



CHAPTER 13

Configuring High Availability

Date: 3/22/13

This chapter describes how to configure high availability for ANM servers and ACE devices.



Note

When naming ACE objects (such as a real server, virtual server, parameter map, class map, health probe, and so on), enter an alphanumeric string of 1 to 64 characters, which can include the following special characters: underscore (_), hyphen (-), dot (.), and asterisk (*). Spaces are not allowed.

If you are using ANM with an ACE module or ACE appliance and you configure a named object at the ACE CLI, keep in mind that ANM does not support all of the special characters that the ACE CLI allows you to use when configuring a named object. If you use special characters that ANM does not support, you may not be able to import or manage the ACE using ANM.

This chapter includes the following sections:

- [Understanding ANM High Availability, page 13-2](#)
- [Understanding ACE Redundancy, page 13-6](#)
- [Configuring ACE High Availability, page 13-14](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Clearing ACE High Availability Pairs, page 13-17](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Displaying High Availability Group Statistics and Status, page 13-21](#)
- [Switching Over an ACE High Availability Group, page 13-22](#)
- [Deleting ACE High Availability Groups, page 13-23](#)
- [ACE High Availability Tracking and Failure Detection Overview, page 13-23](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)
- [Tracking Hosts for High Availability, page 13-25](#)
- [Configuring Host Tracking Probes, page 13-26](#)
- [Configuring ACE Peer Host Tracking Probes, page 13-28](#)
- [Configuring ACE HSRP Groups, page 13-29](#)
- [Synchronizing ACE High Availability Configurations, page 13-30](#)
- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#)

Understanding ANM High Availability

ANM 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 server becomes unresponsive or a critical host or interface fails. High availability uses two ANM nodes, where one node is the *active* node and the other is the *standby* node.

The ANM high availability features are as follows:

- 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.

During normal operation, ANM 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, ANM high availability function stops certain processes on the standby node.

When you log into ANM, you log in using 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 ANM using the same VIP. All ANM functions remain the same.

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 ANM High Availability Processes, page 13-3](#)
- [Configuring ANM High Availability Overview, page 13-3](#)
- [CLI Commands for ANM High Availability Processes, page 13-4](#)
- [Recovering From an HA Database Replication Failure, page 13-6](#)

Understanding ANM High Availability Processes

During normal high availability operation, the active node runs all ANM processes required for normal operation of ANM. The standby node runs only a minimal set of processes. [Table 13-1](#) lists the processes, their descriptions, and on which node they run.



Note

If you are running standalone ANM, all processes show in [Table 13-1](#), with the exception of the heartbeat process, are constantly running.

Table 13-1 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
ANM	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 ANM High Availability Processes, page 13-4](#)
- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability, page 13-14](#)
- [Understanding ACE Redundancy, page 13-6](#)

Configuring ANM High Availability Overview

Configuring ANM high availability depends on whether you are using ANM Virtual Appliance or ANM server.

ANM Virtual Appliance

You can implement redundancy for ANM Virtual Appliance using the high availability feature of the underlying VMware vSphere platform. VMware HA (High Availability) detects faults in the operation of managed virtual machines and provides redundancy in case of a failure.

You implement VMware HA for ANM Virtual Appliance in the same manner as for any VM-based application running on VMware infrastructure; that is, ANM Virtual Appliance does not impose any special requirements for implementing VMware HA.

For additional information about installing ANM Virtual Appliance, see the [Installation Guide for the Cisco Application Networking Manager 5.2 Virtual Appliance](#).

ANM Server

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

- 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**

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 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 Cisco Application Networking Manager 5.2](#).

Related Topics

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Configuring ACE High Availability, page 13-14](#)

CLI Commands for ANM High Availability Processes

You use two commands to view ANM processes:

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

[Table 13-2](#) lists the sub-commands and their descriptions.

