

Regular Firewall Interfaces for Firepower Threat Defense

This chapter includes regular firewall Firepower Threat Defense interface configuration including EtherChannels, VLAN subinterfaces, IP addressing, and more.

Note For initial interface configuration on the Firepower 4100/9300, see Configure Interfaces.

- Requirements and Prerequisites for Regular Firewall Interfaces, on page 1
- Configure EtherChannel and Redundant Interfaces, on page 1
- Configure VLAN Subinterfaces and 802.1Q Trunking, on page 9
- Configure Routed and Transparent Mode Interfaces, on page 11
- Configure Advanced Interface Settings, on page 25
- History for Regular Firewall Interfaces for Firepower Threat Defense, on page 35

Requirements and Prerequisites for Regular Firewall Interfaces

Model Support

FTD

User Roles

- Admin
- Access Admin
- Network Admin

Configure EtherChannel and Redundant Interfaces

This section tells how to configure EtherChannels and redundant interfaces.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel) for more information.

Note Only ASA 5500-X models support redundant interfaces; Firepower models do not support them.

About EtherChannels

This section describes EtherChannels.

About Redundant Interfaces (ASA Platform Only)

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the Firepower Threat Defense device reliability.

You can configure up to 8 redundant interface pairs.

Redundant Interface MAC Address

The redundant interface uses the MAC address of the first physical interface that you add. If you change the order of the member interfaces in the configuration, then the MAC address changes to match the MAC address of the interface that is now listed first. Alternatively, you can assign a manual MAC address to the redundant interface, which is used regardless of the member interface MAC addresses. When the active interface fails over to the standby, the same MAC address is maintained so that traffic is not disrupted.

About EtherChannels

An 802.3ad EtherChannel is a logical interface (called a port-channel interface) consisting of a bundle of individual Ethernet links (a channel group) so that you increase the bandwidth for a single network. A port channel interface is used in the same way as a physical interface when you configure interface-related features.

You can configure up to 48 EtherChannels, depending on how many interfaces your model supports.

Channel Group Interfaces

Each channel group can have up to 8 active interfaces, except for ASA models and the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.

All interfaces in the channel group must be the same type and speed. The first interface added to the channel group determines the correct type and speed.

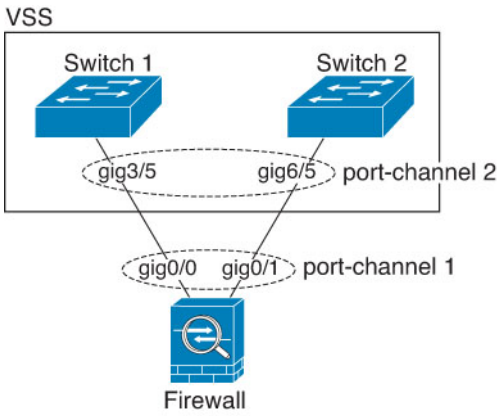
The EtherChannel aggregates the traffic across all the available active interfaces in the channel. The interface is selected using a proprietary hash algorithm, based on source or destination MAC addresses, IP addresses, TCP and UDP port numbers and VLAN numbers.

Connecting to an EtherChannel on Another Device

The device to which you connect the FTD EtherChannel must also support 802.3ad EtherChannels; for example, you can connect to the Catalyst 6500 switch or the Cisco Nexus 7000.

When the switch is part of a Virtual Switching System (VSS) or Virtual Port Channel (vPC), then you can connect FTD interfaces within the same EtherChannel to separate switches in the VSS/vPC. The switch interfaces are members of the same EtherChannel port-channel interface, because the separate switches act like a single switch.

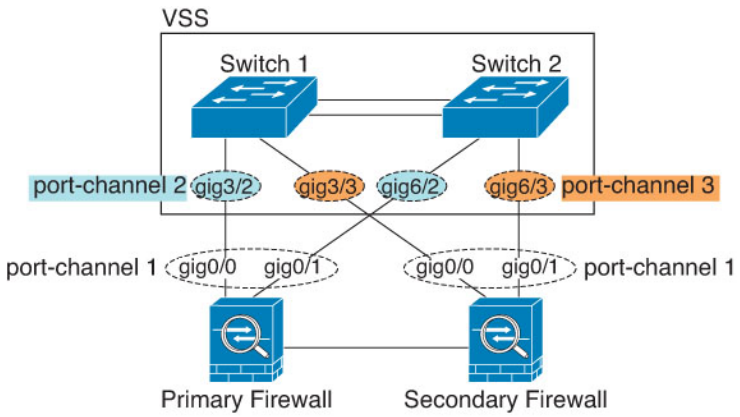
Figure 1: Connecting to a VSS/vPC



Note If the FTD device is in transparent firewall mode, and you place the FTD device between two sets of VSS/vPC switches, then be sure to disable Unidirectional Link Detection (UDLD) on any switch ports connected to the FTD device with an EtherChannel. If you enable UDLD, then a switch port may receive UDLD packets sourced from both switches in the other VSS/vPC pair. The receiving switch will place the receiving interface in a down state with the reason "UDLD Neighbor mismatch".

If you use the FTD device in an Active/Standby failover deployment, then you need to create separate EtherChannels on the switches in the VSS/vPC, one for each FTD device. On each FTD device, a single EtherChannel connects to both switches. Even if you could group all switch interfaces into a single EtherChannel connecting to both FTD devices (in this case, the EtherChannel will not be established because of the separate FTD system IDs), a single EtherChannel would not be desirable because you do not want traffic sent to the standby FTD device.

Figure 2: Active/Standby Failover and VSS/vPC



Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) aggregates interfaces by exchanging the Link Aggregation Control Protocol Data Units (LACPDUs) between two network devices.

You can configure each physical interface in an EtherChannel to be:

- Active—Sends and receives LACP updates. An active EtherChannel can establish connectivity with either an active or a passive EtherChannel. You should use the active mode unless you need to minimize the amount of LACP traffic.
- Passive—Receives LACP updates. A passive EtherChannel can only establish connectivity with an active EtherChannel. Not supported on hardware models.
- On—The EtherChannel is always on, and LACP is not used. An “on” EtherChannel can only establish a connection with another “on” EtherChannel. Not supported on hardware models.

LACP coordinates the automatic addition and deletion of links to the EtherChannel without user intervention. It also handles misconfigurations and checks that both ends of member interfaces are connected to the correct channel group. “On” mode cannot use standby interfaces in the channel group when an interface goes down, and the connectivity and configurations are not checked.

Load Balancing

The FTD device distributes packets to the interfaces in the EtherChannel by hashing the source and destination IP address of the packet (this criteria is configurable). The resulting hash is divided by the number of active links in a modulo operation where the resulting remainder determines which interface owns the flow. All packets with a *hash_value* **mod** *active_links* result of 0 go to the first interface in the EtherChannel, packets with a result of 1 go to the second interface, packets with a result of 2 go to the third interface, and so on. For example, if you have 15 active links, then the modulo operation provides values from 0 to 14. For 6 active links, the values are 0 to 5, and so on.

