CHAPTER 6

Configuring Redundant ACEs

This chapter describes how to configure the Cisco 4700 Series Application Control Engine (ACE) appliance for redundancy, which provides fault tolerance for the stateful switchover of flows. It contains the following major sections:

Information About Redundancy

Redundancy (or fault tolerance) uses a maximum of two ACEs to ensure that your network remains operational even if one of the appliances becomes unresponsive. Redundancy ensures that your network services and applications are always available.

Redundancy provides seamless switchover of flows in case an ACE becomes unresponsive or a critical host or interface fails. Redundancy supports the following network applications that require fault tolerance:

- Mission-critical enterprise applications
- Banking and financial services
- E-commerce
- Long-lived flows such as FTP and HTTP file transfers

This section contains the following topics:

Redundancy Protocol

The ACE uses a proprietary protocol to enable redundant configurations of two ACEs (peers). You can configure a maximum of two ACEs for redundancy. Each peer appliance can contain one or more fault-tolerant (FT) groups. Each FT group consists of two members: one active context and one standby context. For more information about contexts, see the *Cisco 4700 Series Application Control Engine Appliance Virtualization Configuration Guide*. An FT group has a unique group ID that you assign.

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 do not require updating. The ACE selects a VMAC from a pool of virtual MACs available to it. For more information about VMACs, see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.

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, interface, or HSRP group fails (see the "Configuring Tracking and Failure Detection" section).
- You enter the **ft switchover** command to force a switchover (see the "Forcing a Failover" section).

Figure 6-1 shows two possible redundancy configurations, where N is the number of ACEs configured for redundancy. The letters (A, B, C, and D) represent the active contexts in each redundancy group, while the primed letters (A', B', C', and D') are the standby contexts. The contexts are evenly distributed between the two ACEs. You always configure the active and the standby contexts on different ACEs.

Figure 6-1 Even Distribution of Contexts

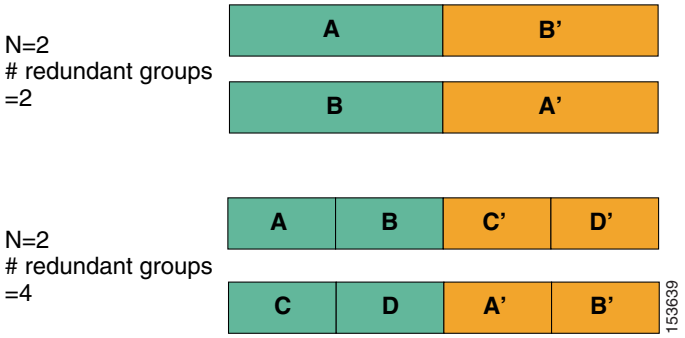

Figure 6-2 shows the uneven distribution of contexts between the two ACEs. As an example, it is possible that the FT groups A,B, C, and D use only half the resources that E and F require.

Figure 6-2 Uneven Distribution of Contexts

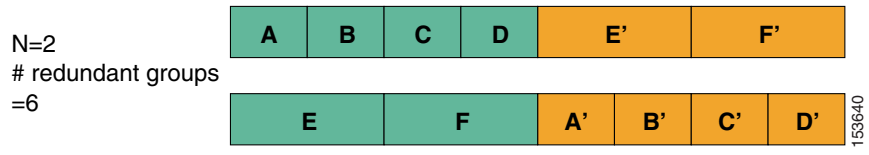

To outside nodes (clients and servers), the active and standby FT group members appear as one node with respect to their IP addresses and associated VMAC. The ACE provides active-active redundancy with multiple-contexts only when there are multiple FT groups configured on each appliance and both appliances 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. For details about configuring contexts, see the *Cisco 4700 Series Application Control Engine Appliance Virtualization Configuration Guide*.

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 (see the "Configuring an FT Peer" section).

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, causing the member with the higher priority always to assert itself and become active (see the "Configuring an FT Group" section).

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.

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 appliance 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
- Sticky table

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.

Note During failover, the ACE sends failover traffic to destination addresses as Layer 3 unicast and Layer 2 broadcast. As a result, you may encounter high CPU utilization in the interrupt context on the switch that connects the two ACEs in the failover setup.

FT VLAN

Redundancy uses a dedicated FT VLAN between redundant ACEs to transmit flow-state information and the redundancy heartbeat. You configure this same VLAN on both peer appliances.

The two redundant appliances constantly communicate over the FT VLAN to determine the operating status of each appliance. 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 FT VLAN resides in the system configuration file. Each FT VLAN on the ACE has one unique MAC address associated with it. The ACE uses these device MAC addresses as the source or destination MACs for sending or receiving redundancy protocol state and configuration replication packets.

Configuration Synchronization

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.

Note In a redundant configuration, with a large configuration on the active ACE, you may encounter a lengthy period of time (sometimes up to 4 hours) for the configuration to be applied and synchronized to the standby ACE.

For information about configuring config sync, see the "Synchronizing Redundant Configurations" section.

Redundancy State for Software Upgrade or Downgrade

The STANDBY_WARM and WARM_COMPATIBLE redundancy states are used when upgrading or downgrading the ACE software. 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 CLI incompatibility.

When the software versions are different while upgrading or downgrading, the STANDBY_WARM and WARM_COMPATIBLE states allows the configuration and state synchronization process to continue on a best-effort basis, which 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 CLI commands or state information. These states allow the standby ACE to come up with best-effort support. In the STANDBY_WARM state, as with the STANDBY_HOT state, the configuration mode is disabled and configuration and state synchronization continues. A failover from the active to the standby based on priorities and preempt can still occur while the standby is in the STANDBY_WARM state.

Guidelines and Limitations

Configuring redundant ACEs has the following guidelines and limitations:

- Redundancy is not supported between an ACE module and an ACE appliance operating as peers. Redundancy must be of the same ACE device type and software release.
- You can configure a maximum of two ACEs (peers) for redundancy.
- Each peer appliance can contain one or more fault-tolerant (FT) groups. Each FT group consists of two members: one active context and one standby context. For more information about contexts, see the *Cisco 4700 Series Application Control Engine Appliance Virtualization Configuration Guide*. An FT group has a unique group ID that you assign.
- 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 do not require updating. The ACE selects a VMAC from a pool of virtual MACs available to it. For more information about VMACs, see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.
- In bridged mode (Layer 2), two contexts cannot share the same VLAN.
- To achieve active-active redundancy, a minimum of two contexts and two FT 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 *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.
- By default, ACE does not replicate IP address sticky table entries on the standby ACE unless you use the **replicate sticky** command in sticky-IP configuration mode. For details on the **replicate sticky** command, see the *Cisco 4700 Series Application Control Engine Appliance Server Load-Balancing Configuration Guide*.
- The ACE does not replicate SSL and other terminated (proxied) connections from the active context to the standby context.
- The ACE does not support the stateful failover of any connections that are proxied. Such connections include Layer 7 connections (including SSL), inspection, and HTTP compression. Also, any connections that are candidates for compression in the VIP but are not being compressed because of the mime type of the data, for example, will remain proxied and will not be supported by stateful failover.
- 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 appliance.

- Do not use this dedicated VLAN for any other network traffic, including HSRP and data.
- Redundancy uses a dedicated FT VLAN between redundant ACEs to transmit flow-state information and the redundancy heartbeat. You must configure this same VLAN on both peer appliances. You also must configure a different IP address within the same subnet on each appliance for the FT VLAN.
- The IP address and the MAC address of the FT VLAN do not change at switchover.
- For redundancy to function properly, both members of an FT group must have identical configurations. Ensure that both ACE appliances include the same bandwidth software license (2G or 1G) and the same virtual context software license. If there is a mismatch in a software license between the two ACE appliances in an FT group, the following operational behavior can occur:
  - If there is a mismatch in the virtual context software license, synchronization between the active ACE and standby ACE may not work properly.
  - If both the active and the standby ACE appliances have the same virtual context software license but have a different bandwidth software license, synchronization will work properly but the standby ACE may experience a potential loss of traffic on switchover from the 2G ACE appliance to the 1G ACE appliance.

  For details about the available ACE software licenses, see Chapter 3, Managing ACE Software Licenses.

Default Settings

Table 6-1 lists the default settings for the ACE redundancy parameters.

Table 6-1 Default Redundancy Parameters

Parameter	Default
Connection replication	Enabled
Heartbeat interval (frequency in milliseconds (ms) at which the active member of the FT group sends the heartbeat packets to the standby member)	300 ms
Heartbeat count (number of missed heartbeats that the standby member must detect before determining that the active member is not available)	10
A member (context) of an FT group becomes the active member through an election process based on the priority that you configure for the group on each peer. The group member with the higher priority becomes the active member.	The group member with the higher priority becomes the active member.
Priority setting of an FT group on the active member.	100
Priority setting of an FT group on the remote standby member.	100
Automatic synchronization of the startup and running configurations between the active and the standby contexts of an FT group.	Enabled
Priority level for multiple probes on the active member.	0
Preempt	Enabled

Configuring Redundant ACEs

This section describes how to configure redundant ACEs and includes the following topics:

Task Flow for Configuring Redundancy

Follow these steps to configure redundancy on the ACE:

Step 1 If you are operating in multiple contexts, observe the CLI prompt to verify that you are operating in the desired context. If necessary, change to the correct context.

```
host1/Admin# changeto C1
host1/C1#
```

The rest of the examples in this table use the Admin context, unless otherwise specified. For details on creating contexts, see the *Cisco 4700 Series Application Control Engine Appliance Virtualization Configuration Guide*.

