redundant interfaces. The Management interface is not supported.
- 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.
- For transparent mode, do not use this procedure for Management interfaces; see Configure a Management Interface for Transparent Mode, on page 10 to configure the Management interface.

### Procedure

**Step 1** Enter interface configuration mode:

**interface** *id*

**Example:**

```
ciscoasa(config)# interface gigabitethernet 0/0
```

The interface ID can be:

- **redundant**
- **port-channel**
- *physical*—For example, **ethernet**, **gigabitethernet**, **tengigabitethernet**. Management interfaces are not supported. Refer to the hardware installation guide for your model for interface names.
- *physical_or_port-channel_or_redundant***.***subinterface*—For example, **gigabitethernet0/0.100**, **port-channel1.100**. or **redundant2.100**.
- **vni**

- **vlan**
- *mapped_name*—For multiple context mode.

**Step 2** Assign the interface to a bridge group:

**bridge-group** *number*

**Example:**

```
ciscoasa(config-if)# bridge-group 1
```

The *number* is an integer between 1 and 250, and must match the BVI interface number. You can assign up to 4 interfaces to a bridge group. You cannot assign the same interface to more than one bridge group.

**Step 3** Name the interface:

**nameif** *name*

**Example:**

```
ciscoasa(config-if)# nameif inside1
```

The *name* 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.

**Step 4** Set the security level:

**security-level** *number*

**Example:**

```
ciscoasa(config-if)# security-level 50
```

The *number* is an integer between 0 (lowest) and 100 (highest)..

**Related Topics**

Configure the MTU and TCP MSS

# Configure a Management Interface for Transparent Mode

In transparent firewall mode, all interfaces must belong to a bridge group. The only exception is 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)) which you can configure as a separate management interface; for the Firepower 9300 chassis, the management interface ID depends on the mgmt-type interface that you assigned to the ASA logical device. You cannot use any other interface types as management interfaces. You can configure one management interface in single mode or per context. For more information see Management Interface for Transparent Mode.

### Before you begin

- Do not assign this interface to a bridge group; a non-configurable bridge group (ID 301) 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 ASASM.) For the Firepower 9300 chassis, the management interface ID depends on the mgmt-type interface that you assigned to the ASA logical device.
- In multiple context mode, you cannot share any interfaces, including the Management interface, across contexts. You must connect to a data interface.
- 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.

### Procedure

**Step 1** Enter interface configuration mode:

**interface** {{**port-channel** *number* | **management** *slot*/*port* | *mgmt-type_interface_id* }[**.** *subinterface*] | *mapped_name*}

**Example:**

```
ciscoasa(config)# interface management 0/0.1
```

The **port-channel** *number* argument is the EtherChannel interface ID, such as **port-channel 1**. The EtherChannel interface must have only Management member interfaces.

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

In multiple context mode, enter the *mapped_name* if one was assigned using the **allocate-interface** command.

For the Firepower 9300 chassis, specify the interface ID for the mgmt type interface (individual or EtherChannel) that you assigned to the ASA logical device.

**Step 2** Name the interface:

**nameif** *name*

**Example:**

```
ciscoasa(config-if)# nameif management
```

The *name* 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.

**Step 3** Set the IP address using one of the following methods.

- Set the IP address manually:

  For use with failover, you must set the IP address and standby address manually; DHCP is not supported.

The *ip_address* and *mask* arguments set the interface IP address and subnet mask.

The standby *ip_address* argument is used for failover.

**ip address** *ip_address* [*mask*] [**standby** *ip_address*]

Example:

```
ciscoasa(config-if)# ip address 10.1.1.1 255.255.255.0 standby 10.1.1.2
```

- Obtain an IP address from a DHCP server:

  **ip address dhcp** [**setroute**]

  Example:

  ```
  ciscoasa(config-if)# ip address dhcp
  ```

  The **setroute** keyword lets the ASA use the default route supplied by the DHCP server.

  Reenter this command to reset the DHCP lease and request a new lease.

  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.

**Step 4** Set the security level:

**security-level** *number*

**Example:**

```
ciscoasa(config-if)# security-level 100
```

The *number* is an integer between 0 (lowest) and 100 (highest).

# Configure IPv6 Addressing

This section describes how to configure IPv6 addressing.

## About IPv6

This section includes information about IPv6.

### 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. For a bridge group, this address needs to be configured for the BVI, and not per member interface. You can also configure a global IPv6 address for the management interface in transparent mode.