If an active interface goes down and is not replaced by a standby interface, then traffic is rebalanced between the remaining links. The failure is masked from both Spanning Tree at Layer 2 and the routing table at Layer 3, so the switchover is transparent to other network devices.

EtherChannel MAC Address

All interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links.

Firepower Hardware

The port-channel interface uses the MAC address of the internal interface Internal-Data 0/1. Alternatively you can manually configure a MAC address for the port-channel interface. All EtherChannel interfaces on a chassis use the same MAC address, so be aware that if you use SNMP polling, for example, multiple interfaces will have the same MAC address.



Note Member interfaces only use the Internal-Data 0/1 MAC address after a reboot. Prior to rebooting, the member interface uses its own MAC address. If you add a new member interface after a reboot, you will have to perform another reboot to update its MAC address.

ASA Hardware

The port-channel interface uses the lowest numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can manually configure a MAC address for the port-channel interface. We recommend manually configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.

Guidelines for EtherChannels

Bridge Group

In routed mode, FMC-defined EtherChannels are not supported as bridge group members. EtherChannels on the Firepower 4100/9300 can be bridge group members.

High Availability

- When you use an EtherChannel interface as a High Availability link, it must be pre-configured on both units in the High Availability pair; you cannot configure it on the primary unit and expect it to replicate to the secondary unit because *the High Availability link itself is required for replication*.
- If you use an EtherChannel interface for the state link, no special configuration is required; the configuration can replicate from the primary unit as normal. For the Firepower 4100/9300 chassis, all interfaces, including EtherChannels, need to be pre-configured on both units.
- You can monitor EtherChannel interfaces for High Availability. When an active member interface fails over to a standby interface, this activity does not cause the EtherChannel interface to appear to be failed when being monitored for device-level High Availability. Only when all physical interfaces fail does the EtherChannel interface appear to be failed (for an EtherChannel interface, the number of member interfaces allowed to fail is configurable).
- If you use an EtherChannel interface for a High Availability or state link, then to prevent out-of-order packets, only one interface in the EtherChannel is used. If that interface fails, then the next interface in

the EtherChannel is used. You cannot alter the EtherChannel configuration while it is in use as a High Availability link. To alter the configuration, you need to temporarily disable High Availability, which prevents High Availability from occurring for the duration.

Model Support

- You cannot add EtherChannels in the FMC for the Firepower 4100/9300 or the FTDv. The Firepower 4100/9300 supports EtherChannels, but you must perform all hardware configuration of EtherChannels in FXOS on the chassis.

General EtherChannel Guidelines

- You can configure up to 48 EtherChannels, depending on how many interfaces are available on your model.
- Each channel group can have up to 8 active interfaces, except for ASA models and the ISA 3000, which supports 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same media type and speed capacity. The media type can be either RJ-45 or SFP; SFPs of different types (copper and fiber) can be mixed. You cannot mix interface capacities (for example 1GB and 10GB interfaces) by setting the speed to be lower on the larger-capacity interface.
- The device to which you connect the FTD EtherChannel must also support 802.3ad EtherChannels.
- The FTD device does not support LACPDUs that are VLAN-tagged. If you enable native VLAN tagging on the neighboring switch using the Cisco IOS **vlan dot1Q tag native** command, then the FTD device will drop the tagged LACPDUs. Be sure to disable native VLAN tagging on the neighboring switch.
- Devices do not support LACP rate fast, except for ASA models and the ISA 3000; LACP always uses the normal rate. This setting is not configurable. Note that the Firepower 4100/9300, which configures EtherChannels in FXOS, has the LACP rate set to fast by default; on these platforms, the rate is configurable.
- In Cisco IOS software versions earlier than 15.1(1)S2, FTD did not support connecting an EtherChannel to a switch stack. With default switch settings, if the FTD EtherChannel is connected cross stack, and if the primary switch is powered down, then the EtherChannel connected to the remaining switch will not come up. To improve compatibility, set the **stack-mac persistent timer** command to a large enough value to account for reload time; for example, 8 minutes or 0 for indefinite. Or, you can upgrade to more a more stable switch software version, such as 15.1(1)S2.
- All the FTD configuration refers to the logical EtherChannel interface instead of the member physical interfaces.

Configure a Redundant Interface (ASA Platform Only)

A logical redundant interface consists of a pair of physical interfaces: an active and a standby interface. When the active interface fails, the standby interface becomes active and starts passing traffic. You can configure a redundant interface to increase the Firepower Threat Defense reliability. By default, redundant interfaces are enabled.

- You can configure up to 8 redundant interface pairs.
- Both member interfaces must be of the same physical type. For example, both must be GigabitEthernet.

Note Redundant interfaces are not supported on the Firepower platform; only ASA 5500-X models support redundant interfaces.

Before you begin

- You cannot add a physical interface to the redundant interface if you configured a name for it. You must first remove the name.

Caution If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the Physical Interface and Configure Ethernet Settings.

Step 3 Click **Add Interfaces** > **Redundant Interface**.

Step 4 On the **General** tab, set the following parameters:

a) **Redundant ID**—Set an integer between 1 and 8.

b) **Primary Interface**—Choose an interface from the drop-down list. After you add the interface, any configuration for it (such as an IP address) is removed.

c) **Secondary Interface**—The second interface must be the same physical type as the first interface.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 (Optional) Add a VLAN subinterface. See Add a Subinterface, on page 10.

Step 8 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 13 or Configure Bridge Group Interfaces, on page 15.

Configure an EtherChannel

This section describes how to create an EtherChannel port-channel interface, assign interfaces to the EtherChannel, and customize the EtherChannel.

Guidelines

- You can configure up to 48 EtherChannels, depending on the number of interfaces for your model.
- Each channel group can have up to 16 active interfaces. For switches that support only 8 active interfaces, you can assign up to 16 interfaces to a channel group: while only 8 interfaces can be active, the remaining interfaces can act as standby links in case of interface failure.
- All interfaces in the channel group must be the same type, speed, and duplex. Half duplex is not supported.

Note For the Firepower 4100/9300, you configure EtherChannels in FXOS. See Add an EtherChannel (Port Channel) for more information.

Before you begin

- You cannot add a physical interface to the channel group if you configured a name for it. You must first remove the name.

Note If you are using a physical interface already in your configuration, removing the name will clear any configuration that refers to the interface.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the member interfaces according to Enable the Physical Interface and Configure Ethernet Settings.

Step 3 Click **Add Interfaces** > **Ether Channel Interface**.

Step 4 On the **General** tab, set the **Ether Channel ID** to a number between 1 and 48.

Step 5 In the **Available Interfaces** area, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members.

Make sure all interfaces are the same type and speed. The first interface you add determines the type and speed of the EtherChannel. Any non-matching interfaces you add will be put into a suspended state. The FMC does not prevent you from adding non-matching interfaces.

Step 6 (Optional) Click the **Advanced** tab to customize the EtherChannel. Set the following parameters on the **Information** sub-tab:

- **Load Balancing**—Select the criteria used to load balance the packets across the group channel interfaces. By default, the Firepower Threat Defense device balances the packet load on interfaces according to the source and destination IP address of the packet. If you want to change the properties on which the packet is categorized, choose a different set of criteria. For example, if your traffic is biased heavily towards the same source and destination IP addresses, then the traffic assignment to interfaces in the EtherChannel will be unbalanced. Changing to a different algorithm can result in more evenly distributed traffic. For more information about load balancing, see Load Balancing, on page 4.

