



7000 and 8000 Series Device High Availability

The following topics describe how to configure high availability for Firepower 7000 Series and 8000 Series devices in the Firepower System:

- [About 7000 and 8000 Series Device High Availability, on page 1](#)
- [Establishing Firepower 7000/8000 Series High Availability, on page 5](#)
- [Editing Device High Availability, on page 6](#)
- [Configuring Individual Devices in a High-Availability Pair, on page 7](#)
- [Configuring Individual Device Stacks in a High-Availability Pair, on page 7](#)
- [Configuring Interfaces on a Device in a High-Availability Pair, on page 8](#)
- [Switching the Active Peer in a Device High-Availability Pair, on page 9](#)
- [Placing a High-Availability Peer into Maintenance Mode, on page 9](#)
- [Replacing a Device in a Stack in a High-Availability Pair, on page 10](#)
- [Device High Availability State Sharing, on page 10](#)
- [Device High Availability State Sharing Statistics for Troubleshooting, on page 13](#)
- [Separating Device High-Availability Pairs, on page 16](#)

About 7000 and 8000 Series Device High Availability

With 7000 and 8000 Series device high availability, you can establish redundancy of networking functionality and configuration data between two peer devices or two peer device stacks.

You achieve configuration redundancy by configuring two peer devices or two peer device stacks into a high-availability pair to act as a single logical system for policy deploys, system updates, and registration. The system automatically synchronizes other configuration data.



Note Static routes, non-SFRP IP addresses, and routing priorities are not synchronized between the peer devices or peer device stacks. Each peer device or peer device stack maintains its own routing intelligence.

Related Topics

[SFRP](#)

[Advanced Virtual Switch Settings](#)

Device High Availability Requirements

Before you can configure a 7000 and 8000 Series device high-availability pair, the following must be true:

- You can only pair single devices with single devices or device stacks with device stacks.
- Both devices or device stacks must have normal health status, be running the same software, and have the same licenses. See [Using the Health Monitor](#) for more information. In particular, the devices cannot have hardware failures that would cause them to enter maintenance mode and trigger a failover.



Note After you pair the devices, you cannot change the license options for individual paired devices, but you can change the license for the entire high-availability pair.

- Interfaces must be configured on each device or each primary device in a stack.
- Both devices or the primary members of the device stacks must be the same model and have identical copper or fiber interfaces.
- Device stacks must have identical hardware configurations, except for an installed malware storage pack. For example, you can pair a Firepower 8290 with another 8290. None, one, or all devices in either stack might have a malware storage pack.



Caution Do not attempt to install a hard drive that was not supplied by Cisco in your device. Installing an unsupported hard drive may damage the device. Malware storage pack kits are available for purchase **only** from Cisco, and are for use **only** with 8000 Series devices. Contact Support if you require assistance with the malware storage pack. See the *Firepower System Malware Storage Pack Guide* for more information.

- If the devices are targeted by NAT policies, both peers must have the same NAT policy.
- In a multidomain deployment, you can only establish 7000 or 8000 Series device high-availability or device stacks within a leaf domain.



Note After failover and recovery, SFRP preempts to the primary node.

Device High Availability Failover and Maintenance Mode

With a 7000 and 8000 Series device high availability, the system fails over either manually or automatically. You manually trigger failover by placing one of the paired devices or stacks in maintenance mode.

Automatic failover occurs after the health of the active device or stack becomes compromised, during a system update, or after a user with Administrator privileges shuts down the device. Automatic failover also occurs after an active device or device stack experiences NMSB failure, NFE failure, hardware failure, firmware failure, critical process failure, a disk full condition, or link failure between two stacked devices. If the health

of the standby device or stack becomes similarly compromised, the system does not fail over and enters a degraded state. The system also does not fail over when one of the devices or device stacks is in maintenance mode. Note that disconnecting the stacking cable from an active stack sends that stack into maintenance mode. Shutting down the secondary device in an active stack also sends that stack into maintenance mode.



Note If the active member of the high-availability pair goes into maintenance mode and the active role fails over to the other pair member, when the original active pair member is restored to normal operation it does not automatically reclaim the active role.

Configuration Deployment and Upgrade Behavior for High-Availability Pairs

This topic describes upgrade and deployment behavior for 7000 and 8000 Series devices (and stacks) in high availability pairs.

