



# CHAPTER 10

## Configuring High Availability

---

**Date:** 4/24/09

High Availability (or fault tolerance) ensures that your network services and applications are always available. High availability (HA) provides seamless switchover of flows in case an ANM becomes unresponsive or a critical host or interface fails. High Availability uses two nodes, where one node is that *active* node and the other is the *standby* node.

The ANM high availability feature includes:

- Automatic determination of node status, whether *active* or *standby*, using heartbeat counts
- Designation of the virtual IP address (VIP), which is associated with the active node
- Near real-time replication of ANM configuration and events after a failover occurs
- Automatic inspection of certificate/key presence on HA peer upon SSL certificate or key import.

For information about ACE redundancy, see [Understanding ACE Redundancy, page 10-20](#).

### Related Topics

- [Understanding High Availability, page 10-1](#)
- [Understanding High Availability Processes, page 10-2](#)
- [Configuring High Availability Overview, page 10-4](#)

## Understanding High Availability

During normal operation, high availability performs the following actions:

- The two nodes constantly exchange heartbeat packets over both interfaces.
- Database operations that occur on the active node's database are replicated on the standby node's database.
- The monitor function ensures that the necessary processes are running on both the active and standby node. For example, not all processes necessarily run on the standby node, so after a node changes from active to standby, the ANM high availability function stops certain processes on the standby node.

When you log into the ANM, you log in via a virtual IP address (VIP) that associates with the active node. The VIP is the only IP address you need to remember. If the current active node fails, the standby node takes over as the active node and the VIP automatically associates with the node that has just become active. When a failover occurs and the standby node becomes the active node, all existing web sessions are lost. In addition, there is a slight delay while the standby node takes over as the active node. After the switchover is complete and the ANM fully initializes, you can log into the ANM using the same VIP. All ANM functions remain the same.

The ANM uses heartbeat counts to determine when a failover should occur. Because both nodes are constantly sending and receiving heartbeat packets, if heartbeat packets are no longer being received on a node, its peer node is determined to be dead. If this peer node was the active node, then the standby node takes over as the active node. The VIP automatically associates with the newly active node, and the monitoring process starts any necessary processes on the newly active node that were not already running.

Similarly, if you manually issue a failover to cause the active node to become the standby node, the heartbeat process disassociates the VIP from the node and tells the monitoring function to stop processes that are not normally run on the standby node.

#### Related Topics

- [Understanding High Availability Processes, page 10-2](#)
- [Configuring High Availability Overview, page 10-4](#)

## Understanding High Availability Processes

During normal high availability operation, the active node runs all ANM processes required for normal operation of the ANM. The standby node runs only a minimal set of processes. [Table 10-3](#) lists the processes, their descriptions, and on which node they run.



#### Note

If you are running standalone ANM, all processes show in [Table 10-3](#), with the exception of the heartbeat process, are constantly running.

**Table 10-1 ANM High Availability Processes**

Process	Description	Node on Which Process Runs
Monit	Starts, stops, restarts, and monitors local ANM processes	Active and standby
Heartbeat	Provides UDP-based heartbeat between nodes, helps determine active vs. standby states, and associates the VIP	Active and standby
Mysql	Provides persistent storage and implements database replication between active and standby nodes	Active and standby
DCM	Java process	Active node only
DAL	Java process	Active node only
Ip-disc	Java process	Active node only
Licman	Java process for license management	Active and standby

**Related Topics**

- [CLI Commands for High Availability Processes, page 10-3](#)
- [Understanding High Availability, page 10-1](#)
- [Configuring High Availability Overview, page 10-4](#)

## CLI Commands for High Availability Processes

You can use two commands to view ANM processes:

- You can use the `/opt/CSCOanm/bin/anm-tool` command to start and stop the ANM processes and to view the status of the ANM processes.
- You can use the `/opt/CSCOanm/bin/anm-ha` command to check high availability configuration or to force a node to become standby or active.

[Table 10-2](#) lists the sub-commands and their descriptions.

**Table 10-2** CLI Sub-commands for Processes

Command	Sub-command	Description
/opt/CSCOanm/bin/anm-tool	info-services	Indicates the state of all ANM processes. This command does not return process status if <i>monit</i> is not running.
	stop-services	Stops all ANM processes, including <i>monit</i> . <b>Note</b> <i>Monit</i> must be running in order for the info-services command to provide status information.
	start-services	Starts the relevant ANM processes.
	restart-services	Restarts the relevant ANM processes.
	info	Provides additional information (state, whether running or stopped, start time, and PID) regarding the Java processes. <i>Monit</i> need not be running for this command to return information.
/opt/CSCOanm/bin/anm-ha	check	Checks the local node's high availability configuration. If errors are returned, it's likely that HA will not function correctly until you fix the errors. <b>Note</b> You must run this command on both the active and standby node.  While errors might indicate a problem, they could also simply indicate a known condition. For example, you receive a warning if the ANM cannot ping the peer node via either of the specified IP addresses; however, if the peer is down, the warning can be ignored because this is a known issue. It is also possible that no error might be returned even though there is a configuration problem. For example, the configuration of the two nodes must match; however the check sub-command cannot validate that the configurations match.
	active	Forces the local node to become <i>active</i> and the peer node to become the <i>standby</i> node.
	standby	Forces the local node to become <i>standby</i> and the peer node to become the <i>active</i> node.

