CHAPTER 25

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 and 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:

About FSPF

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 25-1 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 25-1 Fault Tolerant Fabric

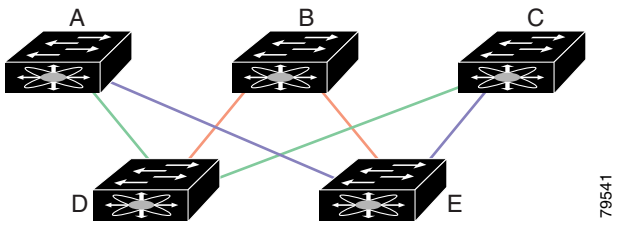

For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Links

To further improve on the topology in Figure 25-1, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. Figure 25-2 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 25-2 Fault Tolerant Fabric with Redundant Links

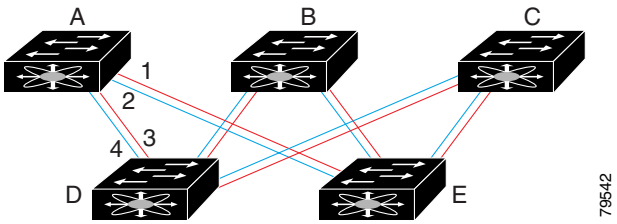

For example, if all links are of equal speed and no PortChannels exist, the FSPF calculates four equal paths from A to C: A1-E-C, A2-E-C, A3-D-C, and A4-D-C. If PortChannels exist, these paths are reduced to two.

Fail-Over Scenarios for PortChannels and FSPF Links

The SmartBits traffic generator was used to evaluate the scenarios displayed in Figure 25-3. 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 at 1 Gbps in two scenarios:

- Disabling the traffic link by physically removing the cable (see Table 25-1).
- Shutting down either switch 1 or switch 2 (see Table 25-2).

Figure 25-3 Fail-Over Scenario Using Traffic Generators

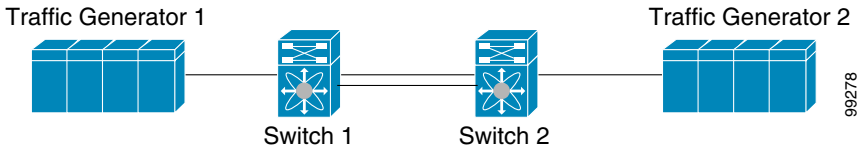

Table 25-1 Physically Removing the Cable for the SmartBits Scenario

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

Table 25-2 Shutting Down the Switch for the SmartBits Scenario

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

FSPF Global Configuration

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

Some FSPF features can be globally configured in each VSAN. By configuring a feature for the entire VSAN, you do not 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.

This section includes the following topics:

About SPF Computational Hold Times

The SPF computational hold time sets the minimum time between two consecutive SPF computations on the VSAN. Setting this to a small value means that FSPF reacts faster to any fabric changes by recomputing paths on the VSAN. A small SPF computational hold time uses more switch CPU time.

About Link State Records

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 25-3 displays the default settings for switch responses.

Table 25-3 LSR Default Settings

LSR Option	Default	Description
Acknowledgment interval (RxmtInterval)	5 seconds	The time a switch waits for an acknowledgment 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.

The LSR minimum arrival time is the period between receiving LSR updates on this VSAN. Any LSR updates that arrive before the LSR minimum arrival time are discarded.

The LSR minimum interval time is the frequency at which this switch sends LSR updates on a VSAN.

Configuring FSPF on a VSAN

To configure an FSPF feature for the entire VSAN using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane as shown in Figure 25-4.