Step 2 Enter configuration mode.

```
host1/Admin# config
host1/Admin(config)#
```

Step 3 Configure one of the Ethernet ports on the ACE for fault tolerance using a dedicated fault-tolerant (FT) VLAN for communication between the members of an FT group.

```
host1/Admin(config-if)# ft-port vlan 200
```

Step 4 Configure a dedicated FT VLAN for communication between the members of the FT group. This FT VLAN is global and is shared by all contexts. Specify the IP address and netmask of the FT VLAN and the IP address and netmask of the remote peer.

```
host1/Admin(config)# ft interface vlan 200
host1/Admin(config-ft-intf)# ip address 192.168.12.1 255.255.255.0
host1/Admin(config-ft-intf)# peer ip address 192.168.12.15 255.255.255.0
host1/Admin(config-ft-intf)# no shutdown
host1/Admin(config-ft-intf)# exit
```

Step 5 Configure a VLAN with an alias IP address that floats between the active and standby ACEs and serves as a shared gateway for the two devices.

```
host1/Admin(config)# interface vlan 100
host1/Admin(config-if)# alias 192.168.1.1 255.255.255.0
host1/Admin(config-if)# exit
```

Step 6 Configure the local redundancy peer appliance, associate the FT VLAN with the peer, configure the heartbeat interval and count, and configure a query interface VLAN.

```
host1/Admin(config)# ft peer 1
host1/Admin(config-ft-peer)# ft-interface vlan 200
host1/Admin(config-ft-peer)# heartbeat count 20
host1/Admin(config-ft-peer)# heartbeat interval 300
host1/Admin(config-ft-peer)# query-interface vlan 400
host1/Admin(config-ft-intf)# exit
```

Step 7 Create at least one FT group on each ACE.

```
host1/Admin(config)# ft group 1
host1/Admin(config-ft-group)#
```

Step 8 Associate a context with each FT group. You must associate the local context and the corresponding peer context with the same FT group.

```
host1/Admin(config-ft-group)# associate-context C1
```

Step 9 Associate the peer context with the FT group.

```
host1/Admin(config-ft-group)# peer 1
```

Step 10 (Optional) Configure the priority of the FT group on the local appliance.

```
host1/Admin(config-ft-group)# priority 100
```

Step 11 (Optional) Configure the priority of the FT group on the peer appliance.

```
host1/Admin(config-ft-group)# peer priority 200
```

Step 12 Place the FT group in service.

```
host1/Admin(config-ft-group)# inservice
host1/Admin(config-ft-group)# exit
```

Step 13 (Optional) Configure one or more critical objects (gateways or hosts, or interfaces) to track for switchover. For example, to configure a critical interface for tracking, enter:

```
host1/Admin(config)# ft track interface VLAN100
host1/Admin(config-ft-track-intf)# track-interface vlan 100
host1/Admin(config-ft-track-intf)# peer track-interface vlan 100
host1/Admin(config-ft-track-intf)# priority 50
host1/Admin(config-ft-track-intf)# peer priority 150
host1/Admin(config-ft-track-intf)# ctrl-z
```

Step 14 (Optional) Enable autosynchronization of the running- and/or startup-configuration file from the active to the standby context.

```
host1/Admin(config)# ft auto-sync running-config
host1/Admin(config)# ft auto-sync startup-config
```

Step 15 (Optional) Save your configuration changes to Flash memory.

```
host1/Admin(config)# exit
host1/Admin# copy running-config startup-config
```

Step 16 (Recommended) Verify your redundancy configuration by using the following commands in Exec mode:

```
host1/Admin# show running-config ft
host1/Admin# show running-config interface
```

Configuring Redundancy

This section describes how to configure redundancy on the ACE and contains the following topics:

- Configuring an FT VLAN
- Configuring an Alias IP Address
- Configuring an FT Peer

- Configuring an FT Group
- Specifying the Peer Hostname
- Specifying the MAC Address Banks for a Shared VLAN
- Forcing a Failover
- Synchronizing Redundant Configurations

Requirements

You must configure the **ft interface**, **ft peer**, and **ft group** commands on all ACEs that participate in the redundancy configuration.

Configuring an FT VLAN

This section describes how to configure an FT VLAN. Peer ACEs communicate with each other over a dedicated FT VLAN. These redundant peers use the FT VLAN to transmit and receive heartbeat packets and state and configuration replication packets. You must configure the same VLAN on each peer appliance.

Prerequisites

To configure one of the Ethernet ports or a port-channel interface on the ACE for fault tolerance using a dedicated FT VLAN for communication between the members of an FT group, use the **ft-port vlan** command in interface configuration mode (see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*).

Note When you specify the **ft-port vlan** command, the ACE modifies the associated Ethernet port or port-channel interface to a trunk port.

When you specify an Ethernet port as a dedicated FT VLAN, you have the option to either configure the dedicated VLAN as the only VLAN associated with the Ethernet port or to include it as part of a VLAN trunk link (see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*). Note that the ACE automatically includes the FT VLAN in the VLAN trunk link. If you choose to configure VLAN trunking, it is not necessary for you to assign the FT VLAN in the trunk link along with the other VLANs.

We recommend that you enable Quality of Service (QoS) on the FT VLAN port to provide higher priority for FT traffic. See the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide* for details.

Restrictions

Do not use this dedicated VLAN for any other network traffic, including HSRP and data.

This topic includes the following restrictions:

- Do not use this dedicated VLAN for any other network traffic, including data.
- On both peer ACE appliances, you must configure the same Ethernet port or port-channel interface as the FT VLAN port. For example:
  - If you configure ACE appliance 1 to use Ethernet port 4 as the FT VLAN port, then be sure to configure ACE appliance 2 to use Ethernet port 4 as the FT VLAN port.
  - If you configure ACE appliance 1 to use port-channel interface255 as the FT VLAN port, then be sure to configure ACE appliance 2 to use port-channel interface 255 as the FT VLAN.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft interface vlan** *vlan_id* **Example:** host1/Admin(config)# ft interface vlan 200 host1/Admin(config-ft-intf)#	Creates an FT VLAN. The *vlan_id* argument specifies a unique identifier for the FT VLAN. Enter an integer from 2 to 4094. This command enters the FT interface configuration mode.
	no ft interface vlan *vlan_id* **Example:** host1/Admin(config)# no ft interface vlan 200	(Optional) Removes an FT VLAN from the redundancy configuration. **Note** To remove an FT VLAN, first remove it from the FT peer by using the **no ft-interface vlan** command in FT peer configuration mode.
Step 3	**ip address** *ip_address netmask* **Example:** host1/Admin(config-ft-intf)# ip address 192.168.12.1 255.255.255.0	Assigns an IP address to the VLAN. The keyword and arguments of this command are: • **address** *ip_address*—Specifies the IP address of the FT VLAN. • *netmask*—Subnet mask of the FT VLAN. Enter a subnet mask in dotted-decimal notation.
	no ip address *ip_address netmask* **Example:** host1/Admin(config-ft-intf)# no ip address 192.168.12.1 255.255.255.0	(Optional) Removes an IP address from an FT VLAN.

	Command	Purpose
Step 4	**peer ip address** *ip_address netmask* **Example:** host1/Admin(config-ft-intf)# peer ip address 192.168.12.15 255.255.255.0	Allows the local member to communicate with the remote peer. The keyword and arguments of this command are as follows: • **address** *ip_address*—Specifies the IP address of the remote peer. • *netmask*—Subnet mask of the remote peer. Enter a subnet mask in dotted-decimal notation.
	no peer ip address *ip_address netmask* **Example:** host1/Admin(config-ft-intf)# no peer ip address 192.168.12.15 255.255.255.0	(Optional) Removes an IP address from the remote peer.
Step 5	**no shutdown** **Example:** host1/Admin(config-ft-intf)# no shutdown	Enables the FT VLAN.
	shutdown **Example:** host1/Admin(config-ft-intf)# shutdown	(Optional) Disables the FT VLAN after you have enabled it.
Step 6	**exit** **Example:** host1/Admin(config-ft-intf)# exit host1/Admin(config)#	(Optional) Exits the fault-tolerant interface configuration mode.
Step 7	**do copy running-config startup-config** **Example:** host1/Admin(config)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring an Alias IP Address

This section describes how to configure an alias IP address. When you configure redundancy, configure a VLAN interface that has an alias IP address that floats between the active and standby appliances. The alias IP address serves as a shared gateway for the two ACE appliances.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**interface vlan** *vlan_id* **Example:** host1/Admin(config)# interface vlan 100 host1/Admin(config-if)#	Enters interface configuration mode. The *vlan_id* argument specifies a unique identifier for the VLAN. This command enters the FT interface configuration mode.

