



Setting Up Aggregate Interfaces

You can group multiple physical Ethernet interfaces into a single logical link on Series 3 managed devices configured in either a Layer 2 deployment that provides packet switching between networks, or a Layer 3 deployment that routes traffic between interfaces. This single aggregate logical link provides higher bandwidth, redundancy, and load-balancing between two endpoints.

You create aggregate links by creating a switched or routed link aggregation group, or LAG. When you create an aggregation group, a logical interface called an aggregate interface is created. To an upper layer entity a LAG looks like a single logical link and data traffic is transmitted through the aggregate interface. The aggregate link provides increased bandwidth by adding the bandwidth of multiple links together. It also provides redundancy by load-balancing traffic across all available links. If one link fails, the system automatically load-balances traffic across all remaining links.



The endpoints in a LAG can be two FirePOWER managed devices, as shown in the illustration above, or a FirePOWER managed device connected to a third-party access switch or router. The two devices do not have to match, but they must have the same physical configuration and they must support the IEEE 802.3ad link aggregation standard. A typical deployment for a LAG might be to aggregate access links between two managed devices, or to create a point-to-point connection between a managed device and an access switch or a router.

Note that you cannot configure aggregate interfaces on a virtual managed device, a Cisco ASA with FirePOWER Services device, or a Cisco NGIPS for Blue Coat X-Series device.

For more information about setting up aggregate interfaces, see [Configuring LAGs, page 8-1](#).

Configuring LAGs

License: Control

Supported Devices: Series 3

There are two types of aggregate interfaces: switched, which are Layer 2 aggregate interfaces, and routed, which are Layer 3 aggregate interfaces. You implement link aggregation through the use of link aggregation groups (LAGs). You configure a LAG by creating an aggregate switched or routed interface and then associating a set of physical interfaces with the link. All of the physical interfaces must be of the same speed and medium.

You create aggregate links either dynamically or statically. Dynamic link aggregation uses Link Aggregation Control Protocol (LACP), a component of the IEEE 802.3ad link aggregation standard, while static link aggregation does not. LACP enables each device on either end of the LAG to exchange link and system information to determine which links will be actively used in the aggregation. A static LAG configuration requires you to manually maintain link aggregations and apply load-balancing and link selection policies.

When you create a switched or routed aggregate interface, a link aggregation group of the same type is created and numbered automatically. For example, when you create your first LAG (switched or routed), the aggregate interface can be identified by the **lag0** label in the Interfaces tab for your managed device. When you associate physical and logical interfaces with this LAG, they appear nested below the primary LAG in a hierarchical tree menu. Note that a switched LAG can only contain switched physical interfaces, and a routed LAG can only contain routed physical interfaces.

Consider the following requirements when you configure a LAG:

- The FireSIGHT system supports a maximum of 14 LAGs, and assigns a unique ID to each LAG interface in the range of 0 to 13. The LAG ID is not configurable.
- You must configure the LAG on both sides of the link, and you must set the interfaces on either side of the link to the same speed.
- You must associate at least two physical interfaces per LAG, up to a maximum of eight. A physical interface cannot belong to more than one LAG.
- Physical interfaces in a LAG cannot be used in any other mode of operation, either as inline or passive, or be used as part of another logical interface for tagged traffic.
- Physical interfaces in a LAG can span multiple NetMods, but cannot span multiple sensors (i.e. all physical interfaces must reside on the same device).
- A LAG cannot contain a stacking NetMod.


Note

Link aggregation is not supported on device clusters.

See the following sections for more information:

- [Specifying a Load-Balancing Algorithm, page 8-2](#)
- [Specifying a Link Selection Policy, page 8-3](#)
- [Configuring LACP, page 8-4](#)
- [Adding Aggregate Switched Interfaces, page 8-5](#)
- [Adding Aggregate Routed Interfaces, page 8-7](#)
- [Adding Logical Aggregate Interfaces, page 8-10](#)
- [Viewing Aggregate Interface Statistics, page 8-12](#)
- [Deleting Aggregate Interfaces, page 8-12](#)

Specifying a Load-Balancing Algorithm

License: Control

Supported Devices: Series 3

You assign an egress load-balancing algorithm to the LAG that determines how to distribute traffic to the LAG bundle's member links. The load-balancing algorithm makes hashing decisions based on values in various packet fields, such as Layer 2 MAC addresses, Layer 3 IP addresses, and Layer 4 port numbers (TCP/UDP traffic). The load-balancing algorithm you select applies to all of the LAG bundle's member links.