- **LACP Mode**—Choose Active, Passive, or On. We recommend using Active mode (the default).
- **Active Physical Interface: Range**—From the left drop-down list, choose the minimum number of active interfaces required for the EtherChannel to be active, between 1 and 16. The default is 1. From the right drop-down list, choose the maximum number of active interfaces allowed in the EtherChannel, between 1 and 16. The default is 16. If your switch does not support 16 active interfaces, be sure to set this command to 8 or fewer.
- **Active Mac Address**—Set a manual MAC address if desired. The mac_address is in H.H.H format, where H is a 16-bit hexadecimal digit. For example, the MAC address 00-0C-F1-42-4C-DE is entered as 000C.F142.4CDE.

Step 7 (Optional) Click the **Hardware Configuration** tab and set the Duplex and Speed to override these settings for all member interfaces. This method provides a shortcut to set these parameters because these parameters must match for all interfaces in the channel group.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 10 (Optional) Add a VLAN subinterface. See Add a Subinterface, on page 10.

Step 11 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 13 or Configure Bridge Group Interfaces, on page 15.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical, redundant, or EtherChannel interface into multiple logical interfaces that are tagged with different VLAN IDs. An interface with one or more VLAN subinterfaces is automatically configured as an 802.1Q trunk. Because VLANs let you keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces or devices.

Guidelines and Limitations for VLAN Subinterfaces

High Availability

You cannot use a subinterface for the failover or state link.

Additional Guidelines

- Preventing untagged packets on the physical interface—If you use subinterfaces, you typically do not also want the physical interface to pass traffic, because the physical interface passes untagged packets. This property is also true for the active physical interface in a redundant interface pair and for EtherChannel links. Because the physical, redundant, or EtherChannel interface must be enabled for the subinterface to pass traffic, ensure that the physical, redundant, or EtherChannel interface does not pass traffic by not configuring a name for the interface. If you want to let the physical, redundant, or EtherChannel interface pass untagged packets, you can configure the name as usual.

- You cannot configure subinterfaces on the interface.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- The FTD does not support the Dynamic Trunking Protocol (DTP), so you must configure the connected switch port to trunk unconditionally.
- You might want to assign unique MAC addresses to subinterfaces defined on the FTD, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD.

Maximum Number of VLAN Subinterfaces by Device Model

The device model limits the maximum number of VLAN subinterfaces that you can configure. Note that you can configure subinterfaces on data interfaces only, you cannot configure them on the management interface.

The following table explains the limits for each device model.

Model	Maximum VLAN Subinterfaces
Firepower 4100	1024
Firepower 9300	1024
FTDv	50
ASA 5506-X ASA 5506W-X ASA 5506H-X	30
ASA 5508-X	50
ASA 5512-X	100
ASA 5515-X	100
ASA 5516-X	100
ASA 5525-X	200
ASA 5545-X	300
ASA 5555-X	500

Add a Subinterface

Add one or more subinterfaces to a physical, redundant, or port-channel interface.

Note The parent physical interface passes untagged packets. You may not want to pass untagged packets, so be sure not to include the parent interface in your security policy.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Enable the parent interface according to Enable the Physical Interface and Configure Ethernet Settings.

Step 3 Click **Add Interfaces** > **Sub Interface**.

Step 4 On **General**, set the following parameters:

a) **Interface**—Choose the physical, redundant, or port-channel interface to which you want to add the subinterface.

b) **Sub-Interface ID**—Enter the subinterface ID as an integer between 1 and 4294967295. The number of subinterfaces allowed depends on your platform. You cannot change the ID after you set it.

c) **VLAN ID**—Enter the VLAN ID between 1 and 4094 that will be used to tag the packets on this subinterface.

This VLAN ID must be unique.

Step 5 Click **OK**.

Step 6 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 7 Configure the routed or transparent mode interface parameters. See Configure Routed Mode Interfaces, on page 13 or Configure Bridge Group Interfaces, on page 15.

Configure Routed and Transparent Mode Interfaces

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

About Routed and Transparent Mode Interfaces

Firewall mode interfaces subject traffic to firewall functions such as maintaining flows, tracking flow states at both IP and TCP layers, IP defragmentation, and TCP normalization. You can also optionally configure IPS functions for this traffic according to your security policy.

The types of firewall interfaces you can configure depends on the firewall mode set for the device: routed or transparent mode. See Transparent or Routed Firewall Mode for Firepower Threat Defense for more information.

- Routed mode interfaces (routed firewall mode only)—Each interface that you want to route between is on a different subnet.

- Bridge group interfaces (routed and transparent firewall mode)—You can group together multiple interfaces on a network, and the Firepower Threat Defense device 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. In routed mode, the Firepower Threat Defense device routes between BVIs and regular routed interfaces. In transparent mode, each bridge group is separate and cannot communicate with each other.

Dual IP Stack (IPv4 and IPv6)

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

Guidelines and Limitations for Routed and Transparent Mode Interfaces

High Availability, Clustering

- Do not configure failover links with the procedures in this chapter. See the High Availability chapter for more information.
- For cluster interfaces, see the clustering chapter for requirements.
- When you use High Availability, you must set the IP address and standby address for data interfaces manually; DHCP and PPPoE are not supported. Set the standby IP addresses on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. See the High Availability chapter for more information.

IPv6

- IPv6 is supported on all interfaces.
- You can only configure IPv6 addresses manually in transparent mode.
- The FTD device does not support IPv6 anycast addresses.

Model Guidelines

- For the FTDv on VMware with bridged ixgbevf interfaces, bridge groups are not supported.

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 FTD device 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 FTD 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).
- Management interfaces are not supported as bridge group members.
- For the FTDv on VMware with bridged ixgbevf interfaces, transparent mode is not supported, and bridge groups are not supported in routed mode.
- 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 FTD 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 interface.
- Transparent mode is not supported on threat defense virtual instances deployed on Amazon Web Services, Microsoft Azure, Google Cloud Platform, and Oracle Cloud Infrastructure.
- In routed mode, to route between bridge groups and other routed interfaces, you must name the BVI.
- In routed mode, FTD-defined EtherChannel 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 FTD when using bridge group members. If there are two neighbors on either side of the FTD running BFD, then the FTD will drop BFD echo packets because they have the same source and destination IP address and appear to be part of a LAND attack.

Additional Guidelines and Requirements

- The FTD supports only one 802.1Q header in a packet and does not support multiple headers (known as Q-in-Q support) for firewall interfaces.**Note:** For inline sets and passive interfaces, the FTD supports Q-in-Q up to two 802.1Q headers in a packet, with the exception of the Firepower 4100/9300, which only supports one 802.1Q header.

Configure Routed Mode Interfaces

This procedure describes how to set the name, security zone, and IPv4 address.