	Command	Purpose
Step 3	**alias** *ip_address netmask* **Example:** host1/Admin(config-if)# alias 192.168.1.1 255.255.255.0	Configures an alias IP address. The *ip_address netmask* arguments specify the IP address and netmask for the VLAN interface. Enter the IP address and subnet mask in dotted-decimal notation.
	no alias *ip_address netmask* **Example:** host1/Admin(config-if)# no alias 192.168.1.1 255.255.255.0	(Optional) Removes an alias IP address.
Step 4	**do copy running-config startup-config** **Example:** host1/Admin(config-if)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring an FT Peer

This section describes how to configure an FT peer definition on both peer ACEs.

Restrictions

This topic includes the following restrictions:

- You must create FT peers in the admin context only.
- You can configure a maximum of two ACEs as redundancy peers.
- Before you can remove an FT peer from the configuration by using the **no** form of the command, you must remove the peer from the FT group (see the "Configuring an FT Group" section).
- You cannot delete a query interface if it is associated with a peer. You must disassociate the interface from the peer first, and then you can delete the interface.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft peer** *peer_id* **Example:** host1/Admin(config)# ft peer 1 host1/Admin(config-ft-peer)	Creates an FT peer. The *peer_id* argument specifies a unique identifier for the peer. You can only enter 1. This command enters the FT peer configuration mode.
	no ft peer *peer_id* **Example:** host1/Admin(config)# no ft peer 1	(Optional) Removes the FT peer from the configuration.

	Command	Purpose
Step 3	**ft-interface vlan** *vlan_id* **Example:** host1/Admin(config-ft-peer) ft-interface vlan 200	Associates an FT VLAN with a peer. The *vlan_id* argument specifies the identifier of an existing VLAN. Enter an integer from 2 to 4094.
	no ft-interface vlan *vlan_id* **Example:** host1/Admin(config-ft-peer) no ft-interface vlan 200	(Optional) Removes the FT VLAN from the peer configuration.
Step 4	**heartbeat** {**count** *number* | **interval** *frequency*} **Example:** host1/Admin(config-ft-peer) heartbeat interval 500	Configures the heartbeat interval and count. The keywords and arguments are: • **count** *number*—Specifies the number of heartbeat intervals that must transpire with no heartbeat packet received by the standby member before the standby member determines that the active member is not available. Enter an integer from 10 to 50. The default is 10 heartbeat intervals. If the standby member of the FT group does not receive a heartbeat packet from the active member, a time period equal to **count** *number* times **interval** *frequency* must elapse before a switchover can occur. For example, in the default case, where the heartbeat frequency is 300 ms and the heartbeat count is 10, if the standby member does not receive a heartbeat packet from the active member for 3000 ms (3 seconds), a switchover occurs. • **interval** *frequency*—Specifies the interval in milliseconds (ms) between heartbeats. Enter an integer from 100 to 1000 ms. The default is 300 ms.
	no heartbeat {**count** *number* | **interval** *frequency*} **Example:** host1/Admin(config-ft-peer) no heartbeat interval 500	(Optional) Resets either the heartbeat count to the default of 10 or the heartbeat interval to the default of 100 ms.
Step 5	**query-interface vlan** *vlan-id* **Example:** host1/Admin(config-ft-peer)# query-interface vlan 400	Configures a query interface to allow the standby member to determine whether the active member is down or if there is a connectivity problem with the FT VLAN. A query interface helps prevent two redundant contexts from becoming active at the same time for the same FT group. Before triggering a switchover, the ACE pings the active member to make sure that it is down. Configuring a query interface allows you to assess the health of the active member, but it increases switchover time. The *vlan_id* argument specifies the identifier of an existing VLAN. Enter an integer from 2 to 4094.

	Command	Purpose
	no query-interface vlan *vlan-id* **Example:** host1/Admin(config-ft-peer)# no query-interface vlan 400	(Optional) Removes a query interface from the peer configuration. **Note** You cannot delete a query interface if it is associated with a peer. You must disassociate the interface from the peer first, and then you can delete the interface.
Step 6	**do copy running-config startup-config** **Example:** host1/Admin(config-ft-peer)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Configuring an FT Group

This section describes how to configure multiple FT groups on each ACE.

Prerequisites

Before you place an FT group in service, be sure that you have associated one context with the FT group and that you have properly configured the two peers.

Restrictions

This topic includes the following restrictions:

- You must configure the same group ID on both peer appliances.
- The maximum number of FT groups that you can create is 64 groups (20 user contexts and 1 Admin context).
- Each FT group consists of a maximum of two members (contexts): one active context on one appliance and one standby context on the peer appliance
- Before you can remove a context from an FT group, you must first take the group out of service by using the **no inservice** command.
- The ACE does not perform bulk config synchronization (sync) on the **peer priority** command value in the FT group associated with the Admin context to the peer. Therefore, you may observe a peer priority value in the running-configuration file that is different from the actual operating value. For information on bulk config sync, see the "Synchronizing Redundant Configurations" section.
- If you disable preemption by using the **no preempt** command and a member with a higher priority is found after the other member has become active, the electing member becomes the standby member even though it has a higher priority.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft group** *group_id* **Example:** host1/Admin(config) ft group 1 host1/Admin(config-ft-group)#	Creates an FT group. The *group_id* argument specifies a unique identifier of the group. Enter an integer from 1 to 64. This command enters the FT group configuration mode.
	no ft group *group_id* **Example:** host1/Admin(config) no ft group 1	(Optional) Removes the FT group from the configuration.
Step 3	**associate-context** *name* **Example:** host1/Admin(config-ft-group)# associate-context C1	Associates a context with an FT group.
	no associate-context *name* **Example:** host1/Admin(config-ft-group)# no associate-context C1	(Optional) Removes a context from an FT group.
Step 4	**peer** *peer_id* **Example:** host1/Admin(config-ft-group)# peer 1	Associates a peer ACE with an FT group. For the *peer_id* argument, enter 1 as the identifier of an existing peer appliance. You can only enter 1.
	no peer *peer_id* **Example:** host1/Admin(config-ft-group)# no peer 1	(Optional) Removes the peer association with the FT group.
Step 5	**priority** *number* **Example:** host1/Admin(config-ft-group)# priority 150	Configures the priority of an FT group on the active member. Configure a higher priority on the FT group member that you want to be the active member. The *number* argument specifies the priority of the FT group on the local peer. Enter an integer from 1 to 255. The default is 100.
	no priority **Example:** host1/Admin(config-ft-group)# no priority	(Optional) Restores the default priority of 100.
Step 6	**peer priority** *number* **Example:** host1/Admin(config-ft-group)# peer priority 150	Configures the priority of an FT group on the remote standby member. Configure a lower priority on the FT group member that you want to be the standby member. The *number* argument specifies the priority of the FT group on the standby member. Enter an integer from 1 to 255. The default is 100.

	Command	Purpose
	no peer priority **Example:** host1/Admin(config-ft-group)# no priority	(Optional) Restores the default priority of 100.
Step 7	**preempt** **Example:** host1/Admin(config-ft-group)# preempt	Configures preemption after it has been disabled. Preemption ensures that the group member with the higher priority always asserts itself and becomes the active member. By default, preemption is enabled.
	no preempt **Example:** host1/Admin(config-ft-group)# no preempt	(Optional) Disables preemption.
Step 8	**inservice** **Example:** host1/Admin(config-ft-group)# inservice	Places an FT group in service.
	no inservice **Example:** host1/Admin(config-ft-group)# no inservice	(Optional) Takes the FT group out of service.
Step 9	**do copy running-config startup-config** **Example:** host1/Admin(config-ft-group)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Modifying an FT Group

This section describes how to modify an FT group.

Note You can modify the **priority**, **peer priority**, and **preempt** command values without taking the FT group out of service.

Details

Follow these steps to modify an FT group:

Step 1 Remove the FT group from service by using the **no inservice** command.

Step 2 Make the necessary modifications to the FT group.

Step 3 Place the FT group back in service by using the **inservice** command.

Specifying the Peer Hostname

This section describes how to specify the peer hostname.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**peer hostname** *name* **Example:** host1/Admin(config)# peer hostname ACE_2	Specifies the hostname of a peer ACE. For details about this command, see the "Assigning a Name to the ACE" section.

Specifying the MAC Address Banks for a Shared VLAN

This section describes how to specify the MAC address banks to be used by the local ACE and the peer ACE with a shared VLAN (FT VLAN). You configure these commands to prevent MAC address conflicts between the two peer ACEs. For details about these commands, see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.