Table 13-2 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> . Note <i>Monit</i> must be running in order for the info-services command to provide status information. Note When ANM is running in HA mode and the standby ANM is just starting up, the active ANM copies its entire database to the standby ANM. During the copy process, the active ANM cannot be stopped or restarted using the anm-tool command. Check the Admin > ANM Management page for the HA Replication Status and wait until the status is set to OK before attempting to stop ANM.
	start-services	Starts the relevant ANM processes.
	restart-services	Restarts the relevant ANM processes. Note When ANM is running in HA mode and the standby ANM is just starting up, the active ANM copies its entire database to the standby ANM. During the copy process, the active ANM cannot be stopped or restarted using the anm-tool command. Check the Admin > ANM Management page for the HA Replication Status and wait until the status is set to OK before attempting to restart ANM.
	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, HA might not function correctly until you fix the errors. Note 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 ANM High Availability Processes, page 13-3](#)

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability, page 13-14](#)
- [Understanding ACE Redundancy, page 13-6](#)


Recovering From an HA Database Replication Failure

This section provides an overview of the database replication process that occurs between ANM HA active and standby nodes and how to recover from a replication failure.

When the active ANM is running and the standby ANM is just starting up, the active ANM copies its entire database to the standby ANM. This process normally takes from a few seconds to a few minutes depending on the size of the configuration data and monitoring data. During the replication process, the active ANM database is locked and the active ANM cannot be stopped or restarted using the **anm-tool** command nor can it perform a failover.

It is possible for the database replication process to fail if the standby ANM is stopped or powered down, the connectivity is down, or the active ANM is powered down. The failure of the replication process does not affect the integrity of the active ANM database. The procedure in this section describes what to do if you encounter a replication failure.

Procedure

-
- Step 1** Check the standby ANM and make sure that it has stopped.
- If the standby ANM is still running, stop it because its database might be incomplete due to the replication failure.
- Step 2** Check the connectivity between the active ANM and standby ANM and make sure that both links are up and connected.
- Step 3** Do one of the following:
- If the active ANM is still running, login and check to see that its configuration is normal.
 - If the active ANM has stopped or powered down, restart it now.
- Step 4** After the active ANM is running normally, restart the standby ANM.
-
-  **Caution** Do not restart the standby ANM before the active ANM is running and operating normally.
-
- Step 5** From the standby ANM GUI, choose **Admin > ANM Management** to display the ANM Server window and make sure that the HA Replication Status is set to OK before performing any daily management tasks.
-

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 version.

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

This section includes the following topics:

- [ACE High Availability Polling, page 13-7](#)
- [ACE Redundancy Protocol, page 13-8](#)
- [ACE Stateful Failover, page 13-9](#)
- [ACE Fault-Tolerant VLAN, page 13-10](#)
- [ACE Configuration Synchronization, page 13-11](#)
- [ACE Redundancy Configuration Requirements and Restrictions, page 13-12](#)
- [ACE High Availability Troubleshooting Guidelines, page 13-12](#)

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**.

**Note**

To display statistics and status information for a particular high availability group displayed in the High Availability (HA) Setup window (Config > Devices > admin_context > High Availability (HA) > Setup), see the “[Displaying High Availability Group Statistics and Status](#)” section on page 13-21.

The possible HA states are as follows:

- **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.

- **Inconclusive**—Indicates that ANM was able to determine that the given ACE was configured in HA, however ANM was able to find more than one ACE module or ACE appliance that appeared to be a peer. In this case, ANM was unable to conclusively find a unique HA peer for the given ACE module or ACE appliance. For additional details on addressing this state, see the [“ANM Requirements for ACE High Availability” section on page 5-8](#) for details.

Inconclusive is not shown in the HA State field but is shown in the HA Peer field. It is possible that a context HA peer is inconclusive, but its HA State and HA Peer state are still shown normally because these states are from context polling from the ACE device.

**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 13-7](#)
- [ACE Redundancy Protocol, page 13-8](#)

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 the [“Configuring Virtual Contexts” section on page 6-8](#).

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 the [“Switching Over an ACE High Availability Group” section on page 13-22](#)).

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 the “[Configuring ACE High Availability Peers](#)” section on page 13-15.

