



Configuring Fibre Channel Routing Services and Protocols

Fabric Shortest Path First (FSPF) is the standard path selection protocol used by Fibre Channel fabrics. The FSPF feature is enabled by default on all Fibre Channel switches. Except in configurations that require special consideration, you do not need to configure any FSPF services. FSPF automatically calculates the best path between any two switches in a fabric. Specifically, FSPF is used to:

- Dynamically compute routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Select an alternative path in the event of the failure of a given path.
 - FSPF supports multiple paths.
 - FSPF automatically computes an alternative path around a failed link. It provides a preferred route when two equal paths are available.

This chapter provides details on Fibre Channel routing services and protocols. It includes the following sections:

- [FSPF Features, page 16-2](#)
- [FSPF Examples, page 16-2](#)
- [Configuring FSPF Globally, page 16-4](#)
- [Configuring FSPF for a Specific Interface, page 16-6](#)
- [Configuring Fibre Channel Routes, page 16-8](#)
- [Clearing FSPF Counters, page 16-9](#)
- [Broadcast Routing, page 16-10](#)
- [In-Order Delivery, page 16-10](#)
- [Configuring Flow Statistics, page 16-14](#)
- [Displaying Routing and Forwarding Information, page 16-16](#)
- [Default Settings, page 16-20](#)

FSPF Features

FSPF is the protocol currently standardized by the T11 committee for routing in Fibre Channel networks. The FSPF protocol has the following characteristics and features:

- Supports multipath routing.
- Bases path status on a link state protocol.
- Routes hop by hop, based only on the domain ID.
- Runs only on E ports or TE ports and provides a loop free topology.
- Runs on a per VSAN basis. Connectivity in a given VSAN in a fabric is guaranteed only for the switches configured in that VSAN.
- Uses a topology database to keep track of the state of the links on all switches in the fabric and associates a cost with each link.
- Guarantees a fast reconvergence time in case of a topology change. Uses the standard Dijkstra's algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra's algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.

FSPF Examples

This section provides examples of topologies and applications that demonstrate the benefits of FSPF.




Note

The FSPF feature can be used on any topology.

Fault Tolerant Fabric

Figure 16-1 depicts a fault tolerant fabric using a partial mesh topology. If a link goes down anywhere in the fabric, any switch can still communicate with all others in the fabric. In the same way, if any switch goes down, the connectivity of the rest of the fabric is preserved.

Figure 16-1 Fault Tolerant Fabric




For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in [Figure 16-1](#), each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. [Figure 16-2](#) shows this arrangement. Because switches in the Cisco MDS 9000 Family support PortChanneling, each pair of physical links can appear to the FSPF protocol as one single logical link.

By bundling pairs of physical links, FSPF efficiency is considerably improved by the reduced database size and the frequency of link updates. Once physical links are aggregated, failures are not attached to a single link but to the entire PortChannel. This configuration also improves the resiliency of the network. The failure of a link in a PortChannel does not trigger a route change, thereby reducing the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 16-2 Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in [Figure 16-3](#) and summarized in [Table 16-3](#) and [Table 16-4](#). Two links between switch 1 and switch 2 exist as either equal-cost ISLs or PortChannels. There is one flow from traffic generator 1 to traffic generator 2. The traffic was tested at 100 utilization of 1G in two scenarios:

- Disabling the traffic link by either physically removing the cable (see [Table 16-3](#)).
- Shutting down either switch 1 or switch 2 (see [Table 16-4](#)).

Figure 16-3 Fail-over Scenario Using Traffic Generators




Table 16-1 Physically Removing the Cable for the SmartBits Scenario

| PortChannel Scenario | | FSPF Scenario (Equal cost ISL) | |
|---|----------|--------------------------------|----------|
| Switch 1 | Switch 2 | Switch 1 | Switch 2 |
| 110ms (~2K frame drops) | | 130ms+ (~4k frame drops) | |
| 100 ms hold time when a signal loss is reported as mandated by the standard | | | |

Table 16-2 Shutting Down the Switch for the SmartBits Scenario

| PortChannel Scenario | | FSPF Scenario (Equal cost ISL) | |
|-----------------------|-------------------------|--------------------------------|-------------------------|
| Switch 1 | Switch 2 | Switch 1 | Switch 2 |
| ~0ms (~8 frame drops) | 110ms (~2K frame drops) | 130ms+ (~4K frame drops) | |
| No hold time needed. | Signal loss on switch 1 | No hold time needed. | Signal loss on switch 1 |

Configuring FSPF Globally

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you don't have to specify the VSAN number for every command. This global configuration feature also reduces the chance of typing errors or other minor configuration errors.