**Related Topics**

- [Understanding High Availability Processes, page 10-2](#)
- [Understanding High Availability, page 10-1](#)
- [Configuring High Availability Overview, page 10-4](#)

## Configuring High Availability Overview

The ANM high availability consists of two nodes, which both run the ANM software. Each node must have at least two network interfaces:

- A primary interface, normally used to access the node.
- A heartbeat interface, which is used to provide additional redundancy. The heartbeat interfaces of the two nodes must be connected via a crossover Ethernet connection.
- The two Ethernet interfaces used on one of the hosts should match the two interfaces used on the other host, with regard to the subnets they participate in. For example, if HA Node 1 uses eth0 for the primary interface and eth1 for the heartbeat interface, then HA Node 2 should also use eth0 for the primary interface and eth1 for the heartbeat interface.

**Note**

The ANM does not configure the primary and heartbeat IP addresses of the nodes' interfaces. You must manually configure the node's interfaces.

When you installed the ANM, you provided values for high availability parameters, determined the node IDs of the two nodes designated as *Node 1* and *Node 2*. For additional information about the installation parameters, see the *Installation Guide for the Cisco Application Networking Manager 2.1*.

**Related Topics**

- [Understanding High Availability, page 10-1](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Configuring ACE High Availability, page 10-4](#)

## Configuring ACE High Availability

The tasks involved with configuring high availability on ACE devices are described in [Table 10-3](#).

**Table 10-3** High Availability Task Overview

	Task	Reference
Step 1	Create a fault-tolerant VLAN and identify peer IP addresses and configure peer devices for heartbeat count and interval.	<a href="#">Configuring ACE High Availability Peers, page 10-5</a>

**Table 10-3 High Availability Task Overview (continued)**

	Task	Reference
Step 2	Reconcile SSL certificates and keys, create a fault-tolerant group, assign peer priorities, associate the group with a context, place the group in service, and enable automatic synchronization.	<a href="#">Configuring High Availability Groups, page 10-8</a>
Step 3	Configure tracking for switchover.	<a href="#">High Availability Tracking and Failure Detection Overview, page 10-13</a>

**Related Topics**

- [Understanding ACE Redundancy, page 10-20](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [High Availability Tracking and Failure Detection Overview, page 10-13](#)

## Configuring ACE High Availability Peers

**Note**

This functionality is available for only Admin contexts.

Fault-tolerant peers transmit and receive heartbeat packets and state and configuration replication packets. The standby member uses the heartbeat packet to monitor the health of the active member, while the active member uses the heartbeat packet to monitor the health of the standby member. When the heartbeat packets are not received from the active member when expected, switchover occurs and the standby member assumes all active communications previously on the active member.

Use this procedure to:

- Identify the two members of a high availability pair.
- Assign IP addresses to the peer ACEs.
- Assign a fault-tolerant VLAN to high availability peers and bind a physical gigabit Ethernet interface to the FT VLAN.
- Configure heartbeat frequency and count on the ACEs in a fault-tolerant VLAN.

**Note**

For ANM to properly manage high availability peers, ensure that the combination of FT interface VLAN along with IP and peer IP address always be unique across every pair of ACE devices in high availability when those devices are imported into ANM. For details, see [“ANM Requirements for ACE High Availability” section on page 2-8](#).

**Assumption**

- At least one fault-tolerant VLAN has been configured.

**Note**

A fault-tolerant VLAN cannot be used for other network traffic.

**Procedure**

- Step 1** Select **Config > Devices > admin\_context > High Availability (HA) > Setup**. The HA Management window appears with two columns: One for the selected ANM and one for a peer ANM.
- Step 2** Click **Edit**, then enter the information for the primary ACE and the peer ACE as described in [Table 10-4](#).

**Table 10-4 High Availability Management Configuration Attributes**

Field	This Device	Peer Device
Module	Name of the ACE	
VLAN	Specify a fault-tolerant VLAN to be used for this high availability pair. Valid entries are integers from 1 to 4094.  <b>Note</b> This VLAN cannot be used for other network traffic.	Not applicable.
IP Address	Enter an IP address for the fault-tolerant VLAN in dotted-decimal format, such as 192.168.11.2.	Enter the IP address of the peer interface in dotted-decimal format so that the peer ACE can communicate on the fault-tolerant VLAN.
Netmask	Select the subnet mask that is to be used for the fault-tolerant VLAN.	Not applicable.
Query VLAN	Select the VLAN that the standby ACE is to use to determine whether the active ACE is down or if there is a connectivity problem with the fault-tolerant VLAN.	Select the VLAN that the standby ACE is to use to determine whether the active ACE is down or if there is a connectivity problem with the fault-tolerant VLAN.
Heartbeat Count	Enter the number of heartbeat intervals that must occur with no heartbeat packet received by the standby ACE before the standby ACE determines that the active member is not available. Valid entries are integers from 10 to 50.	Not applicable.
Heartbeat Interval	Enter the number of milliseconds that the active ACE is to wait between each heartbeat it sends to the standby ACE. Valid entries are integers from 100 to 1000.	Not applicable.
Interface Enabled	Select the Interface Enabled check box to enable the high availability interface. Clear the check box to disable the high availability interface.	Not applicable.
Shared VLAN Host ID	Enter a specific bank of MAC addresses that the ACE uses. Enter a number from 1 to 16. Be sure to configure different bank numbers for multiple ACEs.	Not applicable.

**Table 10-4 High Availability Management Configuration Attributes (continued)**

Field	This Device	Peer Device
Peer Shared VLAN Host ID	Enter a specific bank of MAC addresses for the same ACE in a redundant configuration. Valid entries are integers from 1 to 16. Be sure to configure different bank numbers for multiple ACEs.	Not applicable.
HA State	This is a read-only field with the current state of high availability on the ACE.	Not applicable.

**Step 3** Click:

- **Deploy Now** to save your entries and to continue with configuring high availability groups. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom. See [Configuring High Availability Groups, page 10-8](#) to configure a high availability group.
- **Cancel** to exit this procedure without saving your entries and to view the HA Management screen.