Before you begin

- **Firepower 4100/9300**
  1. Configure a Physical Interface
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (Port Channel)

- Add a Subinterface, on page 10 in FMC

- (Optional) **All other models:**
  - Configure a Redundant Interface (ASA Platform Only), on page 6
  - Configure an EtherChannel, on page 7
  - Add a Subinterface, on page 10

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

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

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The routed interface is a Routed-type interface, and can only belong to Routed-type zones.

Step 9 See Configure the MTU, on page 29 for information about the **MTU**.

Step 10 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP and PPPoE are not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availabilty, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. 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.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.
- **Use 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:

- **VPDN Group Name**—Specify a group name of your choice to represent this connection.
- **PPPoE User Name**—Specify the username provided by your ISP.
- **PPPoE Password/Confirm Password**—Specify and confirm the password provided by your ISP.
- **PPP Authentication**—Choose **PAP**, **CHAP**, or **MSCHAP**.

  PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.
- **PPPoE route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.
- **Store Username and Password in Flash**—Stores the username and password in flash memory.

  The Firepower Threat Defense device stores the username and password in a special location of NVRAM.

Step 11 (Optional) See Configure IPv6 Addressing, on page 20 to configure IPv6 addressing on the **IPv6** tab.

Step 12 (Optional) See Configure the MAC Address, on page 30 to manually configure the MAC address on the **Advanced** tab.

Step 13 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default for RJ-45 interfaces. You cannot select Auto for SFP interfaces.
- **Speed**—Choose **Auto** to have the interface negotiate the speed, link status, and flow control (Auto is only available for RJ-45 interfaces), or pick a specific speed: **10**, **100**, **1000** Mbps. For SFP interfaces, setting the speed enables auto-negotiation of link status and flow control. For SFP interfaces, depending on your hardware, you can select **No Negotiate** to set the speed to 1000 and disable link negotiation.

Step 14 Click **OK**.

Step 15 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Bridge Group Interfaces

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

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

Configure General Bridge Group Member Interface Parameters

This procedure describes how to set the name and security zone for each bridge group member interface. The same bridge group can include different types of interfaces: physical interfaces, VLAN subinterfaces, EtherChannels, and redundant interfaces. The interface is not supported. In routed mode, EtherChannels are not supported.

Before you begin

- **Firepower 4100/9300**
  1. Configure a Physical Interface
  2. (Optional) Configure any special interfaces.
     - Add an EtherChannel (Port Channel)
     - Add a Subinterface, on page 10 in FMC
- (Optional) **All other models:**
  - Configure a Redundant Interface (ASA Platform Only), on page 6
  - Configure an EtherChannel, on page 7
  - Add a Subinterface, on page 10

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

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

Step 4 Enable the interface by checking the **Enabled** check box.

Step 5 (Optional) Set this interface to **Management Only** to limit traffic to management traffic; through-the-box traffic is not allowed.

Step 6 (Optional) Add a description in the **Description** field.

The description can be up to 200 characters on a single line, without carriage returns.

Step 7 In the **Mode** drop-down list, choose **None**.

Regular firewall interfaces have the mode set to None. The other modes are for IPS-only interface types. After you assign this interface to a bridge group, the mode will show as **Switched**.

Step 8 From the **Security Zone** drop-down list, choose a security zone or add a new one by clicking **New**.

The bridge group member interface is a Switched-type interface, and can only belong to Switched-type zones. Do not configure any IP address settings for this interface. You will set the IP address for the Bridge Virtual Interface (BVI) only. Note that the BVI does not belong to a zone, and you cannot apply access control policies to the BVI.

Step 9 See Configure the MTU, on page 29 for information about the **MTU**.

Step 10 (Optional) Set the duplex and speed by clicking **Hardware Configuration**.

- **Duplex**—Choose **Full**, **Half**, or **Auto**. Auto is the default for RJ-45 interfaces. You cannot select Auto for SFP interfaces.
- **Speed**—Choose **Auto** to have the interface negotiate the speed, link status, and flow control (Auto is only available for RJ-45 interfaces), or pick a specific speed: **10**, **100**, **1000** Mbps. For SFP interfaces, setting the speed enables auto-negotiation of link status and flow control. For SFP interfaces, depending on your hardware, you can select **No Negotiate** to set the speed to 1000 and disable link negotiation.

Step 11 (Optional) See Configure IPv6 Addressing, on page 20 to configure IPv6 addressing on the **IPv6** tab.

Step 12 (Optional) See Configure the MAC Address, on page 30 to manually configure the MAC address on the **Advanced** tab.

Step 13 Click **OK**.

Step 14 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure the Bridge Virtual Interface (BVI)

Each bridge group requires a BVI for which you configure an IP address. The Firepower Threat Defense 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.

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

Before you begin

You cannot add the BVI to a security zone; therefore, you cannot apply Access Control policies to the BVI. You must apply your policy to the bridge group member interfaces based on their zones.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Choose **Add Interfaces** > **Bridge Group Interface**.

Step 3 (Routed Mode) In the **Name** field, enter a name up to 48 characters in length.

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 not case-sensitive.

Step 4 In the **Bridge Group ID** field, enter the bridge group ID between 1 and 250.

Step 5 In the **Description** field, enter a description for this bridge group.

Step 6 On the **Interfaces** tab, click an interface and then click **Add** to move it to the **Selected Interfaces** area. Repeat for all interfaces that you want to make members of the bridge group.

Step 7 (Transparent Mode) Click the **IPv4** tab. In the **IP Address** field, enter the IPv4 address and subnet mask.

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 transparent firewall) such as a /30 subnet (255.255.255.252). The Firepower Threat Defense device drops all ARP packets to or from the first and last addresses in a subnet. For example, if you use a /30 subnet and assign a reserved address from that subnet to the upstream router, then the Firepower Threat Defense device drops the ARP request from the downstream router to the upstream router.

For High Availabilty, set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. 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.

Step 8 (Routed Mode) Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

High Availability and clustering interfaces only support static IP address configuration; DHCP is not supported.

- **Use Static IP**—Enter the IP address and subnet mask. For High Availabilty, you can only use a static IP address. Set the standby IP address on the **Devices** > **Device Management** > **High Availability** tab in the **Monitored Interfaces** area. 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.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.

Step 9 (Optional) See Configure IPv6 Addressing, on page 20 to configure IPv6 addressing.

Step 10 (Optional) See Add a Static ARP Entry, on page 31 and Add a Static MAC Address and Disable MAC Learning for a Bridge Group, on page 32 (for transparent mode only) to configure the **ARP** and **MAC** settings.

Step 11 Click **OK**.

Step 12 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure a Diagnostic (Management) Interface for Transparent Mode

In transparent firewall mode, all interfaces must belong to a bridge group. The only exception is the Diagnostic *slot*/*port* interface. For the Firepower 4100/9300 chassis, the diagnostic interface ID depends on the mgmt-type

interface that you assigned to the Firepower Threat Defense logical device. You cannot use any other interface types as diagnostic interfaces. You can configure one diagnostic interface.

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.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface.

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

Step 4 Click the **IPv4** tab. To set the IP address, use one of the following options from the **IP Type** drop-down list.