Figure 25-4 FSPF General Information

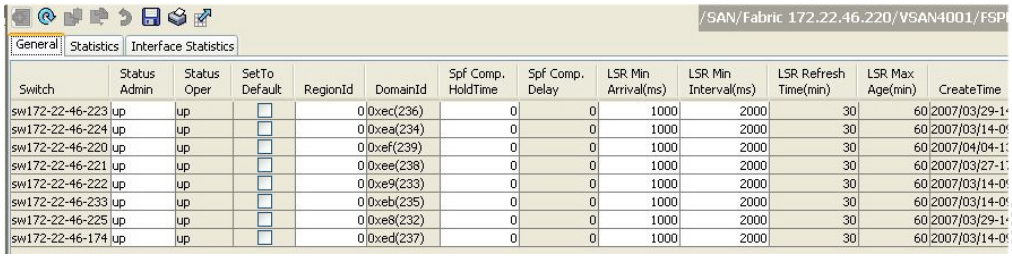

Step 2 The RegionID, **Spf Comp Holdtime**, LSR Min Arrival, and LSR Min Interval field values are applied across all interfaces on the VSAN. You can change them here or, if they don't exist, create them here.

Step 3 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Resetting FSPF to the Default Configuration

To return the FSPF VSAN global configuration to its factory default using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane as shown in Figure 25-4.

Step 2 Check the **SetToDefault** check box for a switch.

Step 3 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Enabling or Disabling FSPF

To enable or disable FSPF using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN and select **FSPF** for a VSAN that you want to configure for FSPF.

You see the FSPF configuration in the Information pane as shown in Figure 25-4.

Step 2 Set the Status Admin drop-down menu to **up** to enable FSPF or to **down** to disable FSPF.

Step 3 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

FSPF Interface Configuration

Several FSPF commands are available on a per interface basis. These configuration procedures apply to an interface in a specific VSAN.

This section includes the following topics:

About FSPF Link 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 1 Gbps is 1000 and for 2 Gbps is 500.

Configuring FSPF Link Cost

To configure FSPF link cost using Fabric Manager, follow these steps:

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane as shown in Figure 25-5.

Figure 25-5 Fibre Channel Physical FSPF Interface

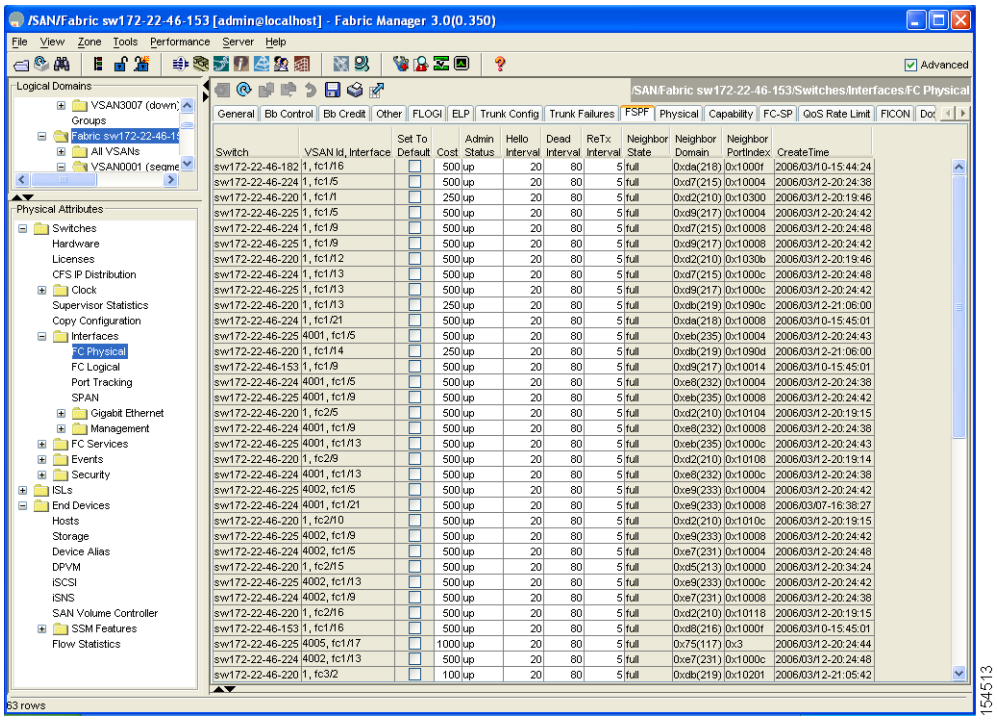

Step 3 Double-click in the Cost field of a switch and change the value.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About 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.