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 the “[Configuring ACE High Availability Groups](#)” section on page 13-17.

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

Related Topics

- [Understanding ACE Redundancy, page 13-6](#)
- [ACE High Availability Polling, page 13-7](#)

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 ACE 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 the *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

Related Topics

- [Understanding ACE Redundancy, page 13-6](#)

ACE Fault-Tolerant VLAN

ACE redundancy uses a dedicated fault-tolerant VLAN between redundant ACEs 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 ACEs 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 either the *Cisco Application Control Engine Module Administration Guide* or the *Cisco 4700 Series Application Control Engine Appliance Administration Guide*.

Related Topics

[Understanding ACE Redundancy, page 13-6](#)

ACE Configuration Synchronization

For redundancy to function properly, both members of a 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 the “[Configuring ACE High Availability Peers](#)” section on [page 13-15](#).

**Note**

The Application Networking Manager manages local configurations only.

When ANM detects a pair of ACE peers operating in high availability (HA), ANM allows you to make configuration changes on either the active or standby ACE. ANM then automatically (and seamlessly) pushes the configuration to the active ACE and locally replicates the configuration on the standby imported into ANM. This action is similar to what is performed by the ACE to the peers.

**Note**

Keep in mind that the configuration pushed while the standby ACE has been selected does not mean that ANM pushed the configuration to the standby ACE. Typically, with auto-sync turned off, configuration changes are disabled on the standby ACE. In this case, ANM tries to push the configuration to the active ACE in the HA device pair.

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 Topics

- [Understanding ACE Redundancy, page 13-6](#)
- [Synchronizing ACE High Availability Configurations, page 13-30](#)
- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#)

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 the [“Configuring Virtual Context VLAN Interfaces”](#) section on page 12-6.
- When importing an ACE HA pair into ANM, follow one of the configuration requirements outlined below for ANM to uniquely identify the ACE HA pair:
 - Use a unique combination of FT interface VLAN and FT IP address/peer IP address for every ACE HA pair imported into ANM. For HA, it is critical that the combination of FT interface VLAN and IP address/peer IP address always be unique across every pair of ACE peer devices.
 - Define a peer IP address in the management interface, using the management IP address of the peer ACE (module or appliance). Note that the management IP address and management peer IP address used for this definition should be the management IP address used to import both ACE devices into ANM.

For more information about the use of multiple HA pairs imported into ANM, see the [“ANM Requirements for ACE High Availability”](#) section on page 5-8

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

Related Topics

- [Understanding ANM High Availability, page 13-2](#)

ACE High Availability Troubleshooting Guidelines

This section provides the following set of guidelines for troubleshooting an ACE high availability (or redundancy) configuration in ANM:

- If the high availability setup of two ACE devices is successful, the HA State field of the ACE HA Management table should indicate no errors. If the HA State field does not read Compatible, verify that both ACE devices are the same type of hardware. ACE modules cannot be synchronized with ACE appliances.
- If the high availability setup of two ACE devices is successful, the License Compatibility and SRG Compatibility fields of the **show ft peer** CLI command output on the ACE (module or appliance) should indicate no errors. See either the *Cisco Application Control Engine Module Administration Guide* or *Cisco 4700 Series Application Control Engine Appliance Administration Guide* for details on the **show ft peer** CLI command.
 - If the SRG Compatibility field indicates a problem, this means that the versions of the ACE software running on the devices are not compatible with each other. One or both of the devices will need to have an appropriate version of the ACE software installed before they can be synchronized.

- If the License Compatibility field indicates a licensing problem, go to the Licenses page of ACE Hardware Setup (see the [“Using ACE Hardware Setup” section on page 3-5](#)) and make sure each ACE device has a valid license installed. Licenses must be installed on each device separately because each license is only valid for one hardware device.

For proper HA functionality, the licenses on both ACEs in the pair must be also compatible with each other. This means both licenses must permit the same bandwidth and the same number of virtual contexts.