- **Use Static IP**—Enter the IP address and subnet mask.
- **Use DHCP**—Configure the following optional parameters:
  - **Obtain default route using DHCP**—Obtains the default route from the DHCP server.
  - **DHCP route metric**—Assigns an administrative distance to the learned route, between 1 and 255. The default administrative distance for the learned routes is 1.
- **Use PPPoE**—Configure the following parameters:
  - **VPDN Group Name**—Specify a group name.
  - **PPPoE User Name**—Specify the username provided by your ISP.
  - **PPPoE Password/Confirm Password**—Specify and confirm the password provided by your ISP.
  - **PPP Authentication**—Choose **PAP**, **CHAP**, or **MSCHAP**.

    PAP passes a cleartext username and password during authentication and is not secure. With CHAP, the client returns the encrypted [challenge plus password], with a cleartext username in response to the server challenge. CHAP is more secure than PAP, but it does not encrypt data. MSCHAP is similar to CHAP but is more secure because the server stores and compares only encrypted passwords rather than cleartext passwords as in CHAP. MSCHAP also generates a key for data encryption by MPPE.
  - **PPPoE route metric**—Assign an administrative distance to the learned route. Valid values are from 1 to 255. By default, the administrative distance for the learned routes is 1.
  - **Enable Route Settings**—To manually configure the PPPoE IP address, check this box and then enter the **IP Address**.
  - **Store Username and Password in Flash**—Stores the username and password in flash memory.

    The Firepower Threat Defense device stores the username and password in a special location of NVRAM.

Step 5 (Optional) See Configure IPv6 Addressing, on page 20 to configure **IPv6** addressing.

Step 6 (Optional) On the **Advanced** tab, configure optional settings.

- See Configure the MAC Address, on page 30.
- See Add a Static ARP Entry, on page 31.
- See Set Security Configuration Parameters, on page 33.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure IPv6 Addressing

This section describes how to configure IPv6 addressing in routed and transparent mode.

About IPv6

This section includes information about IPv6.

IPv6 Addressing

You can configure two types of unicast addresses for IPv6:

- Global—The global address is a public address that you can use on the public network. For a bridge group, this address needs to be configured for the BVI, and not per member interface. You can also configure a global IPv6 address for the management interface in transparent mode.
- Link-local—The link-local address is a private address that you can only use on the directly-connected network. Routers do not forward packets using link-local addresses; they are only for communication on a particular physical network segment. They can be used for address configuration or for the Neighbor Discovery functions such as address resolution. In a bridge group, only member interfaces have link-local addresses; the BVI does not have a link-local address.

At a minimum, you need to configure a link-local address for IPv6 to operate. If you configure a global address, a link-local address is automatically configured on the interface, so you do not also need to specifically configure a link-local address. For bridge group member interfaces, when you configure the global address on the BVI, the FTD device automatically generates link-local addresses for member interfaces. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Modified EUI-64 Interface IDs

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

When this feature is enabled on an interface, the source addresses of IPv6 packets received on that interface are verified against the source MAC addresses to ensure that the interface identifiers use the Modified EUI-64

format. If the IPv6 packets do not use the Modified EUI-64 format for the interface identifier, the packets are dropped and the following system log message is generated:

```
325003: EUI-64 source address check failed.
```

The address format verification is only performed when a flow is created. Packets from an existing flow are not checked. Additionally, the address verification can only be performed for hosts on the local link.

Configure a Global IPv6 Address

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

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

For subinterfaces defined on the FTD, 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 FTD. See Configure the MAC Address, on page 30.

Before you begin

For IPv6 neighbor discovery for bridge groups, you must explicitly allow Neighbor Solicitation (ICMPv6 type 135) and Neighbor Advertisement (ICMPv6 type 136) packets through the FTD bridge group member interfaces using a bidirectional access rule.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **IPv6** page.

For routed mode, the **Basic** page is selected by default. For transparent mode, the **Address** page is selected by default.

Step 4 On the **Basic** page, check **Enable IPv6**.

Step 5 Configure the global IPv6 address using one of the following methods.

- (Routed interface) Stateless autoconfiguration—Check the **Autoconfiguration** check box.

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

Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the Firepower Threat Defense device does send Router Advertisement messages in this case. Uncheck the **IPv6** > **Settings** > **Enable RA** check box to suppress messages.

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

  a. Click the **Address** page, and click **Add Address**.

  The **Add Address** dialog box appears.

  b. In the **Address** field, enter either a full global IPv6 address, including the interface ID, or enter the IPv6 prefix, along with the IPv6 prefix length. (Routed Mode) If you only enter the prefix, then be sure to check the **Enforce EUI 64** check box to generate the interface ID using the Modified EUI-64 format. For example, 2001:0DB8::BA98:0:3210/48 (full address) or 2001:0DB8::/48 (prefix, with EUI 64 checked).

  For High Availabilty (if you did not set **Enforce EUI 64**), set the standby IP address on the **Devices** > **Device Management** > **High Availability** page in the **Monitored Interfaces** area. 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.

Step 6 For Routed interfaces, you can optionally set the following values on the **Basic** page:

- To enforce the use of Modified EUI-64 format interface identifiers in IPv6 addresses on a local link, check the **Enforce EUI-64** check box.
- To manually set the link-local address, enter an address in the **Link-Local address** field.

  A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. If you do not want to configure a global address, and only need to configure a link-local address, you have the option of manually defining the link-local address. Note that we recommend automatically assigning the link-local address based on the Modified EUI-64 format. For example, if other devices enforce the use of the Modified EUI-64 format, then a manually-assigned link-local address may cause packets to be dropped.
- Check the **Enable DHCP for address config** check box to set the Managed Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.
- Check the **Enable DHCP for non-address config** check box to set the Other Address Config flag in the IPv6 router advertisement packet.

  This flag in IPv6 router advertisements informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.

Step 7 For Routed interfaces, see Configure IPv6 Neighbor Discovery, on page 23 to configure settings on the **Prefixes** and **Settings** pages. For BVI interfaces, see the following parameters on the **Settings** page:

- **DAD attempts**—The maximum number of DAD attempts, between 1 and 600. Set the value to 0 to disable duplicate address detection (DAD) processing. This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses. 1 attempt is the default.
- **NS Interval**—The interval between IPv6 neighbor solicitation retransmissions on an interface, between 1000 and 3600000 ms. The default value is 1000 ms.

- **Reachable Time**—The amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred, between 0 and 3600000 ms. The default value is 0 ms. 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.

Step 8 Click **OK**.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

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.

Before you begin

Supported in Routed mode only. For IPv6 neighbor settings supported in transparent mode, see Configure a Global IPv6 Address, on page 21.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click **IPv6**, and then **Prefixes**.

Step 4 (Optional) To configure which IPv6 prefixes are included in IPv6 router advertisements, perform the following steps:

a) Click **Add Prefix**.

b) In the **Address** field, enter the IPv6 address with the prefix length or check the **Default** check box to use the default prefix.

c) (Optional) Uncheck the **Advertisement** check box to indicate that the IPv6 prefix is not advertised.

d) Check the **Off Link** check box to indicate that the specified prefix is assigned to the link. Nodes sending traffic to addresses that contain the specified prefix consider the destination to be locally reachable on the link. This prefix should not be used for on-link determination.

e) To use the specified prefix for autoconfiguration, check the **Autoconfiguration** check box.

f) For the **Prefix Lifetime**, click **Duration** or **Expiration Date**.

