

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, on page 1
- Guidelines and limitations for routed and transparent mode interfaces, on page 3
- Configure Routed Mode Interfaces, on page 5
- Configure Bridge Group Interfaces, on page 10
- Configure IPv6 Addressing, on page 15
- Monitoring Routed and Transparent Mode Interfaces, on page 29
- Examples for Routed and Transparent Mode Interfaces, on page 34
- History for Routed and Transparent Mode Interfaces, on page 37

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 supports both routed and bridged interfaces, and you can route between routed interfaces and BVIs. 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.

Whether you assign a security level to a BVI depends on the firewall mode. In transparent mode, the BVI interface does not have a security level because it does not participate in routing between interfaces. In routed mode, BVI interfaces have a security level if you choose to route between the BVIs and other interfaces. For routed mode, the security level on a bridge group member interface only applies for communication within the bridge group. Similarly, the BVI security level only applies for inter-BVI/Layer 3 interface communication.

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)

A dual IP stack is a network configuration that enables both IPv6 and IPv4 addresses on the same interface on ASA. Make sure you configure a default route for both IPv4 and IPv6.

31-bit subnet masks

A 31-bit subnet mask includes only two addresses for point-to-point connections.

31-bit subnet mask implementation

For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. Normally, the first and last address in a 2-address subnet is reserved for the network and broadcast, so a 2-address subnet is not usable. However, if you have a point-to-point connection and do not need network or broadcast addresses, a 31-bit subnet is a useful way to preserve addresses in IPv4.

31-bit subnet mask use cases

The failover link between two ASAs only requires two addresses. Any packet that is transmitted by one end of the link is always received by the other, and broadcasting is unnecessary. You can also have a directly-connected management station running SNMP or Syslog.

Note: 31-bit subnet and clustering compatibility

You can use a 31-bit subnet mask in Spanned clustering mode, excluding the management interface and the cluster control link.

You cannot use a 31-bit subnet mask in individual clustering mode on any interface.

31-bit subnet and failover

A ASA using a 31-bit subnet for its interface IP address can't configure a standby IP address due to insufficient addresses.

Interface monitoring limitations

When using a 31-bit subnet for the ASA interface IP address, you cannot configure a standby IP address for the interface because there are not enough addresses. Normally, an interface for failover should have a standby IP address so the active unit can perform interface tests to ensure standby interface health. Without a standby IP address, the ASA cannot perform any network tests; only the link state can be tracked.

For the failover and optional separate state link, which are point-to-point connections, you can also use a 31-bit subnet.

31-bit subnet and management

If you have a directly-connected management station, you can use a point-to-point connection for SSH or HTTP on the ASA, or for SNMP or Syslog on the management station.

Recommendation: 31-bit subnet unsupported features

The 31-bit subnet does not support certain network features that require multiple host addresses.

- BVI interfaces for bridge groups—The bridge group requires at least 3 host addresses: the BVI, and two hosts connected to two bridge group member interfaces. you must use a /29 subnet or smaller.
- Multicast Routing

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.
- DHCPv6 and prefix delegation options are not supported with multiple context mode.
- In routed firewall mode, bridge group interfaces are not supported in multiple context mode.

Failover, Clustering

Do not configure failover links with the procedures in this chapter. See the Failover chapter for more information.

- For cluster interfaces, see the clustering chapter for requirements.
- 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.
- DHCPv6 and prefix delegation options are not supported with multiple context mode, transparent mode, clustering, or Failover.

Model guidelines

- For the ASAv50, bridge groups are not supported in either transparent or routed mode.

VLAN IDs for the ASASM

- 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 64 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.
- For the ASAv50 on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- For the 200/ 1010/ 1210// 1220, you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.
- 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.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, ASA-defined EtherChannel and VNI interfaces are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.
- 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.

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).

- Interface problems, such as frequent up/down status changes, can prevent the floating connection timer from applying correctly to the connections going through the interface. If you have problems with an interface's status, consider clearing all connections after the status becomes stable to clear invalid connections.

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.

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:

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

Note For the Firepower 1010, you cannot configure switch ports as routed mode interfaces.

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, and for loopback interfaces, 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. For point-to-point connections, you can specify a 31-bit subnet mask (255.255.255.254). In this case, no IP addresses are reserved for the network or broadcast addresses. You cannot set the standby IP address in this case.

Example:

```
ciscoasa(config-if)# ip address 10.1.1.0 255.255.255.254
```

- 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)..

Note For loopback interfaces, you do not set the security level because the interface is only supported for to/from the device traffic.

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.

Note For loopback interfaces, you do not set the management mode because the interface is only supported for to/from the device traffic.

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 15
Enable the Physical Interface and Configure Ethernet Parameters
Configure PPPoE, on page 9

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 Bridge Group Interfaces

A bridge group is a group of interfaces that the ASA bridges instead of routes. Bridge groups are supported in both transparent and routed firewall mode. For more information about bridge groups, see 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.

For routed mode, if you provide a name for the BVI, then the BVI participates in routing. Without a name, the bridge group remains isolated as in transparent firewall mode.

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 (Transparent Mode) 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.

Step 3 (Routed Mode) Set the IP address using one of the following methods.

