Interfaces

The following topics explain how to configure the interfaces on your FTD device.

About FTD Interfaces

FTD includes data interfaces as well as a Management/Diagnostic interface.

When you attach a cable to an interface connection, you need to configure the interface. At minimum, you need to enable the physical interface and give it an IP address. If you intend to create VLAN subinterfaces rather than a single physical interface on a given port, you would typically configure the IP addresses on the subinterface, not on the physical interface. VLAN subinterfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN IDs.

The interface list shows the available interfaces, their names, addresses, and states. You can change the state of an interface, on or off, directly in the list of interfaces. The list shows the interface characteristics based on your configuration.

The following topics explain the limitations of configuring interfaces through the FDM as well as other interface management concepts.

Routed Interfaces

In routed firewall mode, each interface is a Layer 3 routed interface for which you need to set an IP address on a unique subnet.

You can configure both IPv6 and IPv4 addresses on an interface. Make sure you configure a default route for both IPv4 and IPv6.

Management/Diagnostic Interface and Network Deployment

The physical management interface is shared between the Diagnostic logical interface and the Management logical interface.

Management Interface

The Management logical interface is separate from the other interfaces on the device. It is used to run the configuration interface, allow access to the device command line interface (CLI), and to obtain updates for various features. Configure the address on the **System Settings** > **Device Management IP** page. You can configure additional settings at the CLI using the **configure network** command.

Diagnostic Interface

The Diagnostic logical interface can be configured along with the rest of the data interfaces. Using the Diagnostic interface is optional. For example, configure an IP address if you do not want to send system log messages to a remote syslog server through a data interface. The Diagnostic interface only allows management traffic, and does not allow through traffic.

Routed Mode Deployment

We recommend that you do *not* configure an IP address for the Diagnostic interface if you do not have an inside router. The benefit to leaving the IP address off of the Diagnostic interface is that you can place the Management interface on the same network as any other data interfaces. If you configure the Diagnostic interface, its IP address is typically on the same network as the Management IP address, and it counts as a regular interface that cannot be on the same network as any other data interfaces. Because the Management interface requires Internet access for updates, putting Management on the same network as an inside interface means you can deploy the FTD device with only a switch on the inside and point to the inside interface as its gateway. See the following deployment that uses an inside switch:

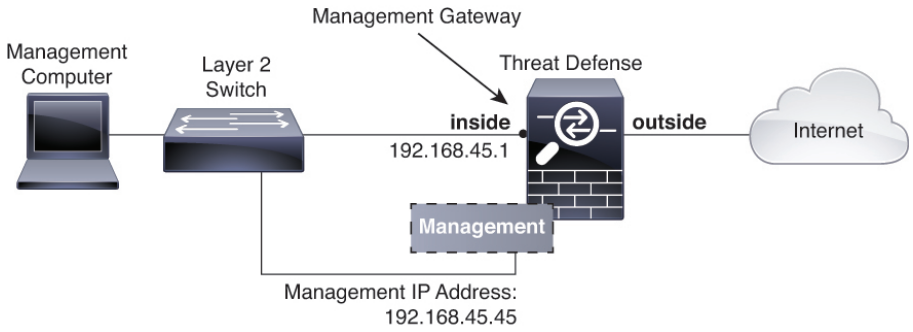

To cable the above scenario on the ASA 5508-X, or ASA 5516-X, see the following:

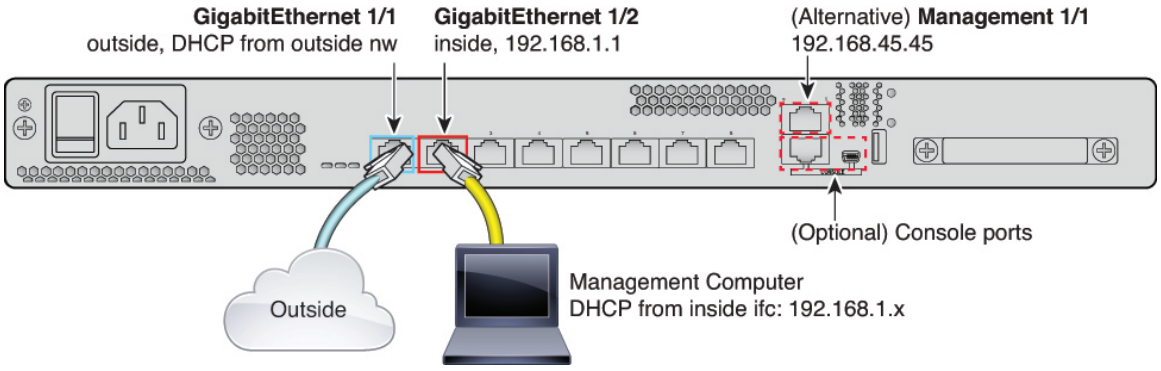

If you configure the Diagnostic IP address, then you need an inside router:

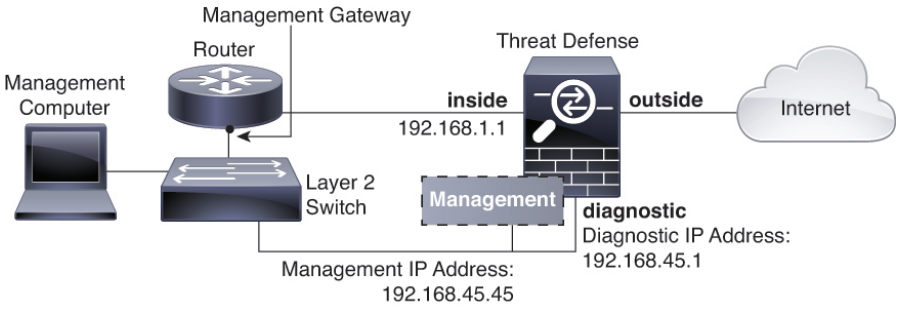

Security Zones

Each interface can be assigned to a single security zone. You then apply your security policy based on zones. For example, you can assign the inside interface to the inside zone; and the outside interface to the outside zone. You can configure your access control policy to enable traffic to go from inside to outside, but not from outside to inside, for example.

You do not include the Management/Diagnostic interface in a zone. Zones apply to data interfaces only.

You can create security zones on the **Objects** page.

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. You cannot specify any of the following as a global address.
  - Internally reserved IPv6 addresses: fd00::/56 (from=fd00:: to= fd00:0000:0000:00ff:ffff:ffff:ffff:ffff)
  - An unspecified address, such as ::/128
  - The loopback address, ::1/128
  - multicast addresses, ff00::/8
  - Link-local addresses, fe80::/10

- 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 Network Discovery functions such as address resolution and neighbor discovery.

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. If you do not configure a global address, then you need to configure the link-local address, either automatically or manually.

Auto-MDI/MDIX Feature

For RJ-45 interfaces, the default auto-negotiation setting also includes the Auto-MDI/MDIX feature. Auto-MDI/MDIX eliminates the need for crossover cabling by performing an internal crossover when a straight cable is detected during the auto-negotiation phase. Either the speed or duplex must be set to auto-negotiate to enable Auto-MDI/MDIX for the interface. If you explicitly set both the speed and duplex to a fixed value, thus disabling auto-negotiation for both settings, then Auto-MDI/MDIX is also disabled. For Gigabit Ethernet, when the speed and duplex are set to 1000 and full, then the interface always auto-negotiates; therefore Auto-MDI/MDIX is always enabled and you cannot disable it.

Guidelines and Limitations for Interfaces

The following topics cover some of the limitations for interfaces.

Limitations for Interface Configuration

When you use the FDM to configure the device, there are several limitations to interface configuration. If you need any of the following features, you must use the FMC to configure the device.