- **Duration**—Enter a **Preferred Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being valid. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default is 2592000 (30 days). Enter a **Valid Lifetime** for the prefix in seconds. This setting is the amount of time that the specified IPv6 prefix is advertised as being preferred. The maximum value represents infinity. Valid values are from 0 to 4294967295. The default setting is 604800 (seven days). Alternatively, check the **Infinite** checkbox to set an unlimited duration.
- **Expiration Date**—Choose a **Valid** and **Preferred** date and time.

g) Click **OK**.

Step 5 Click **Settings**.

Step 6 (Optional) Set the maximum number of **DAD attempts**, between 1 and 600. 1 attempt is the default. Set the value to 0 to disable duplicate address detection (DAD) processing.

This setting configures the number of consecutive neighbor solicitation messages that are sent on an interface while DAD is performed on IPv6 addresses.

During the stateless autoconfiguration process, Duplicate Address Detection verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces.

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

```
325002: Duplicate address ipv6_address/MAC_address on interface
```

If the duplicate address is the link-local address of the interface, the processing of IPv6 packets is disabled on the interface. If the duplicate address is a global address, the address is not used.

Step 7 (Optional) Configure the interval between IPv6 neighbor solicitation retransmissions in the **NS Interval** field, between 1000 and 3600000 ms.

The default value is 1000 ms.

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

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

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

Step 8 (Optional) Configure the amount of time that a remote IPv6 node is considered reachable after a reachability confirmation event has occurred in the **Reachable Time** field, between 0 and 3600000 ms.

The default value is 0 ms. 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.

Step 9 (Optional) To suppress the router advertisement transmissions, uncheck the **Enable RA** check box. If you enable router advertisement transmissions, you can set the RA lifetime and interval.

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 Firepower Threat Defense device to supply the IPv6 prefix (for example, the outside interface).

- **RA Lifetime**—Configure the router lifetime value in IPv6 router advertisements, between 0 and 9000 seconds.

  The default is 1800 seconds.

- **RA Interval**—Configure the interval between IPv6 router advertisement transmissions, between 3 and 1800 seconds.

  The default is 200 seconds.

Step 10 Click **OK**.

Step 11 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Configure Advanced Interface Settings

This section describes how to configure MAC addresses for regular firewall mode interfaces, how to set the maximum transmission unit (MTU), and how to set other advanced parameters.

About Advanced Interface Configuration

This section describes advanced interface settings.

About MAC Addresses

You can manually assign MAC addresses to override the default.

Note You might want to assign unique MAC addresses to subinterfaces defined on the FTD, because they use the same burned-in MAC address of the parent interface. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD device.

Default MAC Addresses

Default MAC address assignments depend on the type of interface.

- Physical interfaces—The physical interface uses the burned-in MAC address.
- EtherChannels (Firepower Models)—For an EtherChannel, all interfaces that are part of the channel group share the same MAC address. This feature makes the EtherChannel transparent to network applications and users, because they only see the one logical connection; they have no knowledge of the individual links. The port-channel interface uses a unique MAC address from a pool; interface membership does not affect the MAC address.
- EtherChannels (ASA Models)—The port-channel interface uses the lowest-numbered channel group interface MAC address as the port-channel MAC address. Alternatively you can configure a MAC address for the port-channel interface. We recommend configuring a unique MAC address in case the group channel interface membership changes. If you remove the interface that was providing the port-channel MAC address, then the port-channel MAC address changes to the next lowest numbered interface, thus causing traffic disruption.
- Subinterfaces—All subinterfaces of a physical interface use the same burned-in MAC address. You might want to assign unique MAC addresses to subinterfaces. For example, your service provider might perform access control based on the MAC address. Also, because IPv6 link-local addresses are generated based on the MAC address, assigning unique MAC addresses to subinterfaces allows for unique IPv6 link-local addresses, which can avoid traffic disruption in certain instances on the FTD.

About the MTU

The MTU specifies the maximum frame *payload* size that the FTD device can transmit on a given Ethernet interface. The MTU value is the frame size *without* Ethernet headers, VLAN tagging, or other overhead. For example, when you set the MTU to 1500, the expected frame size is 1518 bytes including the headers, or 1522 when using VLAN. Do not set the MTU value higher to accommodate these headers.

Path MTU Discovery

The FTD device supports Path MTU Discovery (as defined in RFC 1191), which lets all devices in a network path between two hosts coordinate the MTU so they can standardize on the lowest MTU in the path.

Default MTU

The default MTU on the FTD device is 1500 bytes. This value does not include the 18-22 bytes for the Ethernet header, VLAN tagging, or other overhead.

MTU and Fragmentation

For IPv4, if an outgoing IP packet is larger than the specified MTU, it is fragmented into 2 or more frames. Fragments are reassembled at the destination (and sometimes at intermediate hops), and fragmentation can

cause performance degradation. For IPv6, packets are typically not allowed to be fragmented at all. Therefore, your IP packets should fit within the MTU size to avoid fragmentation.

For TCP packets, the endpoints typically use their MTU to determine the TCP maximum segment size (MTU - 40, for example). If additional TCP headers are added along the way, for example for site-to-site VPN tunnels, then the TCP MSS might need to be adjusted down by the tunneling entity. See About the TCP MSS, on page 27.

For UDP or ICMP, the application should take the MTU into account to avoid fragmentation.

Note The FTD device can receive frames larger than the configured MTU as long as there is room in memory.

MTU and Jumbo Frames

A larger MTU lets you send larger packets. Larger packets might be more efficient for your network. See the following guidelines:

- Matching MTUs on the traffic path—We recommend that you set the MTU on all FTD interfaces and other device interfaces along the traffic path to be the same. Matching MTUs prevents intermediate devices from fragmenting the packets.
- Accommodating jumbo frames—You can set the MTU 9000 bytes or higher when you enable jumbo frames. The maximum depends on the model.

About the TCP MSS

The TCP maximum segment size (MSS) is the size of the TCP payload *before* any TCP and IP headers are added. UDP packets are not affected. The client and the server exchange TCP MSS values during the three-way handshake when establishing the connection.

You can set the TCP MSS on the FTD device for through traffic using the Sysopt_Basic object in FlexConfig; see FlexConfig Policies for FTD; by default, the maximum TCP MSS is set to 1380 bytes. This setting is useful when the FTD device needs to add to the size of the packet for IPsec VPN encapsulation. However, for non-IPsec endpoints, you should disable the maximum TCP MSS on the FTD device.

If you set a maximum TCP MSS, if either endpoint of a connection requests a TCP MSS that is larger than the value set on the FTD device, then the FTD device overwrites the TCP MSS in the request packet with the FTD device maximum. If the host or server does not request a TCP MSS, then the FTD device assumes the RFC 793-default value of 536 bytes (IPv4) or 1220 bytes (IPv6), but does not modify the packet. For example, you leave the default MTU as 1500 bytes. A host requests an MSS of 1500 minus the TCP and IP header length, which sets the MSS to 1460. If the FTD device maximum TCP MSS is 1380 (the default), then the FTD device changes the MSS value in the TCP request packet to 1380. The server then sends packets with 1380-byte payloads. The FTD device can then add up to 120 bytes of headers to the packet and still fit in the MTU size of 1500.

