



Configuring Redundant ACE Modules

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

- [Overview of Redundancy](#)
- [Redundancy Configuration Quick Start](#)
- [Configuring Redundancy](#)
- [Configuring Tracking and Failure Detection](#)
- [Example of a Redundancy Configuration](#)
- [Displaying Redundancy Information](#)
- [Clearing Redundancy Statistics](#)

Overview of Redundancy

Redundancy (or fault tolerance) uses a maximum of two ACEs in the same Catalyst 6500 series switch or in separate switches to ensure that your network remains operational even if one of the modules becomes unresponsive. Redundancy ensures that your network services and applications are always available.

**Note**

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.

Redundancy provides seamless switchover of flows in case an ACE becomes unresponsive or a critical host, interface, or HSRP group 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](#)
- [Stateful Failover](#)
- [FT VLAN](#)
- [Configuration Synchronization](#)
- [Configuration Requirements and Restrictions](#)

Redundancy Protocol

You can configure a maximum of two ACEs (peers) in the same Catalyst 6500 series switch or in different chassis for redundancy. Each peer module 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 Application Control Engine Module 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 does not require updating. The ACE selects a VMAC from a pool of virtual MACs available to it. You can specify the pool of MAC addresses that the local ACE and the peer ACE use by configuring the **shared-vlan-hostid** command and the **peer shared-vlan-hostid** command, respectively. To avoid MAC address conflicts, be sure that the two pools are different on the two ACEs. For more information about VMACs and MAC address pools, see the *Cisco Application Control Engine Module 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

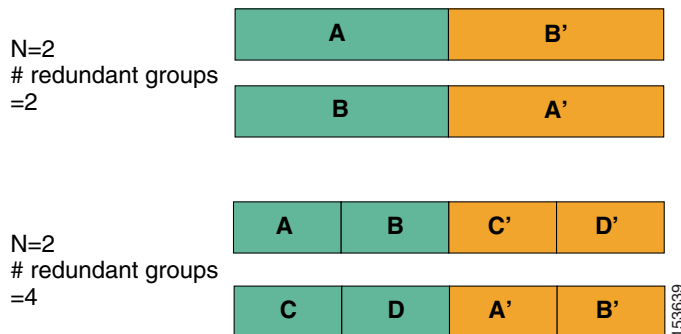



Figure 6-2 shows the uneven distribution of contexts between the two ACEs. As an example, it is possible that the FT groups A,B, C, and D use only half the resources that E and F require.

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 module and both modules 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 Application Control Engine Module 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. For details about configuring the heartbeat, 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. To disable preemption, use the **preempt** command. Entering this command causes 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.



Note

By default, connection replication is enabled in the ACE.

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



Note

The ACE does not replicate SSL and other terminated (proxied) connections from the active context to the standby context.

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

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

**Note**

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

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 must configure this same VLAN on both peer modules. You also must configure a different IP address within the same subnet on each module for the FT VLAN.

**Note**

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

The two redundant modules constantly communicate over the FT VLAN to determine the operating status of each module. 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.

**Note**

The IP address and the MAC address of the FT VLAN do not change at switchover.

Configuration Synchronization

For redundancy to function properly, both members of an FT group must have identical configurations. Ensure that both ACE modules include the same bandwidth software license (4 Gbps, 8 Gbps, or 16 Gbps) and the same virtual context software license. If there is a mismatch in a software license between the two ACE modules 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 modules have the same virtual content 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, for example, an 8-Gbps ACE module to a 4-Gbps ACE module.

For details about the available ACE software licenses, see [Chapter 3, Managing ACE Software Licenses](#).

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.

For information about configuring config sync, see the [“Synchronizing Redundant Configurations”](#) section.

Configuration Requirements and Restrictions

Follow these requirements and restrictions when configuring the redundancy feature.

- 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.
- 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 Application Control Engine Module Routing and Bridging Configuration Guide*.

Redundancy Configuration Quick Start

[Table 6-1](#) provides a quick overview of the steps required to configure redundancy for *each* ACE in the redundancy configuration. Each step includes the CLI command or a reference to the procedure required to complete the task. For a complete description of each feature and all the options associated with the CLI commands, see the sections following [Table 6-1](#).

Table 6-1 Redundancy Configuration Quick Start**Task and Command Example**

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 Application Control Engine Module Virtualization Configuration Guide*.

2. Enter configuration mode.

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

3. 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 60
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
```

4. Configure the local redundancy peer module, associate the FT VLAN with the peer, and configure the heartbeat interval and count.

```
host1/Admin(config)# ft peer 1
host1/Admin(config-ft-peer)# ft-interface vlan 60
host1/Admin(config-ft-peer)# heartbeat count 20
host1/Admin(config-ft-peer)# heartbeat interval 300
host1/Admin(config-ft-intf)# exit
```

5. Create at least one FT group on each ACE.

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

6. 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
```

Table 6-1 Redundancy Configuration Quick Start (continued)

Task and Command Example

7. Associate the peer context with the FT group.

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

8. (Optional) Configure the priority of the FT group on the local module.

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

9. (Optional) Configure the priority of the FT group on the peer module.

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

10. Place the FT group in service.

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

11. (Optional) Configure one or more critical objects (gateways or hosts, interfaces, or HSRP groups) 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
```

12. (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
```

13. (Optional) Save your configuration changes to Flash memory.