Configuring Hello Time Intervals

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

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane as shown in Figure 25-5.

Step 3 Change the Hello Interval field for a switch.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About Dead Time 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.

Configuring Dead Time Intervals

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

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane as shown in Figure 25-5.

Step 3 Double-click the Dead Interval field for a switch and provide a new value.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About 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.

Configuring Retransmitting Intervals

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

Step 1 Expand **Switches**, expand **Interfaces**, and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane as shown in Figure 25-5.

Step 3 Double-click the ReTx Interval field and enter a value.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

About 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.

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.

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

Step 1 Expand **Switches**, expand **Interfaces** and then select **FC Physical**.

You see the interface configuration in the Information pane.

Step 2 Click the **FSPF** tab.

You see the FSPF interface configuration in the Information pane shown in Figure 25-5.

Step 3 Set a switch Admin Status drop-down menu to **down**.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

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.

Displaying the FSPF Database

The FSPF database for a specified VSAN includes the following information:

- Link State Record (LSR) type
- Domain ID of the LSR owner
- Domain ID of the advertising router
- LSR age
- LSR incarnation member
- Number of links

To display the FSPF database using Device Manager, follow these steps:

Step 1 Choose **FC > Advanced > FSPF**.

You see the FSPF dialog box shown in Figure 25-6.

Figure 25-6 FSPF Dialog Box in Device Manager

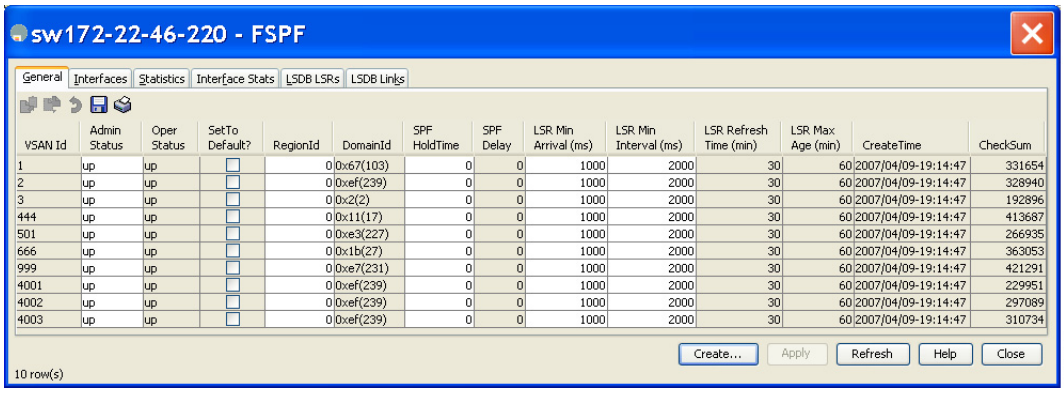

Step 2 Click the **LSDB LSRs** tab.

You see the FSPF database information shown in Figure 25-7.

Figure 25-7 FSPF Database Information in the LSDB LSRs Tab

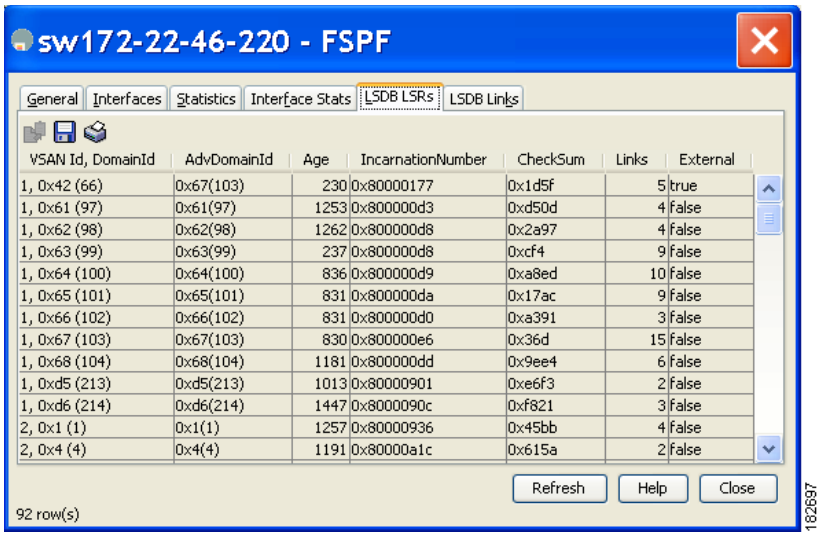

Step 3 Click **Close** to close the dialog box.