- 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 Neighbor Discovery functions such as address resolution. In a bridge group, only member interfaces have link-local addresses; the BVI does not have a link-local address.

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 the interface, so you do not also need to specifically configure a link-local address. For bridge group member interfaces, when you configure the global address on the BVI, the ASA automatically generates link-local addresses for member interfaces. 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.

## 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:

```
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.

# Configure a Global IPv6 Address

To configure a global IPv6 address for any routed mode interface and for the transparent mode BVI, perform the following steps.

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

For subinterfaces, we recommend that you also set the MAC address manually, because they use the same burned-in MAC address of the parent interface. IPv6 link-local addresses are generated based on the MAC address, so assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the ASA. See Manually Configure the MAC Address.

### Before you begin

- 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.

### Procedure

**Step 1** Enter interface configuration mode:

**interface** *id*

**Example:**

```
ciscoasa(config)# interface gigabithethernet 0/0
```

In transparent mode, specify the BVI:

**Example:**

```
ciscoasa(config)# interface bvi 1
```

In transparent mode, in addition to the BVI, you can also specify a Management interface:

**Example:**

```
ciscoasa(config)# interface management 1/1
```

**Step 2** (Routed interface) Set the IP address using one of the following methods.

- Enable stateless autoconfiguration on the interface:

  **ipv6 address autoconfig**

  Enabling stateless autoconfiguration on the interface configures IPv6 addresses based on prefixes received in Router Advertisement messages. A link-local address, based on the Modified EUI-64 interface ID, is automatically generated for the interface when stateless autoconfiguration is enabled.

  **Note** Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the ASA does send Router Advertisement messages in this case. See the **ipv6 nd suppress-ra** command to suppress messages.

- Manually assign a global address to the interface:

  **ipv6 address** *ipv6_address*/*prefix-length* [**standby** *ipv6_address*]

  Example:

  ```
  ciscoasa(config-if)# ipv6 address 2001:0DB8:BA98::3210/64 standby 2001:0DB8:BA98::3211
  ```

  When you assign a global address, the link-local address is automatically created for the interface.

  **standby** specifies the interface address used by the secondary unit or failover group in a failover pair.

- Assign a global address to the interface by combining the specified prefix with an interface ID generated from the interface MAC address using the Modified EUI-64 format:

  **ipv6 address** *ipv6-prefix*/*prefix-length* **eui-64**

  Example:

  ```
  ciscoasa(config-if)# ipv6 address 2001:0DB8:BA98::/64 eui-64
  ```

  When you assign a global address, the link-local address is automatically created for the interface.

  You do not need to specify the standby address; the interface ID will be generated automatically.

**Step 3** (BVI interface) Manually assign a global address to the BVI. For a management interface in Transparent mode, use this method as well.

**ipv6 address** *ipv6_address*/*prefix-length* [**standby** *ipv6_address*]

Example:

```
ciscoasa(config-if)# ipv6 address 2001:0DB8::BA98:0:3210/48
```

When you assign a global address, the link-local address is automatically created for the interface.

**standby** specifies the interface address used by the secondary unit or failover group in a failover pair.

**Step 4** (Optional) Enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link:

**ipv6 enforce-eui64** *if_name*

**Example:**

```
ciscoasa(config)# ipv6 enforce-eui64 inside
```

The *if_name* argument is the name of the interface, as specified by the **nameif** command, on which you are enabling the address format enforcement.

# Configure IPv6 Neighbor Discovery

The IPv6 neighbor discovery process uses ICMPv6 messages and solicited-node multicast addresses to determine the link-layer address of a neighbor on the same network (local link), verify the readability of a neighbor, and keep track of neighboring routers.

Nodes (hosts) use neighbor discovery to determine the link-layer addresses for neighbors known to reside on attached links and to quickly purge cached values that become invalid. Hosts also use neighbor discovery to find neighboring routers that are willing to forward packets on their behalf. In addition, nodes use the protocol to actively keep track of which neighbors are reachable and which are not, and to detect changed link-layer addresses. When a router or the path to a router fails, a host actively searches for functioning alternates.

-

### Procedure

**Step 1** Specify the IPv6 interface you want to configure.

**interface** *name*

**Example:**

```
ciscoasa(config)# interface gigabitethernet 0/0
ciscoasa(config-if)#
```

**Step 2** Specify the number Duplicate Address Detection (DAD) attempts.

**ipv6 nd dad attempts** *value*

Valid values for the *value* argument range from 0 to 600. A 0 value disables DAD processing on the specified interface. The default is 1 message.

DAD ensures the uniqueness of new unicast IPv6 addresses before they are assigned, and ensures that duplicate IPv6 addresses are detected in the network on a link basis. The ASA uses neighbor solicitation messages to perform DAD.

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:

```
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.