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

14. (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

To configure redundancy on the ACE, use the commands in the following sections. You must configure the **ft interface**, **ft peer**, and **ft group** commands on all ACEs that participate in the redundancy configuration. This section 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](#)

Configuring 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 module.

**Note**

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

This section contains the following topics:

- [Creating an FT VLAN](#)
- [Configuring an FT VLAN IP Address](#)
- [Configuring the Peer IP Address](#)
- [Enabling the FT VLAN](#)

Creating an FT VLAN

To create an FT VLAN, use the **ft interface** command in configuration mode. The syntax of this command is as follows:

```
ft interface vlan vlan_id
```

The *vlan_id* argument specifies a unique identifier for the FT VLAN. Enter an integer from 2 to 4094.

For example, enter:

```
host1/Admin(config)# ft interface vlan 200  
host1/Admin(config-ft-intf)#
```

**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. See the [“Associating the FT VLAN with the Local Peer”](#) section.

To remove the FT VLAN from the redundancy configuration, enter:

```
host1/Admin(config)# no ft interface vlan 200
```

Configuring an FT VLAN IP Address

After you create the FT VLAN, you must assign an IP address to the VLAN. To assign an IP address to the VLAN, use the **ip** command in FT interface configuration mode. The syntax of this command is as follows:

```
ip address ip_address netmask
```

The keyword and arguments of this command are:

- **address** *ip_address*—Specifies the IP address of the FT VLAN. Enter an IP address in dotted-decimal notation (for example, 192.168.12.1).
- *netmask*—Subnet mask of the FT VLAN. Enter a subnet mask in dotted-decimal notation (for example, 255.255.255.0).

For example, to configure an IP address for the FT VLAN, enter:

```
host1/Admin(config-ft-intf)# ip address 192.168.12.1 255.255.255.0
```

To remove the IP address from an FT VLAN, enter:

```
host1/Admin(config-ft-intf)# no ip address
```

Configuring the Peer IP Address

The local member of the FT group communicates with the remote peer over the FT VLAN. To allow the local member to communicate with the remote peer, use the **peer ip address** command in FT interface configuration mode. The syntax of this command is as follows:

```
peer ip address ip_address netmask
```

The keyword and arguments of this command are:

- **address** *ip_address*—Specifies the IP address of the remote peer. Enter an IP address in dotted-decimal notation (for example, 192.168.12.15).
- *netmask*—Subnet mask of the remote peer. Enter a subnet mask in dotted-decimal notation (for example, 255.255.255.0).

For example, to configure an IP address on the remote peer, enter:

```
host1/Admin(config-ft-intf)# peer ip address 192.168.12.15  
255.255.255.0
```

To remove an IP address from the remote peer, enter:

```
host1/Admin(config-ft-intf)# no peer ip address 192.168.12.15  
255.255.255.0
```

Enabling the FT VLAN

To enable the FT VLAN, use the **no shutdown** command in FT interface configuration mode. The syntax of this command is as follows:

```
no shutdown
```

For example, to enable the FT VLAN, enter:

```
host1/Admin(config-ft-intf)# no shutdown
```

To disable the FT VLAN after you have enabled it, enter:

```
host1/Admin(config-ft-intf)# shutdown
```

Configuring 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 modules. The alias IP address serves as a shared gateway for the two ACE modules.

To configure an alias IP address, use the **alias** command in interface configuration mode. The syntax of this command is as follows:

```
alias ip_address netmask
```

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 (for example, 192.168.1.1 255.255.255.0).

For example, to configure an alias IP address, enter:

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

To remove an alias IP address, enter:

```
host1/Admin(config-if)# no alias 192.168.1.1 255.255.255.0
```

Configuring an FT Peer

On both peer ACEs, configure an FT peer definition in the Admin context only. You can configure a maximum of two ACEs as redundancy peers.

To create an FT peer, use the **ft peer** command in configuration mode. The syntax of this command is as follows:

```
ft peer peer_id
```

The *peer_id* argument specifies a unique identifier for the peer. You can only enter 1.

For example, enter:

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



Note

Before you can remove an FT peer from the configuration, remove the peer from the FT group. See the [“Associating a Peer with an FT Group”](#) section.

To remove the FT peer from the configuration, enter:

```
host1/Admin(config)# no ft peer 1
```

After you create an FT peer, configure the peer attributes as described in the following topics:

- [Associating the FT VLAN with the Local Peer](#)
- [Configuring the Heartbeat Interval and Count](#)
- [Configuring a Query Interface](#)

Associating the FT VLAN with the Local Peer

After you create an FT peer, associate the existing FT VLAN with that local peer so that it can communicate with the remote peer. The redundancy peers use this dedicated FT VLAN to exchange heartbeat packets and flow-state information. For information about configuring an FT VLAN, see the “[Configuring an FT VLAN](#)” section.

To associate an FT VLAN with a peer, use the **ft-interface vlan** command in FT peer configuration mode. The syntax of this command is as follows:

```
ft-interface vlan vlan_id
```

The *vlan_id* argument specifies the identifier of an existing VLAN. Enter an integer from 2 to 4094.