Select the load-balancing algorithm that supports your deployment scenario from the following options when you configure a LAG:

- Destination IP
- Destination MAC
- Destination Port
- Source IP
- Source MAC
- Source Port
- Source and Destination IP
- Source and Destination MAC
- Source and Destination Port

**Note**

You should configure both ends of the LAG to have the same load-balancing algorithm. Higher layer algorithms will back off to lower layer algorithms as necessary (such as a Layer 4 algorithm backing off to Layer 3 for ICMP traffic).

Specifying a Link Selection Policy

License: Control

Supported Devices: Series 3

Link aggregation requires the speed and medium of each link to be the same at both endpoints. Because link properties can change dynamically, the link selection policy helps determine how the system manages the link selection process. A link selection policy that maximizes the highest port count supports link redundancy, while a link selection policy that maximizes total bandwidth supports overall link speed. A stable link selection policy attempts to minimize excessive changes in link states.

**Note**

You should configure both ends of the LAG to have the same link selection policy.

Select the link selection policy that supports your deployment scenario from the following options when you configure a LAG:

- **Highest Port Count** — select this option for the highest total active port count to provide added redundancy.
- **Highest Total Bandwidth** — select this option to provide the highest total bandwidth for the aggregated link.
- **Stable** — select this option if your primary concern is link stability and reliability. Once you configure a LAG, the active links change only when absolutely necessary (such as link failure) rather than doing so for added port count or bandwidth.

- **LACP Priority** — select this option to use the LACP algorithm to determine which links are active in the LAG. This setting is appropriate if you have undefined deployment goals, or if the device at the other end of the LAG is a non-FirePOWER device.

When LACP is enabled, a link selection policy based on LACP priority uses two properties of LACP, the system priority and link priority, which are described as follows:

- **LACP system priority.** You configure this value on each partnered device running LACP to determine which one is superior in link aggregation. The system with the lower value has the higher system priority. In dynamic link aggregation, the system with the higher LACP system priority sets the selected state of member links on its side first, then the system with the lower priority sets its member links accordingly. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768.
- **LACP link priority.** You configure this value on each link belonging to the aggregation group. The link priority determines the active and standby links in the LAG. Links with lower values have higher priority. If an active link goes down, the standby link with the highest priority is selected to replace the downed link. However, if two or more links have the same LACP link priority, the link with the lowest physical port number is selected as the standby link. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768.

LACP is a key aspect of automating the link selection method that supports dynamic link aggregation. For more information, see [Configuring LACP, page 8-4](#).

Configuring LACP

License: Control

Supported Devices: Series 3

Link Aggregation Control Protocol (LACP), a component of IEEE 802.3ad, is a method of exchanging system and port information to create and maintain LAG bundles. When you enable LACP, each device on either end of the LAG uses LACP to determine which links will be actively used in the aggregation. LACP provides availability and redundancy by exchanging LACP packets (or control messages) between links. It learns the capabilities of the links dynamically and informs the other links. Once LACP identifies correctly matched links, it facilitates grouping the links into the LAG. If a link fails, traffic continues on the remaining links. LACP must be enabled at both ends of the LAG for the link to be operational.

When you enable LACP, you need to select a transmission mode for each end of the LAG that determines how LACP packets are exchanged between partnered devices. There are two options for LACP mode:

- **Active** — select this mode to place a device into an active negotiating state, in which the device initiates negotiations with remote links by sending LACP packets.
- **Passive** — select this mode to place a device into a passive negotiating state, in which the device responds to LACP packets it receives but does not initiate LACP negotiation.



Note

Both modes allow LACP to negotiate between links to determine if they can form a link bundle based on criteria such as port speed. However, you should avoid a passive-passive configuration, which essentially places both ends of the LAG in listening mode.

LACP has a timer which defines how often LACP packets are sent between devices. LACP exchanges packets at these rates:

- **Slow** — 30 seconds

- Fast — 1 second

The device where this option is applied expects to receive LACP packets with this frequency from the partner device on the other side of the LAG.

**Note**

When a LAG is configured on a managed device that is part of a device stack, only the primary device participates in LACP communication with the partner system. All secondary devices forward LACP messages to the primary device. The primary device relays any dynamic LAG modifications to the secondary devices.