You can also configure the minimum TCP MSS; if a host or server requests a very small TCP MSS, the FTD device can adjust the value up. By default, the minimum TCP MSS is not enabled.

For to-the-box traffic, including for SSL VPN connections, this setting does not apply. The FTD device uses the MTU to derive the TCP MSS: MTU - 40 (IPv4) or MTU - 60 (IPv6).

Default TCP MSS

By default, the maximum TCP MSS on the FTD device is 1380 bytes. This default accommodates IPv4 IPsec VPN connections where the headers can equal up to 120 bytes; this value fits within the default MTU of 1500 bytes.

Suggested Maximum TCP MSS Setting

The default TCP MSS assumes the FTD device acts as an IPv4 IPsec VPN endpoint and has an MTU of 1500. When the FTD device acts as an IPv4 IPsec VPN endpoint, it needs to accommodate up to 120 bytes for TCP and IP headers.

If you change the MTU value, use IPv6, or do not use the FTD device as an IPsec VPN endpoint, then you should change the TCP MSS setting using the Sysopt_Basic object in FlexConfig.



Note Even if you explicitly set an MSS, if a component such as TLS/SSL decryption or server discovery needs a particular MSS, it will set that MSS based on the interface MTU and ignore your MSS setting.

See the following guidelines:

- Normal traffic—Disable the TCP MSS limit and accept the value established between connection endpoints. Because connection endpoints typically derive the TCP MSS from the MTU, non-IPsec packets usually fit this TCP MSS.
- IPv4 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 120. For example, if you use jumbo frames and set the MTU to 9000, then you need to set the TCP MSS to 8880 to take advantage of the new MTU.
- IPv6 IPsec endpoint traffic—Set the maximum TCP MSS to the MTU - 140.

ARP Inspection for Bridge Group Traffic

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection.

ARP inspection prevents malicious users from impersonating other hosts or routers (known as ARP spoofing). ARP spoofing can enable a “man-in-the-middle” attack. For example, a host sends an ARP request to the gateway router; the gateway router responds with the gateway router MAC address. The attacker, however, sends another ARP response to the host with the attacker MAC address instead of the router MAC address. The attacker can now intercept all the host traffic before forwarding it on to the router.

ARP inspection ensures that an attacker cannot send an ARP response with the attacker MAC address, so long as the correct MAC address and the associated IP address are in the static ARP table.

When you enable ARP inspection, the FTD device compares the MAC address, IP address, and source interface in all ARP packets to static entries in the ARP table, and takes the following actions:

- If the IP address, MAC address, and source interface match an ARP entry, the packet is passed through.
- If there is a mismatch between the MAC address, the IP address, or the interface, then the FTD device drops the packet.
- If the ARP packet does not match any entries in the static ARP table, then you can set the FTD device to either forward the packet out all interfaces (flood), or to drop the packet.

Note The dedicated interface never floods packets even if this parameter is set to flood.

MAC Address Table

When you use bridge groups, the FTD learns and builds a MAC address table in a similar way as a normal bridge or switch: when a device sends a packet through the bridge group, the FTD adds the MAC address to its table. The table associates the MAC address with the source interface so that the FTD knows to send any packets addressed to the device out the correct interface. Because traffic between bridge group members is subject to the FTD security policy, if the destination MAC address of a packet is not in the table, the FTD does not flood the original packet on all interfaces as a normal bridge does. Instead, it generates the following packets for directly-connected devices or for remote devices:

- Packets for directly-connected devices—The FTD generates an ARP request for the destination IP address, so that it can learn which interface receives the ARP response.
- Packets for remote devices—The FTD generates a ping to the destination IP address so that it can learn which interface receives the ping reply.

The original packet is dropped.

Default Settings

- If you enable ARP inspection, the default setting is to flood non-matching packets.
- The default timeout value for dynamic MAC address table entries is 5 minutes.
- By default, each interface automatically learns the MAC addresses of entering traffic, and the FTD device adds corresponding entries to the MAC address table.

Note Firepower Threat Defense device generates a reset packet to reset a connection that is denied by a stateful inspection engine. Here, the destination MAC address of the packet is not determined based on the ARP table lookup but instead it is taken directly from the packets (connections) that are being denied.

Guidelines for ARP Inspection and the MAC Address Table

- ARP inspection is only supported for bridge groups.
- MAC address table configuration is only supported for bridge groups.

Configure the MTU

Customize the MTU on the interface, for example, to allow jumbo frames.

Caution Changing the highest MTU value on the device for a data interface restarts the Snort process when you deploy configuration changes, temporarily interrupting traffic inspection. Inspection is interrupted on all data interfaces, not just the interface you modified. Whether this interruption drops traffic or passes it without further inspection depends on the model of the managed device and the interface type. This caution does not apply to the Diagnostic interface or management-only interfaces. See Snort® Restart Traffic Behavior for more information.

Before you begin

- Changing the MTU above 1500 bytes automatically enables jumbo frames; for ASA models, you must reload the system before you can use jumbo frames.
- If you use an interface in an inline set, the MTU setting is not used. However, the jumbo frame setting *is* relevant to inline sets; jumbo frames enable the inline interfaces to receive packets up to 9000 bytes. To enable jumbo frames, you must set the MTU of *any* interface above 1500 bytes.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 On the **General** tab, set the **MTU** between 64 and 9198 bytes; the maximum is 9000 for the Firepower Threat Defense Virtual and 9184 for the Firepower Threat Defense on the Firepower 4100/9300 chassis.

The default is 1500 bytes.

Step 4 Click **OK**.

Step 5 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Step 6 For ASA models, if you set the MTU above 1500 bytes, reload the system to enable jumbo frames.

Configure the MAC Address

You might need to manually assign a MAC address. You can also set the Active and Standby MAC addresses on the **Devices** > **Device Management** > **High Availability** tab. If you set the MAC address for an interface on both screens, the addresses on the **Interfaces** > **Advanced** tab take precedence.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab.
The **Information** tab is selected.

Step 4 In the **Active MAC Address** field, enter a MAC address in H.H.H format, where H is a 16-bit hexadecimal digit.

For example, the MAC address 00-0C-F1-42-4C-DE would be entered as 000C.F142.4CDE. The MAC address must not have the multicast bit set, that is, the second hexadecimal digit from the left cannot be an odd number.

Step 5 In the **Standby MAC Address** field, enter a MAC address for use with High Availability.

If the active unit fails over and the standby unit becomes active, the new active unit starts using the active MAC addresses to minimize network disruption, while the old active unit uses the standby address.

Step 6 Click **OK**.

Step 7 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static ARP Entry

By default, all ARP packets are allowed between bridge group members. You can control the flow of ARP packets by enabling ARP inspection (see Configure ARP Inspection). ARP inspection compares ARP packets with *static* ARP entries in the ARP table.