For failover and clustering, you must set the IP address manually; DHCP is 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.

  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.

  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 (Routed Mode) Name the interface:

nameif *name*

Example:

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

You must name the BVI if you want to route traffic outside the bridge group members, for example, to the outside interface or to members of other bridge groups. 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 5 (Routed Mode) 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).

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

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, and EtherChannels. The Management interface is not supported. In routed mode, EtherChannels and VNIs are 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 14 to configure the Management interface.

Procedure

Step 1 Enter interface configuration mode:

interface *id*

Example:

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

The interface ID can be:

- **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.subinterface*—For example, **gigabitethernet0/0.100** or **port-channel1.100**.
- **vni**
- **vlan**
- *mapped_name*—For multiple context mode.



Note For the Firepower 1010, you cannot configure switch ports as bridge group members.

You cannot mix logical VLAN interfaces and physical router interfaces in the same bridge group.



Note In routed mode, the **port-channel** and **vni** interfaces are not supported as bridge group members.

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 64 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 4100/9300 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.

The Management interface is for to- and from-the-box traffic only and cannot pass through traffic.

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.
- For the Firepower 4100/9300 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.

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

For the Firepower 4100/9300 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

IPv6 addressing provides 128 bits of network addressing, which provides more than enough globally unique IP addresses. It supports both public global addresses and private link-local addresses, and

IPv6 address types

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.

Address verification process

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 this 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.

IPv6 prefix delegation client configuration

The ASA can act as a DHPCv6 Prefix Delegation client so that the client interface, for example the outside interface connected to a cable modem, can receive one or more IPv6 prefixes that the ASA can then subnet and assign to its inside interfaces.

IPv6 prefix delegation

IPv6 prefix delegation enables the ASA to function as a Dynamic Host Control Protocol for IPv6 (DHCPv6) prefix delegation client and receive one or more IPv6 prefixes. This feature allows the ASA to subnet and assign these prefixes to its inside interfaces. It enable the hosts connected to inside interfaces to use Stateless Address Auto Configuration (SLAAC) to obtain global IPv6 addresses.

Client and server behavior

The client interface, such as the outside interface connected to a cable modem, receives IPv6 prefixes that the ASA can then subnet and assign to its inside interfaces. The inside ASA interfaces do not in turn act as Prefix Delegation servers; the ASA can only provide global IP addresses to SLAAC clients.

If a router is connected to the ASA, it can act as a SLAAC client to obtain its IP address. However, if you want to use a subnet of the delegated prefix for the networks behind the router, you must manually configure those addresses on the router's inside interfaces.

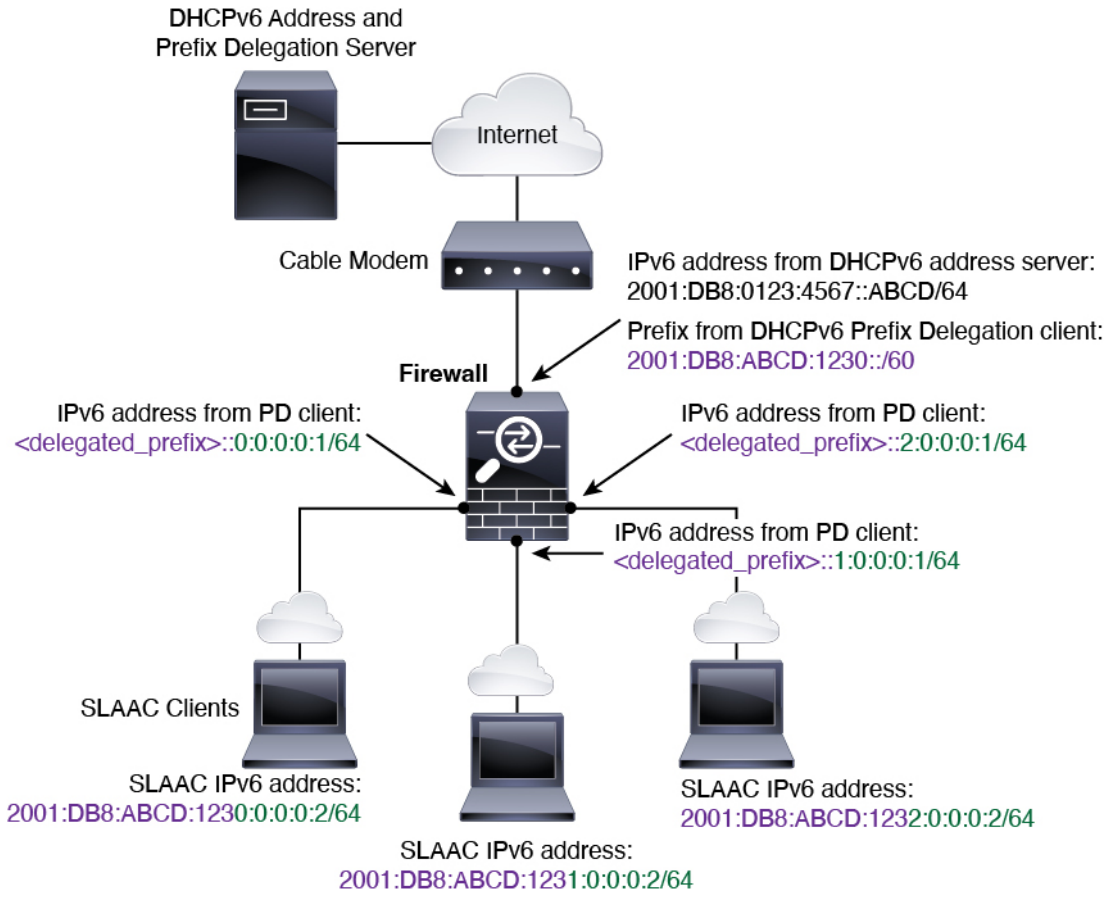
The ASA includes a light DHCPv6 server. This server provides information, such as the DNS server and domain name, to SLAAC clients when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients. You will configure the client to generate its own IPv6 address by enabling IPv6 autoconfiguration on the client. Enabling stateless autoconfiguration on a client configures IPv6 addresses based on prefixes received in Router Advertisement messages. In other words, the addresses are based on the prefix that the ASA received using Prefix Delegation.