Adding Aggregate Switched Interfaces

License: Control

Supported Devices: Series 3

You can combine between two and eight physical ports on a managed device to create a switched LAG interface. You must assign a switched LAG interface to a virtual switch before it can handle traffic. A managed device can support up to 14 LAG interfaces.

**Caution**

Changing any (Series 2) or the highest (Series 3) MTU value for a sensing interface or inline set temporarily interrupts traffic inspection on all sensing interfaces on the device, not just the interface you changed, when you apply your changes. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and the interface type. See [How Snort Restarts Affect Traffic, page 1-9](#).

To edit an existing switched LAG interface, click the edit icon (✎) next to the interface.

To configure a switched LAG interface:

Access: Admin/Network Admin

-
- Step 1** Select **Devices > Device Management**.
The Device Management page appears.
- Step 2** Next to the device where you want to configure the switched LAG interface, click the edit icon (✎).
The Interfaces tab appears.
- Step 3** From the **Add** drop-down menu, select **Add Aggregate Interface**.
- Step 4** Click **Switched** to display the switched LAG interface options.
- Step 5** Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
- Step 6** From the **Virtual Switch** drop-down list, you must select an existing virtual switch or select **New** to add a new virtual switch.

**Note**

If you add a new virtual switch, you must configure it on the Virtual Switches tab of the Device Management page (**Devices > Device Management > Virtual Switches**) after you set up the switched interface. See [Adding Virtual Switches, page 6-6](#).

- Step 7** Select the **Enabled** check box to allow the switched LAG interface to handle traffic. If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.
- Step 8** From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the interface is configured to auto negotiate speed and duplex settings. Note that mode settings are available only for copper interfaces.





Note Interfaces on 8000 Series appliances do not support half-duplex options. When links auto negotiate speed, all active links are selected for the LAG based on the same speed setting.



- Step 9** From the **MDI/MDIX** drop-down list, select an option to designate whether the interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.

By default, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.

- Step 10** In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed.
- The range within which you can set the MTU can vary depending on the FireSIGHT System device model and the interface type. See [MTU Ranges for Managed Devices, page 4-64](#) for more information.

- Step 11** Under **Link Aggregation**, you have two options for selecting physical interfaces to add to the LAG bundle:
- Next to **Available Interfaces**, select one or more interfaces, then click the add selected icon (). Use **Ctrl** or **Shift** to select multiple physical interfaces.
 - To add all interface pairs to the LAG bundle, click the add all icon (.



Tip To remove physical interfaces from the LAG bundle, select one or more physical interfaces and click the remove selected icon (). To remove all physical interfaces from the LAG bundle, click the remove all icon (). Deleting the LAG interface from the Interfaces tab also removes the interfaces.

- Step 12** From the **Load-Balancing Algorithm** drop-down list, select the option that supports your deployment scenario. See [Specifying a Load-Balancing Algorithm, page 8-2](#) for more information.
- Step 13** From the **Link Selection Policy** drop-down list, select the option that supports your deployment scenario: Highest Port Count (redundancy), Highest Total Bandwidth (speed), Stable (no excessive change in maintain link state), or LACP Priority (automatic link aggregation).
- If you select **LACP Priority**, you need to assign a value for **System Priority**. You then need to click the **Configure Interface Priority** link to assign a priority value for each interface in the LAG. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768. See [Specifying a Link Selection Policy, page 8-3](#) for more information.



Note Select **LACP Priority** when you configure an aggregate interface between a FireSIGHT System device and a third-party network device.

- Step 14** From the **Tunnel Level** drop-down list, select the option that supports your deployment scenario, either **Inner** or **Outer**.

Note that the tunnel level only applies to IPv4 traffic when Layer 3 load balancing is configured. The outer tunnel is always used for Layer 2 and IPv6 traffic. If the **Tunnel Level** is not explicitly set, the default is **Outer**.

Step 15 Under **LACP**, select the **Enabled** check box to allow the switched LAG interface to handle traffic using the Link Aggregation Control Protocol. See [Configuring LACP, page 8-4](#) for more information.

If you clear the check box, the LAG interface becomes a static configuration and the FireSIGHT System will use all of the physical interfaces selected for the aggregation.

Step 16 Select a **Rate** radio button to set the frequency that determines how often LACP control messages are received from the partner device.

- Select **Slow** to receive packets every 30 seconds.
- Select **Fast** to receive packets every 1 second.

Step 17 Select a **Mode** radio button to establish the listening mode of the device.