Restrictions

This topic includes the following restrictions:

- Perform this task from the Admin context only.
- Select a bank of MAC addresses for the peer that is different from that used by the local ACE.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**shared-vlan-hostid** *number* **Example:** host1/Admin(config)# shared-vlan-hostid 3	Configures the 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. The *number* argument is the 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. For details about this command, see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.

	Command	Purpose
	no shared-vlan-hostid **Example:** host1/Admin(config)# no shared-vlan-hostid	(Optional) Removes a configured bank of MAC addresses.
Step 3	**peer shared-vlan-hostid** *number* **Example:** host1/Admin(config)# peer shared-vlan-hostid 3	Configures a specific bank of MAC addresses for a peer ACE in a redundant configuration. The *number* argument is the 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. For details about this command, see the *Cisco 4700 Series Application Control Engine Appliance Routing and Bridging Configuration Guide*.
	no peer shared-vlan-hostid **Example:** host1/Admin(config)# no peer shared-vlan-hostid	(Optional) Removes the configured bank of MAC addresses.
Step 4	**do copy running-config startup-config** **Example:** host1/Admin(config)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Forcing a Failover

This section describes how to force a failover (switchover). 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 becoming the active member of the FT group, a switchover occurs.

Note During failover, the ACE sends failover traffic to destination addresses as Layer 3 unicast and Layer 2 broadcast. As a result, you may encounter high CPU utilization in the interrupt context on the switch that connects the two ACEs in the failover setup.

The switchover process exhibits the following behavior, depending on whether you perform the task from the Admin context or a user context:

- Admin context—If you specify an FT group ID, then the FT group specified by the group ID switches over. If you do not specify a group ID, then the Admin context switches over.
- User context—Because you cannot specify an FT group ID in a user context, the context in which you enter the command switches over.

Note When you specify the **ft switchover** command to force a switchover, there may be brief periods of time when the configuration mode is enabled on the new active group member to allow the administrator to make configuration changes. However, any configuration changes made during this time are not synchronized with the standby group member and will exist only on the active group member. We recommend that you refrain from making any configuration changes after you enter the **ft switchover** command until the FT states stabilize to ACTIVE and STANDBY_HOT. Once the FT group reaches the

steady state of ACTIVE and STANDBY_HOT, any configuration changes performed on the active group member will be dynamically synchronized to the standby group member, assuming that configuration synchronization is enabled.

Prerequisites

To use the **ft switchover** command, you must disable preemption by using the **no preempt** command. For information on the **preempt** command, see the "Configuring an FT Group" section.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft group** *group_id* **Example:** host1/Admin(config) ft group 1 host1/Admin(config-ft-group)#	Enters the FT group configuration mode.
Step 3	**no preempt** **Example:** host1/Admin(config-ft-group)# no preempt	Disables preemption.
Step 4	**Ctrl-z** **Example:** host1/Admin(config-ft-group)# Ctrl-z host1/Admin#	Returns to the Exec mode prompt.
Step 5	**ft switchover** [**all** [**force**] | **force** | [*group_id* [**force**]]] **Example:** host1/Admin# ft switchover 1 This command will cause card to switchover (yes/no)? [no] yes	Causes a switchover. The keywords, arguments, and options are: • **all**—(Optional) Causes a switchover of all FT groups configured in the ACE simultaneously. This keyword is available in the Admin context only. • **force**—(Optional) Causes a switchover while ignoring the state of the standby member. Use this option only when the FT VLAN is down. This keyword is available in the Admin context only. • *group_id*—(Optional) FT group that you want to switch over. Enter the ID of an existing FT group as an integer from 1 to 255. This argument is available in the Admin context only.

Synchronizing Redundant Configurations

This section describes how to synchronize redundant configurations. To ensure that the running configurations on both the active and the standby contexts of an FT group are identical, the ACE automatically synchronizes the running configurations between the two contexts. After the active

context has accepted either a new configuration or modifications to an existing configuration, the ACE automatically applies the new configuration or configuration changes to the standby context and disables configuration mode in the standby context.

The ACE supports the following two types of configuration synchronizations:

- Bulk config sync—Synchronizes the entire active context configuration to the standby context when the peer comes up or when autosynchronization is enabled
- Dynamic incremental sync—Synchronizes the configuration applied to the active context to the standby context if the peer is already up

Note When you upgrade from one major release of ACE software to another major release (for example, from), dynamic incremental sync is disabled while the active ACE is running A4(1.0) and the standby is running the earlier release (split mode). We recommend that you do not make any configuration changes during this time and that you do not keep the ACEs in this state for an extended period of time. However, if you must make configuration changes while the ACEs are in split mode, ensure that you manually synchronize to the standby ACE any configuration changes that you make on the active ACE. After you complete the software upgrade of both ACEs, a bulk sync occurs automatically and dynamic incremental sync will be enabled again.

You can enable automatic synchronization of the running-configuration and the startup-configuration files after they have been explicitly disabled.

Caution Toggling **ft auto-sync running-config** in the Admin context may have undesirable side effects if the same command is also disabled in an active user context. If **ft auto-sync running-config** is disabled in the active Admin context and in an active user context, and you subsequently enable **ft auto-sync running-config** in the active Admin context first, the entire configuration of the standby user context will be lost. Always enable **ft auto-sync running-config** in the active user context first, and then enable the command in the active Admin context.

Restrictions

This topic includes the following restrictions:

- The configurations on both the active context and the standby context must be identical. If there is a mismatch between configuration objects, then configuration synchronization may fail.
- In a redundant configuration, with a large configuration on the active ACE, you may encounter a lengthy period of time (sometimes up to 4 hours) for the configuration to be applied and synchronized to the standby ACE.
- If the standby ACE has reached the maximum resource limit for a configuration object even if some of the configuration objects are not in the redundant context and you configure one more object of the same type in the redundant context of the active ACE, configuration synchronization will fail. For example, suppose that you have configured two contexts on each ACE (Admin and C1) and the C1 context is the only one in the FT group. On the standby ACE, you have configured 8,192 **match source-address** statements in the Admin context and in the C1 context for a total of 16,384 **match source-address** statements (the ACE limit). When you configure one new **match source-address** statement on the active ACE in C1, configuration synchronization will fail, the new match statement will not be replicated to the standby, and syslog ACE-1-727005 is generated.
- If you operate the active ACE with config sync disabled for a prolonged period of time, you must manually duplicate any changes that you make to the active ACE on the standby ACE to ensure that connection replication works properly.

- If a license mismatch occurs between the two ACEs in a redundant configuration, the **ft auto-sync** command is automatically disabled and a syslog message is generated.
- If you temporarily disable **ft auto-sync running-config** on the active ACE (for example, to test changes to your configuration), when you subsequently reenable config sync, any changes that you made to the active ACE are duplicated on the standby ACE. Note that the standby ACE remains in the STANDBY_HOT state even when config sync is disabled on the active ACE.
- If the configuration synchronization fails, the running-configuration file reverts to the startup-configuration file.
- The ACE does not copy or write changes in the running-configuration file to the startup-configuration file unless you enter the **copy running-config startup-config** command or the **write memory** command for the current context. To write the contents of the running-configuration file to the startup-configuration file for all contexts, use the **write memory all** command. At this time, if the **ft auto-sync startup-config** command is enabled, the ACE synchronizes the startup-configuration file on the active ACE to the standby ACE.
- The ACE does not synchronize the SSL certificates and key pairs that are present in the active context with the standby context of an FT group. If the ACE performs a configuration synchronization and does not find the necessary certificates and keys in the standby context, config sync fails and the standby context enters the STANDBY_COLD state.

Caution Do not enter the **no inservice** command followed by the **inservice** command on the active context of an FT group when the standby context is in the STANDBY_COLD state. Doing so may cause the standby context running-configuration file to overwrite the active context running-configuration file.

To copy the certificates and keys to the standby context, you must export the certificates and keys from the active context to an FTP or TFTP server using the **crypto export** command, and then import the certificates and keys to the standby context using the **crypto import** command. For more information about importing and exporting certificates and keys, see the *Cisco 4700 Series Application Control Engine Appliance SSL Configuration Guide*.

To return the standby context to the STANDBY_HOT state in this case, ensure that you have imported the necessary SSL certificates and keys to the standby context, and then perform a bulk sync of the active context configuration by entering the following commands in configuration mode in the active context of the FT group:

1. **no ft auto-sync running-config**
2. **ft auto-sync running-config**

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/C1# config host1/C1#(config)#	Enters global configuration mode.
Step 2	**ft auto-sync** {**running-config** | **startup-config**} **Example:** host1/C1(config) ft auto-sync running-config	Enables automatic synchronization of the running-configuration and the startup-configuration files after they have been explicitly disabled. The keywords are: • **running-config**—Enables autosynchronization of the running-configuration file. The default is enabled. • **startup-config**—Enables autosynchronization of the startup-configuration file. The default is enabled.
	no ft auto-sync {**running-config** | **startup-config**} **Example:** host1/C1(config) no ft auto-sync running-config	(Optional) Disables automatic synchronization of the running-configuration and the startup-configuration files.