Viewing FSPF Statistics

To view FSPF statistics using Fabric Manager, follow these steps:

Step 1 Expand a Fabric, expand a VSAN, and then select **FSPF** in the Logical Domains pane.

You see the FSPF configuration dialog box.

Step 2 Click the **Statistics** tab.

You see the FSPF VSAN statistics in the Information pane (see Figure 25-8).

Figure 25-8 FSPF VSAN Statistics

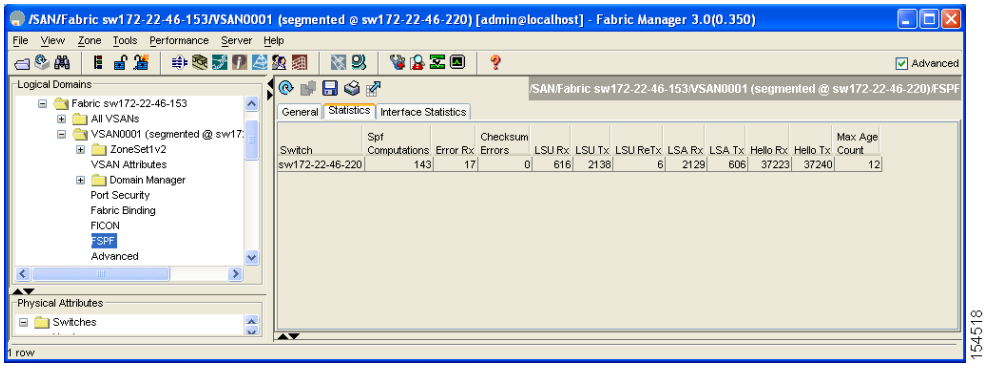

Step 3 Click the **Interface Statistics** tab.

You see the FSPF interface statistics in the Information pane.

FSPF Routes

FSPF routes traffic across the fabric, based on entries in the FSPF database. These routes can be learned dynamically, or configured statically.

This section includes the following topics:

About Fibre Channel Routes

Each port implements forwarding logic, which forwards frames based on its FC ID. Using 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 25-9).

Figure 25-9 Fibre Channel Routes

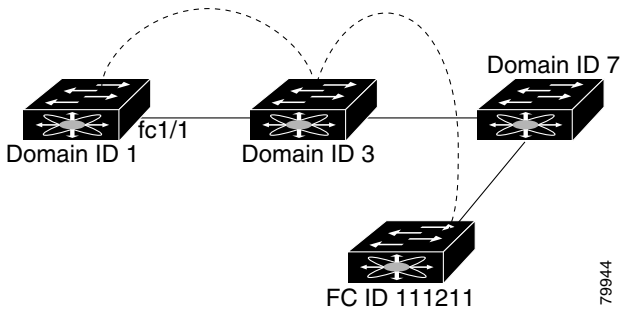

Note Other than in VSANs, runtime checks are not performed on configured and suspended static routes.

Configuring Fibre Channel Routes

If you disable FSPF, you can manually configure a Fibre Channel route.

To configure a Fibre Channel route using Device Manager, follow these steps:

Step 1 Click **FC > Advanced > Routes**.

You see the FC Static Route Configuration dialog box shown in Figure 25-10.

Figure 25-10 Fibre Channel Static Route Configuration Dialog Box

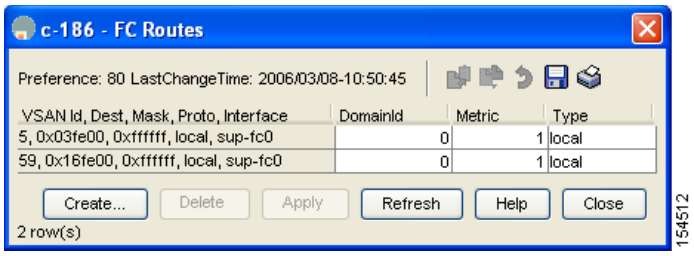

Step 2 Click **Create** to create a static route.