- Routed firewall mode only is supported. You cannot configure transparent firewall mode interfaces.
- You cannot configure passive or ERSPAN interfaces.
- You cannot configure interfaces to be inline (in an inline set), or inline tap, for IPS-only processing. IPS-only mode interfaces bypass many firewall checks and only support IPS security policy. In comparison, 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 firewall mode traffic according to your security policy.
- You cannot configure EtherChannel or redundant interfaces.
- You can only add one bridge group.
- You cannot configure PPPoE for IPv4. If the Internet interface is connected to a DSL, cable modem, or other connection to your ISP, and your ISP uses PPPoE to provide your IP address, you must use the FMC instead of the FDM.
- For the ASA 5512-X, 5515-X, 5525-X, 5545-X, and 5555-X, you can install an optional network interface module. Modules are only discovered during bootstrap (that is, during installation, when switching between local/remote management, and during a major/minor release upgrade, but not patch or hot fix upgrades). For a module that includes SFP interfaces, the FDM sets the speed and duplex to auto; however,

the SFP interfaces do not support the speed and duplex set to auto. You must set the speed and duplex manually. Set the speed to 1000 and the duplex to Full and then deploy the configuration. If the link does not come up, try a different speed.

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
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

Configure a Physical Interface

At minimum, you must enable a physical interface to use it. You would also typically name it and configure IP addressing. You would not configure IP addressing if you intend to create VLAN subinterfaces.

You can disable an interface to temporarily prevent transmission on the connected network. You do not need to remove the interface's configuration.

Procedure

Step 1 Click the name of the device in the menu, then click the link in the **Interfaces** summary.

The interface list shows the available interfaces, their names, addresses, and states.

Step 2 Click the edit icon ( ) for the physical interface you want to edit.

Step 3 Set the following:

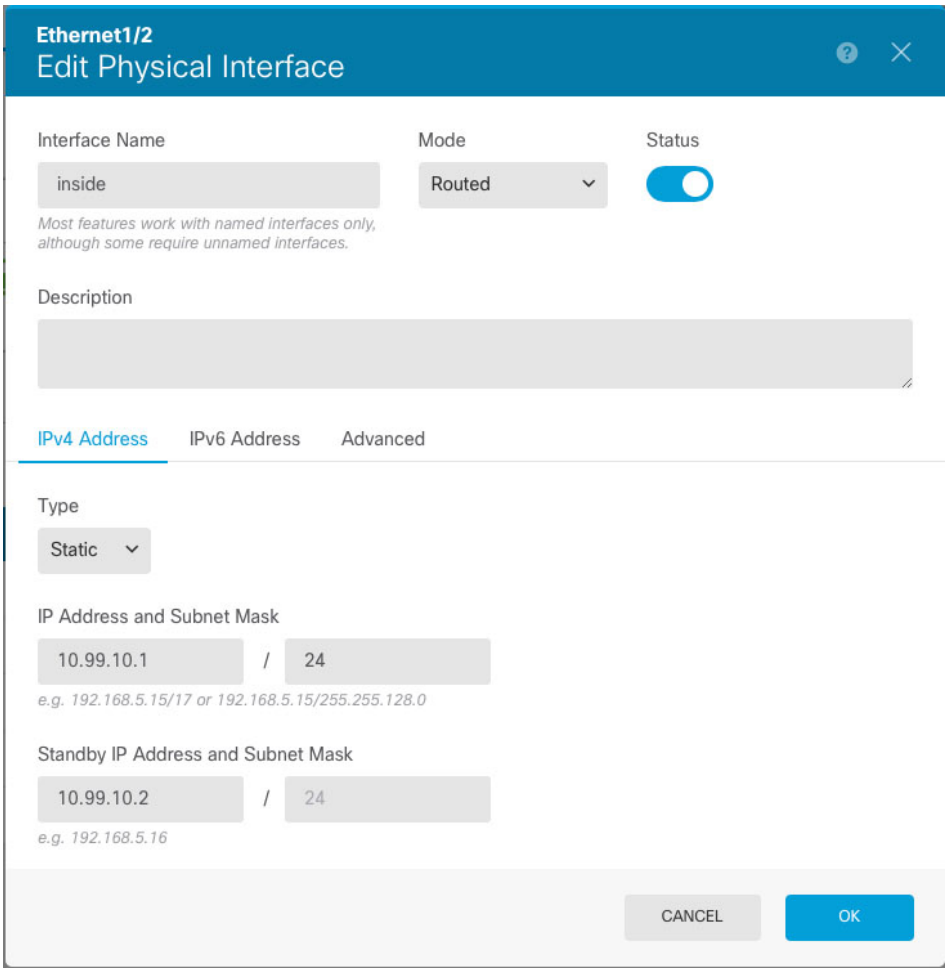

a) Set the **Interface Name**.