Note

FSPF is enabled by default. Generally, you do not need to configure these advanced features.



Caution

The default for the backbone region is 0 (zero). You do not need to change this setting unless your region is different from the default. If you are operating with other vendors using the backbone region, you can change this default to be compatible with those settings.

To configure a FSPF feature for the entire VSAN, follow these steps:

| | Command | Purpose |
|---------------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fspf config vsan 1 switch-config-(fspf-config)# | Enters FSPF global configuration mode for the specified VSAN. |
| Step 3 | switch-config-(fspf-config)# spf static switch-config-(fspf-config)# | Forces static SPF computation for the dynamic (default) incremental VSAN. |
| Step 4 | switch-config-(fspf-config)# spf hold-time 10 switch-config-(fspf-config)# | Configures the hold time between two route computations in milliseconds for the entire VSAN. The default value is 0. Note If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly. |
| Step 5 | switch-config-(fspf-config)# region 7 switch-config-(fspf-config)# | Configures the autonomous region for this VSAN and specifies the region ID (7). |

Deleting the Entire FSPF Configuration

To delete FSPF configuration for the entire VSAN, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf config vsan 3 switch(config)# | Deletes the FSPF configuration for VSAN 3. |

Disabling FSPF Routing Protocols

By default, FSPF is enabled on switches in the Cisco MDS 9000 Family.

To enable or disable FSPF routing protocols, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# no fspf enable vsan 5 switch(config)# | Disables FSPF routing protocol in VSAN 5. |
| | switch(config)# fspf enable vsan 7 switch(config)# | Enables FSPF routing protocol in VSAN 7. |

Link State Record Defaults

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches, and then flooded throughout the fabric. [Table 16-3](#) displays the default settings for switch responses.

Table 16-3 LSR Default Settings

| LSR Option | Default | Description |
|---|------------|--|
| Acknowledgement interval (RxmtInterval) | 5 seconds | The time a switch waits for an acknowledgement from the LSR before retransmission. |
| Refresh time (LSRefreshTime) | 30 minutes | The time a switch waits before sending an LSR refresh transmission. |
| Maximum age (MaxAge) | 60 minutes | The time a switch waits before dropping the LSR from the database. |

Configuring FSPF for a Specific Interface

Several FSPF commands are available on a per interface basis. The following configuration procedures apply to an interface in a specific VSAN and are described in this section.

- [Computing Route Cost, page 16-6](#)
- [Specifying Hello Time Intervals, page 16-6](#)
- [Specifying Dead Intervals, page 16-7](#)
- [Disabling FSPF for Specific Interfaces, page 16-7](#)
- [Retransmitting Intervals, page 16-8](#)

Computing Route Cost

FSPF tracks the state of links on all switches in the fabric, associates a cost with each link in its database, and then chooses the path with a minimal cost. The cost associated with an interface can be administratively changed to implement the FSPF route selection. The integer value to specify cost can range from 1 to 65,535. The default cost for 1Gbps is 1000 and 2Gbps is 500.

To configure FSPF link cost, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf cost 5 vsan 90 switch(config-if)# | Configures the cost for the selected interface in VSAN 90. |

Specifying Hello Time Intervals

You can set the FSPF hello time interval to specify the interval between the periodic hello messages sent to verify the health of the link. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.

To configure the FSPF Hello time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf hello-interval 15 vsan 175 switch(config-if)# | Specifies the hello message interval (15 seconds) to verify the health of the link in VSAN 175. The default is 20 seconds. |

Specifying Dead Intervals

You can set the FSPF dead time interval to specify the maximum interval for which a hello message must be received before the neighbor is considered lost and removed from the database. The integer value can range from 1 to 65,535 seconds.



Note

This value must be the same in the ports at both ends of the ISL.