IPv6 prefix delegation /64 subnet example

Here is an IPv6 prefix delegation /64 subnet example that demonstrates how a device subnets delegated prefixes into /64 networks and assigns global IPv6 addresses to interfaces.

IPv6 prefix delegation /64 subnet configuration

This example shows the ASA receiving an IP address on the outside interface using the DHCPv6 address client. It also gets a delegated prefix using the DHCPv6 Prefix Delegation client. The ASA subnets the delegated prefix into /64 networks and assigns global IPv6 addresses to its inside interfaces dynamically using the delegated prefix plus a manually configured subnet (::0, ::1, or ::2) and IPv6 address (0:0:0:1) per interface. SLAAC clients connected to those inside interfaces obtain IPv6 addresses on each /64 subnet.

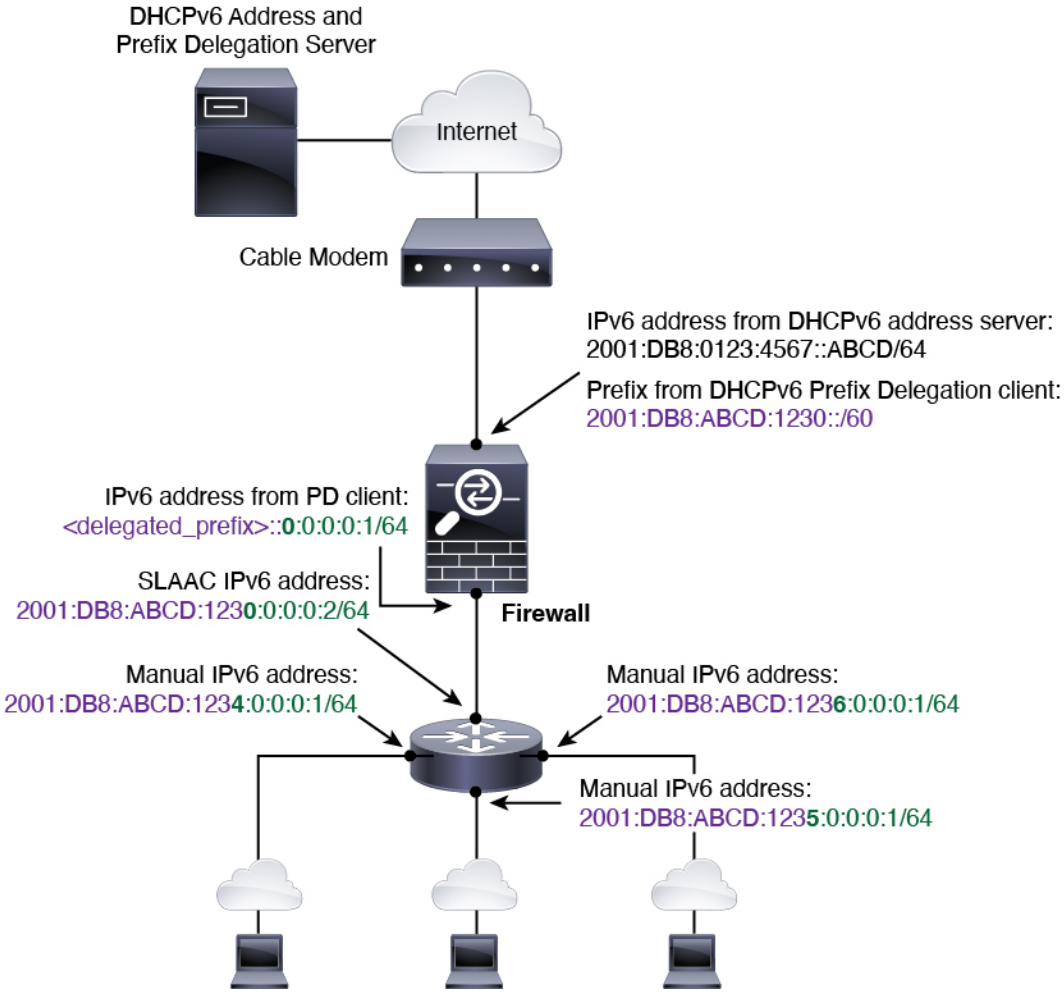


IPv6 prefix delegation /62 subnet examples

IPv6 prefix delegation /62 subnet enables subnetting of prefixes into 4 /62 subnets for downstream router delegation. It allows manual configuration of /64 subnets on router interfaces when dynamic prefix assignment is not available.