**Note**

If the licenses' bandwidth limits do not match, configuration synchronization may appear to work (although Admin context synchronization may actually not be functional), and the License Compatibility field may not show an error. However, failover from the higher bandwidth ACE to a lower bandwidth ACE could result in loss of traffic.

Configuring ACE High Availability

The tasks involved with configuring high availability on ACE devices are described in [Table 13-3](#).

Table 13-3 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.	Configuring ACE High Availability Peers, page 13-15
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.	Configuring ACE High Availability Groups, page 13-17
Step 3	Configure tracking for switchover.	ACE High Availability Tracking and Failure Detection Overview, page 13-23

Related Topics

- [Understanding ACE Redundancy, page 13-6](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [ACE High Availability Tracking and Failure Detection Overview, page 13-23](#)
- [Synchronizing ACE High Availability Configurations, page 13-30](#)
- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#)

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 do the following tasks:

- 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 the [“ANM Requirements for ACE High Availability” section on page 5-8](#).

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 Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.

The HA Management window appears with two columns; one for the selected ACE and one for a peer ACE.

Step 2 Click **Edit** and enter the information for the primary ACE and the peer ACE as described in [Table 13-4](#).

Table 13-4 High Availability Management Configuration Attributes

Field	This Device	Peer Device
Module	Name of the ACE	Not applicable.
VLAN	Fault-tolerant VLAN to be used for this high availability pair. Valid entries are from 1 to 4094. Note This VLAN cannot be used for other network traffic.	Not applicable.

Table 13-4 High Availability Management Configuration Attributes (continued)

Field	This Device	Peer Device
IP Address	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	Subnet mask that is to be used for the fault-tolerant VLAN.	Not applicable.
Query VLAN	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.	Choose 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	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 from 10 to 50.	Not applicable.
Heartbeat Interval	Number of milliseconds that the active ACE is to wait between each heartbeat it sends to the standby ACE. Valid entries are from 100 to 1000.	Not applicable.
Interface Enabled	Interface Enabled check box that enables the high availability interface. Uncheck the check box to disable the high availability interface.	Not applicable.
Shared VLAN Host ID	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.
Peer Shared VLAN Host ID	Specific bank of MAC addresses for the same ACE in a redundant configuration. Valid entries are from 1 to 16. Be sure to configure different bank numbers for multiple ACEs.	Not applicable.
HA State	Read-only field with the current state of high availability on the ACE.	Not applicable.

Step 3 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. Continue with configuring high availability groups. The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom. See the “[Configuring ACE High Availability Groups](#)” section on page 13-17 to configure a high availability group.
- Click **Cancel** to exit this procedure without saving your entries and to view the HA Management window.

Related Topics

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability, page 13-14](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Synchronizing ACE High Availability Configurations, page 13-30](#)

- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers](#), page 13-32
- [Tracking ACE VLAN Interfaces for High Availability](#), page 13-24

Clearing ACE High Availability Pairs

**Note**

This functionality is available for only Admin contexts.

You can remove a high availability link between two ACEs.

Procedure

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

Related Topics

- [Understanding ANM High Availability](#), page 13-2
- [Configuring ACE High Availability Peers](#), page 13-15
- [Editing High Availability Groups](#), page 13-19
- [ACE High Availability Tracking and Failure Detection Overview](#), page 13-23
- [Tracking ACE VLAN Interfaces for High Availability](#), page 13-24
- [Tracking Hosts for High Availability](#), page 13-25

Configuring ACE High Availability Groups

**Note**

This functionality is available for only Admin contexts.

You can configure a high availability group, or fault-tolerant group, which 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).

**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 the [“Configuring ACE High Availability Peers”](#) section on page 13-15).

Procedure

Step 1 Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.

The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

Step 2 In the HA Groups table of the HA Management window, 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.

Uncheck the Enabled check box to disable the high availability group.

Step 4 In the Context field, choose the virtual context to associate with this high availability group.

Step 5 In the Priority (Actual) field, enter the priority that you want to assign to the first device in the group.

Valid entries are 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 specify that the group member with the higher priority is to always assert itself and become the active member.