Caution

An error is reported at the command prompt if the configured dead time interval is less than the Hello time interval.

To configure the FSPF dead time interval, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf dead-interval 25 vsan 7 switch(config-if)# | Specifies the maximum interval for VSAN 7 before which a Hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds. |

Disabling FSPF for Specific Interfaces

You can disable the FSPF protocol for selected interfaces. By default, FSPF is enabled on all E ports and TE ports. This default can be disabled by setting the interface as passive.



Note

FSPF must be enabled at both ends of the interface for the protocol to work.

To disable FSPF for a specific interface, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures a specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf passive vsan 1 switch(config-if)# | Disables the FSPF protocol for the specified interface in the specified VSAN. |
| | switch(config-if)# no fspf passive vsan 1 switch(config-if)# | Reenables the FSPF protocol for the specified interface in the specified VSAN. |

Retransmitting Intervals

You can specify the time after which an unacknowledged link state update should be transmitted on the interface. The integer value to specify retransmit intervals can range from 1 to 65,535 seconds.


Note

This value must be the same on the switches on both ends of the interface.


To configure the FSPF retransmit time interval, follow these steps:

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# interface fc1/4 switch(config-if)# | Configures the specified interface, or if already configured, enters configuration mode for the specified interface. |
| Step 3 | switch(config-if)# fspf retransmit-interval 15 vsan 12 switch(config-if)# | Specifies the retransmit time interval for unacknowledged link state updates in VSAN 12. The default is 5 seconds. |

Configuring Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. To configure the FC ID for the specified interface and domain, you can configure the specified route (for example FC ID 111211 and domain ID 3) in the switch with domain ID 1 (see Figure 16-4).

Figure 16-4 Fibre Channel Routes


Note

Other than in VSANs, run time checks are not performed on configured and suspended static routes.

To configure an FC route, follow these steps:

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t | Enters configuration mode. |
| Step 2 | switch(config)# fcroute 0x111211 interface fc1/1 domain 3 vsan 2 switch(config)# | Configures the route for the specified Fibre Channel interface and domain. In this example, interface fc1/1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch. |
| | switch(config)# fcroute 0x111211 interface port-channel 1 domain 3 vsan 4 switch(config)# | Configures the route for the specified PortChannel interface and domain. In this example, interface port-channel 1 is assigned an FC ID (0x111211) and a domain ID (3) to the next hop switch. |
| | switch(config)# fcroute 0x031211 interface fc1/1 domain 3 metric 1 vsan 1 switch(config-if)# | Configures the static route for a specific FC ID and next hop domain ID and also assigns the cost of the route. If the remote destination option is not specified, the default is direct. |
| Step 3 | switch(config)# fcroute 0x111112 interface fc1/1 domain 3 metric 3 remote vsan 3 | Adds a static route to the RIB. If this is an active route and the FIB ¹ records are free, it is also added to the FIB. If the cost (metric) of the route is not specified, the default is 10. |
| | switch(config)# fcroute 0x610000 0xff0000 interface fc 1/1 domain 1 vsan 2 switch(config)# | Configures the netmask for the specified route in interface fc1/1 (or PortChannel). You can specify one of three routes: ff0000 matches only the domain, ffff00 matches the domain and the area, fffff matches the domain, area, and port. |

1. FIB = Forwarding Information Base

Clearing FSPF Counters

To clear the FSPF statistics counters for one interface or for the entire VSAN, follow this step:

| | Command | Purpose |
|--------|--|---|
| Step 1 | switch# clear fspf counters vsan 1 switch# | Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared. |
| | switch# clear fspf counters vsan 200 interface fc1/1 switch# | Clears the FSPF statistics counters for the specified interface in VSAN 200. |

Broadcast Routing

Broadcast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric (for broadcast traffic).

FSPF provides the topology information to compute the distribution tree. Fibre Channel defines 256 multicast groups and one broadcast address for each VSAN. Switches in the Cisco MDS 9000 Family only use broadcast routing. By default, they use the principal switch as the root node to derive the distribution tree information. The protocols create a loop-free broadcast distribution tree.



Caution

All switches in the fabric should run the same multicast and broadcast distribution tree algorithm to ensure the same distribution tree.