For example, enter:

```
host1/Admin(config-ft-peer)# ft-interface vlan 200
```

To remove the FT VLAN from the peer configuration, enter:

```
host1/Admin(config-ft-peer)# no ft-interface vlan 200
```

Configuring the Heartbeat Interval and Count

The heartbeat interval determines the frequency in milliseconds (ms) at which the active member of the FT group sends the heartbeat packets to the standby member. The heartbeat count is the number of missed heartbeats that the standby member must detect before determining that the active member is not available. To configure the heartbeat interval and count, use the **heartbeat** command in peer configuration mode. The syntax of this command is as follows:

```
heartbeat {count number | interval frequency}
```

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.

For example, to set the heartbeat count to 20, enter:

```
host1/Admin(config-ft-peer)# heartbeat count 20
```

To reset the heartbeat count to the default of 10, enter:

```
host1/Admin(config-ft-peer)# no heartbeat count
```

For example, to set the heartbeat interval to 500 ms, enter:

```
host1/Admin(config-ft-peer)# heartbeat interval 500
```

To reset the heartbeat interval to the default of 100 ms, enter:

```
host1/Admin(config-ft-peer)# no heartbeat interval
```

Configuring a Query Interface

Configure 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.

To configure a query interface, use the **query-interface vlan** command in FT peer configuration mode. The syntax of this command is as follows:

```
query-interface vlan vlan_id
```

The *vlan_id* argument specifies the identifier of an existing VLAN. Enter an integer from 2 to 4094.

For example, to configure a query interface, enter:

```
host1/Admin(config-ft-peer) # query-interface vlan 400
```

To remove the query interface from the peer configuration, enter:

```
host1/Admin(config-ft-peer) # no query-interface vlan 400
```



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.

Configuring an FT Group

On each ACE, you can create multiple FT groups, up to a maximum of 251 groups (250 user contexts and 1 Admin context). Each group consists of a maximum of two members (contexts): one active context on one module and one standby context on the peer module.

To create an FT group, use the **ft group** command in configuration mode. You must configure the same group ID on both peer modules. The syntax of this command is as follows:

```
ft group group_id
```

The *group_id* argument specifies a unique identifier of the group. Enter an integer from 1 to 255.

For example, enter:

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

To remove the group from the configuration, enter:

```
host1/Admin(config) # no ft group 1
```

After you create an FT group, configure the FT group attributes as described in the following topics:

- [Associating a Context with an FT Group](#)
- [Associating a Peer with an FT Group](#)
- [Assigning a Priority to the Active FT Group Member](#)
- [Assigning a Priority to the Standby FT Group Member](#)
- [Configuring Preemption](#)
- [Placing an FT Group in Service](#)
- [Modifying an FT Group](#)

Associating a Context with an FT Group

An FT group consists of two members (contexts) with the same name, each residing on a different ACE. To associate a context with an FT group, use the **associate-context** command in FT group configuration mode. You need to make this association for both redundant contexts in an FT group. The syntax of this command is as follows:

```
associate-context name
```

For the *name* argument, enter the unique identifier of the context that you want to associate with the FT group.

For example, enter:

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



Note

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. See the [“Placing an FT Group in Service”](#) section.

To remove a context from an FT group, enter:

```
host1/Admin(config-ft-group)# no associate-context C1
```

Associating a Peer with an FT Group

To associate a peer ACE with an FT group, use the **peer** command in FT group configuration mode. The syntax of this command is as follows:

```
peer peer_id
```

For the *peer_id* argument, enter 1 as the identifier of an existing peer module. You can only enter 1.

For example, enter:

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

To remove the peer association with the FT group, enter:

```
host1/Admin(config-ft-group)# no peer
```

Assigning a Priority to the Active FT Group Member

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. To ensure that the member with the higher priority always becomes the active member, use the **preempt** command, which is enabled by default. For details, see the “[Configuring Preemption](#)” section.

To configure the priority of an FT group on the active member, use the **priority** command in FT group configuration mode. You must configure the priority of an FT group on both modules. Configure a higher priority for the group on the ACE where you want the active member to initially reside. The syntax of this command is as follows:

```
priority number
```

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.

**Tip**

Configure a higher priority on the FT group member that you want to be the active member.

For example, to configure the priority of the FT group on the active member, enter:

```
host1/Admin(config-ft-group)# priority 150
```

To restore the default priority of 100, enter:

```
host1/Admin(config-ft-group)# no priority
```

Assigning a Priority to the Standby FT Group Member

To configure the priority of an FT group on the remote standby member, use the **peer priority** command in FT group configuration mode. You must configure the priority of an FT group on both redundant modules. Configure a lower priority for the FT group on the ACE where you want the standby member to initially reside. The syntax of this command is as follows:

```
peer priority number
```

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.



Tip

Configure a lower priority on the FT group member that you want to be the standby member.



Note

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.

For example, to configure the priority of the FT group member on the remote standby member, enter:

```
host1/Admin(config-ft-group)# peer priority 50
```

To restore the default priority of 100, enter:

```
host1/Admin(config-ft-group)# no peer priority
```

Configuring Preemption

Preemption ensures that the group member with the higher priority always asserts itself and becomes the active member. By default, preemption is enabled. To configure preemption after it has been disabled, use the **preempt** command in FT group configuration mode. The syntax of this command is as follows:

preempt

For example, enter:

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

To disable preemption, enter:

```
host1/Admin(config-ft-group)# no preempt
```



Note

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.

Placing an FT Group in Service



Note

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.

