

Transparent or Routed Firewall Mode

This chapter describes how to set the firewall mode to routed or transparent, as well as how the firewall works in each firewall mode.

You can set the firewall mode independently for each context in multiple context mode.

- Firewall mode, on page 1
- Default settings, on page 10
- Guidelines for firewall mode, on page 10
- Set the Firewall Mode, on page 11
- Examples for Firewall Mode, on page 12
- History for the Firewall Mode, on page 23

Firewall mode

A ASA device supports two firewall modes :

- Routed Firewall mode
- Transparent Firewall mode

Routed firewall mode

In routed mode, the ASA is considered to be a router hop in the network. Each interface that you want to route between is on a different subnet. You can share Layer 3 interfaces between contexts.

Integrated routing and bridging features

With Integrated Routing and Bridging, you can use a "bridge group" where you group together multiple interfaces on a network, and the ASA uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. The ASA routes between BVIs and regular routed interfaces. If you do not need multiple context mode or clustering or EtherChannel or VNI member interfaces, you might consider using routed mode instead of transparent mode. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

Transparent firewall mode

Traditionally, a firewall is a routed hop and acts as a default gateway for hosts that connect to one of its screened subnets. A transparent firewall, on the other hand, is a Layer 2 firewall that acts like a “bump in the wire,” or a “stealth firewall,” and is not seen as a router hop to connected devices. However, like any other firewall, access control between interfaces is controlled, and all of the usual firewall checks are in place.

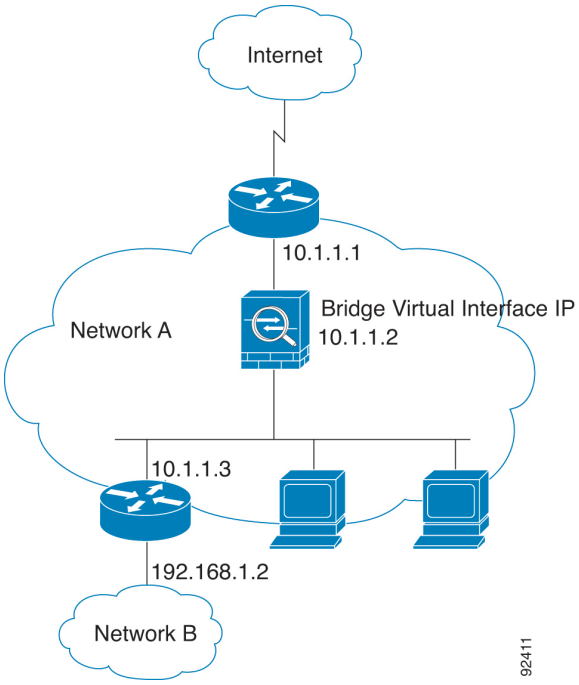
Layer 2 connectivity is achieved by using a "bridge group" where you group together the inside and outside interfaces for a network, and the ASA uses bridging techniques to pass traffic between the interfaces. Each bridge group includes a Bridge Virtual Interface (BVI) to which you assign an IP address on the network. You can have multiple bridge groups for multiple networks. In transparent mode, these bridge groups cannot communicate with each other.

Transparent firewall network usage

The ASA connects the same network between its interfaces. Because the firewall is not a routed hop, you can easily introduce a transparent firewall into an existing network.

The figure shows a typical transparent firewall network where the outside devices are on the same subnet as the inside devices. The inside router and hosts appear to be directly connected to the outside router.

Figure 1: Transparent firewall network



Management Interface

In addition to each Bridge Virtual Interface (BVI) IP address, you can add a separate Management *slot*/*port* interface that is not part of any bridge group, and that allows only management traffic to the ASA. For more information, see Management Interface.

Passing traffic for routed-mode features

For features that are not directly supported on the transparent firewall, you can allow traffic to pass through so that upstream and downstream routers can support the functionality.

Supported traffic types

For example, by using an access rule, you can allow DHCP traffic (instead of the unsupported DHCP relay feature) or multicast traffic such as that created by IP/TV. You can also establish routing protocol adjacencies through a transparent firewall; you can allow OSPF, RIP, EIGRP, or BGP traffic through based on an access rule. Likewise, protocols like HSRP or VRRP can pass through the ASA.

Bridge groups

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. Like any other firewall interfaces, access control between interfaces is controlled, and all of the usual firewall checks are in place.

Bridge virtual interfaces

Each bridge group includes a Bridge Virtual Interface (BVI). The ASA uses the BVI 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 bridge group member interfaces. The BVI does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

BVI functionality by mode

In transparent mode: Only bridge group member interfaces are named and can be used with interface-based features.

In routed mode: The BVI acts as the gateway between the bridge group and other routed interfaces. To route between bridge groups/routed interfaces, you must name the BVI. For some interface-based features, you can use the BVI itself:

- Access rules—You can configure access rules for both bridge group member interfaces and for the BVI; for inbound rules, the member interface is checked first. For outbound rules, the BVI is checked first.
- DHCPv4 server—Only the BVI supports the DHCPv4 server configuration.
- Static routes—You can configure static routes for the BVI; you cannot configure static routes for the member interfaces.
- Syslog server and other traffic sourced from the ASA—When specifying a syslog server (or SNMP server, or other service where the traffic is sourced from the ASA), you can specify either the BVI or a member interface.

If you do not name the BVI in routed mode, then the ASA does not route bridge group traffic. This configuration replicates transparent firewall mode for the bridge group. If you do not need multiple context mode or clustering or EtherChannel or VNI member interfaces, you might consider using routed mode instead. In routed mode, you can have one or more isolated bridge groups like in transparent mode, but also have normal routed interfaces as well for a mixed deployment.