**Related Topics**

- [Understanding High Availability, page 10-1](#)
- [Configuring High Availability Overview, page 10-4](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Synchronizing ACE High Availability Configurations, page 10-19](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

## Clearing High Availability Pairs

**Note**

This functionality is available for only Admin contexts.

Use this procedure to remove a high availability link between two ACEs.

**Procedure**

- 
- Step 1** Select **Config > Devices > *admin\_context* > High Availability (HA) > Setup**. The HA Management screen appears.
- Step 2** Select the ACE pair whose high availability configuration you want to remove, then click **Clear**. A message appears asking you to confirm the clearing of the high availability link.
- Step 3** Click:
- **OK** to confirm the removal of this high availability link and to return to the HA Management screen.
  - **Cancel** to exit this procedure without removing this high availability link and to return to the HA Management screen.
-

**Related Topics**

- [Understanding High Availability, page 10-1](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Editing High Availability Groups, page 10-10](#)
- [High Availability Tracking and Failure Detection Overview, page 10-13](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)
- [Tracking Hosts for High Availability, page 10-14](#)

## Configuring High Availability Groups

**Note**


---

This functionality is available for only Admin contexts.

---

A fault-tolerant group consists of a maximum of two contexts: One active context on one ACE and one standby context on the peer ACE. You can create multiple fault-tolerant groups on each ACE up to a maximum of:

- For the ACE module—251 groups (250 user contexts and 1 Admin context).
- For the ACE appliance—21 groups (20 user contexts and 1 Admin context).

Use this procedure to configure high availability groups.

**Note**


---

For the replication process to function properly and successfully replicate the configuration for a user context when switching from the active context to the standby context, ensure that each user context has been added to the FT group. All applicable user contexts must be part of an FT group for redundancy to function properly.

---

**Assumption**

At least one high availability pair has been configured. (See [Configuring ACE High Availability Peers, page 10-5](#).)

**Procedure**

- 
- Step 1** **Config > Devices > *admin\_context* > High Availability (HA) > Setup.** The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
  - Step 2** In the HA Groups table, click **Add** to add a new high availability group. The table refreshes with the configurable fields.
  - Step 3** Check the Enabled check box to enable the high availability group. Clear the Enabled check box to disable the high availability group.
  - Step 4** In the Context field, select the virtual context to associate with this high availability group.
  - Step 5** In the Priority (Actual) field, enter the priority you want to assign to the first device in the group. Valid entries are integers from 1 to 255.

A member of a fault-tolerant group becomes the active member through a process based on the priority assigned. In this process, the group member with the higher priority becomes the active member. When you set up a fault-tolerant pair, use a higher priority for the group where the active member initially resides.

**Step 6** Check the Preempt check box to indicate that the group member with the higher priority is to always assert itself and become the active member. Clear the Preempt check box to indicate that you do not want the group member with the higher priority to always become the active member.

**Step 7** In the Peer Priority (Actual) field, enter the priority you want to assign to the peer device in the group. Valid entries are integers from 1 to 255.

A member of a fault-tolerant group becomes the active member through a process based on the priority assigned. In this process, the group member with the higher priority becomes the active member. When you set up a fault-tolerant pair, use a higher priority for the group where the active member initially resides.

**Step 8** Leave the Autosync Run check box *unchecked* to enable automatic synchronization of the running configuration files. Clear the Autosync Run check box to disable automatic synchronization of the running configuration files. If you disable automatic synchronization, you need to update the configuration of the standby context manually. See [Synchronizing Virtual Context Configurations, page 3-67](#).



---

**Note** If you check **Autosync Run** for the HA group, you must manually sync the standby context in order for ANM to allow subsequent configuration changes. Until you have done this, the standby context will be marked out of sync. See [Synchronizing Virtual Context Configurations in High Availability Mode, page 10-20](#).

---

**Step 9** Check the Autosync Startup check box to enable automatic synchronization of the startup configuration files. Clear the Autosync Run check box to disable automatic synchronization of the startup configuration files. If you disable automatic synchronization, you need to update the configuration of the standby context manually. See [Synchronizing Virtual Context Configurations, page 3-67](#).

**Step 10** Click:

- **Deploy Now** to accept your entries. The HA Groups table refreshes with the new high availability group.
- **Cancel** to exit this procedure without saving your entries and to return to the HA Management screen and HA Groups table.

---

#### Related Topics

- [Configuring ACE High Availability Peers, page 10-5](#)
- [Editing High Availability Groups, page 10-10](#)
- [Synchronizing Virtual Context Configurations, page 3-67](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)
- [Tracking Hosts for High Availability, page 10-14](#)

## Editing High Availability Groups

Use this procedure to modify the attributes of a high availability group.

**Note**


---

This functionality is available for only Admin contexts.

---

**Note**


---

If you need to modify a fault-tolerant group, take the group out of service before making any other changes (see [Taking a High Availability Group Out of Service, page 10-10](#)). When you finish making all changes, place the group back into service (see [Enabling a High Availability Group, page 10-11](#)).

---

**Procedure**

- 
- Step 1** Select **Config > Devices > admin\_context > High Availability (HA) > Setup**. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, select the high availability group you want to modify, then click **Edit**. The table refreshes with configurable fields.
- Step 3** Modify the fields as desired. For information on these fields, see [Configuring High Availability Groups, page 10-8](#).

**Note**


---

If you leave unchecked **Autosync Run** for the HA group, you must manually sync the standby context in order for ANM to allow subsequent configuration changes. Until you have done this, the standby context will be marked out of sync. See [Synchronizing Virtual Context Configurations in High Availability Mode, page 10-20](#).