In-Order Delivery

In-order delivery of data frames guarantees frame delivery to a destination in the same order that they were sent by the originator.


Some Fibre Channel protocols or applications cannot handle out-of-order frame delivery. In these cases, switches in the Cisco MDS 9000 Family preserve frame ordering in the frame flow. The source ID (SID), destination ID (DID), and optionally the originator exchange ID (OX ID) identify the flow of the frame.

In case of a single switch, all frames received by a specific ingress port and destined to a certain egress port are always delivered in the same order in which they were received.

Reordering Network Frames

When you experience a route change in the network. The new selected path may be faster or less congested than the old route (see [Figure 16-5](#)).

Figure 16-5 Route Change Delivery



In [Figure 16-5](#), the new path from Switch 1 to Switch 4 is faster. Hence, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.


If the in-order guarantee feature is enabled, the frames within the network are treated as specified below:

- Frames in the network are delivered in the order in which they are transmitted.
- Frames which can not be delivered in-order within the network latency drop period are dropped inside the network.
- The number of dropped frames are reduced by slowing down the traffic at the frame source.

Reordering PortChannel Frames

When a link change occurs in a PortChannel, the frames for the same exchange or the same flow can switch from one path to another faster path (see [Figure 16-6](#)).

Figure 16-6 Link Congestion Delivery



In [Figure 16-6](#), the port of the old path (red dot) is congested. Hence Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

When the in-order guarantee feature is enabled, the frames crossing a PortChannel are treated as specified below:

- Frames using the old path are delivered before new frames are accepted.
- Frames which cannot be delivered in-order, through the old path, within the switch latency drop period are dropped.
- The new frames are delivered through the new path after the switch latency drop period has elapsed.

Enabling In-Order Delivery

By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.



Tip

We recommend that you only enable this feature in a switch when devices are present in the switch that cannot handle any out-of-order frames. Load-balancing algorithms within the Cisco MDS 9000 Family ensure that frames are delivered in order during normal fabric operation. The load-balancing algorithms based on source FC ID, destination FC ID, and exchange ID are enforced in hardware without any performance degradation. However, if the fabric encounters a failure and this feature is enabled, the recovery will be delayed due to an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out of order.

Use the **show in-order-guarantee** command to display the present configuration status:

```
switch# show in-order-guarantee
```

■ In-Order Delivery

inorder delivery is not guaranteed

To enable in-order delivery, follow these steps.

| | Command | Purpose |
|--------|--|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# in-order-guarantee | Enables in-order delivery in the switch. |
| | switch(config)# no in-order-guarantee | Reverts the switch to the factory defaults and disables the in-order delivery feature. |

Configuring the Drop Latency Time

Use this command if you need to change the default latency time for either a network or a switch.

To configure the network and the switch drop latency time, follow these steps.

| | Command | Purpose |
|--------|---|--|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcdroplateny network 5000 | Configures network drop latency time to be 5000 milliseconds for the network. The valid range is 0 to 60000 milliseconds. The default is 2000 milliseconds. Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network |
| | switch(config)# no fcdroplateny network 4500 | Removes the current fcdroplateny network configuration (4500) and reverts the switch to the factory defaults. |
| Step 3 | switch(config)# fcdroplateny switch 4000 | Configures switch drop latency time to be 4000 milliseconds for the switch. The valid range is 0 to 60000 milliseconds. The default is 500 milliseconds. Note The switch drop latency parameter should have the same value in all the switches in the network |
| | switch(config)# no fcdroplateny switch 4500 | Removes the current fcdroplateny switch configuration (4500) and reverts the switch to the factory defaults. |

Displaying Latency Information

You can view the configured latency parameters using the **show fcdroplateny** command (see [Example 16-1](#)).