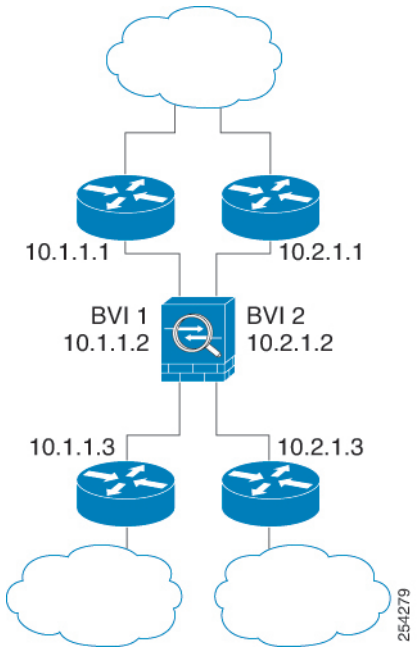
Bridge groups in transparent firewall mode

Bridge group traffic is isolated from other bridge groups; traffic is not routed to another bridge group within the ASA , and traffic must exit the ASA before it is routed by an external router back to another bridge group in the ASA. Although the bridging functions are separate for each bridge group, many other functions are shared between all bridge groups. For example, all bridge groups share a syslog server or AAA server configuration. For complete security policy separation, use security contexts with one bridge group in each context.

You can include multiple interfaces per bridge group. See Guidelines for firewall mode, on page 10 for the exact number of bridge groups and interfaces supported. If you use more than 2 interfaces per bridge group, you can control communication between multiple segments on the same network, and not just between inside and outside. For example, if you have three inside segments that you do not want to communicate with each other, you can put each segment on a separate interface, and only allow them to communicate with the outside interface. Or you can customize the access rules between interfaces to allow only as much access as desired.

This figure shows two networks connected to the ASA, which has two bridge groups.

Figure 2: Transparent firewall network with two bridge groups



Bridge groups in routed firewall mode

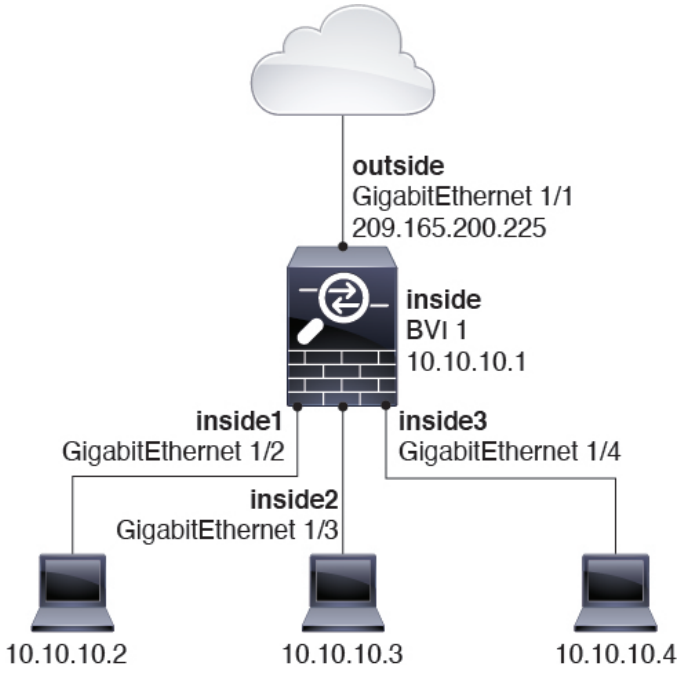
Bridge group traffic can be routed to other bridge groups or routed interfaces. You can choose to isolate bridge group traffic by not assigning a name to the BVI interface for the bridge group. If you name the BVI, then the BVI participates in routing like any other regular interface.

Bridge group usage and configuration

One use for a bridge group in routed mode is to use extra interfaces on the ASA instead of an external switch. For example, the default configuration for some devices include an outside interface as a regular interface, and then all other interfaces assigned to the inside bridge group. Because the purpose of this bridge group is to replace an external switch, you need to configure an access policy so all bridge group interfaces can freely

communicate. For example, as in the default configuration, set all the interfaces to the same security level, and then enable same-security interface communication; no access rule is required.

Figure 3: Routed firewall network with an inside bridge group and an outside routed interface



Passing Traffic Not Allowed in Routed Mode

In routed mode, some types of traffic cannot pass through the ASA even if you allow it in an access rule. The bridge group, however, can allow almost any traffic through using either an access rule (for IP traffic) or an EtherType rule (for non-IP traffic):

- IP traffic—In routed firewall mode, broadcast and multicast traffic is blocked even if you allow it in an access rule, including unsupported dynamic routing protocols and DHCP (unless you configure DHCP relay). Within a bridge group, you can allow this traffic with an access rule (using an extended ACL).
- Non-IP traffic—AppleTalk, IPX, BPDUs, and MPLS, for example, can be configured to go through using an EtherType rule.

Note The bridge group does not pass CDP packets packets, or any packets that do not have a valid EtherType greater than or equal to 0x600. An exception is made for BPDUs and IS-IS, which are supported.

Allow layer 3 traffic

This topic describes how different types of Layer 3 traffic are handled through bridge groups and the access rule requirements for each traffic type.

- Unicast IPv4 and IPv6 traffic is allowed through the bridge group automatically from a higher security interface to a lower security interface, without an access rule.

- For Layer 3 traffic traveling from a low to a high security interface, an access rule is required on the low security interface.
- ARPs are allowed through the bridge group in both directions without an access rule. ARP traffic can be controlled by ARP inspection.
- IPv6 neighbor discovery and router solicitation packets can be passed using access rules.
- Broadcast and multicast traffic can be passed using access rules.

Allowed MAC addresses

These destination MAC addresses are allowed through the bridge group if allowed by your access policy (For more information, refer to Allow layer 3 traffic, on page 5). Any MAC address not on this list is dropped.