To place an FT group in service, use the **inservice** command in FT group configuration mode. The syntax of this command is as follows:

inservice

For example, to place an FT group in service, enter:

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

To take the FT group out of service, enter:

```
host1/Admin(config-ft-group)# no inservice
```

Modifying an FT Group

If you need to modify an FT group, perform the following steps in FT group configuration mode:

-
- 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.
-



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

Specifying the Peer Hostname

To specify the hostname of a peer ACE, use the **peer hostname** command in configuration mode in the Admin context. For details about this command, see [Chapter 1, Setting Up the ACE](#).

Specifying the MAC Address Banks for a Shared VLAN

To specify the MAC address banks to be used by the local ACE and the peer ACE with a shared VLAN (FT VLAN), use the **shared-vlan-hostid** command and the **peer shared-vlan-hostid** command, respectively, in configuration mode in the Admin context. You configure these commands to prevent MAC address conflicts between the two peer ACEs. Be sure to select a bank of MAC addresses for the peer that is different from that used by the local ACE. For details about this command, see the *Cisco Application Control Engine Module Routing and Bridging Configuration Guide*.

Forcing a Failover

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

To cause a switchover, use the **ft switchover** command in Exec mode. To use this command, you must disable preemption by using the **no preempt** command. For information on the **preempt** command, see the “[Configuring Preemption](#)” section.

The syntax of the **ft switchover** command is:

```
ft switchover [all [force] | force | [group_id [force]]]
```

The keywords, arguments, and options are:

- **all**—(Optional) Causes a switchover of all FT groups configured in the ACE simultaneously.
- **force**—(Optional) Causes a switchover while ignoring the state of the standby member. Use this option only when the FT VLAN is down.
- *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.

The **ft switchover** command exhibits the following behavior, depending on whether you enter the command 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.

For example, to cause a failover from the active module to the standby module of FT group1, enter:

```
host1/Admin# ft switchover 1
This command will cause card to switchover (yes/no)? [no] yes
host1/Admin#
```

Synchronizing Redundant Configurations

The configurations on both the active context and the standby context must be identical. 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 config sync—Synchronizes the configuration applied to the active context to the standby context if the peer is already up

To enable automatic synchronization of the running-configuration and the startup-configuration files after they have been explicitly disabled, use the **ft auto-sync** command in configuration mode.



Note

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, if 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 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. (For more information about FT states, see [Table 6-2](#).) 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.

**Note**

If a license mismatch occurs between the two ACEs in a redundant configuration, the **auto-sync** command is automatically disabled and a syslog message is generated.

The syntax of this command is as follows:

```
ft auto-sync {running-config | startup-config}
```

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.

**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.

**Note**

If the config sync 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. For more information about FT states, see [Table 6-2](#).

**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 Application Control Engine Module 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**

For example, to enable autosynchronization of the running-configuration file in the C1 context after it has been disabled, enter:

```
host1/C1(config)# ft auto-sync running-config
```

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.

This section contains the following topics:

- [Overview of Tracking and Failure Detection](#)
- [Configuring Tracking and Failure Detection for a Host or Gateway](#)
- [Configuring Tracking and Failure Detection for an Interface](#)
- [Creating a Tracking and Failure Detection Process for an HSRP Group](#)

Overview of Tracking and Failure Detection

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
- Hot Standby Router Protocol (HSRP) groups

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 before you configure 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 the 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.

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. It contains the following topics:

- [Creating a Tracking and Failure Detection Process for a Host or Gateway](#)
- [Configuring the Gateway or Host IP Address Tracked by the Active Member](#)
- [Configuring a Probe on the Active Member for Host Tracking](#)
- [Configuring a Priority on the Active Member for Multiple Probes](#)
- [Configuring the Gateway or Host IP Address Tracked by the Standby Member](#)
- [Configuring a Probe on the Standby Member for Host Tracking](#)
- [Configuring a Priority on the Standby Member for Multiple Probes](#)
- [Example of a Tracking Configuration for a Gateway](#)

Creating a Tracking and Failure Detection Process for a Host or Gateway

To create a tracking and failure detection process for a gateway or host, use the **ft track host** command in configuration mode. The syntax of this command is as follows:

```
ft track host name
```

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.

For example, to create a tracking process for a gateway, enter:

```
host1/Admin(config)# ft track host TRACK_GATEWAY1  
host1/Admin(config-ft-track-host)#
```

To remove the gateway-tracking process, enter:

```
host1/Admin(config)# no ft track host TRACK_GATEWAY1
```

Configuring the Gateway or Host IP Address Tracked by the Active Member

To allow the active member to track a gateway or host, you need to configure the IP address of the gateway or host. To configure the IP address, use the **track-host** command in FT track host configuration mode. The syntax of this command is as follows:

```
track-host ip_address
```

The *ip_address* argument specifies the IP address of the gateway or host that you want the active FT group member to track. Enter the IP address in dotted-decimal notation (for example, 192.168.12.101).

For example, to track the gateway located at 192.168.12.101, enter:

```
host1/Admin(config-ft-track-host)# track-host 192.168.12.101
```

Configuring a Probe on the Active Member for Host Tracking

Configure one or more probes on the active FT group member to track the health of the gateway or host. For details about creating probes, see the *Cisco Application Control Engine Module Server Load-Balancing Configuration Guide*. To associate an existing probe with a gateway or host for tracking by the active member, use the **probe** command in FT track host configuration mode. The syntax of this command is as follows:

```
probe name priority number
```

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.