---

- Step 4** When you finish modifying this group, click:
- **Deploy Now** to accept your entries and to return to the HA Groups table.
  - **Cancel** to exit this procedure without saving your entries and to return to the HA Management screen.
- 

**Related Topics**

- [Configuring High Availability Groups, page 10-8](#)
- [Taking a High Availability Group Out of Service, page 10-10](#)
- [Enabling a High Availability Group, page 10-11](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [High Availability Tracking and Failure Detection Overview, page 10-13](#)

## Taking a High Availability Group Out of Service

**Note**


---

This functionality is available for only Admin contexts.

---

If you need to modify a fault-tolerant group, you must first take the group out of service before making any other changes. Use this procedure to take a high availability group out of service.

**Procedure**

- 
- Step 1** Select **Config > Devices > admin\_context > High Availability (HA) > Setup**. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, select the high availability group you want to take out of service, then click **Edit**. The table refreshes with configurable fields.
- Step 3** Clear the **Enabled** check box.
- Step 4** Click **Deploy Now** to take the high availability group out of service and to return to the HA Groups table. You can now make the necessary modifications to the high availability group. To put the high availability group back in service, see [Enabling a High Availability Group, page 10-11](#).
- 

**Related Topic**

- [Enabling a High Availability Group, page 10-11](#)

## Enabling a High Availability Group

**Note**

This functionality is available for only Admin contexts.

After you take a high availability group out of service to modify it, you need to reenable the group. Use the following procedure to put a high availability group back in service.

**Procedure**

- 
- Step 1** Select **Config > Devices > admin\_context > High Availability (HA) > Setup**. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, select the high availability group you want to take out of service, then click **Edit**. The table refreshes with configurable fields.
- Step 3** Select the **Enabled** check box.
- Step 4** Click **Deploy Now** to put the high availability group in service and to return to the HA Groups table.
- 

**Related Topic**

- [Taking a High Availability Group Out of Service, page 10-10](#)

## Switching Over a High Availability Group

**Note**

This functionality is available for only Admin contexts.

You may need to cause a switchover when you want to make a particular context the standby (for example, for maintenance or a software upgrade on the currently active context). If the standby group member can statefully become the active member of the high availability group, a switchover occurs.

Use this procedure to force the failover of a high availability group.

**Procedure**

- 
- Step 1** Select **Config > Devices > *admin\_context* > High Availability (HA) > Setup**. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, select the group you want to switch over, then click **Switchover**. The standby group member becomes active, while the previously active group member becomes the standby member.



**Note** You must manually sync the standby context in order for ANM to allow subsequent configuration changes. Until you have done this, the standby context will be marked out of sync. See [Synchronizing Virtual Context Configurations in High Availability Mode, page 10-20](#).

---

**Related Topics**

- [Understanding High Availability, page 10-1](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

## Deleting High Availability Groups



**Note** This functionality is available for only Admin contexts.

---

Use this procedure to remove a high availability group from ANM management.

**Procedure**

- 
- Step 1** Select **Config > Devices > *admin\_context* > High Availability (HA) > Setup**. The HA Management screen appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, select the high availability group that you want to remove, then click **Delete**. A message appears asking you to confirm the deletion.
- Step 3** Click:
- **Deploy Now** to delete the high availability group and to return to the HA Groups table. The selected group no longer appears.
  - **Cancel** to exit this procedure without deleting the high availability group and to return to the HA Groups table.
- 

**Related Topics**

- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

# High Availability Tracking and Failure Detection Overview

The ANM supports the tracking and detection of failures to ensure that switchover occurs as soon as the criteria are met (see [Configuring ACE High Availability Peers, page 10-5](#)). You can track and detect failures on:

- Hosts—See [Tracking Hosts for High Availability, page 10-14](#).
- Interfaces—See [Tracking VLAN Interfaces for High Availability, page 10-13](#).

When the active member of a fault-tolerant group becomes unresponsive, the following occurs:

1. The active member's priority is reduced by 10.
2. If the resulting priority value is less than that of the standby member, the active member switches over and the standby member becomes the new active member. All active flows continue uninterrupted.
3. When the failed member comes back up, its priority is incremented by 10.
4. If the resulting priority value is greater than that of the currently active member, a switchover occurs again, returning the flows to the originally active member.

**Note**

In a user context, the ACE allows a switchover only of the fault-tolerant groups belonging to that context. In an Admin context, the ACE allows a switchover of all fault-tolerant groups on all configured contexts on the ACE.

**Related Topics**

- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)
- [Tracking Hosts for High Availability, page 10-14](#)

## Tracking VLAN Interfaces for High Availability

Use this procedure to configure a tracking and failure detection process for a VLAN interface.

**Procedure**

- Step 1** Select **Config > Devices > *admin\_context* > HA Tracking And Failure Detection > Interfaces**. The Track Interface table appears.
- Step 2** Click **Add** to add a new tracking process to this table, or select an existing entry, then click **Edit** to modify it. The Track Interface configuration screen appears.
- Step 3** In the Track Object Name field, enter a unique identifier for the tracking process. Valid entries are unquoted text strings with no spaces.
- Step 4** In the Priority field, enter the priority for the interface on the active member. Valid entries are integers from 0 to 255 with higher values indicating higher priorities. The values that you enter here and in the Interface Peer Priority field (see [Step 6](#)) reflect the point at which you want switchover to occur. If the tracked interface goes down, the priority of that fault-tolerant group is decremented by the value entered in the Priority field. If the priority of the fault-tolerant group on the active member falls below that of the standby member, a switchover occurs.