- TRUE broadcast destination MAC address equal to FFFF.FFFF.FFFF
- IPv4 multicast MAC addresses from 0100.5E00.0000 to 0100.5EFE.FFFF
- IPv6 multicast MAC addresses from 3333.0000.0000 to 3333.FFFF.FFFF
- BPDU multicast address equal to 0100.0CCC.CCCD
- AppleTalk multicast MAC addresses from 0900.0700.0000 to 0900.07FF.FFFF

BPDU Handling

To prevent loops using the Spanning Tree Protocol, BPDUs are passed by default. To block BPDUs, you need to configure an EtherType rule to deny them. You can also block BPDUs on the external switches. For example, you can block BPDUs on the switch if members of the same bridge group are connected to switch ports in different VLANs. In this case, BPDUs from one VLAN will be visible in the other VLAN, which can cause Spanning Tree Root Bridge election process problems.

If you are using failover, you might want to block BPDUs to prevent the switch port from going into a blocking state when the topology changes. See Bridge Group Requirements for Failover for more information.

MAC address and route lookups

For traffic within a bridge group, the outgoing interface of a packet is determined by performing a destination MAC address lookup instead of a route lookup.

Route lookups, however, are necessary for these situations:

- Traffic originating on the ASA—Add a default/static route on the ASA for traffic destined for a remote network where a syslog server, for example, is located.
- SSL policy enabled in transparent firewall mode—
  - TLS 1.3 traffic passes through the ASA device.
  - ASA acts as a TCP proxy. It terminates the client-side TCP connection and initiates a separate server-side connection. The ASA TCP stack owns the session state.
  - For idle proxied sessions, ASA generates TCP ACK keepalive probes to preserve the connection state. These are not passthrough packets—the firewall originates them and sends them toward the client or server depending on the state of the TLS transit connection.

- Because the ASA originates these keepalives, it must perform an adjacency table lookup (next-hop or egress MAC resolution) even in transparent mode.
- If the adjacency lookup fails for one of these proxy-generated packets, the ASA logs syslog 110003.

  To prevent these syslog messages, configure routes on the ASA for the transit connection endpoints (client and server IP addresses). Proper routing enables successful adjacency lookups for keepalive packets.

  If the ASA successfully sends a TCP keepalive and receives a reply, it keeps the transit connection active. If the device cannot send the keepalives, or receives no reply, it removes the transit connection after approximately 240 seconds.

- Voice over IP (VoIP) and TFTP traffic with inspection enabled, and the endpoint is at least one hop away—Add a static route on the ASA for traffic destined for the remote endpoint so that secondary connections are successful. The ASA creates a temporary "pinhole" in the access control policy to allow the secondary connection; and because the connection might use a different set of IP addresses than the primary connection, the ASA needs to perform a route lookup to install the pinhole on the correct interface.

  Affected applications include:

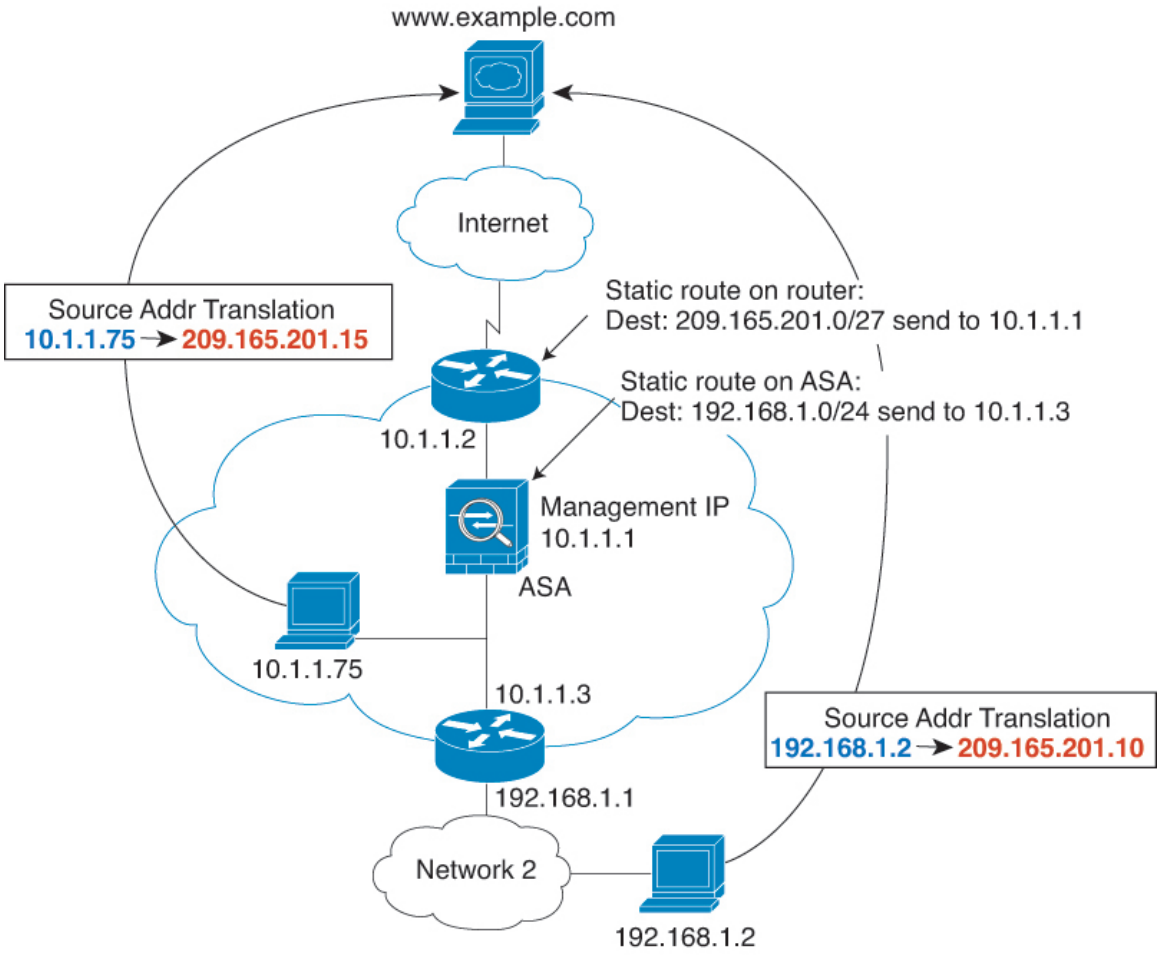
  - CTIQBE
  - GTP
  - H.323
  - MGCP
  - RTSP
  - SIP and Skinny (SCCP)
  - SQL*Net
  - SunRPC
  - TFTP

- Traffic at least one hop away for which the ASA performs NAT—Configure a static route on the ASA for traffic destined for the remote network. You also need a static route on the upstream router for traffic destined for the mapped addresses to be sent to the ASA.