Configuring Tracking and Failure Detection

This section describes the tracking and failure detection feature of the ACE. This feature allows you to designate certain network items as critical so that if one or more items fail, the ACE reduces the priority of the associated active FT group accordingly. If the priority of the active FT group falls below the priority of the corresponding FT group on the standby, a switchover occurs.

The ACE supports the tracking and failure detection of several network items. You can configure an ACE to track and detect failures in the following items in the Admin context and any user context:

- Gateways or hosts
- Interfaces

If one of the items that you configure for tracking and failure detection becomes unresponsive and is associated with the active member of an FT group, by default, the ACE subtracts a value of 0 from the configured priority of the active member. If you configure a nonzero value for the tracking priority and the resulting priority value of the active member 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 that exist at the time of the switchover continue uninterrupted on the new active member of the FT group.

When the failed item comes back up, the ACE increments the priority of the associated group member by a value of 0 by default. If you configure a non-zero value for the tracking priority and the resulting priority of the standby member is greater than the priority of the active member, a switchover occurs back to the original active group member.

You can configure the unit priority associated with tracked items to be greater than 0. This option allows you to fine tune the switchover scenario so that a switchover occurs when either all or any of the tracked objects fails.

Note To prevent an unexpected switchover from occurring, we strongly recommend that you disable preemption while you are configuring tracking. After you configure tracking and before you reenable preemption, ensure that the tracked network objects are up and operating properly. A switchover may occur immediately when you reenable preemption. Preemption must be enabled for a tracking switchover to work. For details about preemption, see the "Configuring an FT Group" section.

For example, suppose that on ACE 1 you configure the active FT group member with a priority of 100 and on ACE 2 you configure the standby FT group member with a priority of 70. Assume that you configure the FT group to track three critical interfaces, each with a unit priority of 15. To trigger a switchover, all three interfaces must fail so that the priority of the active member is less than the priority of the standby member (100 – 45 = 55).

To illustrate the "any" scenario, assume that the active and the standby FT group members have the same individual priorities as in the previous example (100 and 70, respectively). However, this time you configure the three tracked interfaces, each with a unit priority of 40. If any one of the interfaces associated with the active member goes down, then the priority of the active member falls below the priority of the standby member and a switchover occurs. If that failed interface later returns to service, the ACE increments the associated group member priority by 40, and a switchover would occur back to the original active member. To guarantee a switchover if any tracked item goes down, configure the unit priority on each tracked item equal to the group member's priority. In this case, you could configure the unit priority to be 100.

This section contains the following topics:

- Configuring Tracking and Failure Detection for a Host or Gateway
- Configuring Tracking and Failure Detection for an Interface

Configuring Tracking and Failure Detection for a Host or Gateway

This section describes how to configure tracking and failure detection for a gateway or a host.

Restrictions

If you remove a probe from the active FT group member configuration and you have not configured a tracking priority for the FT group, the ACE increments the net FT group priority by the priority value of the deleted probe. You cannot delete a probe from the running-configuration file if the ACE is using the probe for tracking.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft track host** *name* **Example:** host1/Admin(config)# ft track host TRACK_GATEWAY1 host1/Admin(config-ft-track-host)#	Creates a tracking and failure detection process for a gateway or host. For the *name* argument, enter a unique identifier of the tracking process as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters. This commands enters the FT track host configuration mode.
Step 3	**track-host** *ip_address* **Example:** host1/Admin(config-ft-track-host)# track-host 192.168.12.101	Configures the IP address of the gateway or host. The *ip_address* argument specifies the IP address of the gateway or host that you want the active FT group member to track. This command enters the FT group configuration mode.
	no track-host *ip_address* **Example:** host1/Admin(config-ft-track-host)# no track-host 192.168.12.101	(Optional) Removes the IP address of the gateway or host from the tracking process on the standby member configuration.
Step 4	**probe** *name* **priority** *number* **Example:** host1/Admin(config-ft-track-host)# probe TCP_PROBE1 priority 50	Associates an existing probe with a gateway or host for tracking by the active member. For information about creating probes, see the *Cisco 4700 Series Application Control Engine Appliance Server Load-Balancing Configuration Guide*. The keyword and arguments are: • *name*—Identifier of an existing probe that you want to associate with a gateway or host for tracking. • **priority** *number*—Specifies the priority of the probe sent by the active member. Enter 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 gateway or host that the probe is tracking. If the probe goes down, the ACE decrements the priority of the FT group on the active member by the value of the *number* argument. If the resulting priority of the FT group on the active member is less than the priority of the FT group on the standby member, a switchover occurs.
	no probe *name* **Example:** host1/Admin(config-ft-track-host)# no probe TCP_PROBE1	(Optional) Removes the tracking probe from the active member.

	Command	Purpose
Step 5	**priority** *number* **Example:** `host1/Admin(config-ft-track-host)# priority 50`	Assigns a priority for multiple probes on the active member. The *number* argument specifies the priority of the probes on the active member. Enter a priority value 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 gateway or host that the probes are tracking. If all the probes go down, the ACE decrements the priority of the FT group on the active member by the value of the *number* argument. If the resulting priority of the FT group on the active member is less than the priority of the FT group on the standby member, a switchover occurs.
	no priority *number* **Example:** `host1/Admin(config-ft-track-host)# no priority 50`	(Optional) Resets the priority to the default value of 0.
Step 6	**peer track-host** *ip_address* **Example:** `host1/Admin(config-ft-track-host)# peer track-host 172.16.27.1`	Configures the IP address of the gateway or host. The *ip_address* argument specifies the IP address of the gateway or host that you want the standby FT group member to track.
	no peer track-host *ip_address* **Example:** `host1/Admin(config-ft-track-host)# no peer track-host 172.16.27.1`	(Optional) Removes the host tracked by the standby member.
Step 7	**peer probe** *name* **priority** *number* **Example:** `host1/Admin(config-ft-track-host)# peer probe TCP_PROBE1 priority 25`	Associates an existing probe with a gateway or host for tracking by the standby member. The keyword and arguments are: • *name*—Identifier of an existing probe that you want to associate with a gateway or host for tracking. • **priority** *number*—Specifies the priority of the probe sent by the standby member. Enter 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 gateway or host that the probe is tracking. If the probe goes down, the ACE decrements the priority of the FT group on the standby member by the value of the *number* argument.
	no peer probe *name* **Example:** `host1/Admin(config-ft-track-host)# no peer probe TCP_PROBE1`	(Optional) Removes the tracking probe from the standby member.

	Command	Purpose
Step 8	**peer priority** *number* **Example:** host1/Admin(config-ft-track-host)# peer priority 25	Assigns a priority for multiple probes on the standby member. The *number* argument specifies the priority of the probes configured for the gateway or host on the standby member. Enter a priority value 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 gateway or host that the probes are tracking. If all the probes go down, the ACE decrements the priority of the FT group on the standby member by the value of the *number* argument.
	no peer priority *number* **Example:** host1/Admin(config-ft-track-host)# peer priority 25	(Optional) Reset the multiple-probe priority to the default value of 0 on the standby member.
Step 9	**do copy running-config startup-config** **Example:** host1/Admin(config-ft-track-host)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Examples

The following example demonstrates a tracking configuration for a gateway on the active member of an FT group:

```
ft track host TRACK_GATEWAY
  track-host 192.161.100.1
  probe GATEWAY_TRACK1 priority 10
  probe GATEWAY_TRACK2 priority 20
  priority 50
```

In this configuration example, if the GATEWAY_TRACK1 probe goes down, the ACE reduces the priority of the FT group on the active member by 10. If the GATEWAY_TRACK2 probe goes down, the ACE reduces the priority of the FT group on the active member by 20. If both probes go down, the ACE reduces the priority of the FT group on the active member by 50. If at any time 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.

Configuring Tracking and Failure Detection for an Interface

This section describes how to configure tracking and failure detection for an interface.

Restrictions

You cannot delete an interface if the ACE is using the interface for tracking. Also, you cannot configure the FT VLAN for tracking.