- Step 5** In the VLAN Interface field, select the fault-tolerant VLAN that you want the active member to track.
- Step 6** In the Interface Peer Priority field, enter the priority for the interface on the standby member. Valid entries are integers from 0 to 255 with higher values indicating higher priorities. The values that you enter here and in the Priority field (See [Step 4](#)) reflect the point at which you want switchover to occur. If the tracked interface goes down, the priority of that fault-tolerant group is decremented by the value entered in the Interface Peer Priority field. If the priority of the fault-tolerant group on the active member falls below that of the standby member, a switchover occurs.
- Step 7** In the Peer VLAN Interface field, enter the identifier of an existing fault-tolerant VLAN that you want the standby member to track. Valid entries are integers from 1 to 4096.
- Step 8** Click:
- **Deploy Now** to save your entries and to return to the Track Interface table.
  - **Cancel** to exit this procedure without saving your entries and to return to the Track Interface table.
  - **Next** to deploy your entries and to configure the next entry in the Track Interface table.
- 

#### Related Topics

- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking Hosts for High Availability, page 10-14](#)

## Tracking Hosts for High Availability

Use this procedure to configure a tracking and failure detection process for a gateway or host.

#### Procedure

- Step 1** Select **Config > Devices > *admin\_context* > HA Tracking And Failure Detection > Hosts**. The Track Host table appears.
- Step 2** Click **Add** to add a new tracking process to the table, or select an existing entry, then click **Edit** to modify it. The Track Host configuration screen appears.
- Step 3** In the Track Object Name field, enter a unique identifier for the tracking process. Valid entries are unquoted text strings with no spaces.
- Step 4** In the Track Host/IP Address field, enter the IP address or hostname of the gateway or host that you want the active member of the high availability group to track. Enter the IP address in dotted-decimal format, such as 192.168.11.2.
- Step 5** In the Priority field, enter the priority of the probe sent by the active member. Valid entries are integers from 0 to 255. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the host that the probe is tracking. If the probe goes down, the ACE decrements the priority of the fault-tolerant group on the active member by the value in the Priority field.
- Step 6** In the Peer Host/IP Address field, enter the IP address or hostname of the host that you want the standby member to track. Enter the IP address using dotted-decimal notation, such as 192.168.11.2.

- Step 7** In the Peer Priority field, enter the priority of the probe sent by the standby member. Valid entries are integers from 0 to 255. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the host that the probe is tracking. If the probe goes down, the ACE decrements the priority of the fault-tolerant group on the standby member by the value in the Priority field.
- Step 8** Click:
- **Deploy Now** to save your entries and to continue with configuring track host probes. See [Configuring Host Tracking Probes, page 10-15](#).
  - **Cancel** to exit this procedure without saving your entries and to return to the Track Host table.
  - **Next** to deploy your entries and to configure another tracking process.
- 

#### Related Topics

- [Configuring Host Tracking Probes, page 10-15](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

## Configuring Host Tracking Probes

Use this procedure to configure probes on the active high availability group member to track the health of the gateway or host.

#### Assumptions

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability, page 10-14](#).)
- At least one health monitoring probe has been configured (see [Configuring Health Monitoring for Real Servers, page 5-28](#)).

#### Procedure

- 
- Step 1** Select **Config > Devices > admin\_context > HA Tracking And Failure Detection > Hosts**. The Track Host table appears.

- Step 2** Select the tracking process you want to modify, then select the Peer Track Host Probe tab. The Peer Track Host Probes table appears.
- Step 3** In the Peer Track Host Probes table, click **Add** to add a peer host tracking probe, or select an existing peer host tracking probe, then click Edit to modify it. The Peer Track Host Probes configuration screen appears.
- Step 4** In the Probe Name field, select the name of the probe to be used for the peer host tracking process.
- Step 5** In the Priority field, enter a priority for the host you are tracking by the active member. Valid entries are integers from 1 to 255 with higher values indicating higher priorities. Assign a priority value based on the relative importance of the gateway or host that the probes are tracking. If the host goes down, the ACE decrements the priority of the high availability group on the standby member by the value in this Priority field.
- Step 6** Click:
- **Deploy Now** to save your entries and to return to the Track Host Probe table. The table includes the added probe.
  - **Cancel** to exit this procedure without saving your entries and to return to the Track Host Probe table.
  - **Next** to deploy your entries and to configure another track host probe.
- 

#### Related Topics

- [Configuring Peer Host Tracking Probes, page 10-17](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

## Deleting Host Tracking Probes

Use this procedure to remove a high availability host tracking probe.

#### Procedure

- Step 1** Select **Config > Devices > ACE *admin\_context* > HA Tracking And Failure Detection > Hosts**. The Track Host table appears.
- Step 2** Select the tracking process you want to modify, then select the Track Host Probe tab. The Track Host Probe table appears.
- Step 3** In the Track Host table, select the probe you want to remove, then click **Delete**. The probe is deleted and the Track Host Probe table refreshes without the deleted probe.
- 

#### Related Topics

- [Configuring Peer Host Tracking Probes, page 10-17](#)
- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

# Configuring Peer Host Tracking Probes

Use this procedure to configure probes on the standby member of a high availability group to track the health of the gateway or host.

## Assumptions

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability](#), page 10-14.)
- At least one health monitoring probe has been configured (see [Configuring Health Monitoring for Real Servers](#), page 5-28).