  This routing requirement is also true for embedded IP addresses for VoIP and DNS with inspection and NAT enabled, and the embedded IP addresses are at least one hop away. The ASA needs to identify the correct egress interface so it can perform the translation.

NAT within a bridge group

Figure 4: NAT example: NAT within a bridge group



Unsupported features for bridge groups in transparent mode

This table lists the features that are not supported in bridge groups in transparent mode.

Table 1: Unsupported features in transparent mode

Feature	Description
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on bridge group member interfaces.

Feature	Description
DHCP relay	The transparent firewall can act as a DHCPv4 server, but it does not support DHCP relay. DHCP relay is not required because you can allow DHCP traffic to pass through using two access rules: one that allows DCHP requests from the inside interface to the outside, and one that allows the replies from the server in the other direction.
Dynamic routing protocols	You can, however, add static routes for traffic originating on the ASA for bridge group member interfaces. You can also allow dynamic routing protocols through the ASA using an access rule.
Multicast IP routing	You can allow multicast traffic through the ASA by allowing it in an access rule.
QoS	—
VPN termination for through traffic	The transparent firewall supports site-to-site VPN tunnels for management connections only on bridge group member interfaces. It does not terminate VPN connections for traffic through the ASA. You can pass VPN traffic through the ASA using an access rule, but it does not terminate non-management connections.
Unified Communications	—

Unsupported features for bridge groups in routed mode

This table lists the features are not supported in bridge groups in routed mode.

Table 2: Unsupported features in routed mode

Feature	Description
EtherChannel or VNI member interfaces	Only physical interfaces and subinterfaces are supported as bridge group member interfaces. Management interfaces are also not supported.
Clustering	Bridge groups are not supported in clustering.
Dynamic DNS	—
DHCPv6 stateless server	Only the DHCPv4 server is supported on BVIs.
DHCP relay	The routed firewall can act as a DHCPv4 server, but it does not support DHCP relay on BVIs or bridge group member interfaces.
Dynamic routing protocols	You can, however, add static routes for BVIs. You can also allow dynamic routing protocols through the ASA using an access rule. Non-bridge group interfaces support dynamic routing.
Multicast IP routing	You can allow multicast traffic through the ASA by allowing it in an access rule. Non-bridge group interfaces support multicast routing.

Feature	Description
Multiple Context Mode	Bridge groups are not supported in multiple context mode.
QoS	Non-bridge group interfaces support QoS.
VPN termination for through traffic	You cannot terminate a VPN connection on the BVI. Non-bridge group interfaces support VPN. Bridge group member interfaces support site-to-site VPN tunnels for management connections only. It does not terminate VPN connections for traffic through the ASA. You can pass VPN traffic through the bridge group using an access rule, but it does not terminate non-management connections. Clientless SSL VPN is also not supported.
Unified Communications	Non-bridge group interfaces support Unified Communications.

Default settings

Lists the default settings for firewall mode.

Default mode

The default mode is routed mode.

Bridge group defaults

By default, all ARP packets are passed within the bridge group.

Guidelines for firewall mode

Review these guidelines for firewall mode.

Context mode guidelines

Set the firewall mode per context.

Bridge group limits and network requirements

- Each directly-connected network must be on the same subnet.
- The Firewall Threat Defense device does not support traffic on secondary networks; only traffic on the same network as the BVI IP address is supported.

BVI IP address requirements

- 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 Firewall Threat Defense device. 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).

Unsupported bridge group members

- 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 Firepower 2100 series, bridge groups are not supported in routed mode.
- For the 200/ 1010/ , you cannot mix logical VLAN interfaces and physical firewall interfaces in the same bridge group.

Transparent mode requirements and limitations

- 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.

Routed mode requirements and limitations

- 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.

Traffic handling limitation

- 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.

Set the Firewall Mode

This section describes how to change the firewall mode.

Note We recommend that you set the firewall mode before you perform any other configuration because changing the firewall mode clears the running configuration.

Before you begin

When you change modes, the ASA clears the running configuration (see Guidelines for firewall mode, on page 10 for more information).

- If you already have a populated configuration, be sure to back up your configuration before changing the mode; you can use this backup for reference when creating your new configuration. See Back Up and Restore Configurations or Other Files.
- Use the CLI at the console port to change the mode. If you use any other type of session, including the ASDM Command Line Interface tool or SSH, you will be disconnected when the configuration is cleared, and you will have to reconnect to the ASA using the console port in any case.
- Set the mode within the context.

Note To set the firewall mode to transparent and also configure ASDM management access after the configuration is cleared, see Configure ASDM Access.

Procedure

Set the firewall mode to transparent:

firewall transparent

Example:

```
ciscoasa(config)# firewall transparent
```

To change the mode to routed, enter the **no firewall transparent** command.

Note
You are not prompted to confirm the firewall mode change; the change occurs immediately.

Examples for Firewall Mode

This section includes examples of how traffic moves through the ASA in the routed and transparent firewall mode.