Behavior During Deploy

You deploy configuration changes to the members of a high availability pair at the same time. Deploy either succeeds or fails for both peers. The Firepower Management Center deploys to the active device; if that succeeds then changes are deployed to the standby.



Caution When you deploy, resource demands may result in a small number of packets dropping without inspection. Additionally, deploying some configurations restarts the Snort process, which interrupts traffic inspection. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See [Snort® Restart Traffic Behavior](#) and [Configurations that Restart the Snort Process When Deployed or Activated](#).

Behavior During Upgrade

You should not experience interruptions in traffic flow or inspection while upgrading devices (or device stacks) in high availability pairs. To ensure continuity of operations, they upgrade one at a time. Devices operate in maintenance mode while they upgrade.

Which peer upgrades first depends on your deployment:

- Routed or switched—Standby upgrades first. The devices switch roles, then the new standby upgrades. When the upgrade completes, the devices' roles remain switched. If you want to preserve the active/standby roles, manually switch the roles before you upgrade. That way, the upgrade process switches them back.
- Access control only—Active upgrades first. When the upgrade completes, the active and standby maintain their old roles.

Deployment Types and Device High Availability

You determine how to configure 7000 or 8000 Series device high availability depending on your Firepower System deployment: passive, inline, routed, or switched. You can also deploy your system in multiple roles at once. Of the four deployment types, only passive deployments require that you configure devices or stacks using high availability to provide redundancy. You can establish network redundancy for the other deployment

types with or without device high availability. For a brief overview on high availability in each deployment type, see the sections below.



Note You can achieve Layer 3 redundancy without using device high availability by using the Cisco Redundancy Protocol (SFRP). SFRP allows devices to act as redundant gateways for specified IP addresses. With network redundancy, you configure two devices or stacks to provide identical network connections, ensuring connectivity for other hosts on the network.

Passive Deployment Redundancy

Passive interfaces are generally connected to tap ports on central switches, which allows them to analyze all of the traffic flowing across the switch. If multiple devices are connected to the same tap feed, the system generates events from each of the devices. When configured in a high-availability pair, devices act as either active or standby, which allows the system to analyze traffic even in the event of a system failure while also preventing duplicate events.

Inline Deployment Redundancy

Because an inline set has no control over the routing of the packets being passed through it, it must always be active in a deployment. Therefore, redundancy relies on external systems to route traffic correctly. You can configure redundant inline sets with or without 7000 or 8000 Series device high availability.

To deploy redundant inline sets, you configure the network topology so that it allows traffic to pass through only one of the inline sets while preventing circular routing. If one of the inline sets fails, the surrounding network infrastructure detects the loss of connectivity to the gateway address and adjusts the routes to send traffic through the redundant set.

Routed Deployment Redundancy

Hosts in an IP network must use a well-known gateway address to send traffic to different networks. Establishing redundancy in a routed deployment requires that routed interfaces share the gateway addresses so that only one interface handles traffic for that address at any given time. To accomplish this, you must maintain an equal number of IP addresses on a virtual router. One interface advertises the address. If that interface goes down, the standby interface begins advertising the address.

In devices that are not members of a high-availability pair, you use SFRP to establish redundancy by configuring gateway IP addresses shared between multiple routed interfaces. You can configure SFRP with or without 7000 or 8000 Series device high availability. You can also establish redundancy using dynamic routing such as OSPF or RIP.

Switched Deployment Redundancy

You establish redundancy in a switched deployment using the Spanning Tree Protocol (STP), one of the advanced virtual switch settings. STP is a protocol that manages the topology of bridged networks. It is specifically designed to allow redundant links to provide automatic standby for switched interfaces without configuring standby links. Devices in a switched deployment rely on STP to manage traffic between redundant interfaces. Two devices connected to the same broadcast network receive traffic based on the topology calculated by STP.



Note Cisco strongly recommends that you enable STP when configuring a virtual switch that you plan to deploy in a 7000 or 8000 Series device high-availability pair.

7000/8000 Series High Availability Configuration

When establishing 7000 or 8000 Series device high availability, you designate one of the devices or stacks as active and the other as standby. The system applies a merged configuration to the paired devices. If there is a conflict, the system applies the configuration from the device or stack you designated as active.