**Note**

If you remove a probe from the active FT group member configuration and you have not configured a tracking priority for the FT group (see the “[Configuring a Priority on the Active Member for Multiple Probes](#)” section), 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.

For example, enter:

```
host1/Admin(config-ft-track-host)# probe TCP_PROBE1 priority 50
```

To remove the tracking probe from the active member, enter:

```
host1/Admin(config-ft-track-host)# no probe TCP_PROBE1
```

Configuring a Priority on the Active Member for Multiple Probes

You can assign a tracking priority that the active member uses when multiple tracking probes are defined. To assign a priority for multiple probes on the active member, use the **priority** command in FT track host configuration mode. The syntax of this command is as follows:

priority *number*

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.

For example, enter:

```
host1/Admin(config-ft-track-host)# priority 50
```

To reset the priority to the default value of 0, enter:

```
host1/Admin(config-ft-track-host)# no priority 50
```

Configuring the Gateway or Host IP Address Tracked by the Standby Member

To allow the standby member to track a gateway or host, you need to configure the IP address of the gateway or host. To configure the IP address, use the **peer track-host** command in FT track host configuration mode. The syntax of this command is as follows:

```
peer track-host ip_address
```

The *ip_address* argument specifies the IP address of the gateway or host that you want the standby FT group member to track. Enter the IP address in dotted-decimal notation (for example, 172.16.27.1).

For example, to track the gateway located at 172.16.27.1, enter:

```
host1/Admin(config-ft-track-host) # peer track-host 172.16.27.1
```

To remove the host tracked by the standby member, enter:

```
host1/Admin(config-ft-track-host) # no peer track-host 172.16.27.1
```

Configuring a Probe on the Standby Member for Host Tracking

Configure one or more probes on the standby member to track the health of the gateway or host. For details about creating probes, see the *Cisco Application Control Engine Module Server Load-Balancing Configuration Guide*. To associate an existing probe with a gateway or host for tracking by the standby member, use the **peer probe** command in FT track host configuration mode. The syntax of this command is as follows:

```
peer probe name priority number
```

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.

For example, enter:

```
host1/Admin(config-ft-track-host)# peer probe TCP_PROBE1 priority 25
```

To remove the tracking probe from the standby member, enter:

```
host1/Admin(config-ft-track-host)# no peer probe TCP_PROBE1
```

Configuring a Priority on the Standby Member for Multiple Probes

You can configure a tracking priority that the standby member of an FT group uses when multiple tracking probes are defined. To assign a priority for multiple probes on the standby member, use the **peer priority** command in FT track host configuration mode. The syntax of this command is as follows:

```
peer priority number
```

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.

For example, enter:

```
host1/Admin(config-ft-track-host)# peer priority 25
```

To reset the multiple-probe priority to the default value of 0 on the standby member, enter:

```
host1/Admin(config-ft-track-host)# no peer priority 25
```

Example of a Tracking Configuration for a Gateway

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.

To configure tracking on the standby member, use the **peer** commands described in the “[Configuring a Probe on the Standby Member for Host Tracking](#)” and the “[Configuring a Priority on the Standby Member for Multiple Probes](#)” sections.

Configuring Tracking and Failure Detection for an Interface

This section describes the commands that you enter to configure tracking and failure detection for an interface. It contains the following topics:

- [Creating a Tracking and Failure Detection Process for an Interface](#)
- [Configuring the Interface Tracked by the Active Member](#)
- [Configuring a Priority for a Tracked Interface on the Active Member](#)
- [Configuring the Interface Tracked by the Standby Member](#)
- [Configuring a Priority for a Tracked Interface on the Standby Member](#)
- [Example of a Tracking Configuration for an Interface](#)

Creating a Tracking and Failure Detection Process for an Interface

To create a tracking and failure detection process for an interface, use the **ft track interface** command in configuration mode. The syntax of this command is as follows:

```
ft track interface name
```

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.

**Note**

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

For example, enter:

```
host1/Admin(config)# ft track interface TRACK_VLAN100
```

To remove the interface-tracking process, enter:

```
host1/Admin(config)# no ft track interface TRACK_VLAN100
```

Configuring the Interface Tracked by the Active Member

To configure the interface that you want the active member to track, use the **track-interface vlan** command in FT track interface configuration mode. The syntax of this command is as follows:

```
track-interface vlan vlan_id
```

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.

For example, to track the critical interface VLAN 100, enter:

```
host1/Admin(config-ft-track-intf)# track-interface vlan 100
```

To remove VLAN 100 from the tracking process, enter:

```
host1/Admin(config-ft-track-intf)# no track-interface vlan 100
```

Configuring a Priority for a Tracked Interface on the Active Member

To assign a priority to the interface that the active member is tracking, use the **priority** command in FT track interface configuration mode. The syntax of this command is as follows:

```
priority number
```

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.

For example, enter:

```
host1/Admin(config-ft-track-intf)# priority 50
```

To reset the interface priority on the active member to the default value of 0, enter:

```
host1/Admin(config-ft-track-intf)# no priority 50
```

Configuring the Interface Tracked by the Standby Member

To configure the interface that you want the standby member to track, use the **peer track-interface vlan** command in FT track interface configuration mode. The syntax of this command is as follows:

```
peer track-interface vlan vlan_id
```

For the *vlan_id* argument, enter the VLAN ID of an existing VLAN configured on the standby member as an integer from 2 to 4094.