**Example:**

```
ciscoasa(config-if)# ipv6 nd dad attempts 20
```

**Step 3** Set the interval between IPv6 neighbor solicitation retransmissions.

**ipv6 nd ns-interval** *value*

Values for the *value* argument range from 1000 to 3600000 milliseconds.

Neighbor solicitation messages (ICMPv6 Type 135) are sent on the local link by nodes attempting to discover the link-layer addresses of other nodes on the local link. After receiving a neighbor solicitation message, the destination node replies by sending a neighbor advertisement message (ICPMv6 Type 136) on the local link.

After the source node receives the neighbor advertisement, the source node and destination node can communicate. Neighbor solicitation messages are also used to verify the reachability of a neighbor after the link-layer address of a neighbor is identified. When a node wants to verifying the reachability of a neighbor, the destination address in a neighbor solicitation message is the unicast address of the neighbor.

Neighbor advertisement messages are also sent when there is a change in the link-layer address of a node on a local link.

**Example:**

```
ciscoasa(config-if)# ipv6 nd ns-interval 9000
```

**Step 4** Set the amount of time that a remote IPv6 node is reachable.

**ipv6 nd reachable-time** *value*

Values for the *value* argument range from 0 to 3600000 milliseconds. When 0 is used for the value, the reachable time is sent as undetermined. It is up to the receiving devices to set and track the reachable time value.

The neighbor reachable time enables detecting unavailable neighbors. Shorter configured times enable detecting unavailable neighbors more quickly, however, shorter times consume more IPv6 network bandwidth and processing resources in all IPv6 network devices. Very short configured times are not recommended in normal IPv6 operation.

**Example:**

```
ciscoasa config-if)# ipv6 nd reachable-time 1700000
```

**Step 5** Set the interval between IPv6 router advertisement transmissions.

**ipv6 nd ra-interval** [**msec**] *value*

The **msec** keyword indicates that the value provided is in milliseconds. If this keyword is not present, the value provided is in seconds. Valid values for the *value* argument range from 3 to 1800 seconds or from 500 to 1800000 milliseconds if the **msec** keyword is provided. The default is 200 seconds.

The interval value is included in all IPv6 router advertisements that are sent out of this interface.

The interval between transmissions should be less than or equal to the IPv6 router advertisement lifetime if the ASA is configured as a default router. To prevent synchronization with other IPv6 nodes, randomly adjust the actual value used to within 20 percent of the desired value.

**Example:**

```
ciscoasa(config-if)# ipv6 nd ra-interval 201
```

**Step 6** Specify the length of time that nodes on the local link should consider the ASA as the default router on the link.

**ipv6 nd ra-lifetime** [**msec**] *value*

The optional **msec** keyword indicates that the value provided is in milliseconds. Otherwise, the value is in seconds. Values for the *value* argument range from 0 to 9000 seconds. Entering 0 indicates that the ASA should not be considered a default router on the selected interface.

The router lifetime value is included in all IPv6 router advertisements sent out of the interface. The value indicates the usefulness of the ASA as a default router on this interface.

**Example:**

```
ciscoasa(config-if)# ipv6 nd ra-lifetime 2000
```

**Step 7** Suppress router advertisements.

**ipv6 nd suppress-ra**

Router advertisement messages (ICMPv6 Type 134) are automatically sent in response to router solicitation messages (ICMPv6 Type 133). Router solicitation messages are sent by hosts at system startup so that the host can immediately autoconfigure without needing to wait for the next scheduled router advertisement message.

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).

Entering this command causes the ASA to appear as a regular IPv6 neighbor on the link and not as an IPv6 router.

**Step 8** Add a flag to IPv6 router advertisements to inform IPv6 autoconfiguration clients to use DHCPv6 to obtain an IPv6 address, in addition to the derived stateless autoconfiguration address.

**ipv6 nd managed-config-flag**

This option sets the Managed Address Config flag in the IPv6 router advertisement packet.

**Step 9** Add a flag to IPv6 router advertisements to inform IPv6 autoconfiguration clients to use DHCPv6 to obtain the DNS server address, or other information.

**ipv6 nd other-config-flag**

This option sets the Other Address Config flag in the IPv6 router advertisement packet.