IPv6 /62 subnet delegation configuration

This example shows the ASA subnetting the prefix into 4 /62 subnets: 2001:DB8:ABCD:123**0**::/62, 2001:DB8:ABCD:123**4**::/62, 2001:DB8:ABCD:123**8**::/62, and 2001:DB8:ABCD:123**C**::/62. The ASA uses one of 4 available /64 subnets on 2001:DB8:ABCD:1230::/62 for its inside network (::0). You can then manually use additional /62 subnets for downstream routers. The router shown uses 3 of 4 available /64 subnets on 2001:DB8:ABCD:1234::/62 for its inside interfaces (::4, ::5, and ::6). In this case, the inside router interfaces cannot dynamically obtain the delegated prefix, so you need to view the delegated prefix on the ASA, and then use that prefix for your router configuration. Usually, ISPs delegate the same prefix to a given client when the lease expires, but if the ASA receives a new prefix, you will have to modify the router configuration to use the new prefix. The DHCP unique identifier (DUID) is persistent across reboots.



Enable the IPv6 Prefix Delegation Client

Enable the DHCPv6 Prefix Delegation client on one or more interfaces. The ASA obtains one or more IPv6 prefixes that it can subnet and assign to inside networks. Typically, the interface on which you enable the prefix delegation client obtains its IP address using the DHCPv6 address client; only other ASA interfaces use addresses derived from the delegated prefix.

Before you begin

- This feature is only supported in routed firewall mode.
- This feature is not supported in multiple context mode.
- This feature is not supported in clustering.
- You cannot configure this feature on a management-only interface.
- When you use Prefix Delegation, you must set the ASA IPv6 neighbor discovery router advertisement interval to be much lower than the preferred lifetime of the prefix assigned by the DHCPv6 Server to prevent IPv6 traffic interruption. For example, if the DHCPv6 server sets the preferred Prefix Delegation

lifetime to 300 seconds, you should set the ASA RA interval to be 150 seconds. To set the preferred lifetime, use the **show ipv6 general-prefix** command. To set the ASA RA interval, see Configure IPv6 Neighbor Discovery, on page 24; the default is 200 seconds.

Procedure

Step 1 Enter interface configuration mode for the interface connected to the DHCPv6 server network:

interface *id*

Example:

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

Step 2 Enable the DHCPv6 Prefix Delegation client, and name the prefix(es) obtained on this interface:

ipv6 dhcp client pd *name*

Example:

```
ciscoasa(config-if)# ipv6 dhcp client pd Outside-Prefix
```

The *name* can be up to 200 characters.

Step 3 Provide one or more hints about the delegated prefix you want to receive:

ipv6 dhcp client pd hint *ipv6_prefix/prefix_length*

Example:

```
ciscoasa(config-if)# ipv6 dhcp client pd hint 2001:DB8:ABCD:1230::/60
```

Typically you want to request a particular prefix length, such as ::/60, or if you have received a particular prefix before and want to ensure you get it again when the lease expires, you can enter the whole prefix as the hint. If you enter multiple hints (different prefixes or lengths), then it is up to the DHCP server which hint to honor, or whether to honor the hint at all.

Step 4 See Configure a Global IPv6 Address, on page 21 to assign a subnet of the prefix as the global IP address for an ASA interface.

Step 5 (Optional) See Configure the DHCPv6 Stateless Server to provide domain-name and server parameters to SLAAC clients.

Step 6 (Optional) See Configure IPv6 Network Settings to advertise the prefix(es) with BGP.

Example

The following example configures the DHCPv6 address client and prefix delegation client on GigabitEthernet 0/0, then assigns addresses with the prefix on GigabitEthernet 0/1 and 0/2:

```
interface gigabitethernet 0/0
  ipv6 address dhcp default
  ipv6 dhcp client pd Outside-Prefix
  ipv6 dhcp client pd hint ::/60
interface gigabitethernet 0/1
  ipv6 address Outside-Prefix ::1:0:0:0:1/64
interface gigabitethernet 0/2
  ipv6 address Outside-Prefix ::2:0:0:0:1/64
```

Configure a Global IPv6 Address

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

DHCPv6 and prefix delegation options are not supported with multiple context mode.

Note Configuring the global address automatically configures the link-local address, so you do not need to configure it separately. For bridge groups, configuring the global address on the BVI automatically configures link-local addresses on all member interfaces.

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 or for a bridge group in routed 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.

For failover and clustering, and for loopback interfaces, you must set the IP address manually. For clustering, manually configuring the link-local address is not supported.

- Enable StateLess Address Auto Configuration (SLAAC) on the interface:

  ipv6 address autoconfig [**default trust** {**dhcp** | **ignore**}]

  Enabling SLAAC 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 whenSLAAC is enabled.