Detailed Steps

	Command	Purpose
Step 1	**config** **Example:** host1/Admin# config host1/Admin#(config)#	Enters global configuration mode.
Step 2	**ft track interface** *name* **Example:** host1/Admin(config)# ft track interface TRACK_VLAN100 host1/Admin(config-ft-track-intf)#	Creates a tracking and failure detection process for an interface. For the *name* argument, enter a unique identifier for the tracking process as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters. This commands enters the FT track interface configuration mode.
Step 3	**no ft track interface** *name* **Example:** host1/Admin(config)# ft track interface TRACK_VLAN100	(Optional) Removes the interface-tracking process.
Step 4	**track-interface vlan** *vlan_id* **Example:** host1/Admin(config-ft-track-intf)# track-interface vlan 100	Configures the interface that you want the active member to track. For the *vlan_id* argument, enter the VLAN ID of an existing VLAN configured on the active member as an integer from 2 to 4094.
	no track-interface vlan *vlan_id* **Example:** host1/Admin(config-ft-track-intf)# no track-interface vlan 100	(Optional) Removes the VLAN from the tracking process.
Step 5	**priority** *number* **Example:** host1/Admin(config-ft-track-intf)# priority 50	Configures the interface that you want the active member to track. The *number* argument specifies the priority of the interface on the active member. Enter a priority value 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 interface that you are tracking. If the tracked interface goes down, the ACE decrements the priority of the FT group on the active member by the value of the number argument. 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.
	no priority *number* **Example:** host1/Admin(config-ft-track-intf)# no priority 50	(Optional) Resets the interface priority on the active member to the default value of 0.
Step 6	**peer track-interface vlan** *vlan_id* **Example:** host1/Admin(config-ft-track-intf)# peer track-interface vlan 200	Configures the interface that you want the standby member to track. The *vlan_id* argument is a VLAN ID of an existing VLAN configured on the standby member as an integer from 2 to 4094.

	Command	Purpose
	no peer track-interface vlan *vlan_id* **Example:** host1/Admin(config-ft-track-intf)# no peer track-interface vlan 200	(Optional) Removes the VLAN from the tracking process.
Step 7	**peer priority** *number* **Example:** host1/Admin(config-ft-track-intf)# peer priority 25	Assigns a priority to the tracked interface that the standby member is tracking. The *number* argument specifies the priority of the interface on the standby member. Enter a priority value 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 interface that you are tracking.
	no peer priority *number* **Example:** host1/Admin(config-ft-track-intf)# no peer priority 25	(Optional) Resets the interface priority on the standby member to the default value of 0.
Step 8	**do copy running-config startup-config** **Example:** host1/Admin(config-ft-track-intf)# do copy running-config startup-config	(Optional) Copies the running configuration to the startup configuration.

Examples

The following example demonstrates a tracking configuration for an interface on the active member of an FT group and configures the interface that you want the standby member to track:

```
ft track interface TRACK_VLAN100
  track-interface vlan 100
  priority 50
  peer track-interface vlan 200
  peer priority 25
```

In this configuration example, if VLAN 100 goes down, then the ACE reduces the priority of the FT group on the active member by 50. If at any time 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.

Displaying or Clearing Redundancy Information

This section describes how to display or clear information about redundancy and contains the following sections:

- Displaying Redundancy Information
- Clearing Redundancy Statistics

Displaying Redundancy Information

This section describes the **show** commands that display configuration, status, and statistical information for your redundancy configuration and contains the following sections:

Displaying Redundancy Configuration Information

To display the list of redundancy or fault-tolerance (FT) configurations configured for the current context, perform the following task:

Command	Purpose
show running-config ft	Displays the list of redundancy or fault-tolerance (FT) configurations configured for the current context. The ACE also displays configuration information for each ft configuration listed.

Displaying Bulk Synchronization Command Failures on the Standby ACE

To display the configuration commands that fail on the standby ACE appliance during bulk synchronization in a redundant configuration per context, perform the following task:

Command	Purpose
show ft config-error [*context_name*]	Displays the commands that fail on the standby ACE appliance during bulk synchronization in a redundant configuration per context. If all commands succeed on the standby ACE appliance, the command displays the following message: `No bulk config apply errors` In the Admin context, the optional *context_name* argument is the name of a user context. If you do not enter the argument, the command uses the Admin context. In a user context, this argument is not available.

Displaying FT Group Information

To display redundancy statistics per context, perform the following task:

Command	Purpose
show ft group {{[*group_id*] {**detail** | **status** | **summary**}} | **brief**}	Displays redundancy statistics per context. The keywords, arguments, and options are: • **group** *group_id*—Displays FT group statistics for the specified FT group. In the Admin context, this keyword displays statistics for all FT groups in the ACE. Also, in the Admin context, you can specify an FT group number to display statistics for an individual group. In a user context, this keyword displays statistics only for the FT group to which the user context belongs. • **detail**—Displays detailed information for all FT groups or the specified FT group. The **detail** keyword includes the status of autosync and whether it is disabled or enabled for both the running-config and the startup-config. • **status**—Displays the current operating status for all FT groups or the specified FT group. • **summary**—Displays summary information for all FT groups or the specified FT group. • **brief**—Displays the group ID, local state, peer state, context name, context ID, and configuration synchronization status of all the FT groups that are configured in the ACE.

Table 6-2 describes the fields in the **show ft group** command output.

Table 6-2 Field Descriptions for the show ft group Command Output

Field	Description
FT Group	FT group identifier.
No. of Contexts	Number of contexts associated with the FT group.
Context Name	Name of the context associated with the FT group.
Context ID	Identifier of the context associated with the FT group.
Configured Status	Configured state of the FT group. Possible states are the in-service or out-of-service states.
Maintenance Mode	Current maintenance mode of the local context in an FT group. Applications can turn on maintenance mode when there is an inability to communicate with the peer, license mismatches, too many application errors, and so on. Possible states are: • MAINT_MODE_OFF—Maintenance mode is turned off. • MAINT_MODE_PARTIAL— All standby contexts transition to the FSM_FT_STATE_STANDBY_COLD state (see the “My State” field description). The ACE enters this mode if configuration synchronization fails. • MAINT_MODE_FULL—All contexts on the ACE become nonredundant causing their peer contexts to become active. The ACE enters this mode just before you reboot the appliance and is used primarily when you upgrade the ACE software.

Table 6-2 Field Descriptions for the show ft group Command Output **(continued)**

Field	Description
My State	State of the FT group member in the local ACE. Possible states are: • FSM_FT_STATE_INIT—Configuration for the FT group exists but the group is not in service. This is the initial state for each member (local and peer) of an FT group. • FSM_FT_STATE_ELECT—When you configure the **inservice** command for an FT group, the local group member enters this state. Through the election process, the local context negotiates with its peer context in the FT group to determine their states. One member enters the ACTIVE state and the other member enters the STANDBY_CONFIG state. • FSM_FT_STATE_ACTIVE—Local member of the FT group is active and processing flows. • FSM_FT_STATE_STANDBY_COLD—Either the FT VLAN is down, but the peer device 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. • FSM_FT_STATE_STANDBY_CONFIG—Local standby context is waiting to receive configuration information from its active peer context in the FT group. The active peer context receives a notification to send a snapshot of its running-configuration file to the local standby context. FSM_FT_STATE_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.
My State (Cont.)	• FSM_FT_STATE_STANDBY_HOT—Local standby context has all the state information it needs to statefully assume the active state if a switchover occurs. • FSM_FT_STATE_STANDBY_WARM—State used when upgrading or downgrading the ACE software. 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 CLI incompatibility. When the software versions are different while upgrading or downgrading, the STANDBY_WARM state allows the configuration and state synchronization process to continue on a best-effort basis, which 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 CLI commands or state information. This standby state allows the standby ACE to come up with best-effort support. In the STANDBY_WARM state, as with the STANDBY_HOT state, the configuration mode is disabled and configuration and state synchronization continues. A failover from the active to the standby based on priorities and preempt can still occur while the standby is in the STANDBY_WARM state.
My Config Priority	Priority configured on the FT group in the local ACE.
My Net Priority	Priority of the FT group equal to the configured priority minus the priority of the FT tracking failures if any.
My Preempt	Preemption value of the FT group in the local ACE. Possible values are Enabled or Disabled.
Peer State	State of the FT group in the remote ACE. For possible state values, see the “My State” field description.
Peer Config Priority	Priority configured for the FT group in the remote ACE.
Peer Net Priority	Priority of the FT group in the remote ACE computed from the configured priority and the priority of the FT tracking failures.
Peer Preempt	Preemption value of the FT group in the remote ACE. Possible values are Enabled or Disabled.

Table 6-2 Field Descriptions for the show ft group Command Output **(continued)**

Field	Description
Peer ID	FT peer identifier.
Last State Change Time	Time and date that the peer last changed from the active to standby state, or standby to active state.
Running Cfg Sync Enabled	Configured state of config sync for the running-config. Possible values are Enabled or Disabled.
Running Cfg Sync Status	Current status of config sync for the running-config. For example, Running configuration sync has completed or Config sync disabled when peer is not fully CLI compatible.
Startup Cfg Sync Enabled	Configured state of config sync for the startup-config. Possible states are Enabled or Disabled.
Startup Cfg Sync Status	Current status of config sync for the startup-config. For example, Startup configuration sync is disabled or Config sync disabled when peer is not fully CLI compatible.
Bulk Sync Done for ARP	Number of “bulk synchronization done” messages received on the standby ACE during state synchronization from the ARP module in the control plane.
Bulk Sync Done for LB	Number of “bulk synchronization done” messages received on the standby ACE during state synchronization from the load balancer (LB) module in the data plane.
Bulk Sync Done for ICM	Number of “bulk synchronization done” messages received on the standby ACE during state synchronization from the ICM input connection manager module in the data plane.