Example 16-1 Displays Administrative Distance

```
switch# show fcdroplateny
switch latency value:4000 milliseconds
network latency value:5000 milliseconds
```

Configuring Flow Statistics

Flow statistics count the ingress traffic in the aggregated statistics table. You can collect two kinds of statistics:

- Aggregated flow statistics to count the traffic for a VSAN
- Flow statistics to count the traffic for a source and destination ID pair in a VSAN

If you enable flow counters, you can enable a maximum of 1K entries for aggregate flow and flow statistics. Be sure to assign an unused flow index to a module for each new flow. Flow indexes can be repeated across modules. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

To count the aggregated flow statistics for a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---------------------------------------|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)# | Enables the aggregated flow counter. |
| | switch(config)# no fcflow stats aggregated module 1 index 1005 vsan 1 switch(config)# | Disables the aggregated flow counter. |

To count the flow statistics for a source and destination FC ID in a VSAN, follow these steps:

| | Command | Purpose |
|--------|---|---|
| Step 1 | switch# config t switch(config)# | Enters configuration mode. |
| Step 2 | switch(config)# fcflow stats module 1 index 1 0x145601 0x5601ff ffffff vsan 1 switch(config)# | Enables the flow counter. Note The source ID and the destination ID are specified in the FC ID hex format (for example, 0x123aff). The mask can be one of ff0000 or ffffff. |
| | Step 3 | switch(config)# no fcflow stats aggregated module 2 index 1001 vsan 2 switch(config)# |

Clearing FIB¹ Statistics

To clear the aggregated flow counter, use the **clear fcflow stats** command (see Examples 16-2 and 16-3).

Example 16-2 Clears Aggregated Flow Counters

```
switch# clear fcflow stats aggregated module 2 index 1
```

Example 16-3 Clears Flow Counters for Source and Destination FC IDs

```
switch# clear fcflow stats module 2 index 1
```

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics (see Example 16-4 to 16-6).

Example 16-4 Displays Aggregated fcflow Details for the Specified Module

```
switch# show fcflow stats aggregated module 2
Idx  VSAN # frames # bytes
----  -
0000 4    387,653  674,235,875
0001 6     34,402   2,896,628
```

Example 16-5 Displays fcflow Details for the Specified Module

```
switch# show fcflow stats module 2
Idx  VSAN D ID          S ID          mask          # frames # bytes
----  -
0000 4    032.001.002 007.081.012 ff.ff.ff      387,653  674,235,875
0001 6    004.002.001 019.002.004 ff.00.00     34,402   2,896,628
```

Example 16-6 Displays fcflow Index Usage for the Specified Module

```
switch# show fcflow stats usage module 2
2 flows configured
configured flow : 3,7
```

1. FIB = Forwarding Information Base

Displaying Routing and Forwarding Information

You can view specific information about existing Fibre Channel and FSPF configurations at any time from the EXEC mode. The following **show** commands provide further details on existing Fibre Channel paths and routes (see Examples 16-7 to 16-15).



Note

When the number of routes are displayed in the command output, both visible and hidden routes are included in the total number of routes. While the hidden routes are added to the count, they will not be visible.

Example 16-7 Displays Administrative Distance

```
switch# show fcroute distance

      Route
UUID  Distance      Name
----  -
10    20               RIB
22    40               FCDOMAIN
39    80               RIB-CONFIG
12    100              FSPF
17    120              FLOGI
21    140              TLPM
14    180              MCAST
64    200              RIB-TEST
```

Example 16-8 Displays Multicast Routing Information

```
switch# show fcroute multicast
VSAN FC ID      # Interfaces
----  -
1     0xffffffff 0
2     0xffffffff 1
3     0xffffffff 1
4     0xffffffff 0
5     0xffffffff 0
6     0xffffffff 0
7     0xffffffff 0
8     0xffffffff 0
9     0xffffffff 0
10    0xffffffff 0
```

Example 16-9 Displays FCID Information for a Specified VSAN

```
switch# show fcroute multicast vsan 3

VSAN FC ID      # Interfaces
----  -
3     0xffffffff 1
```

Example 16-10 Displays FCID and interface Information for a Specified VSAN

```
switch# show fcroute multicast 0xffffffff vsan 2
VSAN FC ID      # Interfaces
----  -
2     0xffffffff 1
      fc1/1
```


Example 16-11 Displays Unicast Routing Information

```
switch# show fcroute unicast
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  1    0x010101 0xffffffff 0x00 0x00 D P A 1    10
static  2    0x111211 0xffffffff 0x00 0x00 R P A 1    10
fspf    3    0x610000 0xff0000   0x00 0x00 D P A 4    500
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 16-12 Displays Unicast Routing Information for a Specified VSAN