Note Although RFC 4862 specifies that hosts configured for SLAAC 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.

  If you want to install a default route, specify **default trust dhcp** or **ignore**. **dhcp** specifies the ASA only uses a default route from Router Advertisements that come from a trusted source (in other words, from the same server that provided the IPv6 address). **ignore** specifies that Router Advertisements can be sourced from another network, which can be a riskier method.

- Obtain an address using DHCPv6:

  ipv6 address dhcp [**default**]

  Example:

  ```
  ciscoasa(config-if)# ipv6 address dhcp default
  ```

  The **default** keyword obtains a default route from Router Advertisements.

- 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.

- Use a delegated prefix:

  ipv6 address *prefix_name ipv6_address/prefix_length*

  Example:

  ```
  ciscoasa(config-if)# ipv6 address Outside-Prefix ::1:0:0:0:1/64
  ```

  This feature requires the ASA to have the DHCPv6 Prefix Delegation client enabled *on a different interface*. See Enable the IPv6 Prefix Delegation Client, on page 19. Typically, the delegated prefix will be /60 or smaller so you can subnet to multiple /64 networks. /64 is the supported subnet length if you want to support SLAAC for connected clients. You should specify an address that completes the /60 subnet, for example ::1:0:0:0:1. Enter :: before the address in case the prefix is smaller than /60. For example, if the delegated prefix is 2001:DB8:1234:5670::/60, then the global IP address assigned to this interface is 2001:DB8:1234:5671::1/64. The prefix that is advertised in router advertisements is 2001:DB8:1234:5671::/64. In this example, if the prefix is smaller than /60, the remaining bits of the prefix will be 0's as indicated by the leading ::. For example, if the prefix is 2001:DB8:1234::/48, then the IPv6 address will be 2001:DB8:1234::1:0:0:0:1/64.

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, the ASA randomly adjusts the value that you set (jitter).

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 StateLess Address Auto Configuration (SLAAC) clients to use DHCPv6 to obtain an IPv6 address, in addition to the derived SLAAC 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 SLAAC 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. SLAAC 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 SLAAC 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. For the **default** prefix, this setting only applies to on-link prefixes. Off-link prefixes will still be advertised unless you specify **no-advertise** for a specific off-link prefix.
- **no-autoconfig**—Specifies that the prefix cannot be used for IPv6 SLAAC.
- **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 Set the Recursive DNS Server (RDNSS) option to advertise a list of DNS servers to IPv6 clients.

ipv6 nd ra dns server *dns_server_ipv6_address* [*lifetime_value*]

- *lifetime_value*—The maximum time, in seconds, that the DNS server will be used for name resolution, between 200 and 4294967295. If you set the value to 0, the entry will not be used. If you set the value to 4294967295, the entry will never expire. The default is 3 x the maximum RA interval (see the **ipv6 nd ra-interval** command). The value must be greater than or equal to the RA interval.

This option is useful for clients that use SLAAC instead of DHCPv6 for their addressing. Note that if you enabled the prefix delegation client on the ASA, you can alternatively pass along the DNS server information that the ASA received using the ASA's DHCPv6 stateless server. If you configure both methods, the client will receive both sets of servers.

You can set up to 8 DNS servers. The servers are advertised in the order you add them.

Example:

```
ciscoasa(config-if)# ipv6 nd ra dns server 2620:119:35::35
ciscoasa(config-if)# ipv6 nd ra dns server 2620:119:53::53
ciscoasa(config-if)# ipv6 nd ra dns server 2606:4700:4700::1111
ciscoasa(config-if)# ipv6 nd ra dns server 2606:4700:4700::1001
```

Step 12 Set the DNS Search List (DNSSL) option to advertise a list of search domains to IPv6 clients.

ipv6 nd ra dns-search-list domain *domain_name* [**lifetime** *lifetime_value*

- *lifetime_value*—The maximum time, in seconds, that the domain will be used, between 200 and 4294967295. If you set the value to 0, the entry will not be used. If you set the value to 4294967295, the entry will never expire. The default is 3 x the maximum RA interval (see the **ipv6 nd ra-interval** command). The value must be greater than or equal to the RA interval.

This option is useful for clients that use SLAAC instead of DHCPv6 for their addressing. Note that if you enabled the prefix delegation client on the ASA, you can alternatively pass along the DNS domain that the ASA received using the ASA's DHCPv6 stateless server. If you configure both methods, the client will receive both sets of domains.

You can set up to 5 domains. The domains are advertised in the order you add them.

Example:

```
ciscoasa(config-if)# ipv6 nd ra dns-search-list domain cisco.com
ciscoasa(config-if)# ipv6 nd ra dns-search-list domain example.com
```

Step 13 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 Firepower and Secure Firewall models, 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.