- Select **Active** to initiate negotiations with remote links by sending LACP packets to the partner device.
- Select **Passive** to respond to LACP packets received.

Step 18 Click **Save**.

The switched LAG interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices, page 4-25](#) for more information.

Adding Aggregate Routed Interfaces

License: Control

Supported Devices: Series 3

You can combine between two and eight physical ports on a managed device to create a routed LAG interface. You must assign a routed LAG interface to a virtual router before it can route traffic. A managed device can support up to 14 LAG interfaces.



Caution

Adding a routed interface pair on a Series 3 device restarts the Snort process when you apply your changes, temporarily interrupting traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and how it handles traffic. See [How Snort Restarts Affect Traffic, page 1-9](#) for more information.


You can add static Address Resolution Protocol (ARP) entries to a routed LAG interface. If an external host needs to know the MAC address of the destination IP address it needs to send traffic to on your local network, it sends an ARP request. When you configure static ARP entries, the virtual router responds with an IP address and associated MAC address.

Note that disabling the **ICMP Enable Responses** option for routed LAG interfaces does not prevent ICMP responses in all scenarios. You can add rules to an access control policy to drop packets where the destination IP is the routed interface's IP and the protocol is ICMP; see [Controlling Traffic with Network-Based Rules, page 15-1](#).

If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see [Understanding Advanced Device Settings, page 4-54](#).


**Caution**

Changing any (Series 2) or the highest (Series 3) MTU value for a sensing interface or inline set temporarily interrupts traffic inspection on all sensing interfaces on the device, not just the interface you changed, when you apply your changes. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and the interface type. See [How Snort Restarts Affect Traffic, page 1-9](#).

To edit an existing routed LAG interface, click the edit icon () next to the interface.

To configure a routed LAG interface:

Access: Admin/Network Admin

-
- Step 1** Select **Devices > Device Management**.
- The Device Management page appears.
- Step 2** Next to the device where you want to configure the routed LAG interface, click the edit icon ().
- The Interfaces tab for that device appears.
- Step 3** From the **Add** drop-down menu, select **Add Aggregate Interface**.
- Step 4** Click **Routed** to display the routed LAG interface options.
- Step 5** Optionally, from the **Security Zone** drop-down list, select an existing security zone or select **New** to add a new security zone.
- Step 6** From the **Virtual Router** drop-down list, you must select an existing virtual router or select **New** to add a new virtual router.

**Note**

If you add a new virtual router, you must configure it on the Virtual Routers tab of the Device Management page (**Devices > Device Management > Virtual Routers**) after you set up the routed interface. See [Adding Virtual Routers, page 7-9](#).

- Step 7** Select the **Enabled** check box to allow the routed LAG interface to handle traffic.
- If you clear the check box, the interface becomes disabled so that users cannot access it for security purposes.
- Step 8** From the **Mode** drop-down list, select an option to designate the link mode or select **Autonegotiation** to specify that the LAG interface is configured to auto negotiate speed and duplex settings. Note that mode settings are available only for copper interfaces.

**Note**

Interfaces on 8000 Series appliances do not support half-duplex options. When links auto negotiate speed, all active links are selected for the LAG based on the same speed setting.

- Step 9** From the **MDI/MDIX** drop-down list, select an option to designate whether the LAG interface is configured for MDI (medium dependent interface), MDIX (medium dependent interface crossover), or Auto-MDIX. Note that MDI/MDIX settings are available only for copper interfaces.
- Normally, MDI/MDIX is set to Auto-MDIX, which automatically handles switching between MDI and MDIX to attain link.
- Step 10** In the **MTU** field, type a maximum transmission unit (MTU), which designates the largest size packet allowed. Note that the MTU is the Layer 2 MTU/MRU and not the Layer 3 MTU.

The range within which you can set the MTU can vary depending on the FireSIGHT System device model and the interface type. See [MTU Ranges for Managed Devices, page 4-64](#) for more information.