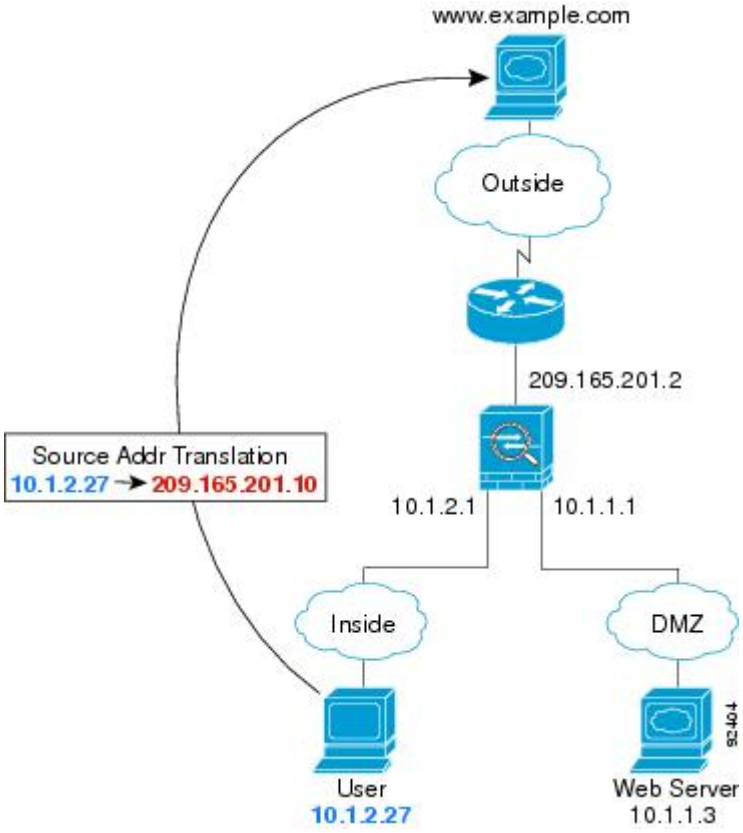
How Data Moves Through the ASA in Routed Firewall Mode

The following sections describe how data moves through the ASA in routed firewall mode in multiple scenarios.

An Inside User Visits a Web Server

The following figure shows an inside user accessing an outside web server.

Figure 5: Inside to Outside



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA translates the real address (10.1.2.27) to the mapped address 209.165.201.10, which is on the outside interface subnet.

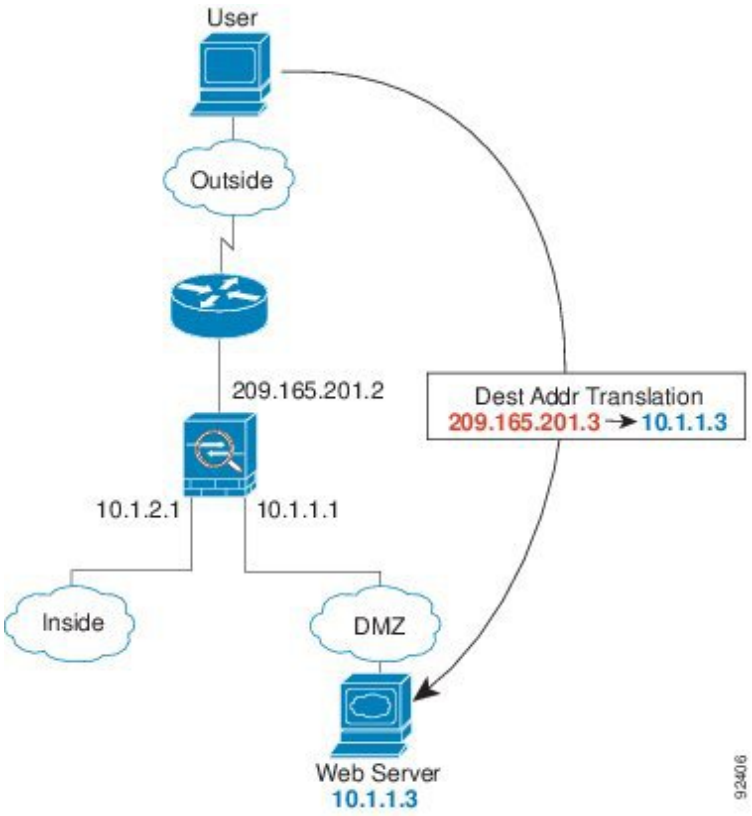
   The mapped address could be on any subnet, but routing is simplified when it is on the outside interface subnet.
4. The ASA then records that a session is established and forwards the packet from the outside interface.

5. When www.example.com responds to the request, the packet goes through the ASA, and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by untranslating the global destination address to the local user address, 10.1.2.27.
6. The ASA forwards the packet to the inside user.

An Outside User Visits a Web Server on the DMZ

The following figure shows an outside user accessing the DMZ web server.

Figure 6: Outside to DMZ



The following steps describe how data moves through the ASA:

1. A user on the outside network requests a web page from the DMZ web server using the mapped address of 209.165.201.3, which is on the outside interface subnet.
2. The ASA receives the packet and untranslates the mapped address to the real address 10.1.1.3.
3. Because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy.

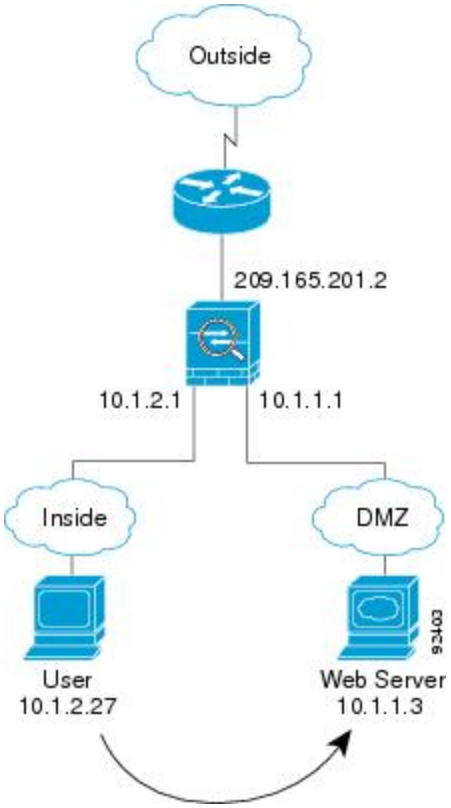
   For multiple context mode, the ASA first classifies the packet to a context.
4. The ASA then adds a session entry to the fast path and forwards the packet from the DMZ interface.