You see the Create Route dialog box shown in Figure 25-11.

Figure 25-11 Create Fibre Channel Route Dialog Box

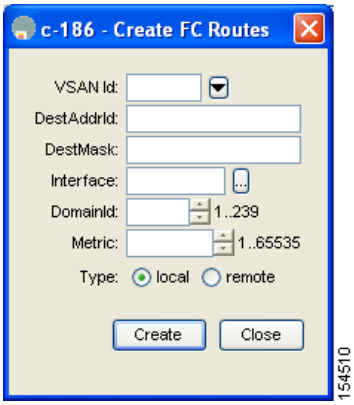

Step 3 Select the VSAN ID that for which you are configuring this route.

Step 4 Fill in the destination address and destination mask for the device you are configuring a route.

Step 5 Select the interface that you want to use to reach this destination.

Step 6 Select the next hop domain ID and route metric.

Step 7 Select either the **local** or **remote** radio button.

Step 8 Click **Create** to save these changes, or click **Close** to discard any unsaved changes.

About Broadcast and Multicast Routing

Broadcast and multicast in a Fibre Channel fabric uses the concept of a distribution tree to reach all switches in the fabric.

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 a loop-free distribution tree for multicast and broadcast routing in a VSAN.

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

To interoperate with other vendor switches (following FC-SW3 guidelines), the SAN-OS software uses the lowest domain switch as the root to compute the multicast tree in interop mode.

About Multicast Root Switch

By default, the **native** (non-interop) mode uses the principal switch as the root. If you change the default, be sure to configure the same mode in all switches in the fabric. Otherwise, multicast traffic could face potential loop and frame-drop problems.

Note The operational mode can be different from the configured interop mode. The interop mode always uses the lowest domain switch as the root.

Setting the Multicast Root Switch

To use the lowest domain switch for the multicast tree computation using Fabric Manager, follow these steps:

Step 1 Expand a fabric, expand a VSAN, and then select **Advanced** for the VSAN that you want to configure FSPF on.

You see the advanced Fibre Channel configuration in the Information pane.

Step 2 Select the **Multicast Root** tab.

You see the multicast root configuration in the Information pane as shown in Figure 25-12.

Figure 25-12 Multicast Root Configuration

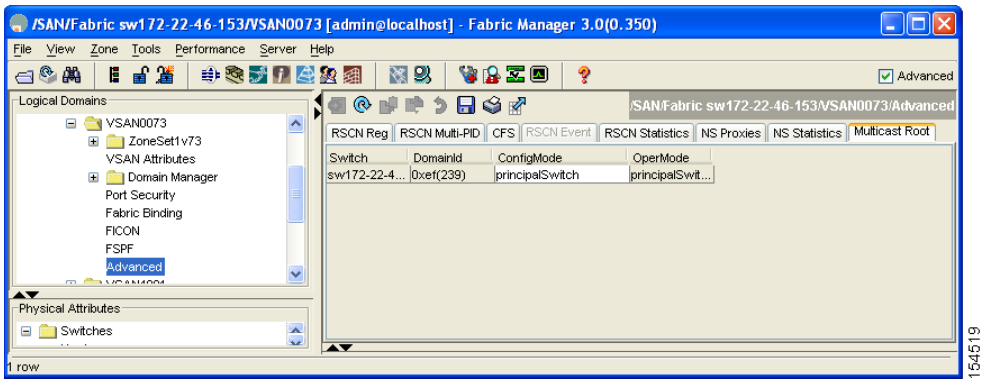

Step 3 Set the Config Mode drop-down menu to **lowestDomainSwitch**.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

In-Order Delivery

In-Order Delivery (IOD) 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.

On any given switch with IOD enabled, 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.

Use IOD only if your environment cannot support out-of-order frame delivery.