Displaying the Redundancy Internal Software History

To display the redundancy internal software history, perform the following task:

Command	Purpose
show ft history {**cfg_cntlr** | **ha_dp_mgr** | **ha_mgr**}	Displays the redundancy internal software history. The keywords are: • **cfg_cntlr**—Displays the configuration controller debug log • **ha_dp_mgr**—Displays the high availability (HA) dataplane manager debug log • **ha_mgr**—Displays the HA manager debug log

Displaying the IDMAP Table

This section describes how to display the IDMAP table. The IDMAP table contains a list of the local ACE to peer (standby) ACE ID mappings for each of the seven object types in the ACE. The local ID and the peer ID for each object type may or may not be the same, but the mappings (local ID to peer ID) should be the same on both the active ACE and the standby ACE. The ACE uses these mappings for configuration synchronization and state replication.

To display the IDMAP table, perform the following task:

Command	Purpose
show ft idmap	Displays the IDMAP table.

Table 6-3 lists the IDMAP table object types available in the ACE.

Table 6-3 ACE Object Types in the IDMAP Table

Object Type	Object Name
0	REAL ID
1	RSERVER ID
2	SERVERFARM ID
3	POLICY ID
4	STICKY GROUP ID
5	IF ID
6	CONTEXT ID

Displaying Memory Statistics

To display redundancy statistics per context, perform the following task:

Command	Purpose
show ft memory [**detail**]	Displays redundancy statistics per context. The optional **detail** keyword displays detailed HA manager memory statistics in the Admin context only.

Displaying Peer Information

To display peer information, perform the following task:

Command	Purpose
show ft peer *peer_id* {**detail** | **status** | **summary**}	Displays redundancy statistics per context. The keywords and arguments are: • *peer_id*—Unique identifier of the remote peer • **detail**—Displays detailed peer information • **status**—Displays the current operating status of the peer • **summary**—Displays summary peer information

Table 6-4 describes the fields in the **show ft peer** command output.

Table 6-4 Field Descriptions for the show ft peer Command Output

Field	Description
Peer ID	Identifier of the remote context in the FT group.
State	Current state of the peer. Possible states are: • FSM_PEER_STATE_INIT—Initial state of the peer after you configure it. • FSM_PEER_STATE_MY_IPADDR—Local ACE IP address is missing. Waiting for the local IP address to be configured. • FSM_PEER_STATE_PEER_IPADDR—Peer IP address is missing. Waiting for the peer IP address to be configured. • FSM_PEER_STATE_START_HB—Peer configuration is complete. Starting the heartbeat to see if there is a peer device.
State (continued)	• FSM_PEER_STATE_TL_SETUP—Heartbeat has detected the presence of the peer device. Redundancy is in the process of establishing a TCP connection to the peer. This connection carries configuration data, application state information, and redundancy protocol packets. • FSM_PEER_STATE_SRG_CHECK—Checking for software version compatibility with the peer device. • FSM_PEER_STATE_LIC_CHECK—Checking for license compatibility with the peer device. • FSM_PEER_STATE_COMPATIBLE—Version and license checks indicate that the peer is compatible for redundancy. • FSM_PEER_STATE_FT_VLAN_DOWN—FT VLAN is down, but, through the query interface, the local ACE has determined that the peer is still alive. • FSM_PEER_STATE_DOWN—Peer device is down. • FSM_PEER_STATE_ERROR—Status of whether an error has occurred with the peer. Possible errors are version mismatch, license mismatch, or failure to establish a TCP connection to the peer. A syslog message appears with more detailed information.
Maintenance Mode	Current maintenance mode of the peer context in an FT group. Applications can turn on maintenance mode when there is an inability to communicate with the peer, license mismatches, too many application errors, and so on. Possible states are: • MAINT_MODE_OFF—Maintenance mode is turned off. • MAINT_MODE_PARTIAL— All standby contexts transition to the STANDBY_COLD state. The ACE enters this mode if configuration synchronization fails. • MAINT_MODE_FULL—All contexts on the ACE become nonredundant causing their peer contexts to become active. The ACE enters this mode just before you reboot the appliance and is used primarily when you upgrade the ACE software.
FT VLAN	Identifier of the interface that is configured as the FT VLAN or Not Configured.
FT VLAN IF State	Current status of the FT VLAN interface. Possible states are UP or DOWN.
My IP Addr	IP address of the local ACE.
Peer IP Addr	IP address of the peer ACE.
Query VLAN	Identifier of the interface that is configured as the query VLAN or Not Configured.

Table 6-4 Field Descriptions for the show ft peer Command Output (continued)

Field	Description
Query VLAN IF State	Current status of the Query VLAN interface (if configured). Possible states are UP or DOWN.
Peer Query IP Addr	IP address of the query interface used to obtain the state of the peer's health when the FT VLAN is down.
Heartbeat interval	Time in seconds that the ACE waits between sending heartbeat packets.
Heartbeat Count	Number of missed heartbeats that an ACE must detect before declaring the peer down.
Tx Packets	Total number of packets that the local ACE sent to the peer.
Tx Bytes	Total number of bytes that the local ACE sent to the peer.
Rx Packets	Total number of packets that the local ACE received from the peer.
Rx Bytes	Total number of bytes that the local ACE received from the peer.
Rx Error Bytes	Total number of error bytes that the local ACE received from the peer.
Tx Keepalive Packets	Total number of keepalive packets that the local ACE sent to the peer.
Rx Keepalive Packets	Total number of keepalive packets that the local ACE received from the peer.
TL_CLOSE Count	Number of Transport Layer close events (TL_CLOSE) received on the redundant TCP connection from the TL driver.
FT_VLAN_ DOWN Count	Number of times that the FT VLAN was unavailable.
PEER_DOWN Count	Number of times that the remote ACE was unavailable.
SRG Compatibility	Status of whether the software version of the local ACE and the software version of the peer ACE are compatible. Possible states are the INIT, COMPATIBLE, or INCOMPATIBLE state.
License Compatibility	Status of whether the license of the local ACE and the license of the peer ACE are compatible. Possible states are the INIT, COMPATIBLE, or INCOMPATIBLE state.
FT Groups	Number of FT groups.

Displaying FT Statistics

To display peer information, perform the following task:

Command	Purpose
show ft stats *group_id*	Displays peer information. The *group_id* argument displays additional load-balancing statistics (LB statistics) for the specified group.

Table 6-5 describes the fields in the **show ft stats** command output.

Table 6-5 Field Descriptions for the show ft stats Command Output

Field	Description
HA Heartbeat Statistics	
Number of Heartbeats Sent	Total number of heartbeat packets sent by the local ACE.
Number of Heartbeats Received	Total number of heartbeat packets received by the local ACE.
Number of Heartbeats Missed	Total number of heartbeat intervals that transpired with no heartbeats received.
Number of Unidirectional HBs Received	Number of heartbeats (HBs) received by the local peer that indicate the remote peer is not receiving HBs. The remote peer is sending heartbeats, but not receiving any. **Note** Both peer appliances send heartbeat packets and each packet indicates whether the other peer has been receiving heartbeats.
Number of HB Timeout Mismatches	Number of times that the local peer received a heartbeat (HB) from the remote peer with a mismatched heartbeat interval. If the heartbeat intervals do not match, a peer adjusts its interval to the lower of the two intervals. **Note** The heartbeat interval should be the same on both peer appliances. Each heartbeat packet contains the configured interval in the packet. When a peer receives a heartbeat packet, it checks to see if the interval in the heartbeat packet matches the interval configured locally.
Num of Peer Up Events Sent	Number of times that the local ACE sent a Peer Up message to the remote ACE.
Num of Peer Down Events Sent	Number of times that the local ACE sent a Peer Down message to the remote ACE.
Successive HBs Miss Intervals Counter	Number of successive heartbeat misses detected by the heartbeat module.
Successive Uni HBs Recv Counter	Number of successive unidirectional heartbeats received by the heartbeat module.
LB Stats for FT Group *N*	
Send-side Stats	
Number of Sticky Entries Shared	Number of sticky database entries that the local ACE sent to the remote ACE.
Number of Replication Packets Sent	Number of packets that contain replication information that the local ACE sent to the remote ACE.
Number of Send Failures	Number of times that the local ACE attempted to send packets to the remote ACE but failed.
Receive-side Stats	
Number of Sticky Entries Dropped	Number of sticky database entries that the remote ACE sent to the local ACE, but the local ACE discarded them.
Number of Replication Packets Received	Number of packets that contain replication information that the local ACE received from the remote ACE.
Number of Receive Failures	Number of times that the remote ACE sent packets to the local ACE, but the local ACE failed to receive them.