Set the name for the interface, up to 48 characters. Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored. Unless you configure subinterfaces, the interface should have a name.

Note If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

b) Set the **Status** slider to the enabled setting ( ).

If you intend to configure subinterfaces for this physical interface, you are probably done. Click **Save** and continue with Configure VLAN Subinterfaces and 802.1Q Trunking, on page 8. Otherwise, continue.

Note Even when configuring subinterfaces, it is valid to name the interface and supply IP addresses. This is not the typical setup, but if you know that is what you need, you can configure it.

c) (Optional) Set the **Description**.

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

Step 4 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. Change the following options if necessary:
  - **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
  - **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.
- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

  Note For an existing interface, your ability to change the address is constrained if you have a DHCP server configured for the interface. The new IP address must be on the same subnet as the DHCP address pool, and it cannot be part of that pool. If you need to configure an address on a different subnet, first delete the DHCP server configuration. See Configuring DHCP Server.

Step 5 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the FTD device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 3.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Suppress RA**—Whether to suppress router advertisements. The FTD can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

  Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the FTD device to supply the IPv6 prefix (for example, the outside interface).

Step 6 (Optional.) Configure Advanced Options, on page 12.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 7 Click **Save**.

What to do next

- Add the interfaces to the appropriate security zones. See Configuring Security Zones.

Configure VLAN Subinterfaces and 802.1Q Trunking

VLAN subinterfaces let you divide a physical 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 allow you to 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.

Create subinterfaces if you attach the physical interface to a trunk port on a switch. Create a subinterface for each VLAN that can appear on the switch trunk port. If you attach the physical interface to an access port on the switch, there is no point in creating a subinterface.

Guidelines and Limitations

- 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. Because the physical interface must be enabled for the subinterface to pass traffic, ensure that the physical interface does not pass traffic by not naming the interface. If you want to let the physical interface pass untagged packets, you can name the interface as usual.
- All subinterfaces on the same parent interface must be either bridge group members or routed interfaces; you cannot mix and match.
- 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 device, 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.

Procedure

Step 1 Click the name of the device in the menu, then click the link in the **Interfaces** summary.

The interface list shows the available interfaces, their names, addresses, and states.

Step 2 Do one of the following:

- Click the + button to create a new subinterface.
- Click the edit icon ( ) for the subinterface you want to edit.

If you no longer need a subinterface, click the delete icon ( ) for the subinterface to delete it.

Step 3 Set the **Status** slider to the enabled setting ( ).

Step 4 Configure the parent interface, name, and description:

a) Choose the **Parent Interface**.

The parent interface is the physical interface to which you want to add the subinterface. You cannot change the parent interface after you create the subinterface.

b) Set the **Subinterface Name**, up to 48 characters.

Alphabetic characters must be lower case. For example, **inside** or **outside**. Without a name, the rest of the interface configuration is ignored.

Note If you change the name, the change is automatically reflected everywhere you used the old name, including security zones, syslog server objects, and DHCP server definitions. However, you cannot remove the name until you first remove all configurations that use the name, because you typically cannot use an unnamed interface for any policy or setting.

c) (Optional) Set a **Description**.

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

d) Set the **VLAN ID**.

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

e) Set the **Subinterface ID**.

Enter the subinterface ID as an integer between 1 and 4294967295. This ID is appended to the interface ID; for example Ethernet1/1.100. You can match the VLAN ID for convenience, but it is not required. You cannot change the ID after you create the subinterface.