DHCP Information

- **show ipv6 dhcp interface** [*ifc_name* [**statistics**]]

  The **show ipv6 dhcp interface** command displays DHCPv6 information for all interfaces. If the interface is configured for DHCPv6 stateless server configuration (see Configure the DHCPv6 Stateless Server), this command lists the DHCPv6 pool that is being used by the server. If the interface has DHCPv6 address client or Prefix Delegation client configuration, this command shows the state of each client and the values received from the server. For a specific interface, you can show message statistics for the DHCP server or client. The following examples show information provided by this command:

```
ciscoasa(config-if)# show ipv6 dhcp interface
GigabitEthernet1/1 is in server mode
  Using pool: Sample-Pool

GigabitEthernet1/2 is in client mode
```

```
  Prefix State is OPEN
  Renew will be sent in 00:03:46
  Address State is OPEN
  Renew for address will be sent in 00:03:47
  List of known servers:
    Reachable via address: fe80::20c:29ff:fe96:1bf4
    DUID: 000100011D9D1712005056A07E06
    Preference: 0
    Configuration parameters:
      IA PD: IA ID 0x00030001, T1 250, T2 400
        Prefix: 2005:abcd:ab03::/48
                preferred lifetime 500, valid lifetime 600
                expires at Nov 26 2014 03:11 PM (577 seconds)
      IA NA: IA ID 0x00030001, T1 250, T2 400
        Address: 2004:abcd:abcd:abcd:abcd:abcd:f2cb/128
                preferred lifetime 500, valid lifetime 600
                expires at Nov 26 2014 03:11 PM (577 seconds)
      DNS server: 2004:abcd:abcd:abcd::2
      DNS server: 2004:abcd:abcd:abcd::4
      Domain name: relay.com
      Domain name: server.com
      Information refresh time: 0
  Prefix name: Sample-PD

Management1/1 is in client mode
  Prefix State is IDLE
  Address State is OPEN
  Renew for address will be sent in 11:26:44
  List of known servers:
    Reachable via address: fe80::4e00:82ff:fe6f:f6f9
    DUID: 000300014C00826FF6F8
    Preference: 0
    Configuration parameters:
      IA NA: IA ID 0x000a0001, T1 43200, T2 69120
        Address: 2308:2308:210:1812:2504:1234:abcd:8e5a/128
                preferred lifetime INFINITY, valid lifetime INFINITY
      Information refresh time: 0


ciscoasa(config-if)# show ipv6 dhcp interface outside statistics

DHCPV6 Client PD statistics:

Protocol Exchange Statistics:


 Number of Solicit messages sent:              1
 Number of Advertise messages received:        1
 Number of Request messages sent:              1
 Number of Renew messages sent:                45
 Number of Rebind messages sent:               0
 Number of Reply messages received:            46
 Number of Release messages sent:              0
 Number of Reconfigure messages received:      0
 Number of Information-request messages sent:  0

Error and Failure Statistics:


 Number of Re-transmission messages sent:                 1
 Number of Message Validation errors in received messages: 0

DHCPV6 Client address statistics:
```

```
Protocol Exchange Statistics:


 Number of Solicit messages sent:              1
 Number of Advertise messages received:        1
 Number of Request messages sent:              1
 Number of Renew messages sent:                45
 Number of Rebind messages sent:               0
 Number of Reply messages received:            46
 Number of Release messages sent:              0
 Number of Reconfigure messages received:      0
 Number of Information-request messages sent:  0

Error and Failure Statistics:


 Number of Re-transmission messages sent:                 1
 Number of Message Validation errors in received messages: 0
```

- **show ipv6 dhcp client** [**pd**] **statistics**

  The **show ipv6 dhcp client statistics** command shows DHCPv6 client statistics and shows the output of the number of messages sent and received. The **show ipv6 dhcp client pd statistics** command shows the Prefix Delegation client statistics. The following examples show information provided by this command:

```
ciscoasa(config)# show ipv6 dhcp client statistics

Protocol Exchange Statistics:
  Total number of Solicit messages sent:              4
  Total number of Advertise messages received:        4
  Total number of Request messages sent:              4
  Total number of Renew messages sent:                92
  Total number of Rebind messages sent:               0
  Total number of Reply messages received:            96
  Total number of Release messages sent:              6
  Total number of Reconfigure messages received:      0
  Total number of Information-request messages sent:  0

Error and Failure Statistics:
  Total number of Re-transmission messages sent:                   8
  Total number of Message Validation errors in received messages: 0


ciscoasa(config)# show ipv6 dhcp client pd statistics

Protocol Exchange Statistics:


 Total number of Solicit messages sent:              1
 Total number of Advertise messages received:        1
 Total number of Request messages sent:              1
 Total number of Renew messages sent:                92
 Total number of Rebind messages sent:               0
 Total number of Reply messages received:            93
 Total number of Release messages sent:              0
 Total number of Reconfigure messages received:      0
 Total number of Information-request messages sent:  0

Error and Failure Statistics:
```

```
 Total number of Re-transmission messages sent:                  1
 Total number of Message Validation errors in received messages: 0
```

- **show ipv6 dhcp ha statistics**

  The **show ipv6 dhcp ha statistics** command shows the transaction statistics between failover units, including how many times the DUID information was synced between the units. The following examples show information provided by this command.

