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. In the Configuration > Device List pane, double-click the context name under the active device IP address.

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, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

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

Step 2 Choose the interface row, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 4 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 5 (Optional) To set this interface as a management-only interface, check the **Dedicate this interface to management-only** check box.

Through traffic is not accepted on a management-only interface.

Note The Channel Group field is read-only and indicates if the interface is part of an EtherChannel.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

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

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.

For failover, set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab. 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.

- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
  - **a.** To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

    Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned.
  - **b.** To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.
  - **c.** (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
  - **d.** (Optional) To assign an administrative distance to the learned route, enter a value between 1 and 255 in the **DHCP Learned Route Metric** field. If this field is left blank, the administrative distance for the learned routes is 1.
  - **e.** (Optional) To enable tracking for DHCP-learned routes, check **Enable Tracking for DHCP Learned Routes**. Set the following values:

    Track ID—A unique identifier for the route tracking process. Valid values are from 1 to 500.

    Track IP Address—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.

    Note Route tracking is only available in single, routed mode.

    SLA ID—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.

    Monitor Options—Click this button to open the **Route Monitoring Options** dialog box. In the **Route Monitoring Options** dialog box you can configure the parameters of the tracked object monitoring process.
  - **f.** (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

    The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.
  - **g.** (Optional) To renew the lease, click **Renew DHCP Lease**.

- (Single mode only) To obtain an IP address using PPPoE, check **Use PPPoE**.
  - **a.** In the **Group Name** field, specify a group name.
  - **b.** In the **PPPoE Username** field, specify the username provided by your ISP.
  - **c.** In the **PPPoE Password** field, specify the password provided by your ISP.
  - **d.** In the **Confirm Password** field, retype the password.
  - **e.** For PPP authentication, click either the **PAP**, **CHAP**, or **MSCHAP** radio button.

    PAP passes cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to

the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.

f. (Optional) To store the username and password in flash memory, check the **Store Username and Password in Local Flash** check box.

The ASA stores the username and password in a special location of NVRAM. If an Auto Update Server sends a **clear configure** 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.

g. (Optional) To display the **PPPoE IP Address and Route Settings** dialog box where you can choose addressing and tracking options, click **IP Address and Route Settings**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Related Topics

Configure IPv6 Addressing, on page 11
Enable the Physical Interface and Configure Ethernet Parameters
Configure PPPoE, on page 6

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 Choose **Configuration** > **Interfaces** > **Add/Edit Interface** > **General**, and then click **PPPoE IP Address and Route Settings**.

Step 2 In the **IP Address** area, choose one of the following:

- **Obtain IP Address using PPP**—Dynamically configure the IP address.
- **Specify an IP Address**—Manually configure the IP address.

Step 3 In the **Route Settings Area**, configure the following:

- **Obtain default route using PPPoE**—Set the default routes when the PPPoE client has not yet established a connection. When using this option, you cannot have a statically defined route in the configuration.
- **PPPoE learned route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. If this field is left blank, the administrative distance for the learned routes is 1.

- **Enable tracking**—Enable route tracking for PPPoE-learned routes. Route tracking is only available in single, routed mode.
- **Primary Track**—Configure the primary PPPoE route tracking.
- **Track ID**—A unique identifier for the route tracking process. Valid values are from 1 to 500.
- **Track IP Address**—Enter the IP address of the target being tracked. Typically, this would be the IP address of the next hop gateway for the route, but it could be any network object available off of that interface.
- **SLA ID**—A unique identifier for the SLA monitoring process. Valid values are from 1 to 2147483647.
- **Monitor Options**—Click this button to open the **Route Monitoring Options** dialog box. In the **Route Monitoring Options** dialog box you can configure the parameters of the tracked object monitoring process.
- **Secondary Track**—Configure the secondary PPPoE route tracking.
- **Secondary Track ID**—A unique identifier for the route tracking process. Valid values are from 1 to 500.

Step 4 Click **OK**.

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 Choose **Configuration** > **Interfaces**, and then choose **Add** > **Bridge Group Interface**.

Step 2 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

You will later assign physical interfaces to this bridge group number.

Step 3 Set the IP address.

a) In the **IP Address** field, enter the IPv4 address.

b) In the **Subnet Mask** field, enter the subnet mask or choose one from the menu.

Do not assign a host address (/32 or 255.255.255.255) to the transparent firewall. 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. For example, 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.