For example, to track the critical interface VLAN 200, enter:

```
host1/Admin(config-ft-track-intf)# peer track-interface vlan 200
```

To remove VLAN 200 from the tracking process, enter:

```
host1/Admin(config-ft-track-intf)# no peer track-interface vlan 200
```

Configuring a Priority for a Tracked Interface on the Standby Member

To assign a priority to the tracked interface that the standby member is tracking, use the **peer priority** command in FT track interface configuration mode. The syntax of this command is as follows:

```
peer priority number
```

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.

If the tracked interface goes down, the ACE decrements the priority of the FT group on the standby 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.

For example, enter:

```
host1/Admin(config-ft-track-intf)# peer priority 25
```

To reset the interface priority on the standby member to the default value of 0, enter:

```
host1/Admin(config-ft-track-intf)# no peer priority 25
```

Example of a Tracking Configuration for an Interface

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

```
ft track interface TRACK_VLAN100
  track-interface vlan 100
  priority 50
```

In the above configuration example, if VLAN 100 goes 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.

To configure tracking on the standby member, use the **peer** commands described in the [“Configuring the Interface Tracked by the Standby Member”](#) and the [“Configuring a Priority for a Tracked Interface on the Standby Member”](#) sections.

Configuring Tracking and Failure Detection for an HSRP Group

This section describes the commands required to configure a tracking and failure detection process for a Hot Standby Router Protocol (HSRP) group that you have previously configured on the Catalyst 6500 supervisor engine. It contains the following topics:

- [Before You Begin](#)
- [Creating a Tracking and Failure Detection Process for an HSRP Group](#)
- [Configuring the HSRP Group to Track on the Active Member](#)
- [Configuring a Priority for the HSRP Group Tracked by the Active Member](#)
- [Configuring the HSRP Group to Track on the Standby Member](#)
- [Configuring a Priority for a Tracked HSRP Group on the Standby Member](#)
- [Example of a Tracking Configuration for an HSRP Group](#)

Before You Begin

**Note**

For best results, observe the following configurational requirements before you attempt to configure HSRP tracking and failure detection on the ACE.

Before you configure an HSRP tracking and failure detection process on the ACE, you must configure the HSRP group on the supervisor engine. For example, if the HSRP group (including the name) is configured on the supervisor engine and it is not in the Active or the Standby state, you will see the following output when you enter the **show ft track detail** command on the ACE:

```
Track type           : TRACK_HSRP
HSRP Group Name     : test
State                : TRACK_DOWN (HSRP Group does not exist
                      on the Supervisor or it is in the INIT
                      State)
Priority              : 20
Transitions          : 1
```

For example, if the HSRP group is in the Standby state, you will see the following output when you enter the **show ft track detail** command on the ACE:

```
Track type           : TRACK_HSRP
HSRP Group Name     : test
State                : TRACK_DOWN (HSRP Group is Standby on
                       the Supervisor)
Priority              : 20
Transitions          : 1
```

For example, if the HSRP group is in the Active state, you will see the following output when you enter the **show ft track detail** command on the ACE:

```
Track type           : TRACK_HSRP
HSRP Group Name     : test
State                : TRACK_UP
Priority              : 20
Transitions          : 2
```

If the HSRP group (including the name) is configured on the supervisor engine after the HSRP tracking process is initially configured on the ACE, you may or may not obtain the expected results when you enter the **show ft track detail** command on the ACE.

If the HSRP group name is changed on the supervisor engine after the HSRP tracking process is configured on the ACE, further state notifications will not be sent to the ACE. You must delete the HSRP tracking process on the ACE after the HSRP group name is changed on the supervisor engine.

Creating a Tracking and Failure Detection Process for an HSRP Group

To create a tracking and failure detection process for an HSRP group, use the **ft track hsrp** command in configuration mode. The syntax of this command is as follows:

```
ft track hsrp tracking_process_name
```

For the *tracking_process_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. The ACE allows you to track up to 250 HSRP groups.

For example, enter:

```
host1/Admin(config)# ft track hsrp HSRP_TRACK_PROCESS1
host1/Admin(config-ft-track-hsrp)#
```

To remove the HSRP group-tracking process, enter:

```
host1/Admin(config)# no ft track hsrp HSRP_TRACK_PROCESS1
```

Configuring the HSRP Group to Track on the Active Member

To track an HSRP group on the active member of an FT group, use the **track-hsrp** command in FT track-HSRP configuration mode. The syntax of this command is as follows:

track-hsrp *name*

For the *name* argument, enter the identifier of an HSRP group previously configured on the Catalyst supervisor that you want to track on the active member. Enter the name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters. The ACE allows you to track up to 250 HSRP groups.

To obtain the correct HSRP group identifier to use for tracking on the ACE, enter the **show standby vlan** command on the Catalyst 6500 series switch or 7600 series router.