5. When the DMZ web server responds to the request, the packet goes through the ASA and because the session is already established, the packet bypasses the many lookups associated with a new connection. The ASA performs NAT by translating the real address to 209.165.201.3.
6. The ASA forwards the packet to the outside user.

An Inside User Visits a Web Server on the DMZ

The following figure shows an inside user accessing the DMZ web server.

Figure 7: Inside to DMZ



The following steps describe how data moves through the ASA:

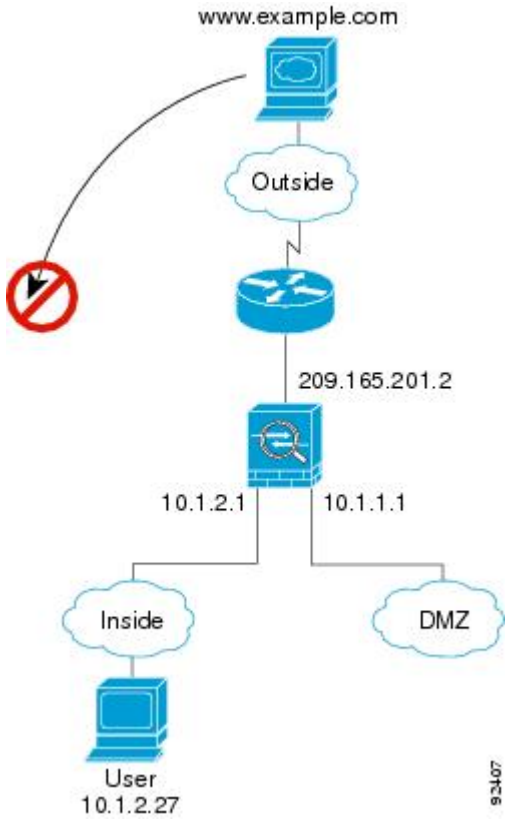
1. A user on the inside network requests a web page from the DMZ web server using the destination address of 10.1.1.3.
2. The ASA receives the packet and because it is a new session, the ASA verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA then records that a session is established and forwards the packet out of the DMZ interface.
4. When the DMZ web server responds to the request, the packet goes through the fast path, which lets the packet bypass the many lookups associated with a new connection.
5. The ASA forwards the packet to the inside user.

An Outside User Attempts to Access an Inside Host

The following figure shows an outside user attempting to access the inside network.

Figure 8: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network attempts to reach an inside host (assuming the host has a routable IP address).

   If the inside network uses private addresses, no outside user can reach the inside network without NAT. The outside user might attempt to reach an inside user by using an existing NAT session.

2. The ASA receives the packet and because it is a new session, it verifies if the packet is allowed according to the security policy.

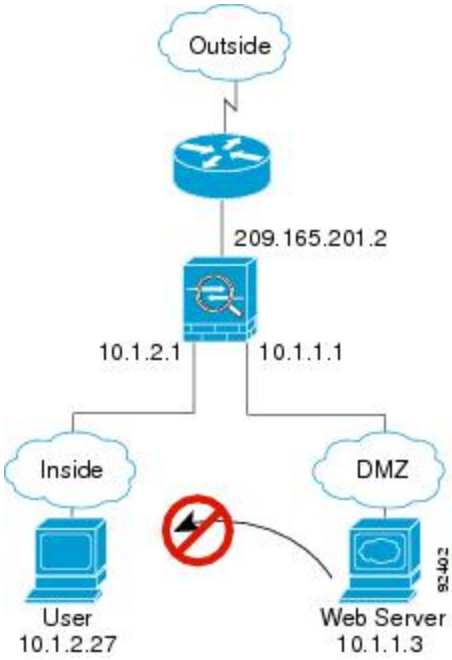
3. The packet is denied, and the ASA drops the packet and logs the connection attempt.

   If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

A DMZ User Attempts to Access an Inside Host

The following figure shows a user in the DMZ attempting to access the inside network.

Figure 9: DMZ to Inside



The following steps describe how data moves through the ASA:

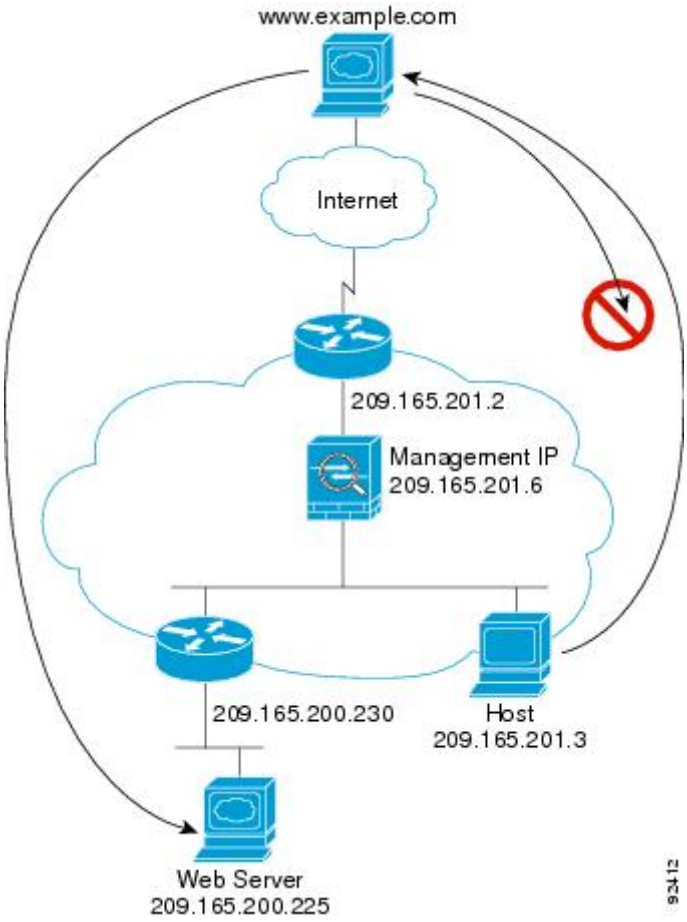
1. A user on the DMZ network attempts to reach an inside host. Because the DMZ does not have to route the traffic on the Internet, the private addressing scheme does not prevent routing.
2. The ASA receives the packet and because it is a new session, it verifies if the packet is allowed according to the security policy.