Step 4 (Optional) In the **Description** field, enter a description for this bridge group.

Step 5 Click **OK**.

A Bridge Virtual Interface (BVI) is added to the interface table, along with the physical and subinterfaces.

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, in the Configuration > Device List pane, double-click the context name under the active device IP address.
- For transparent mode, do not use this procedure for Management interfaces; see Configure a Management Interface for Transparent Mode, on page 9 to configure the Management interface.

Procedure

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

BVIs appear in the table alongside physical interfaces, subinterfaces, redundant interfaces, and EtherChannel port-channel interfaces. In multiple context mode, only interfaces that were assigned to the context in the System execution space appear in the table.

Step 2 Choose the row for a non-BVI interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 In the **Bridge Group** drop-down menu, choose the bridge group to which you want to assign this interface.

Step 4 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 5 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Note The **Channel Group** field is read-only and indicates if the interface is part of an EtherChannel.

Step 7 (Optional) If you install a module, and you want to demonstrate the module functionality on a non-production ASA, check the **Forward traffic to the ASA module for inspection and reporting check box**. See the module chapter or quick start guide for more information.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns. In the case of a failover or state link, the description is fixed as "LAN Failover Interface," "STATE Failover Interface," or "LAN/STATE Failover Interface," for example. You cannot edit this description. The fixed description overwrites any description you enter here if you make this interface a failover or state link.

Step 9 Click **OK**.

Related Topics

Configure the Manual MAC Address, 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, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

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

Step 2 Choose the row for a Management interface, subinterface, or EtherChannel port-channel interface comprised of Management interfaces, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

For the Firepower 9300 chassis, the management interface ID depends on the mgmt-type interface (individual or EtherChannel) that you assigned to the ASA logical device.

Step 3 In the **Bridge Group** drop-down menu, leave the default **--None--**. You cannot assign a management interface to a bridge group.

Step 4 In the **Interface Name** field, enter a name up to 48 characters in length.

Step 5 In the **Security level** field, enter a level between 0 (lowest) and 100 (highest).

Note The **Dedicate this interface to management only** check box is enabled by default and is non-configurable.

Step 6 If the interface is not already enabled, check the **Enable Interface** check box.

Step 7 To set the IP address, use one of the following options.

Note For use with failover, you must set the IP address and standby address manually; DHCP is not supported. Set the standby IP addresses on the **Configuration > Device Management > High Availability > Failover > Interfaces** tab.

- To set the IP address manually, click the **Use Static IP** radio button and enter the IP address and mask.
- To obtain an IP address from a DHCP server, click the **Obtain Address via DHCP** radio button.
  - To force a MAC address to be stored inside a DHCP request packet for option 61, click the **Use MAC Address** radio button.

    Some ISPs expect option 61 to be the interface MAC address. If the MAC address is not included in the DHCP request packet, then an IP address will not be assigned.
  - To use a generated string for option 61, click **Use "Cisco-<MAC>-<interface_name>-<host>"**.
  - (Optional) To obtain the default route from the DHCP server, check **Obtain Default Route Using DHCP**.
  - (Optional) To set the broadcast flag to 1 in the DHCP packet header when the DHCP client sends a discover requesting an IP address, check **Enable DHCP Broadcast flag for DHCP request and discover messages**.

    The DHCP server listens to this broadcast flag and broadcasts the reply packet if the flag is set to 1.
  - (Optional) To renew the lease, click **Renew DHCP Lease**.

Step 8 (Optional) In the **Description** field, enter a description for this interface.

The description can be up to 240 characters on a single line, without carriage returns.

Step 9 Click **OK**.

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.

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 Configure the Manual MAC Address, MTU, and TCP MSS.

Before you begin

- In multiple context mode, complete this procedure in the context execution space. To change from the system to a context configuration, in the Configuration > Device List pane, double-click the context name under the active device IP address.

Procedure

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

Step 2 Choose an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

In transparent mode, select a BVI. For transparent mode, you can also select a management-only interface.

Step 3 Click the **IPv6** tab.

Step 4 Check the **Enable IPv6** check box.

Step 5 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

Step 6 (Routed interface) Configure the global IPv6 address using one of the following methods.

- Stateless autoconfiguration—In the **Interface IPv6 Addresses** area, check the **Enable address autoconfiguration** check box.