For routed interfaces, you can enter static ARP entries, but normally dynamic entries are sufficient. For routed interfaces, the ARP table is used to deliver packets to directly-connected hosts. Although senders identify a packet destination by an IP address, the actual delivery of the packet on Ethernet relies on the Ethernet MAC address. When a router or host wants to deliver a packet on a directly connected network, it sends an ARP request asking for the MAC address associated with the IP address, and then delivers the packet to the MAC address according to the ARP response. The host or router keeps an ARP table so it does not have to send ARP requests for every packet it needs to deliver. The ARP table is dynamically updated whenever ARP responses are sent on the network, and if an entry is not used for a period of time, it times out. If an entry is incorrect (for example, the MAC address changes for a given IP address), the entry needs to time out before it can be updated with the new information.

For transparent mode, the Firepower Threat Defense only uses dynamic ARP entries in the ARP table for traffic to and from the Firepower Threat Defense device, such as management traffic.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP** tab (called **ARP and MAC** for transparent mode).

Step 4 Click **Add ARP Config**.
The **Add ARP Config** dialog box appears.

Step 5 In the **IP Address** field, enter the IP address of the host.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b.

Step 7 To perform proxy ARP for this address, check the **Enable Alias** check box.

If the Firepower Threat Defense device receives an ARP request for the specified IP address, then it responds with the specified MAC address.

Step 8 Click **OK**, and then click **OK** again to exit the Advanced settings.

Step 9 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Add a Static MAC Address and Disable MAC Learning for a Bridge Group

Normally, MAC addresses are added to the MAC address table dynamically as traffic from a particular MAC address enters an interface. You can disable MAC address learning; however, unless you statically add MAC addresses to the table, no traffic can pass through the Firepower Threat Defense device. You can also add static MAC addresses to the MAC address table. One benefit to adding static entries is to guard against MAC spoofing. If a client with the same MAC address as a static entry attempts to send traffic to an interface that does not match the static entry, then the Firepower Threat Defense device drops the traffic and generates a system message. When you add a static ARP entry (see Add a Static ARP Entry, on page 31), a static MAC address entry is automatically added to the MAC address table.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **ARP and MAC** tab.

Step 4 (Optional) Disable MAC learning by unchecking the **Enable MAC Learning** check box.

Step 5 To add a static MAC address, click **Add MAC Config**.
The **Add MAC Config** dialog box appears.

Step 6 In the **MAC Address** field, enter the MAC address of the host; for example, 00e0.1e4e.3d8b. Click **OK**.

Step 7 Click **OK** to exit the Advanced settings.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

Set Security Configuration Parameters

This section describes how to prevent IP spoofing, allow full fragment reassembly, and override the default fragment setting set for at the device level in **Platform Settings** .

Anti-Spoofing

This section lets you enable Unicast Reverse Path Forwarding on an interface. Unicast RPF guards against IP spoofing (a packet uses an incorrect source IP address to obscure its true source) by ensuring that all packets have a source IP address that matches the correct source interface according to the routing table.

Normally, the Firepower Threat Defense device only looks at the destination address when determining where to forward the packet. Unicast RPF instructs the device to also look at the source address; this is why it is called Reverse Path Forwarding. For any traffic that you want to allow through the Firepower Threat Defense device, the device routing table must include a route back to the source address. See RFC 2267 for more information.

For outside traffic, for example, the Firepower Threat Defense device can use the default route to satisfy the Unicast RPF protection. If traffic enters from an outside interface, and the source address is not known to the routing table, the device uses the default route to correctly identify the outside interface as the source interface.

If traffic enters the outside interface from an address that is known to the routing table, but is associated with the inside interface, then the Firepower Threat Defense device drops the packet. Similarly, if traffic enters the inside interface from an unknown source address, the device drops the packet because the matching route (the default route) indicates the outside interface.

Unicast RPF is implemented as follows:

- ICMP packets have no session, so each packet is checked.
- UDP and TCP have sessions, so the initial packet requires a reverse route lookup. Subsequent packets arriving during the session are checked using an existing state maintained as part of the session. Non-initial packets are checked to ensure they arrived on the same interface used by the initial packet.

Fragment per Packet

By default, the Firepower Threat Defense device allows up to 24 fragments per IP packet, and up to 200 fragments awaiting reassembly. You might need to let fragments on your network if you have an application that routinely fragments packets, such as NFS over UDP. However, if you do not have an application that fragments traffic, we recommend that you do not allow fragments through the Firepower Threat Defense device. Fragmented packets are often used as DoS attacks.

Fragment Reassembly

The Firepower Threat Defense device performs the following fragment reassembly processes:

- IP fragments are collected until a fragment set is formed or until a timeout interval has elapsed.
- If a fragment set is formed, integrity checks are performed on the set. These checks include no overlapping, no tail overflow, and no chain overflow.

- IP fragments that terminate at the Firepower Threat Defense device are always fully reassembled.
- If **Full Fragment Reassembly** is disabled (the default), the fragment set is forwarded to the transport layer for further processing.
- If **Full Fragment Reassembly** is enabled, the fragment set is first coalesced into a single IP packet. The single IP packet is then forwarded to the transport layer for further processing.

Before you begin

This screen is only available for named interfaces.

Procedure

Step 1 Select **Devices** > **Device Management** and click **Edit** ( ) for your Firepower Threat Defense device. The **Interfaces** page is selected by default.

Step 2 Click **Edit** ( ) for the interface you want to edit.

Step 3 Click the **Advanced** tab, and then click the **Security Configuration** tab.

Step 4 To enable Unicast Reverse Path Forwarding, check the **Anti-Spoofing** check box.

Step 5 To enable full fragment reassembly, check the **Full Fragment Reassembly** check box.

Step 6 To change the number of fragments allowed per packet, check the **Override Default Fragment Setting** check box, and set the following values:

- **Size**—Set the maximum number of packets that can be in the IP reassembly database waiting for reassembly. The default is 200. Set this value to 1 to disable fragments.
- **Chain**—Set the maximum number of packets into which a full IP packet can be fragmented. The default is 24 packets.
- **Timeout**—Set the maximum number of seconds to wait for an entire fragmented packet to arrive. The timer starts after the first fragment of a packet arrives. If all fragments of the packet do not arrive by the number of seconds specified, all fragments of the packet that were already received will be discarded. The default is 5 seconds.

Step 7 Click **OK**.

Step 8 Click **Save**.

You can now click **Deploy** and deploy the policy to assigned devices. The changes are not active until you deploy them.

History for Regular Firewall Interfaces for Firepower Threat Defense

Feature	Version	Details
Integrated Routing and Bridging	6.2.0	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 Firepower Threat Defense bridges instead of routes. The Firepower Threat Defense is not a true bridge in that the Firepower Threat Defense 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 Firepower Threat Defense 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: clustering. The following features are also not supported on BVIs: dynamic routing and multicast routing. New/Modified screens: • **Devices** > **Device Management** > **Interfaces** > **Edit Physical Interface** • **Devices** > **Device Management** > **Interfaces** > **Add Interfaces** > **Bridge Group Interface** Supported platforms: All except for the Firepower 2100 and the Firepower Threat Defense Virtual