After you pair the devices, you cannot change the license options for individual paired devices, but you can change the license for the entire high availability pair. If there are interface attributes that need to be set on switched interfaces or routed interfaces, the system establishes the high availability pair, but sets it to a pending status. After you configure the necessary attributes, the system completes the high availability pair and sets it to a normal status.

After you establish a high availability pair, the system treats the peer devices or stacks as a single device on the Device Management page. Device high availability pairs display the High Availability icon () in the appliance list. Any configuration changes you make are synchronized between the paired devices. The Device Management page displays which device or stack in the high availability pair is active, which changes after manual or automatic failover.

Removing registration of a device high availability pair from a Firepower Management Center removes registration from both devices or stacks. You remove a device high availability pair from the FMC as you would an individual managed device.

You can then register the high availability pair on another FMC. To register single devices from a high availability pair, you add remote management to the active device in the pair and then add that device to the FMC, which adds the whole pair. To register stacked devices in a high availability pair, you add remote management to the primary device of the either stack and then add that device to the FMC, which adds the whole pair.

After you establish a device high availability pair, you should configure a high-availability link interface.



Note If you plan to set up dynamic NAT, HA state sharing, or VPN using the devices in the high availability pair, you must configure a high-availability link interface. For more information, see [Configuring HA Link Interfaces](#).

Establishing Firepower 7000/8000 Series High Availability

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin



Note This procedure describes establishing a 7000 & 8000 Series device high-availability pair. For information on establishing Firepower Threat Defense high availability, see [Add a Firepower Threat Defense High Availability Pair](#).

When establishing a 7000 & 8000 Series device high-availability pair, you designate one of the devices or stacks as active and the other as standby. The system applies a merged configuration to the paired devices. If there is a conflict, the system applies the configuration from the device or stack you designated as active.

In a multidomain deployment, devices in a high-availability pair must belong to the same domain.

Before you begin

Confirm that all requirements are met; see [Device High Availability Requirements, on page 2](#).

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** From the **Add** drop-down menu, choose **Add High Availability**.
- Step 3** Enter a **Name**.
- Step 4** Under **Device Type**, choose **Firepower**.
- Step 5** Assign roles for the devices or stacks:
- Choose the **Active Peer** device or stack for the high-availability pair.
 - Choose the **Standby Peer** device or stack for the high-availability pair.
- Step 6** Click **Add**. The process takes a few minutes as the system synchronizes data.
-

What to do next

Create an HA Link interface on each of the devices in the high-availability pair if you plan to set up HA state sharing, dynamic NAT, or VPN with the devices. For more information on HA link interfaces, see [Configuring HA Link Interfaces](#).

Editing Device High Availability

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, most changes you make to the device configuration also change the configuration of the whole high-availability pair.

You can view the status of the high-availability pair by hovering your pointer over the status icon in the General section. You can also view which device or stack is the active peer and standby peer in the pair.

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high availability pair where you want to edit the configuration, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 3** Use the sections on the High Availability page to make changes to the high-availability pair configuration as you would a single device configuration.
-

Configuring Individual Devices in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can still configure some attributes for each device within the pair. You can make changes to a paired device just as you would to a single device.

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high-availability pair where you want to edit the configuration, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 3** Click the **Devices** tab.
- Step 4** From the **Selected Device** drop-down list, choose the device you want to modify.
- Step 5** Use the sections on the Devices page to make changes to the individual paired device as you would a single device.
-

Configuring Individual Device Stacks in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	Firepower 8140, Firepower 8200 family, Firepower 8300 family	Leaf only	Admin/Network Admin

After you configure stacked 8000 Series devices into a high-availability pair, the system limits the stack attributes that you can edit. You can edit the name of a stack in a paired stack. In addition, you can edit the

network configuration of the stack, as described in [Configuring Interfaces on a Device in a High-Availability Pair](#), on page 8.

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high-availability pair where you want to edit the configuration, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 3** Click the **Stacks** tab.
- Step 4** From the **Selected Device** drop-down list, choose the stack you want to modify.
- Step 5** Next to the **General** section, click the edit icon (✎).
- Step 6** Enter a **Name**.
- Step 7** Click **Save**.
-

What to do next

- Deploy configuration changes; see [Deploy Configuration Changes](#).