  Enabling stateless autconfiguration on the interface configures IPv6 addresses based upon 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. Check the **Suppress RA** check box to suppress messages.

- Manual configuration—To manually configure a global IPv6 address:

  a. In the **Interface IPv6 Addresses** area, click **Add**.

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

b. In the **Address/Prefix Length** field, the value you enter depends on the method you want to use:

- Full global address—If you want to manually enter the entire address, enter the full address plus the prefix length.
- Modified EUI 64 format—Enter the IPv6 prefix and length, and then check the **EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).

c. Click **OK**.

Step 7 (BVI interface) Manually assign a global address to the BVI. For a management interface in Transparent mode, use this method as well.

a) In the **Interface IPv6 Addresses** area, click **Add**.

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

b) In the **Address/Prefix Length** field, enter the full global IPv6 address along with the IPv6 prefix length.

c) Click **OK**.

Step 8 Click **OK**.

You return to the **Configuration > Device Setup > Interface Settings > Interfaces** pane.

(Optional) Configure the Link-Local Addresses Automatically

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of generating the link-local addresses based on the interface MAC addresses (Modified EUI-64 format. Because MAC addresses use 48 bits, additional bits must be inserted to fill the 64 bits required for the interface ID.)

To automatically configure the link-local addresses for an interface, perform the following steps.

Before you begin

Supported in routed mode only.

Procedure

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

Step 2 Select an interface, and click **Edit**.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **IPv6** tab.

Step 4 In the **IPv6 configuration** area, check the **Enable IPv6** check box.

This option enables IPv6 and automatically generates the link-local address using the Modified EUI-64 interface ID based on the interface MAC address.

Step 5 Click **OK**.

(Optional) Configure the Link-Local Addresses Manually

If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.

To assign a link-local address to an interface, perform the following steps.

Procedure

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

Step 2 Select an interface, and click **Edit**.

For bridge groups, choose a bridge group member interface.

The **Edit Interface** dialog box appears with the **General** tab selected.

Step 3 Click the **IPv6** tab.

Step 4 (Optional) To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.

Step 5 To set the link-local address, enter an address in the **Link-local address** field.

A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. See IPv6 Addresses for more information about IPv6 addressing.

Step 6 Click **OK**.

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 Choose **Configuration** > **Device Setup** > **Interface Settings** > **Interfaces**.

Step 2 Choose the IPv6 interface on which to configure IPv6 neighbor settings, and click **Edit**.

Step 3 Click the **IPv6** tab.

Step 4 Enter the number of allowed **DAD Attempts**.

Values 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.

Step 5 Enter the **NS Interval** in milliseconds to set the interval between IPv6 neighbor solicitation retransmissions.

Valid 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.

Step 6 Enter the **Reachable Time** in seconds to set how long a remote IPv6 node is reachable.

Set the reachable time between 0 to 3600000 milliseconds. When you set the time to 0, then 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.

Step 7 Enter the **RA Lifetime** in seconds to set the length of time that nodes on the local link consider the ASA as the default router on the link.

Values range from 0 to 9000 seconds. Entering 0 indicates that the ASA should not be considered a default router on the selected interface.

Step 8 Check the **Suppress RA** check box to suppress router advertisements.

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).

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

Step 9 Enter the **RA Interval** to set the interval between IPv6 router advertisement transmissions.

Valid values range from 3 to 1800 seconds. The default is 200 seconds.

To add a router advertisement transmission interval value in milliseconds instead, check the **RA Interval in Milliseconds** check box, and enter a value from 500 to 1800000.

Step 10 Check the **Hosts should use DHCP for address config** check box to inform IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.

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

Step 11 Check the **Hosts should use DHCP for non-address config** check box to inform IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

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

Step 12 Configure which IPv6 prefixes are included in IPv6 router advertisements.

a) In the **Interface IPv6 Prefixes** area, click **Add**.
b) Enter the **Address/Prefix Length** or check the **Default** check box to use the default prefix.
c) Check the **No Auto-Configuration** check box to force hosts to configure the IPv6 address manually. Hosts on the local link with the specified prefix cannot use IPv6 autoconfiguration.
d) Check the **No Advertisements** check box to disable prefix advertisement.
e) Check the **Off Link** check box to configure 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.
f) In the **Prefix Lifetime** area, specify a **Lifetime Duration** or **Lifetime Expiration Date**.