  On an active unit:

```
ciscoasa(config)# show ipv6 dhcp ha statistics

DHCPv6 HA global statistics:
  DUID sync messages sent:              1
  DUID sync messages received:          0

DHCPv6 HA error statistics:
  Send errors:                          0
```

  On an standby unit:

```
ciscoasa(config)# show ipv6 dhcp ha statistics

DHCPv6 HA global statistics:
  DUID sync messages sent:              0
  DUID sync messages received:          1

DHCPv6 HA error statistics:
  Send errors:                          0
```

- **show ipv6 general-prefix**

  The **show ipv6 general-prefix** command shows all the prefixes acquired by the DHCPv6 Prefix Delegation clients and the ASA distribution of that prefix to other processes ("Consumer List"). The following example shows information provided by this command:

```
ciscoasa(config)# show ipv6 general-prefix
IPv6 Prefix Sample-PD, acquired via DHCP PD
  2005:abcd:ab03::/48 Valid lifetime 524, preferred lifetime 424
   Consumer List               Usage count
    BGP network command        1
    inside (Address command)   1
```

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 commands.

- **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.

- **show ipv6 nd** {**ra dns server** | **ra dns-search-list** | **detail** | **summary**}

  This command displays IPv6 neighbor discovery parameters that you set in the configuration.

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:

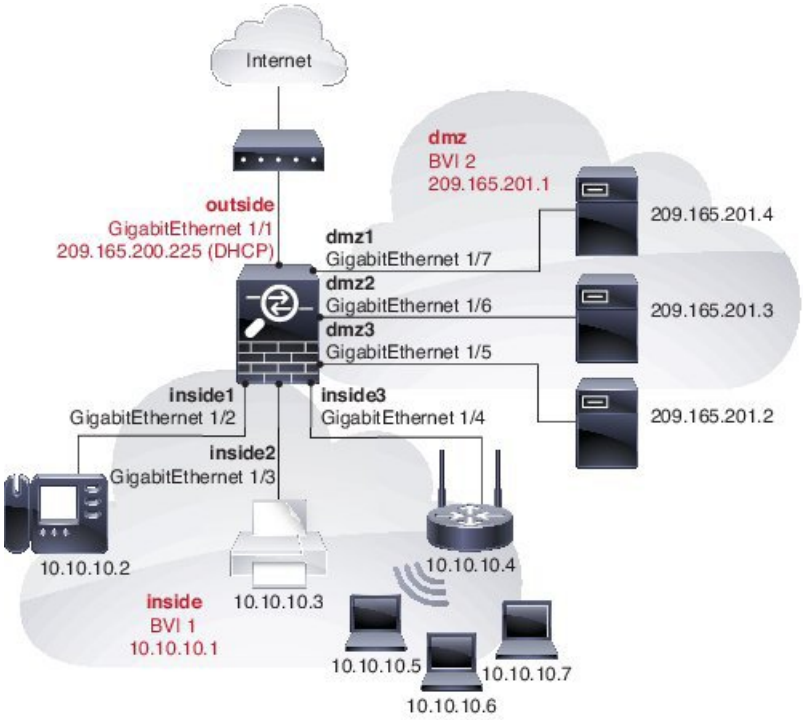
```
interface gigabitethernet 0/0
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 0/1
  nameif outside1
  security-level 0
  bridge-group 1
  no shutdown
interface gigabitethernet 0/2
  nameif dmz1
  security-level 50
  bridge-group 1
  no shutdown
interface 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
```

Switched LAN Segment Example with 2 Bridge Groups

The following example configures 2 bridge groups with 3 interfaces each and one regular routed interface for outside. Bridge group 1 is inside and bridge group 2 is dmz with public web servers. The bridge group member interfaces can communicate freely within the bridge group because each member is at the same security level, and we enabled same security communication. Although the inside member security level is 100 and the dmz member security level is also 100, these security levels do not apply to inter-BVI communications; only the BVI security levels affect inter-BVI traffic. The security levels of the BVIs and outside (100, 50, and 0) implicitly permit traffic from inside to dmz and inside to outside; and from dmz to outside. An access rule is applied to outside to allow traffic to the servers on dmz.




```
interface gigabitethernet 1/1
  nameif outside
  security-level 0
  ip address dhcp setroute
  no shutdown
!
interface gigabitethernet 1/2
  nameif inside1
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 1/3
  nameif inside2
  security-level 100
  bridge-group 1
  no shutdown
interface gigabitethernet 1/4
  nameif inside3
  security-level 100
  bridge-group 1
  no shutdown
```

```
!
interface bvi 1
  nameif inside
  security-level 100
  ip address 10.10.10.1 255.255.255.0
!
interface gigabitethernet 1/5
  nameif dmz1
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/6
  nameif dmz2
  security-level 100
  bridge-group 2
  no shutdown
interface gigabitethernet 1/7
  nameif dmz3
  security-level 100
  bridge-group 2
  no shutdown
!
interface bvi 2
  nameif dmz
  security-level 50
  ip address 209.165.201.1 255.255.255.224
!
same-security-traffic permit inter-interface
!
# Assigns IP addresses to inside hosts
dhcpd address 10.10.10.2-10.10.10.200 inside
dhcpd enable inside
!
# Applies interface PAT for inside traffic going outside
nat (inside1,outside) source dynamic any interface
nat (inside2,outside) source dynamic any interface
nat (inside3,outside) source dynamic any interface
!
# Allows outside traffic to each server for specific applications
object network server1
  host 209.165.201.2
object network server2
  host 209.165.201.3
object network server3
  host 209.165.201.4
!
# Defines mail services allowed on server3
object-group service MAIL
  service-object tcp destination eq pop3
  service-object tcp destination eq imap4
  service-object tcp destination eq smtp
!
# Allows access from outside to servers on the DMZ
access-list SERVERS extended permit tcp any object server1 eq www
access-list SERVERS extended permit tcp any object server2 eq ftp
access-list SERVERS extended permit tcp any object server3 object-group MAIL
access-group SERVERS in interface outside
```