Configuring Interfaces on a Device in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

You can configure interfaces on individual devices in a 7000 or 8000 Series device high-availability pair. However, you must also configure an equivalent interface on the peer device in the pair. For paired stacks, you configure identical interfaces on the primary devices of the stacks. When you configure virtual routers, you select the stack where you want to configure the routers.

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high-availability pair where you want to configure interfaces, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 3** Click the **Interfaces** tab.
- Step 4** From the **Selected Device** drop-down list, choose the device you want to modify.
- Step 5** Configure interfaces as you would on an individual device.
-

Related Topics[Virtual Router Configuration](#)

Switching the Active Peer in a Device High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can manually switch the active and standby peer devices or stacks.

Procedure

-
- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high-availability pair where you want to change the active peer, click the Switch Active Peer icon ()
- Step 3** You can:
- Click **Yes** to immediately make the standby peer the active peer in the high-availability pair.
 - Click **No** to cancel and return to the Device Management page.
-

Placing a High-Availability Peer into Maintenance Mode

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

After you establish a 7000 or 8000 Series device high-availability pair, you can manually trigger failover by placing one of the peers into maintenance mode to perform maintenance on the devices. In maintenance mode, the system administratively takes down all interfaces except for the management interface. After maintenance is completed, you can re-enable the peer to resume normal operation.



Note You should not place both peers in a high-availability pair into maintenance mode at the same time. Doing so will prevent that pair from inspecting traffic.

Procedure

-
- Step 1** Choose **Devices > Device Management**.

- Step 2** Next to the peer you want to place in maintenance mode, click the toggle maintenance mode icon (.
- Step 3** Click **Yes** to confirm maintenance mode.

What to do next

- When maintenance is complete, click the toggle maintenance mode icon () again to bring the peer out of maintenance mode.

Replacing a Device in a Stack in a High-Availability Pair

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	Firepower 8140, 8200 family, 8300 family	Any	Admin/Network Admin

After you place a stack that is a member of a high-availability pair into maintenance mode, you can replace a secondary device in the stack for another device. You can only select devices that are not currently stacked or paired. The new device must follow the same guidelines for establishing a device stack.

Procedure

- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the stack member you want to place into maintenance mode, click the toggle maintenance mode icon (.
- Step 3** Click **Yes** to confirm maintenance mode.
- Step 4** Click the replace device icon (.
- Step 5** Choose the **Replacement Device** from the drop-down list.
- Step 6** Click **Replace** to replace the device.
- Step 7** Click the toggle maintenance mode icon () again to bring the stack immediately out of maintenance mode.

Note You do not need to re-deploy the device configuration.

Device High Availability State Sharing

Device high availability state sharing allows devices or stacks in high-availability pairs to synchronize as much state as necessary, so that if either device or stack fails, the other peer can take over with no interruption to traffic flow. Without state sharing, the following features may not fail over properly:

- Strict TCP enforcement

- Unidirectional access control rules
- Blocking persistence

Note, however, that enabling state sharing slows system performance.

You must configure and enable HA link interfaces on both devices or the primary stacked devices in the high-availability pair before you can configure high availability state sharing. Firepower 82xx Family and 83xx Family devices require a 10G HA link, while other model devices require a 1G HA link.

You must disable state sharing before you can modify the HA link interfaces.



Note If paired devices fail over, the system terminates all existing SSL-encrypted sessions on the active device. Even if you establish high availability state sharing, these sessions must be renegotiated on the standby device. If the server establishing the SSL session supports session reuse and the standby device does not have the SSL session ID, it cannot renegotiate the session.

Strict TCP Enforcement

When you enable strict TCP enforcement for a domain, the system drops any packets that are out of order on TCP sessions. For example, the system drops non-SYN packets received on an unestablished connection. With state sharing, devices in the high-availability pair allow TCP sessions to continue after failover without having to reestablish the connection, even if strict TCP enforcement is enabled. You can enable strict TCP enforcement on inline sets, virtual routers, and virtual switches.

Unidirectional Access Control Rules

If you have configured unidirectional access control rules, network traffic may match a different access control rule than intended when the system reevaluates a connection midstream after failover. For example, consider if you have a policy containing the following two access control rules:

```
Rule 1: Allow from 192.168.1.0/24 to 192.168.2.0/24
Rule 2: Block all
```

Without state sharing, if an allowed connection from 192.168.1.1 to 192.168.2.1 is still active following a failover and the next packet is seen as a response packet, the system denies the connection. With state sharing, a midstream pickup would match the existing connection and continue to be allowed.