```
switch# show fcroute unicast vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040102 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040103 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x040104 0xffffffff 0x00 0x00 R P A 1    103
static  4    0x111211 0xffffffff 0x00 0x00 D P A 1    10
```

Example 16-13 Displays Unicast Routing Information for a Specified FCID

```
switch# show fcroute unicast 0x040101 0xffffffff vsan 4
D:direct R:remote P:permanent V:volatile A:active N:non-active
# Next
Protocol VSAN    FC ID/Mask      RCtrl/Mask  Flags Hops   Cost
-----
static  4    0x040101 0xffffffff 0x00 0x00 R P A 1    103
    fc1/2 Domain 0xa6(166)
```

Example 16-14 Displays Route Database Information

```
switch# show fcroute summary
FC Route Database Created Thu Feb 13 07:21:52 2003
VSAN      Ucast      Mcast      Label      Last Modified Time
-----
1         5          1          0          Thu Feb 13 10:21:06 2003
2         4          1          0          Thu Feb 13 10:21:07 2003
3         4          1          0          Thu Feb 13 10:21:08 2003
4         4          1          0          Thu Feb 13 10:21:09 2003
5         4          1          0          Thu Feb 13 10:21:10 2003
6         4          1          0          Thu Feb 13 10:21:11 2003
7         4          1          0          Thu Feb 13 10:21:12 2003
8         4          1          0          Thu Feb 13 10:21:13 2003
9         4          1          0          Thu Feb 13 10:21:14 2003
10        4          1          0          Thu Feb 13 10:21:15 2003
11        4          1          0          Thu Feb 13 10:21:16 2003
12        4          1          0          Thu Feb 13 10:21:17 2003
13        4          1          0          Thu Feb 13 10:21:18 2003
14        4          1          0          Thu Feb 13 10:21:18 2003
15        4          1          0          Thu Feb 13 10:21:19 2003
-----
Total    61         15         0
```

Example 16-15 Displays Route Database Information for a Specified VSAN

```
switch# show fcroute summary vsan 5
FC Route Database Created Thu Feb 13 07:21:52 2003
```

| VSAN | Ucast | Mcast | Label | Last Modified Time |
|-------|-------|-------|-------|--------------------------|
| 5 | 4 | 1 | 0 | Thu Feb 13 10:21:10 2003 |
| Total | 4 | 1 | 0 | |

Displaying Global FSPF Information

The **show fspf** command (see [Example 16-16](#)) displays global FSPF information for a specific VSAN:

- the domain number of the switch
- the autonomous region for the switch
- Min_LS_arrival: the minimum time that must elapse before the switch accepts LSR updates
- Min_LS_interval: the minimum time that must elapse before the switch can transmit an LSR
- LS_refresh_time: the interval lapse between refresh LSR transmissions
- Max_age: the maximum time aa LSR can stay before being deleted

Example 16-16 Displays FSPF Information for a Specified VSAN

```
switch# show fspf vsan 1
FSPF routing for VSAN 1
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 5000 msec
Local Domain is 0x65(101)
Number of LSRs = 3, Total Checksum = 0x0001288b

Protocol constants :
  LS_REFRESH_TIME = 1800 sec
  MAX_AGE          = 3600 sec

Statistics counters :
  Number of LSR that reached MaxAge = 0
  Number of SPF computations         = 7
  Number of Checksum Errors          = 0
  Number of Transmitted packets :   LSU 65 LSA 55 Hello 474 Retranmsitted LSU 0
  Number of received packets :     LSU 55 LSA 60 Hello 464 Error packets 10
```

Displaying the FSPF Database

The **show fspf database** command displays a summary of the FSPF database for a specified VSAN (see [Example 16-17](#)). If other parameters are not specified, all LSRs in the database are displayed:

- LSR Type
- Domain ID of the LSR owner
- Domain ID of the advertising router

- LSR age
- LS incarnation member
- Number of links

You could narrow the display to obtain specific information by issuing additional parameters for the domain ID of the LSR owner. For each interface, the following information is also available:

- Domain ID of the neighboring switch
- E port index
- Port index of the neighboring switch
- Link type and cost

Example 16-17 Displays FSPF Database Information