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.

- **Lifetime Duration**—Values range from 0 to 4294967295. The default valid lifetime is 2592000 (30 days). The default preferred lifetime is 604800 (7 days). The maximum value represents infinity.
- **Lifetime Expiration Date**—Choose a valid and preferred month and day from the drop-down lists, and then enter a time in hh:mm format.

g) Click **OK** to save your settings.

Step 13 Click **OK**.

Step 14 Configure a static IPv6 neighbor.

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

- This feature is similar to adding a static ARP entry. 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.
- Static entries in the IPv6 neighbor discovery cache are not modified by the neighbor discovery process.

- 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.

See also View and Clear Dynamically Discovered Neighbors, on page 17.

a) Choose **Configuration** > **Device Management** > **Advanced** > **IPv6 Neighbor Discovery Cache**.
b) Click **Add**.

   The **Add IPv6 Static Neighbor** dialog box appears.

c) From the **Interface Name** drop-down list, choose an interface on which to add the neighbor.
d) In the **IP Address** field, enter the IPv6 address that corresponds to the local data-link address, or click the ellipsis (...) to browse for an address.
e) In the **MAC address** field, enter the local data-line (hardware) MAC address.
f) Click **OK**.

Step 15 Click **Apply** to save the running configuration.

View and Clear Dynamically Discovered Neighbors

When a host or node communicates with a neighbor, the neighbor is added to the neighbor discovery cache. The neighbor is removed from the cache when there is no longer any communication with that neighbor.

To view dynamically discovered neighbors and clear these neighbors from the IPv6 neighbor discovery cache, perform the following steps:

Procedure

Step 1 Choose **Monitoring** > **Interfaces** > **IPv6 Neighbor Discovery Cache**.

You can view all static and dynamically discovered neighbors from the IPv6 Neighbor Discovery Cache pane.

Step 2 To clear all dynamically discovered neighbors from the cache, click **Clear Dynamic Neighbor Entries**.

The dynamically discovered neighbor is removed from the cache.

Note This procedure clears only dynamically discovered neighbors from the cache; it does not clear static neighbors.

Monitoring Routed and Transparent Mode Interfaces

You can monitor interface statistics, status, PPPoE, and more.

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

- **Monitoring > Interfaces > Interface Graphs**

  Lets you view interface statistics in graph or table form. If an interface is shared among contexts, the ASA shows only statistics for the current context. The number of statistics shown for a subinterface is a subset of the number of statistics shown for a physical interface.

- **Monitoring > Interfaces > Interface Graphs > Graph/Table**

  Shows a graph for the selected statistics. The Graph window can show up to four graphs and tables at a time. By default, the graph or table displays the real-time statistics. If you enable History Metrics, you can view statistics for past time periods.

Static Route Tracking

- **Monitoring > Interfaces > interface connection > Track Status**

  Displays information about the tracked object.

- **Monitoring > Interfaces > interface connection > Monitoring Statistics**

  Displays statistics for the SLA monitoring process.

PPPoE

- **Monitoring > Interfaces > PPPoE Client > PPPoE Client Lease Information**

  Displays information about current PPPoE connections.

Dynamic ACLs

Monitoring > Interfaces > Dynamic ACLs

Shows a table of the Dynamic ACLs, which are functionally identical to the user-configured ACLs except that they are created, activated and deleted automatically by the ASA. These ACLs do not show up in the configuration and are only visible in this table. They are identified by the "(dynamic)" keyword in the ACL header.

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. We introduced the following screens: Monitoring > Interfaces > IPv6 Neighbor Discovery Cache. Configuration > Device Management > Advanced > IPv6 Neighbor Discovery Cache. Configuration > Device Setup > Interface Settings > Interfaces > IPv6.
IPv6 support for transparent mode	8.2(1)	IPv6 support was introduced for transparent firewall mode.
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 modified or introduced the following screens: **Configuration > Device Setup > Interface Settings > Interfaces** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface**
Address Config Flags for IPv6 DHCP Relay	9.0(1)	We modified the following screen: Configuration > Device Setup > Interfaces > IPv6.

Feature Name	Platform Releases	Feature Information
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. We modified the following screens: **Configuration > Device Setup > Interface Settings > Interfaces** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Bridge Group Interface** **Configuration > Device Setup > Interface Settings > Interfaces > Add/Edit Interface**