Displaying FT Tracking Information

To display tracking information, perform the following task:

Command	Purpose
show ft track {**detail** | **status** | **summary**}	Displays tracking information. The keywords are: • **detail**—Displays detailed tracking information • **status**—Displays the current operating status of the peer plus additional information • **summary**—Displays summary peer information

Table 6-6 describes the fields in the **show ft track** command output.

Table 6-6 Field Descriptions for the show ft track Command Output

Field	Description
FT Group	FT group identifier.
Status	Configured state of the FT group. Possible states are the in-service or out-of-service state.
Maintenance Mode	Current maintenance mode of the local context in an FT group. Applications can turn on maintenance mode when there is an inability to communicate with the peer, license mismatches, too many application errors, and so on. Possible states are: • MAINT_MODE_OFF—Maintenance mode is turned off. • MAINT_MODE_PARTIAL— All standby contexts transition to the FSM_FT_STATE_STANDBY_COLD state (see the "My State" field description). The ACE enters this mode if configuration synchronization fails. • MAINT_MODE_FULL—All contexts on the ACE become nonredundant causing their peer contexts to become active. The ACE enters this mode just before you reboot the appliance and is used primarily when you upgrade the ACE software.

Table 6-6 Field Descriptions for the show ft track Command Output (continued)

Field	Description
My State	State of the FT group member in the local ACE. Possible states are: • FSM_FT_STATE_INIT—Initial state for each member (local and peer) of an FT group. The configuration for the FT group exists but the group is not yet in service. • FSM_FT_STATE_ELECT—State that the local group member enters when you configure the **inservice** command for an FT group. Through the election process, the local context negotiates with its peer context in the FT group to determine their states. One member enters the ACTIVE state and the other member enters the STANDBY_CONFIG state. • FSM_FT_STATE_ACTIVE—State that indicates that the local member of the FT group is active and processing flows. • FSM_FT_STATE_STANDBY_COLD—State that indicates if either the FT VLAN is down but the peer device 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. • FSM_FT_STATE_STANDBY_CONFIG—State that indicates that the local standby context is waiting to receive configuration information from its active peer context in the FT group. The active peer context receives a notification to send a snapshot of its running-configuration file to the local standby context. • FSM_FT_STATE_STANDBY_BULK—State that indicates that the 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. • FSM_FT_STATE_STANDBY_HOT—State that indicates that the local standby context has all the state information it needs to statefully assume the active state if a switchover occurs. • FSM_FT_STATE_STANDBY_WARM—State used when upgrading or downgrading the ACE software. 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 CLI incompatibility. When the software versions are different while upgrading or downgrading, the STANDBY_WARM state allows the configuration and state synchronization process to continue on a best-effort basis, which 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 CLI commands or state information. This standby state allows the standby ACE to come up with best-effort support. In the STANDBY_WARM state, as with the STANDBY_HOT state, the configuration mode is disabled and configuration and state synchronization continues. A failover from the active to the standby based on priorities and preempt can still occur while the standby is in the STANDBY_WARM state.
My Config Priority	Priority configured on the FT group in the local ACE.
My Net Priority	Priority of the FT group equal to the configured priority minus the priority of the FT tracking process failures, if any.
My Preempt	Preemption value of the FT group in the local ACE. Possible values are Enabled or Disabled.
Context Name	Name of the context that is associated with the FT group.
Context ID	Identifier of the context that is associated with the FT group.
Track Type	Type of object being tracked. Possible values are TRACK_HOST or TRACK_INTERFACE.

Table 6-6 Field Descriptions for the show ft track Command Output (continued)

Field	Description
State	State of the tracking process. Possible values are TRACK_UP or TRACK_DOWN.
Priority	Priority of the tracking process.
Transitions	Number of times that the active member of the FT group switched over to the standby member.
Probe Count	Number of probes associated with a TRACK_HOST process.
Probes Down	Number of failed probes.

Clearing Redundancy Statistics

To clear redundancy statistics, use the commands described in the following sections. You must enter all commands in this section in the Admin context unless otherwise indicated.

This section contains the following topics:

- Clearing Transport-Layer Statistics
- Clearing Heartbeat Statistics
- Clearing Tracking-Related Statistics
- Clearing All Redundancy Statistics
- Clearing the Redundancy History

Restrictions

If you configure redundancy on the ACE, then you must explicitly clear statistics on both the active and the standby ACEs. Clearing statistics on the active appliance only does not clear the statistics on the standby appliance.

Clearing Transport-Layer Statistics

To clear all transport layer-related counters that the ACE displays as part of the **show ft peer detail** command output, perform the following task:

Command	Purpose
clear ft ha-stats	Clears all transport layer-related counters that the ACE displays as part of the **show ft peer detail** command output. This command clears the following transport-layer counters: • Tx Packets • Tx Bytes • Rx Packets • Rx Bytes • Rx Error Bytes For an explanation of these fields, see the "Displaying Peer Information" section.

Clearing Heartbeat Statistics

To clear all heartbeat-related statistics, perform the following task:

Command	Purpose
clear ft hb-stats	Clears all heartbeat-related statistics. When you enter this command for the first time, the ACE sets the heartbeat statistics counters to zero and stores a copy of the latest statistics locally. From that point on, when you enter the **show ft hb-stats** command, the ACE displays the difference between the statistics that are stored locally and the current statistics.

Clearing Tracking-Related Statistics

To clear tracking-related statistics for the Admin FT group only, a user context FT group only, or for all FT groups that are configured in the ACE, perform the following task:

Command	Purpose
clear ft track-stats [**all**]	Clears tracking-related statistics for the Admin FT group only, a user context FT group only, or for all FT groups that are configured in the ACE. Use the optional **all** keyword in the Admin context only to clear tracking statistics for all FT groups that are configured in the ACE. If you enter this command in the Admin context without the **all** keyword, it clears the tracking statistics only for the FT group associated with the Admin context. In a user context, you cannot enter the **all** keyword, so you can clear the tracking statistics only for the FT group associated with the user context.

Clearing All Redundancy Statistics

To clear all redundancy statistics, including all TL, heartbeat, and tracking counters, perform the following task in the Admin context only:

Command	Purpose
clear ft all	Clears all redundancy statistics, including all TL, heartbeat, and tracking counters. This command does not affect the redundancy history. To clear the redundancy history, use the **clear ft history** command. For details, see the "Clearing the Redundancy History" section.

Clearing the Redundancy History

To clear the redundancy history, perform the following task in the Admin context only:

Command	Purpose
clear ft history {**cfg_cntlr** | **ha_dp_mgr** | **ha_mgr**}	The keywords are as follows: • **cfg_cntlr**—Clears the Configuration Controller debug log • **ha_dp_mgr**—Clears the HA (redundancy) dataplane manager debug log • **ha_mgr**—Clears the HA (redundancy) manager debug log

Configuration Example of Redundancy

This section shows an example redundancy configuration and illustrates a running-configuration that defines fault tolerance (FT) for a single ACE appliance operating in a redundancy configuration. You must configure a maximum of two ACE appliances (peers) for redundancy to fail over from the active appliance to the standby appliance.

Note All FT parameters are configured in the Admin context.

This configuration addresses the following redundancy components:

- A dedicated FT VLAN for communication between the members of an FT group. You must configure this same VLAN on both peer appliances.
- An FT peer definition.
- An FT group that is associated with the Admin context.
- A critical tracking and failure detection process for an interface.

The redundancy configuration appears in bold in the example.

```
hostname ACE_Appliance_1

interface gigabitEthernet 1/2
  speed 1000M
  duplex FULL
  ft-port vlan 200
  no shutdown

access-list ACL1 line 10 extended permit ip any any

class-map type management match-any L4_REMOTE-MGT_CLASS
  2 match protocol telnet any
  3 match protocol ssh any
  4 match protocol icmp any
  5 match protocol http any
  7 match protocol snmp any
  8 match protocol xml-https any

policy-map type management first-match L4_REMOTE-MGT_POLICY
  class L4_REMOTE-MGT_CLASS
    permit

interface vlan 100
  ip address 192.168.83.219 255.255.255.0
  peer ip address 192.168.83.230 255.255.255.0
  alias 192.168.83.200 255.255.255.0
  access-group input ACL1
```

```
  service-policy input L4_REMOTE-MGT_POLICY
  no shutdown

ft interface vlan 200
  ip address 192.168.1.1 255.255.255.0
  peer ip address 192.168.1.2 255.255.255.0
  no shutdown

ft peer 1
  ft-interface vlan 200
  heartbeat interval 300
  heartbeat count 10

ft group 1
  peer 1
  priority 200
  associate-context Admin
  inservice

ft track interface TRACK_VLAN100
  track-interface vlan 100
  peer track-interface vlan 200
  priority 50
  peer priority 5

ip route 0.0.0.0 0.0.0.0 192.168.83.1
```