Uncheck the Preempt check box to specify 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 that you want to assign to the peer device in the group.

Valid entries are 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 Check the Autosync Run check box to enable automatic synchronization of the running configuration files.

Uncheck 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 the [“Synchronizing Virtual Context Configurations”](#) section on page 6-105.

**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 the “[Synchronizing Virtual Context Configurations in High Availability Mode](#)” section on page 13-31.

Step 9 Check the Autosync Startup check box to enable automatic synchronization of the startup configuration files.

Uncheck 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 the “[Synchronizing Virtual Context Configurations](#)” section on page 6-105.

Step 10 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. The HA Groups table refreshes with the new high availability group.
- Click **Cancel** to exit this procedure without saving your entries and to return to the HA Management window and HA Groups table.

Step 11 (Optional) To display statistics and status information for a particular high availability group, choose the group from the ACE HA Groups table, and click **Details**.

The **show ft group *group_id* detail** CLI command output appears. See the “[Displaying High Availability Group Statistics and Status](#)” section on page 13-21 for details.

Related Topics

- [Configuring ACE High Availability Peers, page 13-15](#)
- [Editing High Availability Groups, page 13-19](#)
- [Synchronizing Virtual Context Configurations, page 6-105](#)
- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)
- [Tracking Hosts for High Availability, page 13-25](#)

Editing High Availability Groups

**Note**

This functionality is available for only Admin contexts.

You can modify the attributes of a high availability group.

**Note**

If you need to modify a fault-tolerant group, take the group out of service before making any other changes (see the “[Taking a High Availability Group Out of Service](#)” section on page 13-20). When you finish making all changes, place the group back into service (see the “[Enabling a High Availability Group](#)” section on page 13-21).

Procedure

-
- Step 1** Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.
- The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, choose the high availability group that you want to modify, and click **Edit**.
- The table refreshes with configurable fields.
- Step 3** Modify the fields as desired. For information on these fields, see the [“Configuring ACE High Availability Groups”](#) section on page 13-17.



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 the [“Synchronizing Virtual Context Configurations in High Availability Mode”](#) section on page 13-31.

- Step 4** When you finish modifying this group, do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. You return to the HA Groups table.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the HA Management window.
-

Related Topics

- [Configuring ACE High Availability Groups, page 13-17](#)
- [Taking a High Availability Group Out of Service, page 13-20](#)
- [Enabling a High Availability Group, page 13-21](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [ACE High Availability Tracking and Failure Detection Overview, page 13-23](#)

Taking a High Availability Group Out of Service



Note This functionality is available for only Admin contexts.

You can take a high availability group out of service, which you must do before you can modify it.

Procedure

-
- Step 1** Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.
- The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

- Step 2** In the HA Groups table, choose the high availability group you want to take out of service, and click **Edit**.
- The table refreshes with configurable fields.
- Step 3** Uncheck 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 the “[Enabling a High Availability Group](#)” section on page 13-21.
-

Related Topics

[Enabling a High Availability Group, page 13-21](#)

Enabling a High Availability Group



Note

This functionality is available for only Admin contexts.

You can put a high availability group back into service after taking it out of service.

Procedure

- Step 1** Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.
- The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.
- Step 2** In the HA Groups table, choose the high availability group you want to take out of service, and click **Edit**.
- The table refreshes with configurable fields.
- Step 3** Check 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 Topics

[Taking a High Availability Group Out of Service, page 13-20](#)

Displaying High Availability Group Statistics and Status

You can display statistics and status information for a particular high availability group by using the **Details** button. ANM accesses the **show ft group group_id detail** CLI command to display detailed ACE HA group information.

Procedure

- Step 1** Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.

The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

Step 2 Choose an ACE HA group from the ACE HA Groups table and click **Details**.

The **show ft group group_id detail** CLI command output appears. For details on the displayed output fields, see either the *Cisco ACE Module Administration Guide* or the *Cisco ACE 4700 Series Appliance Administration Guide*.

Step 3 Click **Update Details** to refresh the output for the **show ft group group_id detail** CLI command.