```
switch# show fspf database vsan 1
```

FSPF Link State Database for VSAN 1 Domain 0x0c(12)

```
LSR Type           = 1
Advertising domain ID = 0x0c(12)
LSR Age            = 1686
LSR Incarnation number = 0x80000024
LSR Checksum       = 0x3caf
Number of links    = 2
```

| NbrDomainId | IfIndex | NbrIfIndex | Link Type | Cost |
|-------------|------------|------------|-----------|------|
| 0x65(101) | 0x0000100e | 0x00001081 | 1 | 500 |
| 0x65(101) | 0x0000100f | 0x00001080 | 1 | 500 |

FSPF Link State Database for VSAN 1 Domain 0x65(101)

```
LSR Type           = 1
Advertising domain ID = 0x65(101)
LSR Age            = 1685
LSR Incarnation number = 0x80000028
LSR Checksum       = 0x8443
Number of links    = 6
```

| NbrDomainId | IfIndex | NbrIfIndex | Link Type | Cost |
|-------------|------------|------------|-----------|------|
| 0xc3(195) | 0x00001085 | 0x00001095 | 1 | 500 |
| 0xc3(195) | 0x00001086 | 0x00001096 | 1 | 500 |
| 0xc3(195) | 0x00001087 | 0x00001097 | 1 | 500 |
| 0xc3(195) | 0x00001084 | 0x00001094 | 1 | 500 |
| 0x0c(12) | 0x00001081 | 0x0000100e | 1 | 500 |
| 0x0c(12) | 0x00001080 | 0x0000100f | 1 | 500 |

FSPF Link State Database for VSAN 1 Domain 0xc3(195)

```
LSR Type           = 1
Advertising domain ID = 0xc3(195)
LSR Age            = 1686
LSR Incarnation number = 0x80000033
LSR Checksum       = 0x6799
Number of links    = 4
```

| NbrDomainId | IfIndex | NbrIfIndex | Link Type | Cost |
|-------------|------------|------------|-----------|------|
| 0x65(101) | 0x00001095 | 0x00001085 | 1 | 500 |
| 0x65(101) | 0x00001096 | 0x00001086 | 1 | 500 |
| 0x65(101) | 0x00001097 | 0x00001087 | 1 | 500 |
| 0x65(101) | 0x00001094 | 0x00001084 | 1 | 500 |

Displaying FSPF Interfaces

The **show fspf** command displays the following information for each selected interface (see [Example 16-18](#)).

- link cost
- timer values
- neighbor's domain ID (if known)
- local interface number
- remote interface number (if known)
- FSPF state of the interface
- interface counters

Example 16-18 Displays FSPF Interface Information

```
switch# show fspf vsan 1 interface fc1/1
FSPF interface fc1/1 in VSAN 1
FSPF routing administrative state is active
Interface cost is 500
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x0c(12), Neighbor Interface index is 0x0f100000
Statistics counters :
  Number of packets received : LSU 8 LSA 8 Hello 118 Error packets 0
  Number of packets transmitted : LSU 8 LSA 8 Hello 119 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

Default Settings

[Table 16-4](#) lists the default settings for FSPF features.

Table 16-4 Default FSPF Settings

| Parameters | Default |
|---|---|
| FSPF | Enabled on all E ports and TE ports. |
| SPF computation | Dynamic. |
| SPF hold time | 0. |
| Backbone region | 0. |
| Acknowledgement interval (RxmtInterval) | 5 seconds. |
| Refresh time (LSRefreshTime) | 30 minutes. |
| Maximum age (MaxAge) | 60 minutes. |
| Hello interval | 20 seconds. |
| Dead interval | 80 seconds. |
| Distribution tree information | Derived from the principal switch (root node). |
| Routing table | FSPF stores up to 16 equal cost paths to a given destination. |

Table 16-4 Default FSPF Settings (continued)

| Parameters | Default |
|---------------------------|---|
| Load balancing | Based on destination ID and source ID on different, equal cost paths. |
| In-order delivery | Disabled. |
| Drop latency | Disabled. |
| Static route cost | If the cost (metric) of the route is not specified, the default is 10. |
| Remote destination switch | If the remote destination switch is not specified, the default is direct. |