Tip If you enable the in-order delivery feature, the graceful shutdown feature is not implemented.

This section includes the following topics:

About 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.

Figure 25-13 Route Change Delivery

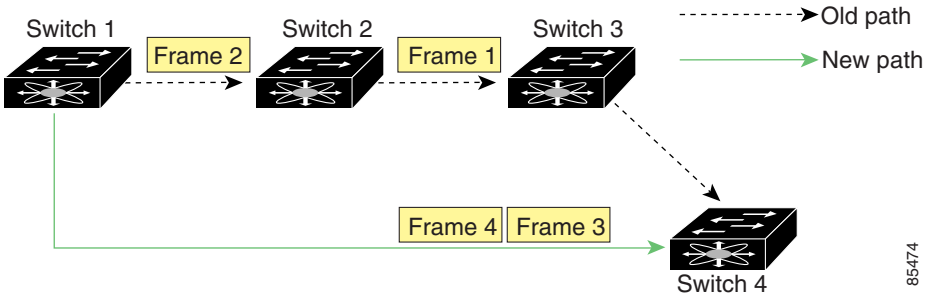

In Figure 25-13, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 may be delivered before Frame 1 and Frame 2.

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

- Frames in the network are delivered in the order in which they are transmitted.
- Frames that cannot be delivered in order within the network latency drop period are dropped inside the network.

About 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.

Figure 25-14 Link Congestion Delivery

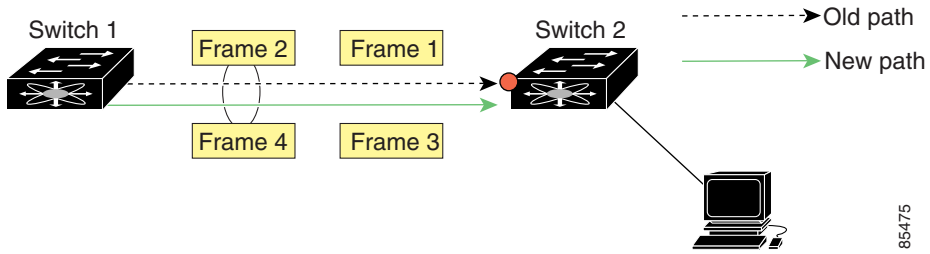

In Figure 25-14, the port of the old path (red dot) is congested. In this scenario, Frame 3 and Frame 4 can be delivered before Frame 1 and Frame 2.

The the in-order delivery feature attempts to minimize the number of frames dropped during PortChannel link changes when the in-order delivery is enabled by sending a request to the remote switch on the PortChannel to flush all frames for this PortChannel.

Note Both switches on the PortChannel must be running Cisco SAN-OS Release 3.0(1) for this IOD enhancement. For earlier releases, IOD waits for the switch latency period before sending new frames.

When the in-order delivery guarantee feature is enabled and a PortChannel link change occurs, the frames crossing the PortChannel are treated as follows:

- Frames using the old path are delivered before new frames are accepted.
- The new frames are delivered through the new path after the switch latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the switch latency drop period are dropped. See the "Configuring the Drop Latency Time" section on page 25-18.

About Enabling In-Order Delivery

You can enable the in-order delivery feature for a specific VSAN or for the entire switch. By default, in-order delivery is disabled on switches in the Cisco MDS 9000 Family.

Tip We recommend that you only enable this feature when devices that cannot handle any out-of-order frames are present in the switch. 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 because of an intentional pausing of fabric forwarding to purge the fabric of resident frames that could potentially be forwarded out-of-order.

Enabling In-Order Delivery Globally

To ensure that the in-order delivery parameters are uniform across all VSANs on an MDS switch, enable in-order delivery globally.

Only enable in-order delivery globally if this is a requirement across your entire fabric. Otherwise, enable IOD only for the VSANs that require this feature.

Note Enable in-order delivery on the entire switch before performing a downgrade to Cisco MDS SAN-OS Release1.3(3) or earlier.

Enabling In-Order Delivery for a VSAN

When you create a VSAN, that VSAN automatically inherits the global in-order-guarantee value. You can override this global value by enabling or disabling in-order-guarantee for the new VSAN.