Step 4 Click **Close** to return to the VLAN Interfaces table.

Switching Over an ACE High Availability Group



Note

This functionality is available for only Admin contexts.

You can force the failover of a high availability group. You may need to force 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.

Procedure

Step 1 Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.

The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

Step 2 In the HA Groups table, choose the group that you want to switch over, and 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 the [“Synchronizing Virtual Context Configurations in High Availability Mode”](#) section on page 13-31.

Related Topics

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Related Topics

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)

Deleting ACE High Availability Groups

**Note**

This functionality is available for only Admin contexts.

You can remove a high availability group from ANM management.

Procedure

Step 1 Choose **Config > Devices > *admin_context* > High Availability (HA) > Setup**.

The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

Step 2 In the HA Groups table, choose the high availability group that you want to remove, and click **Delete**.

A message appears asking you to confirm the deletion.

Step 3 Do one of the following:

- Click **Deploy Now** to delete the high availability group and to return to the HA Groups table. The selected group no longer appears.
- Click **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 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

ACE High Availability Tracking and Failure Detection Overview

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 13-15](#)). You can track and detect failures on the following:

- Hosts—See [Tracking Hosts for High Availability, page 13-25](#).
- Interfaces—See [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#).

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 ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)
- [Tracking Hosts for High Availability, page 13-25](#)

Tracking ACE VLAN Interfaces for High Availability

You can configure a tracking and failure detection process for a VLAN interface.

Procedure

-
- Step 1** Choose **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 choose an existing entry and click **Edit** to modify it.
The Track Interface configuration window appears.
- Step 3** In the Track Object Name field of the Track Interface configuration window, 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 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, choose 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 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 from 1 to 4096.
- Step 8** Do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. You return to the Track Interface table.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Track Interface table.
 - Click **Next** to deploy your entries and to configure the next entry in the Track Interface table.
-

Related Topics

- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking Hosts for High Availability, page 13-25](#)

Tracking Hosts for High Availability

You can configure a tracking and failure detection process for a gateway or host.



Note

Fields and information related to IPv6 require ACE module and ACE appliance software Version A5(1.0) or later.

Procedure

- Step 1** Choose **Config > Devices > admin_context > HA Tracking And Failure Detection > Hosts**.
- The Track Host table appears.
- Step 2** In the Track Host table, click **Add** to add a new tracking process to the table, or choose an existing entry and click **Edit** to modify it.
- The Track Host configuration window appears.
- Step 3** In the Track Object Name field of the Track Host configuration window, enter a unique identifier for the tracking process.
- Valid entries are unquoted text strings with no spaces.
- Step 4** In the IP Address Type field, choose either **IPv4** or **IPv6** for the host address type.
- This field appears only for ACE module and ACE appliance software Version A5(1.0) or later, which supports IPv4 and IPv6.
- Step 5** In the Track Host/IP Address field, enter the IPv4 or IPv6 address or hostname of the gateway or host that you want the active member of the high availability group to track.
- Step 6** In the Priority field, enter the priority of the probe sent by the active member.

Valid entries are 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 7 In the Peer Host/IP Address field, enter the IPv4 or IPv6 address or hostname of the host that you want the standby member to track.

Step 8 In the Peer Priority field, enter the priority of the probe sent by the standby member.

Valid entries are 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 9 Do one of the following:

- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. Continue with configuring track host probes. See [Configuring Host Tracking Probes, page 13-26](#).
- Click **Cancel** to exit this procedure without saving your entries and to return to the Track Host table.
- Click **Next** to deploy your entries and to configure another tracking process.

Related Topics

- [Configuring Host Tracking Probes, page 13-26](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Configuring Host Tracking Probes

You can configure probes on the active high availability group member to track the health of the gateway or host.

Assumptions

This topic assumes the following:

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability, page 13-25](#).)
- At least one health monitoring probe has been configured (see [Configuring Health Monitoring for Real Servers, page 8-51](#)).

Procedure

Step 1 Choose **Config > Devices > admin_context > HA Tracking And Failure Detection > Hosts**.