Step 5 Click the **IPv4 Address** tab and configure the IPv4 address.

Select one of the following options from the **Type** field:

- **DHCP**—Choose this option if the address should be obtained from the DHCP server on the network. Change the following options if necessary:

- **Route Metric**—If you obtain the default route from the DHCP server, the administrative distance to the learned route, between 1 and 255. The default is 1.
- **Obtain Default Route**—Whether to get the default route from the DHCP server. You would normally select this option, which is the default.

- **Static**—Choose this option if you want to assign an address that should not change. Type in the interface's IP address and the subnet mask for the network attached to the interface. For example, if you attach the 10.100.10.0/24 network, you could enter 10.100.10.1/24. Ensure that the address is not already used on the network.

  Note For an existing interface, your ability to change the address is constrained if you have a DHCP server configured for the interface. The new IP address must be on the same subnet as the DHCP address pool, and it cannot be part of that pool. If you need to configure an address on a different subnet, first delete the DHCP server configuration. See Configuring DHCP Server.

Step 6 (Optional.) Click the **IPv6 Address** tab and configure the IPv6 address.

- **State**—To enable IPv6 processing and to automatically configure the link-local address when you do not configure the global address, select **Enabled**. The link local address is generated based on the interface MAC addresses (*Modified* EUI-64 format).

  Note Disabling IPv6 does not disable IPv6 processing on an interface that is configured with an explicit IPv6 address or that is enabled for autoconfiguration.

- **Address Auto Configuration**—Select this option to have the address automatically configured. IPv6 stateless autoconfiguration will generate a global IPv6 address only if the link on which the device resides has a router configured to provide IPv6 services, including the advertisement of an IPv6 global prefix for use on the link. If IPv6 routing services are not available on the link, you will get a link-local IPv6 address only, which you cannot access outside of the device's immediate network link. The link local address is based on the Modified EUI-64 interface ID.

  Although RFC 4862 specifies that hosts configured for stateless autoconfiguration do not send Router Advertisement messages, the FTD device does send Router Advertisement messages in this case. Select **Suppress RA** to suppress messages and conform to the RFC.

- **Static Address/Prefix**—If you do not use stateless autoconfiguration, enter the full static global IPv6 address and network prefix. For example, 2001:0DB8::BA98:0:3210/48. For more information on IPv6 addressing, see IPv6 Addressing, on page 3.

  If you want to use the address as link local only, select the **Link - Local** option. Link local addresses are not accessible outside the local network. You cannot configure a link-local address on a bridge group interface.

  Note A link-local address should start with FE8, FE9, FEA, or FEB, for example fe80::20d:88ff:feee:6a82. 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.

- **Suppress RA**—Whether to suppress router advertisements. The FTD can participate in router advertisements so that neighboring devices can dynamically learn a default router address. By default, router advertisement messages (ICMPv6 Type 134) are periodically sent out each IPv6 configured interface.

Router advertisements are also 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 might want to suppress these messages on any interface for which you do not want the FTD device to supply the IPv6 prefix (for example, the outside interface).

Step 7 (Optional.) Configure Advanced Options, on page 12.

The advanced settings have defaults that are appropriate for most networks. Edit them only if you are resolving network issues.

Step 8 Click **Save**.

What to do next

- Add the subinterfaces to the appropriate security zones. See Configuring Security Zones.

Configure Advanced Interface Options

Advanced options include setting the MTU, hardware settings, management only, MAC address, and other settings.

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.

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 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—A jumbo frame is an Ethernet packet larger than the standard maximum of 1522 bytes (including Layer 2 header and VLAN header), up to 9216 bytes. You can set the MTU to 9000 bytes or higher to accommodate jumbo frames. The maximum depends on the model.

Note Increasing the MTU assigns more memory for jumbo frames, which might limit the maximum usage of other features, such as access rules. If you increase the MTU above the default 1500 on ASA 5500-X series devices , you must reboot the system.

Configure Advanced Options

Advanced interface options have default settings that are appropriate for most networks. Configure them only if you are resolving networking problems.