## Procedure

- 
- Step 1** Select **Config > Devices > ACE admin\_context > HA Tracking And Failure Detection > Hosts**. The Track Host table appears.
- Step 2** Select the tracking process you want to modify, then select the Peer Track Host Probe tab. The Peer Track Host Probes table appears.
- If the Track Host Probe and Peer Track Host Probes tabs do not appear below the Track Host table, click **Show Tabs** below the Track Host table name.
- Step 3** In the Peer Track Host Probes table, click **Add** to add a peer host tracking probe, or select an existing peer host tracking probe, then click **Edit** to modify it. The Peer Track Host Probes configuration screen appears.
- Step 4** In the Probe Name field, select the name of the probe to be used for the peer host tracking process.
- Step 5** In the Priority field, enter a priority for the host you are tracking by the standby member of the high availability group. Valid entries are integers from 0 to 255 with higher values indicating higher priorities. Assign a priority value based on the relative importance of the gateway or host that the probes are tracking. If the host goes down, the ACE decrements the priority of the high availability group on the standby member by the value in this Priority field.
- Step 6** Click:
- **Deploy Now** to save your entries and to return to the Peer Track Host Probes table. The table includes the added probe.
  - **Cancel** to exit this procedure without saving your entries and to return to the Peer Track Host Probes table.
  - **Next** to deploy your entries and to configure another peer track host probe.
- 

## Related Topics

- [Configuring Host Tracking Probes](#), page 10-15
- [Configuring ACE High Availability Peers](#), page 10-5
- [Configuring High Availability Groups](#), page 10-8
- [Tracking VLAN Interfaces for High Availability](#), page 10-13

## Deleting Peer Host Tracking Probes

Use this procedure to remove a high availability peer host tracking probe.

### Procedure

- 
- Step 1** Select **Config > Devices > ACE admin\_context > HA Tracking And Failure Detection > Hosts**. The Track Host table appears.
- Step 2** Select the tracking process you want to modify then, select the Peer Track Host Probe tab. The Peer Track Host Probes table appears.
- If the Track Host Probe and Peer Track Host Probes tabs do not appear below the Track Host table, click **Show Tabs** below the Track Host table name.
- Step 3** In the Peer Track Host Probes table, select the probe you want to remove, then click **Delete**. The probe is deleted and the Peer Track Host Probes table refreshes without the deleted probe.
- 

### Related Topics

- [Configuring Peer Host Tracking Probes, page 10-17](#)
- [Configuring Host Tracking Probes, page 10-15](#)
- [Tracking VLAN Interfaces for High Availability, page 10-13](#)

## Configuring ACE HSRP Groups

This section describes how to add or edit a Hot Standby Router Protocol (HSRP) group.

### Assumptions

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability, page 10-14](#).)
- Before you configure an HSRP tracking and failure detection process on the ACE, you must configure the HSRP group on the Catalyst 6500 Supervisor.

### Procedure

- 
- Step 1** Select **Config > Devices > ACE admin\_context > HA Tracking And Failure Detection > HSRP Groups**. The HSRP Groups table appears.
- Step 2** Click **Add** to add a new HSRP group or select an existing entry, then click **Edit** to modify it. The HSRP Group configuration screen appears.
- Step 3** In the Track Object Name field, enter a unique identifier for the tracking process. Valid entries are unquoted text strings with no spaces.
- Step 4** In the Priority field, enter the priority of the HSRP group as an integer from 0 to 255. The default is 0. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the HSRP group that you are tracking. If the HSRP group goes down, the ACE decrements the priority of the FT group on the active member. If the priority of the FT group on the active member falls below the priority of the FT group on the standby member, a switchover occurs.

- Step 5** In the HSRP Group Name, enter a name for the HSRP group.
- Step 6** In the HSRP Peer Priority field, enter the priority of the HSRP group as an integer from 0 to 255. The default is 0. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the HSRP group you are tracking. If the HSRP group goes down, the ACE decrements the priority of the FT group on the standby member.
- Step 7** In the HSRP Group Name of Peer field, enter a name for the HSRP group on the peer ACE.
- Step 8** Click:
- **Deploy Now** to save your entries and to return to the HSRP Groups table. The table includes the added HSRP group.
  - **Cancel** to exit this procedure without saving your entries and to return to the HSRP Groups table.
- 

## Synchronizing ACE High Availability Configurations

When two ANM devices are configured as high availability peers, their configurations must be synchronized at all times so that the standby member can take over for the active member seamlessly. As they synchronize, however, the configuration on the hot standby ACE can become out of sync with the ANM-maintained configuration data for that ACE.

**Note**

The Application Networking Manager manages local configurations only.

---

**Note**

Although a context might have been configured for syslog notification, changes applied to the standby ACE configuration can change syslog notification configuration so that you are not notified of the out-of-sync configurations. **As a result, it is important for you to manually synchronize the ANM with the standby ACE.**

---

Synchronizing configuration files for the standby ACE requires:

1. Auditing the standby ACE to confirm that its configuration does not agree with the ANM-maintained configuration data for the ACE. See [Synchronizing Virtual Context Configurations, page 3-67](#).
2. Uploading the configuration from the standby ACE to the ANM server. See [Synchronizing Virtual Context Configurations, page 3-67](#).
3. Ensuring that the SSL certificate/keys are imported and identical for the pair. See [Synchronizing SSL Certificate and Key Pairs on Both Peers, page 10-20](#).
4. For an Admin context, uploading configurations on any newly imported user contexts. If new user contexts are not updated, they cannot be managed using ANM.

## Synchronizing Virtual Context Configurations in High Availability Mode

In a high availability pair, the two configured virtual contexts synchronize with each other as part of their ongoing communications. However, their copies do not synchronize in ANM and the configuration on the standby member can become out of sync with the configuration on the ACE appliance.

After the active member of a high availability pair fails and the standby member becomes active, the newly active member detects any out-of-sync virtual context configurations and reports that status in the Virtual Contexts table so that you can synchronize the virtual context configurations.