**Step 10** Configure which IPv6 prefixes are included in IPv6 router advertisements:

**ipv6 nd prefix** {*ipv6_prefix*/*prefix_length* | **default**} [*valid_lifetime preferred_lifetime* | **at** *valid_date preferred_date*] [**no-advertise**] [**no-autoconfig**] [ ] [**off-link**]

The prefix advertisement can be used by neighboring devices to autoconfigure their interface addresses. Stateless autoconfiguration uses IPv6 prefixes provided in router advertisement messages to create the global unicast address from the link-local address.

By default, prefixes configured as addresses on an interface using the **ipv6 address** command are advertised in router advertisements. If you configure prefixes for advertisement using the **ipv6 nd prefix** command, then only these prefixes are advertised.

For stateless autoconfiguration to work correctly, the advertised prefix length in router advertisement messages must always be 64 bits.

- **default**—Indicates that the default prefix is used.
- *valid_lifetime preferred_lifetime* —Specifies the amount of time that the specified IPv6 prefix is advertised as being valid and preferred. An address has no restrictions during the preferred lifetime. After the preferred lifetime expires, the address goes into a deprecated state; while an address is in a deprecated state, its use is discouraged, but not strictly forbidden. After the valid lifetime expires, the address becomes invalid and cannot be used. The valid lifetime must be greater than or equal to the preferred lifetime. Values range from 0 to 4294967295 seconds. The maximum value represents infinity, which can also be specified with the **infinite** keyword. The valid lifetime default is 2592000 (30 days). The preferred lifetime default is 604800 (7 days).
- **at** *valid_date preferred_date*—Indicates a specific date and time at which the prefix expires. Specify the date as the *month_name day hh*:*mm*. For example, enter `dec 1 13:00`.
- **no-advertise**—Disables advertisement of the prefix.
- **no-autoconfig**—Specifies that the prefix cannot be used for IPv6 autoconfiguration.
- **off-link**—Configures the specified prefix as off-link. The prefix will be advertised with the L-bit clear. The prefix will not be inserted into the routing table as a Connected prefix.

When onlink is on (by default), the specified prefix is assigned to the link. Nodes sending traffic to such addresses that contain the specified prefix consider the destination to be locally reachable on the link.

**Example:**

```
ciscoasa(config-if)# ipv6 nd prefix 2001:DB8::/32 1000 900
```

**Step 11** Configure a static entry in the IPv6 neighbor discovery cache.

**ipv6 neighbor** *ipv6_address if_name mac_address*

The following guidelines and limitations apply for configuring a static IPv6 neighbor:

- The **ipv6 neighbor** command is similar to the **arp** command. If an entry for the specified IPv6 address already exists in the neighbor discovery cache—learned through the IPv6 neighbor discovery process—the entry is automatically converted to a static entry. These entries are stored in the configuration when the copy command is used to store the configuration.
- Use the **show ipv6 neighbor** command to view static entries in the IPv6 neighbor discovery cache.
- The **clear ipv6 neighbor** command deletes all entries in the IPv6 neighbor discovery cache except static entries. The **no ipv6 neighbor** command deletes a specified static entry from the neighbor discovery cache; the command does not remove dynamic entries—entries learned from the IPv6 neighbor discovery process—from the cache. Disabling IPv6 on an interface by using the **no ipv6 enable** command deletes all IPv6 neighbor discovery cache entries configured for that interface except static entries (the state of the entry changes to INCMP [Incomplete]).
- Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.
- The **clear ipv6 neighbor** command does not remove static entries from the IPv6 neighbor discovery cache; it only clears the dynamic entries.
- The ICMP syslogs generated are caused by a regular refresh of IPv6 neighbor entries. The ASA default timer for IPv6 neighbor entry is 30 seconds, so the ASA would generate ICMPv6 neighbor discovery and response packets about every 30 seconds. If the ASA has both failover LAN and state interfaces configured with IPv6 addresses, then every 30 seconds, ICMPv6 neighbor discovery and response packets will be generated by both ASAs for both configured and link-local IPv6 addresses. In addition, each packet will generate several syslogs (ICMP connection and local-host creation or teardown), so it may appear that constant ICMP syslogs are being generated. The refresh time for IPV6 neighbor entry is configurable on the regular data interface, but not configurable on the failover interface. However, the CPU impact for this ICMP neighbor discovery traffic is minimal.