   The packet is denied, and the ASA drops the packet and logs the connection attempt.

How Data Moves Through the Transparent Firewall

The following figure shows a typical transparent firewall implementation with an inside network that contains a public web server. The ASA has an access rule so that the inside users can access Internet resources. Another access rule lets the outside users access only the web server on the inside network.

Figure 10: Typical Transparent Firewall Data Path

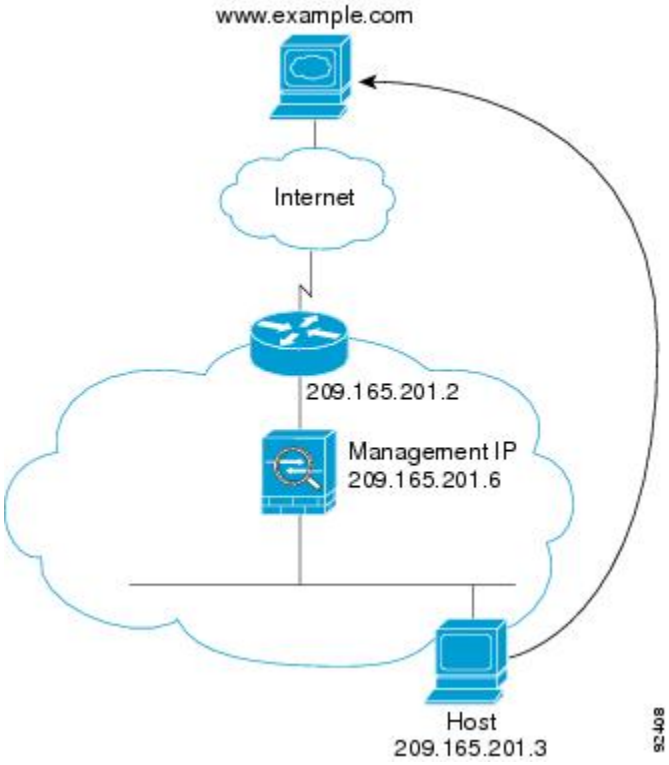


The following sections describe how data moves through the ASA.

An Inside User Visits a Web Server

The following figure shows an inside user accessing an outside web server.

Figure 11: Inside to Outside



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

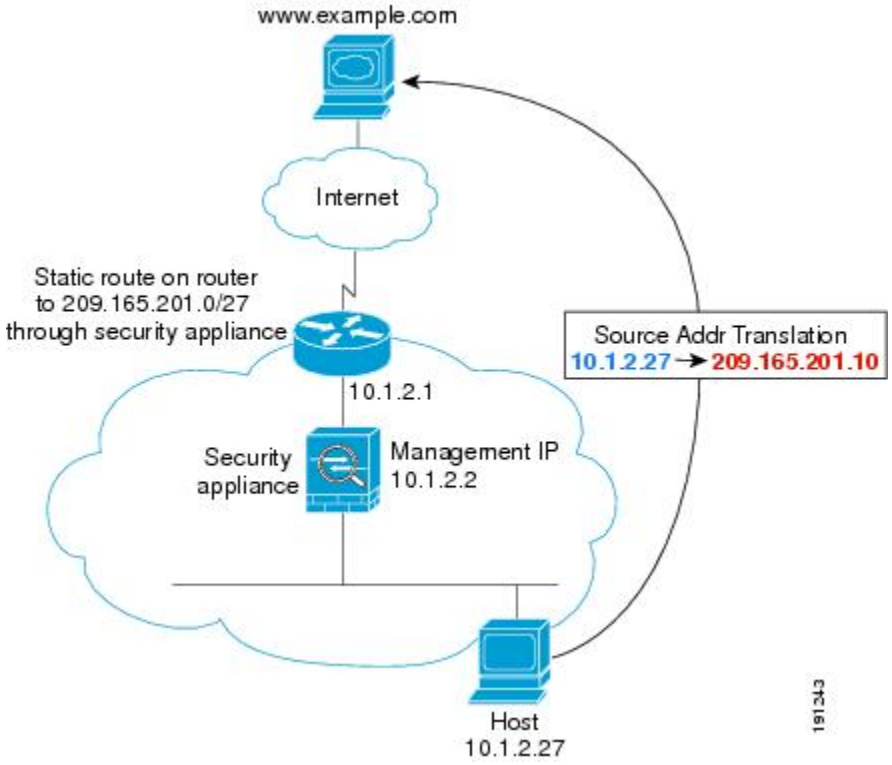
   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 209.165.201.2.

   If the destination MAC address is not in the ASA table, it attempts to discover the MAC address by sending an ARP request or a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
6. The ASA forwards the packet to the inside user.

An Inside User Visits a Web Server Using NAT

The following figure shows an inside user accessing an outside web server.

Figure 12: Inside to Outside with NAT



The following steps describe how data moves through the ASA:

1. The user on the inside network requests a web page from www.example.com.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet according to a unique interface.
3. The ASA translates the real address (10.1.2.27) to the mapped address 209.165.201.10.