The Track Host table appears.

Step 2 Choose the tracking process that you want to modify, and click 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 choose an existing peer host tracking probe and click **Edit** to modify it.
- The Peer Track Host Probes configuration window appears.
- Step 4** In the Probe Name field, choose 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 that you are tracking by the active member.
- Valid entries are 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** Do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. You return to the Track Host Probe table. The table includes the added probe.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Track Host Probe table.
 - Click **Next** to deploy your entries and to configure another track host probe.
-

Related Topics

- [Configuring ACE Peer Host Tracking Probes, page 13-28](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Deleting Host Tracking Probes

You can remove a high availability host tracking probe.

Procedure

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

Related Topics

- [Configuring ACE Peer Host Tracking Probes, page 13-28](#)
- [Configuring ACE High Availability Peers, page 13-15](#)

- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Configuring ACE Peer Host Tracking Probes

You can configure probes on the standby member of a high availability group to track the health of the gateway or host.

Assumptions

This topic assumes the following:

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability, page 13-25](#).)
- At least one health monitoring probe has been configured (see [Configuring Health Monitoring for Real Servers, page 8-51](#)).

Procedure

-
- Step 1** Choose **Config > Devices > ACE admin_context > HA Tracking And Failure Detection > Hosts**.
The Track Host table appears.
- Step 2** In the Track Host table, choose the tracking process that you want to modify, and click 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 choose an existing peer host tracking probe and click **Edit** to modify it.
The Peer Track Host Probes configuration window appears.
- Step 4** In the Probe Name field of the Peer Track Host Probes configuration window, choose 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 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** Do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. You return to the Peer Track Host Probes table. The table includes the added probe.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the Peer Track Host Probes table.
 - Click **Next** to deploy your entries and to configure another peer track host probe.
-

Related Topics

- [Configuring Host Tracking Probes, page 13-26](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Deleting Peer Host Tracking Probes

You can remove a high availability peer host tracking probe.

Procedure

-
- Step 1** Choose **Config > Devices > ACE admin_context > HA Tracking And Failure Detection > Hosts**.
The Track Host table appears.
- Step 2** In the Track Host table, choose the tracking process that you want to modify and click 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, choose the probe that you want to remove, and click **Delete**.
The probe is deleted and the Peer Track Host Probes table refreshes without the deleted probe.
-

Related Topics

- [Configuring ACE Peer Host Tracking Probes, page 13-28](#)
- [Configuring Host Tracking Probes, page 13-26](#)
- [Tracking ACE VLAN Interfaces for High Availability, page 13-24](#)

Configuring ACE HSRP Groups

You can add or edit a Hot Standby Router Protocol (HSRP) group.

Assumptions

This topic assumes the following:

- At least one host tracking process for high availability has been configured (see [Tracking Hosts for High Availability, page 13-25](#).)
- 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** Choose **Config > Devices > ACE admin_context > HA Tracking And Failure Detection > HSRP Groups**.
- The HSRP Groups table appears.
- Step 2** In the HSRP Groups table, click **Add** to add a new HSRP group, or choose an existing entry and click **Edit** to modify it.
- The HSRP Group configuration window appears.
- Step 3** In the Track Object Name field of the HSRP Group configuration window, 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 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 a value 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** Do one of the following:
- Click **Deploy Now** to deploy this configuration on the ACE and save your entries to the running-configuration and startup-configuration files. You return to the HSRP Groups table. The table includes the added HSRP group.
 - Click **Cancel** to exit this procedure without saving your entries and to return to the HSRP Groups table.
-

Synchronizing ACE High Availability Configurations

When two ACE 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**

ANM 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 ANM with the standby ACE.