- Step 11** Next to **ICMP**, select the **Enable Responses** check box to allow the LAG interface to respond to ICMP traffic such as pings and traceroute.
- Step 12** Next to **IPv6 NDP**, select the **Enable Router Advertisement** check box to enable the LAG interface to broadcast router advertisements.
- Step 13** To add an IP address, click **Add**.
The Add IP Address pop-up window appears.
- Step 14** In the **Address** field, type the routed LAG interface's IP address and subnet mask using CIDR notation. Note the following:
- You cannot add network and broadcast addresses, or the static MAC addresses 00:00:00:00:00:00 and FF:FF:FF:FF:FF:FF.
 - You cannot add identical IP addresses, regardless of subnet mask, to interfaces in virtual routers.
- Step 15** Optionally, if your organization uses IPv6 addresses, next to the **IPv6** field, select the **Address Autoconfiguration** check box to set the IP address of the LAG interface automatically.
- Step 16** For **Type**, select either Normal or SFRP.
For SFRP options, see [Configuring SFRP, page 7-7](#) for more information.
- Step 17** Click **OK**.
The IP address is added.
To edit an IP address, click the edit icon (✎). To delete an IP address, click the delete icon (🗑).



Note When adding an IP address to a routed interface of a clustered device, you must add a corresponding IP address to the routed interface on the cluster peer.



- Step 18** To add a static ARP entry, click **Add**.
The Add Static ARP Entry pop-up window appears.
- Step 19** In the **IP Address** field, type an IP address for the static ARP entry.
- Step 20** In the **MAC Address** field, type a MAC address to associate with the IP address. Enter the address using the standard format of six groups of two hexadecimal digits separated by colons (for example, 01:23:45:67:89:AB).
- Step 21** Click **OK**.
The static ARP entry is added.



Tip To edit a static ARP entry, click the edit icon (✎). To delete a static ARP entry, click the delete icon (🗑).

- Step 22** Under **Link Aggregation**, you have two options for selecting physical interfaces to add to the LAG bundle:
- Next to **Available Interfaces**, select one or more interfaces, then click the add selected icon (➡). Use **Ctrl** or **Shift** to select multiple physical interfaces.
 - To add all interface pairs to the LAG bundle, click the add all icon (👉).

**Tip**

To remove physical interfaces from the LAG bundle, select one or more physical interfaces and click the remove selected icon (). To remove all physical interfaces from the LAG bundle, click the remove all icon (). Deleting the LAG interface from the Interfaces tab also removes the interfaces.

- Step 23** From the **Load-Balancing Algorithm** drop-down list, select the option that supports your deployment scenario. See [Specifying a Load-Balancing Algorithm, page 8-2](#) for more information.
- Step 24** From the **Link Selection Policy** drop-down list, select the option that supports your deployment scenario: Highest Port Count (redundancy), Highest Total Bandwidth (speed), Stable (no excessive change in maintain link state), or LACP Priority (automatic link aggregation).
- If you select **LACP Priority**, you need to assign a value for **System Priority**. You then need to click the **Configure Interface Priority** link to assign a priority value for each interface in the LAG. You can specify 0 to 65535. If you do not specify a value, the default priority is 32768. See [Specifying a Link Selection Policy, page 8-3](#) for more information.

**Note**

Select **LACP Priority** when you configure an aggregate interface between a FireSIGHT System device and a third-party network device.

- Step 25** From the **Tunnel Level** drop-down list, select the option that supports your deployment scenario, either **Inner** or **Outer**.
- Note that the tunnel level only applies to IPv4 traffic when Layer 3 load balancing is configured. The outer tunnel is always used for Layer 2 and IPv6 traffic. If the **Tunnel Level** is not explicitly set, the default is **Outer**.
- Step 26** Under **LACP**, select the **Enabled** check box to allow the switched LAG interface to handle traffic using the Link Aggregation Control Protocol. See [Configuring LACP, page 8-4](#) for more information.
- If you clear the check box, the LAG interface becomes a static configuration and the FireSIGHT System will use all of the physical interfaces for the aggregation.
- Step 27** Select a **Rate** radio button to set the frequency that determines how often LACP control messages are received from the partner device.
- Select **Slow** to receive packets every 30 seconds.
 - Select **Fast** to receive packets every 1 second.
- Step 28** Select a **Mode** radio button to establish the listening mode of the device.
- Select **Active** to initiate negotiations with remote links by sending LACP packets to the partner device.
 - Select **Passive** to respond to LACP packets received.
- Step 29** Click **Save**.
- The routed LAG interface is configured. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices, page 4-25](#).

Adding Logical Aggregate Interfaces

License: Control

Supported Devices: Series 3

For each switched or routed aggregate interface, you can add multiple logical interfaces. You must associate each logical LAG interface with a VLAN tag to handle traffic received by the LAG interface with that specific tag. You add logical interfaces to switched or routed aggregate interfaces in the same way you would add them to physical switched or routed interfaces.