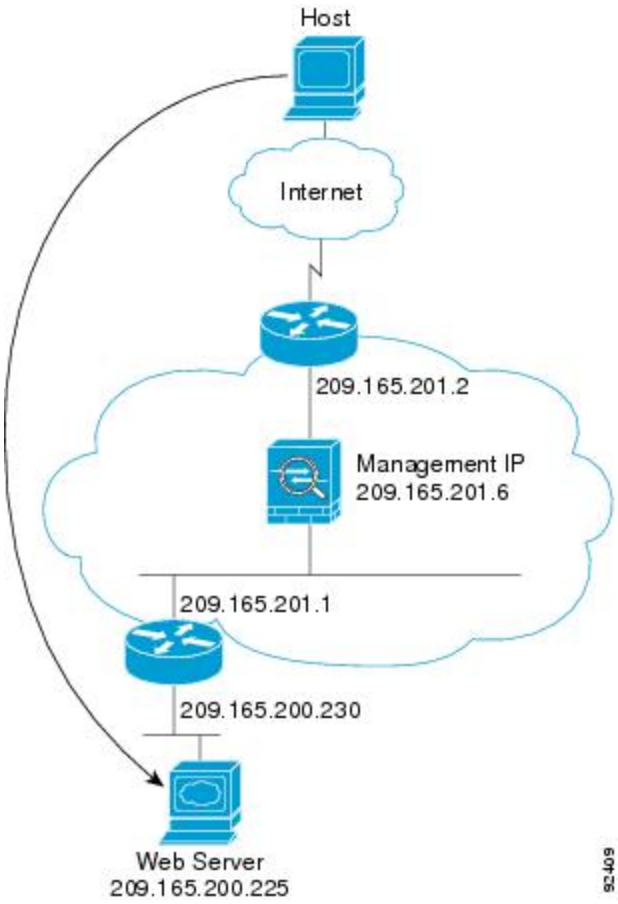
   Because the mapped address is not on the same network as the outside interface, then be sure the upstream router has a static route to the mapped network that points to the ASA.
4. The ASA then records that a session is established and forwards the packet from the outside interface.
5. If the destination MAC address is in its table, the ASA forwards the packet out of the outside interface. The destination MAC address is that of the upstream router, 10.1.2.1.

   If the destination MAC address is not in the ASA table, then it attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
6. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.
7. The ASA performs NAT by untranslating the mapped address to the real address, 10.1.2.27.

An Outside User Visits a Web Server on the Inside Network

The following figure shows an outside user accessing the inside web server.

Figure 13: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network requests a web page from the inside web server.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies that the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The ASA records that a session is established.
4. If the destination MAC address is in its table, the ASA forwards the packet out of the inside interface. The destination MAC address is that of the downstream router, 209.165.201.1.

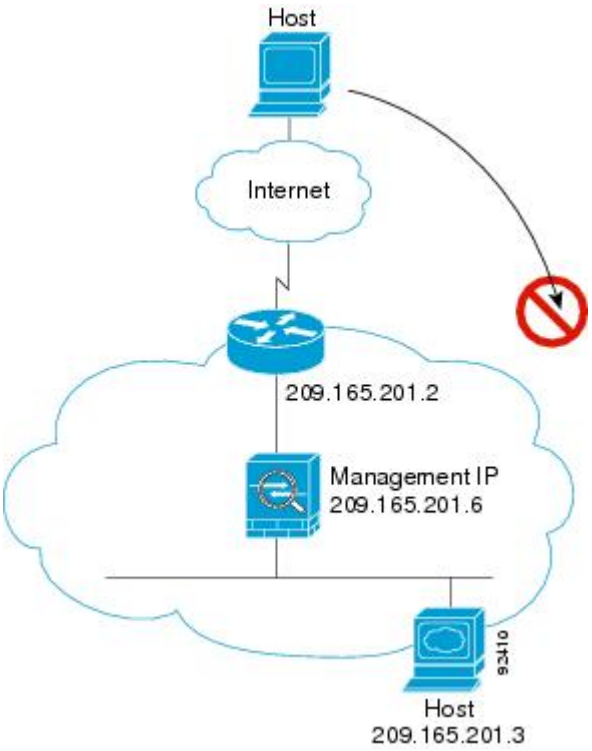
   If the destination MAC address is not in the ASA table, then it attempts to discover the MAC address by sending an ARP request and a ping. The first packet is dropped.
5. The web server responds to the request; because the session is already established, the packet bypasses the many lookups associated with a new connection.

6. The ASA forwards the packet to the outside user.

An Outside User Attempts to Access an Inside Host

The following figure shows an outside user attempting to access a host on the inside network.

Figure 14: Outside to Inside



The following steps describe how data moves through the ASA:

1. A user on the outside network attempts to reach an inside host.
2. The ASA receives the packet and adds the source MAC address to the MAC address table, if required. Because it is a new session, it verifies if the packet is allowed according to the terms of the security policy.

   For multiple context mode, the ASA first classifies the packet to a context.
3. The packet is denied because there is no access rule permitting the outside host, and the ASA drops the packet.
4. If the outside user is attempting to attack the inside network, the ASA employs many technologies to determine if a packet is valid for an already established session.

History for the Firewall Mode

Table 3: Feature History for Firewall Mode

Feature Name	Platform Releases	Feature Information
Transparent Firewall Mode	7.0(1)	A transparent firewall is a Layer 2 firewall that acts like a "bump in the wire," or a "stealth firewall," and is not seen as a router hop to connected devices. We introduced the following commands: **firewall transparent**, **show firewall**.
Transparent firewall bridge groups	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 8 bridge groups in single mode or per context in multiple mode, with 4 interfaces maximum per bridge group. **Note** Although you can configure multiple bridge groups on the ASA 5505, the restriction of 2 data interfaces in transparent mode on the ASA 5505 means you can only effectively use 1 bridge group. We introduced the following commands: **interface bvi**, **bridge-group**, **show bridge-group**.
Mixed firewall mode support in multiple context mode	8.5(1)/9.0(1)	You can set the firewall mode independently for each security context in multiple context mode, so some can run in transparent mode while others run in routed mode. We modified the following command: **firewall transparent**.
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
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**
Support for transparent mode deployment for a Firepower 4100/9300 ASA logical device	9.10(1)	You can now specify transparent or routed mode when you deploy the ASA on a Firepower 4100/9300. New/Modified FXOS commands: **enter bootstrap-key FIREWALL_MODE**, **set value routed**, **set value transparent**