For example, enter the following command:

```
sh-ace-6k-1# show standby vlan 120
Vlan120 - Group 120
  Local state is Active, priority 200, may preempt
  Hellotime 3 sec, holdtime 10 sec
  Next hello sent in 2.022
  Virtual IP address is 192.168.120.254 configured
  Active router is local
  Standby router is 192.168.120.252 expires in 8.360
  Virtual mac address is 0000.0c07.ac78
  7 state changes, last state change 21:54:53
  IP redundancy name is "hsrp-v1120-120" (default)
  Priority tracking 1 interface or object, 1 up:
    Interface or object      Decrement State
    GigabitEthernet4/35      110       Up
```

Use the IP redundancy name (shown in bold in the above output example) as the HSRP group name. The switch or router automatically assigns this name to the HSRP group.

For example, enter:

```
host1/Admin(config-ft-track-hsrp)# track-hsrp hsrp-v1120-120
```

To remove the HSRP group from the tracking process, enter:

```
host1/Admin(config-ft-track-hsrp)# no track-hsrp
```

Configuring a Priority for the HSRP Group Tracked by the Active Member

To assign a priority to the HSRP group that you are tracking on the active member of an FT group, use the **priority** command in FT track HSRP configuration mode. The syntax of this command is as follows:

priority *number*

For the number argument, enter the priority of the HSRP group as an integer from 0 to 255. The default is 0. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the HSRP group that you are tracking. If the HSRP group goes down, the ACE decrements the priority of the FT group on the active member 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.



Note

When you configure HSRP tracking on the FT group member and the HSRP group does not exist on the supervisor engine, the ACE marks the tracking process as TRACK_DOWN and automatically decrements the net priority of the FT group by the tracking priority value.

For example, enter:

```
host1/Admin(config-ft-track-hsrp)# priority 50
```

To reset the priority to the default value of 0, enter:

```
host1/Admin(config-ft-track-hsrp)# no priority 50
```

Configuring the HSRP Group to Track on the Standby Member

To track an HSRP group on the standby member of an FT group, use the **peer track-hsrp** command in FT track HSRP configuration mode. The syntax of this command is as follows:

peer track-hsrp *name*

For the *name* argument, enter the identifier of an HSRP group previously configured on the supervisor engine that you want to track on the standby member of an FT group. Enter the name as an unquoted text string with no spaces and a maximum of 64 alphanumeric characters.

For example, enter:

```
host1/Admin(config-ft-track-hsrp)# peer track-hsrp HSRP_GRP1
```

To remove the HSRP group from the tracking process, enter:

```
host1/Admin(config-ft-track-hsrp)# no peer track-hsrp HSRP_GRP1
```

Configuring a Priority for a Tracked HSRP Group on the Standby Member

To assign a priority to the HSRP group that you are tracking on the standby member of an FT group, use the **peer priority** command in FT track HSRP configuration mode. The syntax of this command is as follows:

peer priority *number*

For the *number* argument, enter the priority of the HSRP group as an integer from 0 to 255. The default is 0. Higher values indicate higher priorities. Assign a priority value based on the relative importance of the HSRP group that you are tracking. If the HSRP group goes down, the ACE decrements the priority of the FT group on the standby member by the value of the *number* argument.

For example, enter:

```
host1/Admin(config-ft-track-hsrp)# peer priority 25
```

To reset the priority to the default value of 0, enter:

```
host1/Admin(config-ft-track-hsrp)# no peer priority 25
```

Example of a Tracking Configuration for an HSRP Group

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

```
ft track hsrp TRACK_HSRP_GRP1
 track-hsrp HSRP_GRP1
 priority 50
```

In the configuration example, if the HSRP_GRP1 group goes 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.

To configure tracking on the standby member, use the **peer** commands described in the “[Configuring the HSRP Group to Track on the Standby Member](#)” and the “[Configuring a Priority for a Tracked HSRP Group on the Standby Member](#)” sections.

Example of a Redundancy Configuration

The following example illustrates a running-configuration that defines fault tolerance (FT) for a single ACE module operating in a redundancy configuration. You must configure a maximum of two ACE modules (peers) for redundancy to fail over from the active module to the standby module.



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 modules.
- 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_Module_1

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 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
```

Displaying Redundancy Information

This section describes how you can use the **show** commands to display configuration information and statistics for your redundancy configuration and contains the following topics:

- [Displaying Redundancy Configurations](#)
- [Displaying FT Group Information](#)
- [Displaying the Redundancy Internal Software History](#)
- [Displaying the IDMAP Table](#)
- [Displaying Memory Statistics](#)
- [Displaying Peer Information](#)
- [Displaying FT Statistics](#)
- [Displaying FT Tracking Information](#)

Displaying Redundancy Configurations

To display redundancy configurations, use the **show running-config ft** command in Exec mode. The syntax of this command is as follows:

```
show running-config ft
```

For example, enter:

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

Displaying FT Group Information

To display redundancy statistics per context, use the **show ft group** command in Exec mode. The syntax of this command is as follows:

```
show ft group {brief | {[group_id]{detail | status | summary}}}
```

The keywords, arguments, and options are:

- **brief**—Displays the group ID, local state, peer state, context name, and context ID of all the FT groups that are configured in the ACE.
- **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.

For example, enter:

```
host1/Admin# show ft group 1 detail
```

[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.

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

Field	Description
Maintenance Mode	<p>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:</p> <ul style="list-style-type: none"> • 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 module 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	<p>State of the FT group member in the local ACE. Possible states are:</p> <ul style="list-style-type: none"> • 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_SELECT—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.

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

Field	Description
My State (Cont.)	<ul style="list-style-type: none"> 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.
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.
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.
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.

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

Field	Description
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, use the **show ft history** command in Exec mode. The syntax of this command is as follows:

```
show ft history {cfg_cntlr | ha_dp_mgr | ha_mgr}
```

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

For example, enter:

```
host1/Admin# show ft history cfg_cntlr
```

Displaying 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, use the **show ft idmap** command in Exec mode. The syntax of this command is as follows:

```
show ft idmap
```

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

For example, enter:

```
host1/Admin# show ft idmap
```

Displaying Memory Statistics

To display redundancy statistics per context, use the **show ft memory** command in Exec mode. The syntax of this command is as follows:

```
show ft memory [detail]
```

The optional **detail** keyword displays detailed HA manager memory statistics in the Admin context only.

For example, enter:

```
host1/Admin# show ft memory detail
```

Displaying Peer Information

To display peer information, use the **show ft peer** command in Exec mode. The syntax of this command is as follows:

```
show ft peer peer_id {detail | status | summary}
```

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

For example, enter:

```
host1/Admin# show ft peer 1
```

[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	<p>Current state of the peer. Possible states are:</p> <p>FSM_PEER_STATE_INIT—Initial state of the peer after you configure it.</p> <p>FSM_PEER_STATE_MY_IPADDR—Local ACE IP address is missing. Waiting for the local IP address to be configured.</p> <p>FSM_PEER_STATE_PEER_IPADDR—Peer IP address is missing. Waiting for the peer IP address to be configured.</p> <p>FSM_PEER_STATE_START_HB—Peer configuration is complete. Starting the heartbeat to see if there is a peer device.</p>

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

Field	Description
State (continued)	<p>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.</p> <p>FSM_PEER_STATE_SRG_CHECK—Checking for software version compatibility with the peer device.</p> <p>FSM_PEER_STATE_LIC_CHECK—Checking for license compatibility with the peer device.</p> <p>FSM_PEER_STATE_COMPATIBLE—Version and license checks indicate that the peer is compatible for redundancy.</p> <p>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.</p> <p>FSM_PEER_STATE_DOWN—Peer device is down.</p> <p>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.</p>

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

Field	Description
Maintenance Mode	<p>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:</p> <ul style="list-style-type: none"> • 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 module 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.
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.

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

Field	Description
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, use the **show ft stats** command in Exec mode. The syntax of this command is as follows:

```
show ft stats group_id
```

The *group_id* argument displays additional load-balancing statistics (LB statistics) for the specified group.

For example, enter:

```
host1/Admin# show ft stats 1
```

[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 modules send heartbeat packets and each packet indicates whether the other peer has been receiving heartbeats.

Table 6-5 *Field Descriptions for the show ft stats Command Output (continued)*

Field	Description
Number of HB Timeout Mismatches	<p>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.</p> <p>Note The heartbeat interval should be the same on both peer modules. 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.</p>
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.

Table 6-5 *Field Descriptions for the show ft stats Command Output (continued)*

Field	Description
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, use the **show ft track** command in Exec mode. The syntax of this command is as follows:

```
show ft track {detail | status | summary}
```

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

For example, enter:

```
host1/Admin# show ft track detail
```

For more information about the **show ft track** command output for HSRP tracking, see the “[Before You Begin](#)” section.

[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	<p>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:</p> <ul style="list-style-type: none"> • 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 module 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	<p>State of the FT group member in the local ACE. Possible states are:</p> <ul style="list-style-type: none"> • 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.

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

Field	Description
My State (Cont.)	<ul style="list-style-type: none"> 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.
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, TRACK_HSRP, or TRACK_INTERFACE.
HSRP Group name	Identifier of the HSRP group that is configured on the Catalyst 6500 series switch that you are tracking.
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.

**Note**

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 module only does not clear the statistics on the standby module.

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](#)

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, use the **clear ft ha-stats** command in Exec mode. The syntax of this command is as follows:

```
clear ft ha-stats
```

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.

For example, enter:

```
host1/Admin# clear ft ha-stats
```

Clearing Heartbeat Statistics

To clear all heartbeat-related statistics, use the **clear ft hb-stats** command in Exec mode. 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. The syntax of this command is as follows:

```
clear ft hb-stats
```

For example, enter:

```
host1/Admin# clear ft hb-stats
```

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, use the **clear ft-track stats** command in Exec mode. The syntax of this command is as follows:

```
clear ft track-stats [all]
```

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.

For example, to clear tracking statistics for all FT groups that are configured in the ACE, enter:

```
host1/Admin# clear ft track-stats all
```

Clearing All Redundancy Statistics

To clear all redundancy statistics, including all TL, heartbeat, and tracking counters, use the **clear ft all** command in Exec mode in the Admin context only. The syntax of this command is as follows:

```
clear ft all
```

**Note**

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.

For example, enter:

```
host1/Admin# clear ft all
```

Clearing the Redundancy History

To clear the redundancy history, use the **clear ft history** command in Exec mode in the Admin context only. The syntax of this command is as follows:

```
clear ft history {cfg_cntlr | ha_dp_mgr | ha_mgr}
```

The keywords are:

- **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

For example, enter:

```
host1/Admin# clear ft history cfg_cntlr
```