For information on synchronizing virtual context configurations, see [Synchronizing Virtual Context Configurations, page 3-67](#).

### Related Topics

- [Configuring ACE High Availability Peers, page 10-5](#)
- [Configuring High Availability Groups, page 10-8](#)
- [Synchronizing Virtual Context Configurations, page 3-67](#)

## Synchronizing SSL Certificate and Key Pairs on Both Peers

When SSL certificate/key import is attempted on a peer that is configured in HA, ANM automatically detects the HA state and also imports the same cert/key into the other HA peer. In addition, when you are configuring two peers in HA from ANM, a warning message appears asking you to perform certificate/key reconciliation and offers the appropriate screen enabling you to do this.

## Understanding ACE Redundancy

ACE module redundancy (or fault tolerance) uses a maximum of two ACEs in the same Catalyst 6500 switch or in separate switches to ensure that your network remains operational even if one of the modules becomes unresponsive.

ACE appliance redundancy uses a maximum of two ACEs to ensure that your network remains operational even if one of the ACE appliances becomes unresponsive.



### Note

Redundancy is supported between ACEs of the same type only. Redundancy is not supported between an ACE appliance and an ACE module operating as peers. Redundancy must be of the same ACE device type and software release.

For additional information about ACE redundancy, see either *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

### Related Topics

- [ACE High Availability Polling, page 10-21](#)
- [ACE Redundancy Protocol, page 10-21](#)

## ACE High Availability Polling

Approximately every two minutes, the ANM issues the **show ft group** command to the ACE to gather the redundancy statistics of each virtual context. The state information is displayed in the HA State and HA Autosync fields when you click **Config > Devices > virtual context**.

The possible HA states are:

- Active—Local member of the FT group is active and processing flows.
- Standby Cold—Indicates if the FT VLAN is down but the peer ACE is still alive, or the configuration or application state synchronization failed. When a context is in this state and a switchover occurs, the transition to the ACTIVE state is stateless.
- Standby Bulk—Local standby context is waiting to receive state information from its active peer context. The active peer context receives a notification to send a snapshot of the current state information for all applications to the standby context.
- Standby Hot—Local standby context has all the state information it needs to statefully assume the active state if a switchover occurs.
- Standby Warm—Allows the configuration and state synchronization process to continue on a best-effort basis when you upgrade or downgrade the ACE software.

**Note**

When you upgrade or downgrade the ACE from one software version to another, there is a point in the process when the two ACEs have different software versions and, therefore, a software incompatibility. When the Standby Warm state appears, this means that the active ACE will continue to synchronize configuration and state information to the standby even though the standby may not recognize or understand the software commands or state information. This standby state allows the standby ACE to come up with best-effort support.

**Related Topics**

- [ACE High Availability Polling, page 10-21](#)
- [ACE Redundancy Protocol, page 10-21](#)

## ACE Redundancy Protocol

You can configure a maximum of two ACEs of the same type (peers) for redundancy in the same Catalyst 6500 switch or in different chassis for redundancy. Each peer ACE can contain one or more fault-tolerant (FT) groups. Each FT group consists of two members: one active context and one standby context. An FT group has a unique group ID that you assign.

**Note**

For the replication process to function properly and successfully replicate the configuration for a user context when switching from the active context to the standby context, ensure that each user context has been added to the FT group. All applicable user contexts must be part of an FT group for redundancy to function properly.

One virtual MAC address (VMAC) is associated with each FT group. The format of the VMAC is: 00-0b-fc-fe-1b-*groupID*. Because a VMAC does not change upon switchover, the client and server ARP tables does not require updating. The ACE selects a VMAC from a pool of virtual MACs available to it. For more information, see [Configuring Virtual Contexts, page 3-5](#).

Each FT group acts as an independent redundancy instance. When a switchover occurs, the active member in the FT group becomes the standby member and the original standby member becomes the active member. A switchover can occur for the following reasons:

- The active member becomes unresponsive.
- A tracked host or interface fails.
- You force a switchover for a high availability group by clicking **Switchover** in the HA Groups table (see [Switching Over a High Availability Group, page 10-11](#)).

To outside nodes (clients and servers), the active and standby FT group members appear as one node with respect to their IP addresses and associated VMAC. ACE provides active-active redundancy with multiple contexts only when there are multiple FT groups configured on each ACE and both devices contain at least one active group member (context). With a single context, the ACE supports active-backup redundancy and each group member is an Admin context.

The ACE sends and receives all redundancy-related traffic (protocol packets, configuration data, heartbeats, and state replication packets) on a dedicated FT VLAN. You cannot use this dedicated VLAN for normal traffic.

To optimize the transmission of heartbeat packets for multiple FT groups and to minimize network traffic, the ACE sends and receives heartbeat messages using a separate process. The ACE uses the heartbeat to probe the peer ACE, rather than probe each context. When an ACE does not receive a heartbeat from the peer ACE, all the contexts in the standby state become active. The ACE sends heartbeat packets over UDP. You can set the frequency with which the ACE sends heartbeat packets as part of the FT peer configuration. For details about configuring the heartbeat, see [Configuring ACE High Availability Peers, page 10-5](#).

The election of the active member within each FT group is based on a priority scheme. The member configured with the higher priority is elected as the active member. If a member with a higher priority is found after the other member becomes active, the new member becomes active because it has a higher priority. This behavior is known as preemption and is enabled by default. You can override this default behavior by disabling preemption. To disable preemption, use the `Preempt` parameter. Enabling `Preempt` causes the member with the higher priority to assert itself and become active. For details about configuring preemption, see [Configuring High Availability Groups, page 10-8](#).