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 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.
Bridge groups for transparent mode	8.4(1)	If you do not want the overhead of security contexts, or want to maximize your use of security contexts, you can group interfaces together in a bridge group, and then configure multiple bridge groups, one for each network. Bridge group traffic is isolated from other bridge groups. You can configure up to eight bridge groups of four interfaces each in single mode or per context. We introduced the following commands: **interface bvi**, **show bridge-group**
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. We modified the following commands: **interface bvi**, **bridge-group**
Transparent mode maximum interfaces per bridge group increased to 64	9.6(2)	The maximum interfaces per bridge group was increased from 4 to 64. We did not modify any commands.

Feature Name	Platform Releases	Feature Information
IPv6 DHCP	9.6(2)	The ASA now supports the following features for IPv6 addressing: • DHCPv6 Address client—The ASA obtains an IPv6 global address and optional default route from the DHCPv6 server. • DHCPv6 Prefix Delegation client—The ASA obtains delegated prefix(es) from a DHCPv6 server. The ASA can then use these prefixes to configure other ASA interface addresess so that StateLess Address Auto Configuration (SLAAC) clients can autoconfigure IPv6 addresses on the same network. • BGP router advertisement for delegated prefixes • DHCPv6 stateless server—The ASA provides other information such as the domain name to SLAAC clients when they send Information Request (IR) packets to the ASA. The ASA only accepts IR packets, and does not assign addresses to the clients. We added or modified the following commands: **clear ipv6 dhcp statistics, domain-name, dns-server, import, ipv6 address autoconfig, ipv6 address dhcp, ipv6 dhcp client pd, ipv6 dhcp client pd hint, ipv6 dhcp pool, ipv6 dhcp server, network, nis address, nis domain-name, nisp address, nisp domain-name, show bgp ipv6 unicast, show ipv6 dhcp, show ipv6 general-prefix, sip address, sip domain-name, sntp address**
Integrated Routing and Bridging	9.7(1)	Integrated Routing and Bridging provides the ability to route between a bridge group and a routed interface. A bridge group is a group of interfaces that the ASA bridges instead of routes. The ASA is not a true bridge in that the ASA continues to act as a firewall: access control between interfaces is controlled, and all of the usual firewall checks are in place. Previously, you could only configure bridge groups in transparent firewall mode, where you cannot route between bridge groups. This feature lets you configure bridge groups in routed firewall mode, and to route between bridge groups and between a bridge group and a routed interface. The bridge group participates in routing by using a Bridge Virtual Interface (BVI) to act as a gateway for the bridge group. Integrated Routing and Bridging provides an alternative to using an external Layer 2 switch if you have extra interfaces on the ASA to assign to the bridge group. In routed mode, the BVI can be a named interface and can participate separately from member interfaces in some features, such as access rules and DHCP server. The following features that are supported in transparent mode are not supported in routed mode: multiple context mode, ASA clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. We modified the following commands: **access-group, access-list ethertype, arp-inspection, dhcpd, mac-address-table static, mac-address-table aging-time, mac-learn, route, show arp-inspection, show bridge-group, show mac-address-table, show mac-learn**

Feature Name	Platform Releases	Feature Information
31-bit Subnet Mask	9.7(1)	For routed interfaces, you can configure an IP address on a 31-bit subnet for point-to-point connections. The 31-bit subnet includes only 2 addresses; normally, the first and last address in the subnet is reserved for the network and broadcast, so a 2-address subnet is not usable. However, if you have a point-to-point connection and do not need network or broadcast addresses, a 31-bit subnet is a useful way to preserve addresses in IPv4. For example, the failover link between 2 ASAs only requires 2 addresses; any packet that is transmitted by one end of the link is always received by the other, and broadcasting is unnecessary. You can also have a directly-connected management station running SNMP or Syslog. This feature is not supported for BVIs for bridge groups or with multicast routing. We modified the following commands: **ip address, http, logging host, snmp-server, ssh**
Recursive DNS Server (RDNSS) and DNS Search List (DNSSL) options to advertise a list of DNS servers and domains to IPv6 clients	9.24(1)	You can now configure Recursive DNS Server (RDNSS) and DNS Search List (DNSSL) options to provide DNS servers and domains to SLAAC clients using router advertisements. New/modified commands: **ipv6 nd ra dns-search-list domain**, **ipv6 nd ra dns server**, **show ipv6 nd detail**, **show ipv6 nd ra dns-search-list**, **show ipv6 nd ra dns server**, **show ipv6 nd summary**