Blocking Persistence

While many connections are blocked on the first packet based on access control rules or other factors, there are cases where the system allows some number of packets through before determining that the connection should be blocked. With state sharing, the system immediately blocks the connection on the peer device or stack as well.

When establishing state sharing for a high-availability pair, you can configure the following options:

Enabled

Click the check box to enable state sharing. Clear the check box to disable state sharing.

Minimum Flow Lifetime

Specify the minimum time (in milliseconds) for a session before the system sends any synchronization messages for it. You can use any integer from 0 to 65535. The system does not synchronize any sessions that have not met the minimum flow lifetime, and the system synchronizes only when a packet is received for the connection.

Minimum Sync. Interval

Specify the minimum time (in milliseconds) between update messages for a session. You can use any integer from 0 to 65535. The minimum synchronization interval prevents synchronization messages for a given connection from being sent more frequently than the configured value after the connection reaches the minimum lifetime.

Maximum HTTP URL Length

Specify the maximum characters for the URL the system synchronizes between the paired devices. You may use any integer from 0 to 225.

Related Topics

[Configuring HA Link Interfaces](#)

Establishing Device High-Availability State Sharing

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

Device high-availability state sharing allows 7000 or 8000 Series devices or stacks in high-availability pairs to synchronize as much state as necessary, so that if either device or stack fails, the other peer can take over with no interruption to traffic flow.

**Caution**

Modifying a high-availability state sharing option on a 7000 or 8000 Series device restarts the Snort process on the primary and secondary devices, temporarily interrupting traffic inspection on both devices. Whether traffic drops during this interruption or passes without further inspection depends on how the target device handles traffic. See [Snort® Restart Traffic Behavior](#) for more information.

Procedure

- Step 1** Configure HA link interfaces for each device in the device high-availability pair; see [Configuring HA Link Interfaces](#).
- Step 2** Choose **Devices > Device Management**.
- Step 3** Next to the device high-availability pair you want to edit, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 4** In the **State Sharing** section, click the edit icon (✎).

- Step 5** Decrease the state sharing values to improve paired peer readiness, or increase the values to allow better performance.
We recommend you use the default values, unless your deployment presents a good reason to change them.
- Step 6** Click **OK**.
-

What to do next

Deploy configuration changes; see [Deploy Configuration Changes](#).

Related Topics

[Configuring HA Link Interfaces](#)

[Snort® Restart Scenarios](#)

Device High Availability State Sharing Statistics for Troubleshooting

The sections below describe the statistics you can view for each device and how you can use them to troubleshoot your state sharing configuration for 7000 and 8000 Series device high-availability pairs.

Messages Received (Unicast)

Messages received are the number of high availability synchronization messages received from the paired peer.

The value should be close to the number of messages sent by the peer. During active use, the values may not match, but should be close. If traffic stops, the values should become stable and the messages received will match the messages sent.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close. The sent value on each peer should be incrementing at approximately the same rate as the received value on the opposite peer.

Contact Support if the received messages stop incrementing or increment slower than the messages sent by the peer.

Packets Received

The system batches multiple messages into single packets in order to decrease overhead. The Packets Received counter displays the total number of these data packets, as well as other control packets that have been received by a device.

The value should be close to the number of packets sent by the peer device. During active use, the values may not match, but should be close. Because the number of messages received should be close and incrementing at the same rate as the number of messages sent by the peer, the number of packets received should have the same behavior.

For troubleshooting, you should view both the packets received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the paired peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received packets stop incrementing or increment slower than the messages sent by the peer.

Total Bytes Received

Total bytes received are the number of bytes that make up the packets received by the peer.

The value should be close to the number of bytes sent by the other peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the total bytes received and the messages sent, compare the rate of increase, and make sure the values are increasing at the same rate. If the sent value on the paired peer is incrementing, the received value on the device should also increase at the same rate.

Contact Support if the received bytes stop incrementing or increment slower than the messages sent by the peer.