**Example:**

```
ciscoasa(config)# ipv6 neighbor 3001:1::45A inside 002.7D1A.9472
```

# Monitoring Routed and Transparent Mode Interfaces

You can monitor interface statistics, status, PPPoE.

**Note** For the Firepower 9300, some statistics are not shown using the ASA commands. You must view more detailed interface statistics using FXOS commands.

- /eth-uplink/fabric# **show interface**
- /eth-uplink/fabric# **show port-channel**
- /eth-uplink/fabric/interface# **show stats**

See the FXOS troubleshooting guide for more information.

# Interface Statistics and Information

- **show interface**

  Displays interface statistics.

- **show interface ip brief**

  Displays interface IP addresses and status.

- **show bridge-group**

  Displays bridge group information such as interfaces assigned, MAC addresses, and IP addresses.

# PPPoE

- **show ip address** *interface_name* **pppoe**

  Displays the current PPPoE client configuration information.

- **debug pppoe** {**event** | **error** | **packet**}

  Enables debugging for the PPPoE client.

- **show vpdn session** [**l2tp** | **pppoe**] [**id** *sess_id* | **packets** | **state** | **window**]

  Views the status of PPPoE sessions.

  The following examples show information provided by this command:

```
ciscoasa# show vpdn

Tunnel id 0, 1 active sessions
     time since change 65862 secs
     Remote Internet Address 10.0.0.1
    Local Internet Address 199.99.99.3
     6 packets sent, 6 received, 84 bytes sent, 0 received
Remote Internet Address is 10.0.0.1
     Session state is SESSION_UP
       Time since event change 65865 secs, interface outside
       PPP interface id is 1
       6 packets sent, 6 received, 84 bytes sent, 0 received
ciscoasa#
ciscoasa# show vpdn session
PPPoE Session Information (Total tunnels=1 sessions=1)
```

```
Remote Internet Address is 10.0.0.1
  Session state is SESSION_UP
    Time since event change 65887 secs, interface outside
    PPP interface id is 1
    6 packets sent, 6 received, 84 bytes sent, 0 received
ciscoasa#
ciscoasa# show vpdn tunnel
PPPoE Tunnel Information (Total tunnels=1 sessions=1)
Tunnel id 0, 1 active sessions
   time since change 65901 secs
   Remote Internet Address 10.0.0.1
   Local Internet Address 199.99.99.3
   6 packets sent, 6 received, 84 bytes sent, 0 received
ciscoasa#
```

## IPv6 Neighbor Discovery

To monitor IPv6 neighbor discovery parameters, enter the following command:

- **show ipv6 interface**

  This command displays the usability status of interfaces configured for IPv6, including the interface name, such as “outside,” and displays the settings for the specified interface. However, it excludes the name from the command and displays the settings for all interfaces that have IPv6 enabled on them. Output for the command shows the following:

  - The name and status of the interface.
  - The link-local and global unicast addresses.
  - The multicast groups to which the interface belongs.
  - ICMP redirect and error message settings.
  - Neighbor discovery settings.
  - The actual time when the command is set to 0.
  - The neighbor discovery reachable time that is being used.

# Examples for Routed and Transparent Mode Interfaces

## Transparent Mode Example with 2 Bridge Groups

The following example for transparent mode 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
```

# History for Routed and Transparent Mode Interfaces

| Feature Name | Platform Releases | Feature Information |
|---|---|---|
| IPv6 Neighbor Discovery | 7.0(1) | We introduced this feature.<br>We introduced the following commands: **ipv6 nd ns-interval**, **ipv6 nd ra-lifetime**, **ipv6 nd suppress-ra**, **ipv6 neighbor**, **ipv6 nd prefix**, **ipv6 nd dad-attempts**, **ipv6 nd reachable-time**, **ipv6 address**, **ipv6 enforce-eui64**. |
| IPv6 support for transparent mode | 8.2(1) | IPv6 support was introduced for transparent firewall mode. |

| Feature Name | Platform Releases | Feature Information |
|---|---|---|
| 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.<br>We introduced the following commands: **interface bvi**, **show bridge-group** |
| Address Config Flags for IPv6 DHCP Relay | 9.0(1) | We introduced the following commands: **ipv6 nd managed-config-flag**, **ipv6 nd other-config-flag**. |
| Transparent mode bridge group maximum increased to 250 | 9.3(1) | The bridge group maximum was increased from 8 to 250 bridge groups. You can configure up to 250 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group.<br>We modified the following commands: **interface bvi**, **bridge-group** |