The following procedure assumes the interface is already defined. You can also edit these settings while initially editing or creating the interface.

Procedure

Step 1 Click the name of the device in the menu, then click the link in the **Interfaces** summary.

The interface list shows the available interfaces, their names, addresses, and states.

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

Step 3 Click **Advanced Options**.

Step 4 To make a data interface management only, select **Management Only**.

A management only interface does not allow through traffic, so there is very little value in setting a data interface as management only. You cannot change this setting for the Management/Diagnostic interface, which is always management only.

Step 5 Change the **MTU** (maximum transmission unit) to the desired value.

The default MTU is 1500 bytes. The minimum and maximum depend on your platform. Set a high value if you typically see jumbo frames on your network.

Note If you increase MTU above 1500 on ASA 5500-X series devices, , you must reboot the device.

Step 6 (Physical interface only.) Modify the speed and duplex settings.

The default is that the interface negotiates the best duplex and speed with the interface at the other end of the wire, but you can force a specific duplex or speed if necessary. Before setting these options for interfaces on a network module, please read Limitations for Interface Configuration, on page 4.

- **Duplex**—Choose **Auto**, **Half**, or **Full**. Auto is the default.
- **Speed**—Choose **10**, **100**, **1000** Mbps, or **Auto**. Auto is the default.

Step 7 Modify the **IPv6 Configuration** settings.

- **Enable DHCP for IPv6 address configuration**—Whether to set the Managed Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain addresses, in addition to the derived stateless autoconfiguration address.
- **Enable DHCP for IPv6 non-address configuration**—Whether to set the Other Address Configuration flag in the IPv6 router advertisement packet. This flag informs IPv6 autoconfiguration clients that they should use DHCPv6 to obtain additional information from DHCPv6, such as the DNS server address.
- **DAD Attempts**—How often the interface performs Duplicate Address Detection (DAD), from 0 - 600. The default is 1. During the stateless autoconfiguration process, DAD verifies the uniqueness of new unicast IPv6 addresses before the addresses are assigned to interfaces. 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. The interface uses neighbor solicitation messages to perform Duplicate Address Detection. Set the value to 0 to disable duplicate address detection (DAD) processing.

Step 8 Click **OK**.

Monitoring Interfaces

You can view some basic information about interfaces in the following areas:

- **Device**. Use the port graphic to monitor the current state of the interfaces. Mouse over a port to see its IP addresses and enabled and link statuses. The IP addresses can be statically assigned or obtained using DHCP.

  Interface ports use the following color coding:

  - Green—The interface is configured, enabled, and the link is up.
  - Gray—The interface is not enabled.
  - Orange/Red—The interface is configured and enabled, but the link is down. If the interface is wired, this is an error condition that needs correction. If the interface is not wired, this is the expected status.

- **Monitoring** > **System**. The **Throughput** dashboard shows information on traffic flowing through the system. You can view information on all interfaces, or you can select a specific interface to examine.
- **Monitoring** > **Ingress Zones** and **Egress Zones**. These dashboards show statistics based on zones, which are composed of interfaces. You can drill into this information for more detail.

Monitoring Interfaces in the CLI

You can also log into the device CLI and use the following commands to get more detailed information about interface-related behavior and statistics.

- **show interface** displays interface statistics and configuration information. This command has many keywords you can use to get to the information you need. Use ? as a keyword to see the available options.
- **show ipv6 interface** displays IPv6 configuration information about the interfaces.
- **show bridge-group** displays information about Bridge Virtual Interfaces (BVI), including member information and IP addresses.
- **show conn** displays information about the connections currently established through the interfaces.
- **show traffic** displays statistics about traffic flowing through each interface.
- **show ipv6 traffic** displays statistics about IPv6 traffic flowing through the device.
- **show dhcpd** displays statistics and other information about DHCP usage on the interfaces, particularly about the DHCP servers configured on interfaces.

Examples for Interfaces

The use case chapter includes the following interface-related examples:

- How to Configure the Device in FDM
- How to Add a Subnet