Protocol Bytes Received

Protocol bytes received are the number of bytes of protocol overhead received, which includes everything but the payload of session state synchronization messages.

The value should be close to the number of bytes sent by the peer. During active use, the values may not match, but should be close.

For troubleshooting, you should view the total bytes received to discover how much actual state data is being shared in comparison to protocol data. If the protocol data is a large percentage of the data being sent, you can adjust the minimum sync interval.

Contact Support if the protocol bytes received increment at a similar rate to the total bytes received. Protocol bytes received should be minimal in relation to the total bytes received.

Messages Sent

Messages sent are the number of high availability synchronization messages sent to the paired peer.

This data is useful in comparison to the number of messages received. During active use, the values may not match, but should be close.

For troubleshooting, you should view both the messages received and the messages sent, compare the rate of increase, and make sure the values are close.

Contact Support if the messages sent increment at a similar rate to the total bytes received.

Bytes Sent

Bytes sent are the total number of bytes sent that make up the high availability synchronization messages sent to the peer.

This data are useful in comparison to the number of messages received. During active use, the values may not match, but should be close. The number of bytes received on the peer should be close to, but not more than this value.

Contact Support if the total bytes received is not incrementing at about the same rate as the bytes sent.

Tx Errors

Tx errors are the number of memory allocation failures the system encounters when trying to allocate space for messages to be sent to the paired peer.

This value should be zero at all times on both peers. Contact Support if this number is not zero or if the number steadily increases, which indicates the system has encountered an error where it cannot allocate memory.

Tx Overruns

Tx overruns are the number of times the system attempts and fails to place a message into the transit queue.

This value should be zero at all times on both peers. When the value is not zero or is steadily increasing, it indicates that the system is sharing too much data across the HA link that cannot be sent quickly enough.

You should increase the HA link MTU if it was previously set below the default value (9918 or 9922). You can change the minimum flow lifetime and minimum synchronization interval settings to reduce the amount of data shared across the HA link to prevent the number from incrementing.

Contact Support if this value persists or continues to increase.

Recent Logs

The system log displays the most recent high availability synchronization messages. The log should not display any ERROR or WARN messages. It should remain comparable between the peers, such as the same number of sockets being connected.

However, the data displayed may be opposite in some instances, for example, one peer reports that it received a connection from the other peer and references different IP addresses. The log provides a comprehensive view of the high availability state sharing connection, and any errors within the connection.

Contact Support if the log displays an ERROR or WARN message, or any message that does not appear to be purely informational.

Viewing Device High Availability State Sharing Statistics

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Leaf only	Admin/Network Admin

After you establish state sharing, you can view the following information about the configuration in the **State Sharing** section of the High Availability page:

- The HA link interface that is being used and its current link state
- Detailed synchronization statistics for troubleshooting issues

The state sharing statistics are primarily counters for different aspects of the high availability synchronization traffic sent and received, along with some other error counters. In addition, you can view the latest system logs for each device in the high-availability pair.

Procedure

- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the device high-availability pair you want to edit, click the edit icon (✎).
In a multidomain deployment, if you are not in a leaf domain, the system prompts you to switch.
- Step 3** In the **State Sharing** section, click the view statistics icon (📊).
- Step 4** Choose a **Device** to view if your high-availability pair is composed of device stacks.
- Step 5** You can:
- Click **Refresh** to update the statistics.
 - Click **View** to view the latest data log for each device in the high-availability pair.
-

Separating Device High-Availability Pairs

Smart License	Classic License	Supported Devices	Supported Domains	Access
N/A	Control	7000 & 8000 Series	Any	Admin/Network Admin

When you separate, or "break," a 7000 or 8000 Series device high-availability pair:

- The active peer (device or stack) retains full deployment functionality
- The standby peer (device or stack) loses its interface configurations and fails over to the active peer, unless you choose to leave the interface configurations active, in which case the standby peer resumes normal operation.
- The standby peer always loses the configuration of passive interfaces.
- Any peer in maintenance mode resumes normal operation.

Procedure

- Step 1** Choose **Devices > Device Management**.
- Step 2** Next to the high-availability pair you want to break, click the Break HA icon (🔪).
- Step 3** Optionally, check the check box to remove the interface configurations on the standby peer.
This step administratively takes down all interfaces except for the management interface.
- Step 4** Click **Yes**.
-