To use the lowest domain switch for the multicast tree computation using Fabric Manager, follow these steps:

Step 1 Expand a fabric and select **All VSANS**.

Step 2 Select the **Attributes** tab.

You see the general VSAN attributes in the Information pane shown in Figure 25-15.

Figure 25-15 General VSAN Attributes

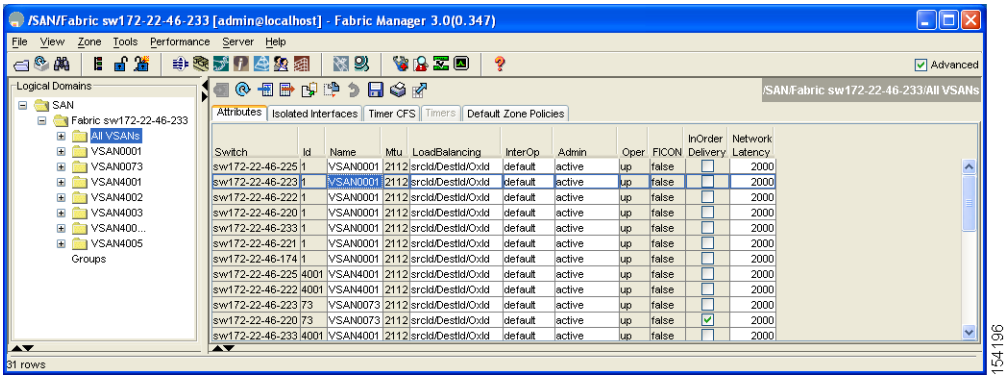

Step 3 Check the **InOrder Delivery** check box to enable IOD for the switch.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Configuring the Drop Latency Time

You can change the default latency time for either the entire switch or a specified VSAN in a switch.

To configure the drop latency time for a switch using Fabric Manager, follow these steps:

Step 1 Expand a fabric and select **All VSANS**.

You see the VSAN configuration in the Information pane.

Step 2 Select the **Attributes** tab.

You see the general VSAN attributes in the Information pane shown in Figure 25-16.

Figure 25-16 General VSAN Attributes

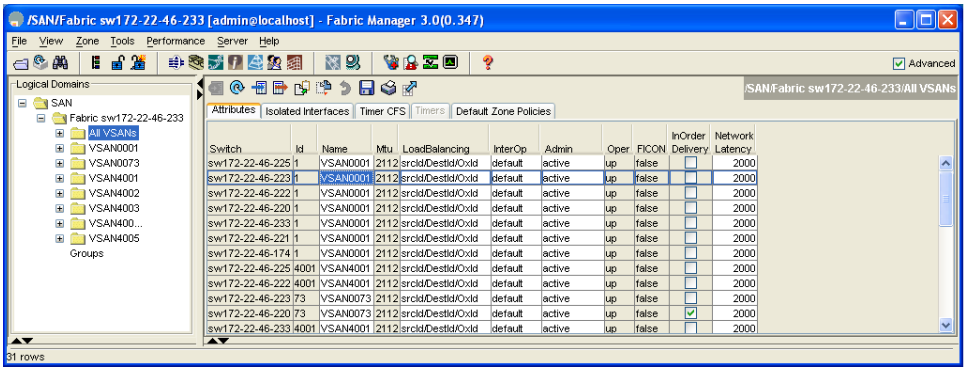

Step 3 Double-click the Network Latency field and change the value.

Step 4 Click **Apply Changes** to save these changes, or click **Undo Changes** to discard any unsaved changes.

Default Settings

Table 25-4 lists the default settings for FSPF features.

Table 25-4 Default FSPF Settings

Parameters	Default
FSPF	Enabled on all E ports and TE ports.
SPF computation	Dynamic.
SPF hold time	0.
Backbone region	0.
Acknowledgment 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).

Table 25-4 Default FSPF Settings (continued)

Parameters	Default
Routing table	FSPF stores up to 16 equal cost paths to a given destination.
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.
Multicast routing	Uses the principal switch to compute the multicast tree.