For additional information about ACE redundancy, see either *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

#### Related Topics

- [Understanding ACE Redundancy, page 10-20](#)
- [ACE High Availability Polling, page 10-21](#)

## ACE Stateful Failover

The ACE replicates flows on the active FT group member to the standby group member per connection for each context. The replicated flows contain all the flow-state information necessary for the standby member to take over the flow if the active member becomes unresponsive. If the active member becomes unresponsive, the replicated flows on the standby member become active when the standby member assumes mastership of the context. The active flows on the former active member transition to a standby state to fully back up the active flows on the new active member.

**Note**

For the replication process to function properly and successfully replicate the configuration for a user context when switching from the active context to the standby context, ensure that the user context has been added to the FT group. All applicable user contexts must be part of an FT group for redundancy to function properly.

**Note**

By default, connection replication is enabled in the ACE.

After a switchover occurs, the same connection information is available on the new active member. Supported end-user applications do not need to reconnect to maintain the same network session.

The state information passed to the standby ANM includes the following data:

- Network Address Translation (NAT) table based on information synchronized with the connection record
- All Transmission Control Protocol (TCP) and User Datagram Protocol (UDP) connections not terminated by the ACE
- HTTP connection states (Optional)
- Sticky table

**Note**

In a user context, the ACE allows a switchover only of the FT group that belongs to that context. In the Admin context, the ACE allows a switchover of all FT groups in all configured contexts in the ACE.

To ensure that bridge learning occurs quickly upon a switchover in a Layer 2 configuration in the case where a VMAC moves to a new location, the new active member sends a gratuitous ARP on every interface associated with the active context. Also, when there are two VLANs on the same subnet and servers need to send packets to clients directly, the servers must know the location of the gateway on the client-side VLAN. The active member acts as the bridge for the two VLANs. In order to initiate learning of the new location of the gateway, the new active member sends an ARP request to the gateway on the client VLAN and bridges the ARP response onto the server VLAN.

For additional information about ACE redundancy, see either *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

**Related Topic**

- [Understanding ACE Redundancy, page 10-20](#)

## ACE Fault-Tolerant VLAN

ACE redundancy uses a dedicated fault-tolerant VLAN between redundant ANMs of the same type to transmit flow-state information and the redundancy heartbeat. Do not use this dedicated VLAN for normal network traffic. You must configure this same VLAN on both peers. You also must configure a different IP address within the same subnet on each ACE for the fault-tolerant VLAN.

The two redundant ANMs constantly communicate over the fault-tolerant VLAN to determine the operating status of each ACE. The standby member uses the heartbeat packet to monitor the health of the active member. The active member uses the heartbeat packet to monitor the health of the standby member. Communications over the switchover link include the following data:

- Redundancy protocol packets
- State information replication data
- Configuration synchronization information
- Heartbeat packets

For multiple contexts, the fault-tolerant VLAN resides in the system configuration data. Each fault-tolerant VLAN on the ACE has one unique MAC address associated with it. The ACE uses these ACE MAC addresses as the source or destination MACs for sending or receiving redundancy protocol state and configuration replication packets.

**Note**


---

The IP address and the MAC address of the fault-tolerant VLAN do not change at switchover.

---

For additional information about ACE redundancy, see the *Cisco Application Control Engine Module Administration Guide*.

**Related Topic**

[Understanding ACE Redundancy, page 10-20](#)

## ACE Configuration Synchronization

For redundancy to function properly, both members of an fault-tolerant group must have identical configurations. The ACE automatically replicates the active configuration on the standby member using a process called *configuration synchronization* (config sync). Config sync automatically replicates any changes made to the configuration of the active member to the standby member. After the ACE synchronizes the redundancy configuration from the active member to the standby peer, it disables configuration mode on the standby. See [Configuring ACE High Availability Peers, page 10-5](#).

**Note**


---

The Application Networking Manager manages local configurations only.

---

For additional information about ACE redundancy, see either the *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

**Related Topic**

[Understanding ACE Redundancy, page 10-20](#)

## ACE Redundancy Configuration Requirements and Restrictions

Follow these requirements and restrictions when configuring the ACE redundancy feature.

- In bridged mode (Layer 2), two contexts cannot share the same VLAN.
- To achieve active-active redundancy, a minimum of two contexts and two fault-tolerant groups are required on each ACE.
- When you configure redundancy, the ACE keeps all interfaces that do not have an IP address in the Down state. The IP address and the peer IP address that you assign to a VLAN interface should be in the same subnet, but different IP addresses. For more information about configuring VLAN interfaces, see [Configuring VLAN Interfaces, page 9-2](#).

- In a high availability pair, the two configured virtual contexts synchronize with each other as part of their ongoing communications. However, their copies do not synchronize in ANM and the configuration on the standby member can become out of sync with the configuration on the ACE. After the active member of a high availability pair fails and the standby member becomes active, ANM on the newly active member detects any out-of-sync virtual context configurations and reports that status in the All Virtual Contexts table so that you can synchronize the virtual context configurations.
- When a virtual context is in either the Standby Hot or Standby Warm state (see [ACE High Availability Polling, page 10-21](#)), the virtual context may receive configuration changes from its ACE peer without updating the ANM GUI. As a result, the ANM GUI will be out of synchronization with the CLI configuration. If you need to check configuration on a standby virtual context using the tracking and failure detection process (see [Tracking VLAN Interfaces for High Availability, page 10-13](#)), we recommend that you first perform a manual synchronization using either the CLI Sync or CLI Sync All buttons before checking the configuration values.

For additional information about ACE redundancy, see either *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

**Related Topic**

[Understanding ACE Redundancy, page 10-20](#)