Synchronizing configuration files for the standby ACE requires the following:

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 6-105](#).
2. Uploading the configuration from the standby ACE to the ANM server. See [Synchronizing Virtual Context Configurations, page 6-105](#).
3. Ensuring that the SSL certificate/keys are imported and identical for the pair. See [Synchronizing SSL Certificate and Key Pairs on Both ACE Peers, page 13-32](#).
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

When configuration changes are made from ANM on any of the ACE devices in a HA pair, ANM automatically detects the active HA peer and deploys the configuration changes to the active ACE alone. ANM does not attempt to deploy a configuration to a standby ACE even if you selected the standby ACE from the ANM device tree. ANM detects the active ACE and will always deploy configuration changes only to the active ACE. In addition, if ACE HA auto-sync is enabled, after the deployment is successful, ANM will locally replicate the configuration in the ANM database on the standby as well to ensure that the ANM configuration is in synchronization with that of the two ACE peers.

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 may 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, 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.

**Note**

If a context is put into an out-of-sync state, this context will be automatically synchronized by the backend ANM. It is not necessary for you to perform an explicit synchronization to take care of the out-of-sync state.

For information on synchronizing virtual context configurations, see [Synchronizing Virtual Context Configurations, page 6-105](#).

Related Topics

- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Synchronizing Virtual Context Configurations, page 6-105](#)

Synchronizing SSL Certificate and Key Pairs on Both ACE Peers

You can reconcile the SSL certificates and key pairs. When SSL certificate/key import is attempted on a peer that is configured in HA, ANM detects the HA state and also imports the same certificate/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 window enabling you to do this.

Guidelines and Restrictions

The certificate/key reconciliation feature is available from the Admin context only; however, executing this feature from the Admin context also reconciles the SSL certificates and key pairs on all the virtual contexts associated with the ACE peers.

Procedure

Step 1 Choose **Config > Devices > admin_context > High Availability (HA) > Setup**.

The HA Management window appears at the top of the content area and the HA Groups table appears at the bottom.

Step 2 In the HA Groups table, choose the group that you want to reconcile the SSL certificates and key pairs on the two HA pairs after a switchover occurs, and click **SSL Certificate/Key Reconcile**.

The SSL Certificate/Key Reconciliation popup window appears. Information appears in this popup window for the primary ACE and the peer ACE as described in [Table 13-5](#).

Table 13-5 SSL Certificate/Key Reconciliation Popup Window Attributes

Field	Description
This Device	IP address for the fault-tolerant VLAN.
Peer Device	Fault-tolerant VLAN to be used for this high availability pair. Valid entries are from 1 to 4094. Note This VLAN cannot be used for other network traffic.
Context Name	Unique name for the virtual context
Matched State	Feature that indicates a match between the SSL certificates and key pairs on the active ACE and the standby ACE peer.
Not Matched State	Feature that indicates that there is not a match between the SSL certificates and key pairs on the active ACE and the standby ACE peer.
SSL Certificates/Keys On Both HA Peers	
File Type	Format of the file: PEM, DER, or PKCS12.
Name	Name of the file that contains the certificate or key pair.
Exportable	Field that indicates whether or not you can export the file from the ACE. Choices are as follows: <ul style="list-style-type: none"> Yes—You can export the file to an FTP, SFTP, or TFP server (see Chapter 11, “Configuring SSL”). No—You cannot export the file as it is protected.
Matched	Field that indicates that the SSL certificate and key pair is a match on the peer ACE.
Available On	Field that identifies the ACE devices that contain the SSL certificate and key pair.

- Step 3** To copy an SSL certificate and key pair to the ACE peer device, choose it from the SSL Certificates/Keys On Both HA Peers list, and then click **Copy To Peer** (or click **Cancel** to close the SSL Certificate/Key Reconciliation popup window without performing the copy).
- Step 4** To delete an SSL certificate and key pair from the ACE HA pair, choose it from the SSL Certificates/Keys On Both HA Peers list, and click **Delete** (or click **Cancel** to close the SSL Certificate/Key Reconciliation popup window without performing the deletion).
-

Related Topics

- [Understanding ANM High Availability, page 13-2](#)
- [Configuring ACE High Availability Peers, page 13-15](#)
- [Configuring ACE High Availability Groups, page 13-17](#)
- [Synchronizing ACE High Availability Configurations, page 13-30](#)