**Note**

When you create a LAG interface you also create an “untagged” logical interface by default, which is identified by the **lag n .0** label, where n is an integer from 0 to 13. To be operational, each LAG requires this one logical interface at a minimum. You can associate additional logical interfaces with any LAG to handle VLAN-tagged traffic. Each additional logical interface requires a unique VLAN tag. The FireSIGHT system supports VLAN tags in the range of 1 through 4094.


You can also configure SFRP on a logical routed interface. See [Configuring SFRP, page 7-7](#) for more information.

Note that disabling the **ICMP Enable Responses** option for logical routed interfaces does not prevent ICMP responses in all scenarios. You can add rules to an access control policy to drop packets where the destination IP is the routed interface’s IP and the protocol is ICMP; see [Controlling Traffic with Network-Based Rules, page 15-1](#).

If you have enabled the **Inspect Local Router Traffic** option on the managed device, it drops the packets before they reach the host, thereby preventing any response. For more information about inspecting local router traffic, see [Understanding Advanced Device Settings, page 4-54](#).


**Caution**

Changing any (Series 2) or the highest (Series 3) MTU value for a sensing interface or inline set temporarily interrupts traffic inspection on all sensing interfaces on the device, not just the interface you changed, when you apply your changes. Whether traffic drops during this interruption or passes without further inspection depends on the model of the managed device and the interface type. See [How Snort Restarts Affect Traffic, page 1-9](#).

To edit an existing logical LAG interface, click the edit icon () next to the interface.

To add a logical LAG interface:

Access: Admin/Network Admin

-
- Step 1** Select **Devices > Device Management**.
The Device Management page appears.
 - Step 2** Next to the device where you want to add the logical LAG interface, click the edit icon ().
The Interfaces tab appears.
 - Step 3** From the **Add** drop-down menu, select **Add Logical Interface**.
The Add Interface pop-up window appears.
 - Step 4** Click **Switched** to display the switched interface options, or select **Routed** to display the routed interface options.

When you create a logical interface for a LAG, you select an available LAG from the **Interface** drop-down list. The aggregate interface is identified by the **lag n** label, where n is an integer from 0 to 13.

See [Adding Logical Switched Interfaces, page 6-3](#) for more information on adding a logical interface to a switched interface.

See [Adding Logical Routed Interfaces, page 7-4](#) for more information on adding a logical interface to a routed interface.

**Note**

When an aggregate interface is disabled, the logical interface associated with the aggregate interface is also disabled.

Viewing Aggregate Interface Statistics

License: Control



Supported Devices: Series 3

You can view protocol and traffic statistics for each aggregate interface. The statistics show LACP protocol information such as LACP key and partner information, packets received, packets transmitter, and packets dropped. Statistics are further refined per member interface to show traffic and link information on a per-port basis.

Aggregate interface information is also presented to the dashboard via predefined dashboard widgets. The Current Interface Status widget shows the status of all interfaces on the appliance, enabled or unused. The Interface Traffic widget shows the rate of traffic received (Rx) and transmitted (Tx) on the appliance's interfaces over the dashboard time range. See [Understanding the Predefined Widgets, page 55-7](#).

To view aggregate interface statistics:

Access: Admin/Network Admin

-
- Step 1** Select **Devices > Device Management**.
The Device Management page appears.
 - Step 2** Next to the device where you want to view the logical aggregate interface statistics, click the edit icon ().
The Interfaces tab for that device appears.
 - Step 3** Next to the interface where you want to view the interface statistics, click the view icon ().
The Statistics pop-up window appears.
 - Step 4** Click **OK** to close the window.
-

Deleting Aggregate Interfaces

License: Control

Supported Devices: Series 3

The following procedure explains how to delete an aggregate interface.

To delete an aggregate interface:

Access: Admin/Network Admin

-
- Step 1** Select **Devices > Device Management**.
The Device Management page appears.
- Step 2** Next to the device where you want to delete the aggregate interface, click the edit icon (✎).
The Interfaces tab for that device appears.
- Step 3** Next to the aggregate interface you want to delete, click the delete icon (🗑️).
The aggregate interface can be identified by the **lag n** label, where **n** can be an integer from 0 to 13.
- Step 4** When prompted, confirm that you want to delete the aggregate interface.
The interface is deleted. Note that your changes do not take effect until you apply the device configuration; see [Applying Changes to Devices, page 4-25](#).
-

