



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, page 25-2](#)
- [FSPF Global Configuration, page 25-4](#)
- [FSPF Interface Configuration, page 25-6](#)
- [FSPF Routes, page 25-12](#)
- [In-Order Delivery, page 25-15](#)
- [Default Settings, page 25-19](#)

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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



79541

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).


Send documentation comments to mdsfeedback-doc@cisco.com

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)			

Send documentation comments to mdsfeedback-doc@cisco.com

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, page 25-4](#)
- [About Link State Records, page 25-4](#)
- [Configuring FSPF on a VSAN, page 25-5](#)
- [Resetting FSPF to the Default Configuration, page 25-5](#)
- [Enabling or Disabling FSPF, page 25-6](#)

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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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

Switch	Status Admin	Status Oper	SetTo Default	RegionId	DomainId	Spf Comp. HoldTime	Spf Comp. Delay	LSR Min Arrival(ms)	LSR Min Interval(ms)	LSR Refresh Time(min)	LSR Max Age(min)	CreateTime
sw172-22-46-223	up	up	<input type="checkbox"/>	00xec(236)	0	0	0	1000	2000	30	60	2007/03/29-11:00:00
sw172-22-46-224	up	up	<input type="checkbox"/>	00xea(234)	0	0	0	1000	2000	30	60	2007/03/14-00:00:00
sw172-22-46-220	up	up	<input type="checkbox"/>	00xef(239)	0	0	0	1000	2000	30	60	2007/04/04-11:00:00
sw172-22-46-221	up	up	<input type="checkbox"/>	00xee(238)	0	0	0	1000	2000	30	60	2007/03/27-11:00:00
sw172-22-46-222	up	up	<input type="checkbox"/>	00xe9(233)	0	0	0	1000	2000	30	60	2007/03/14-00:00:00
sw172-22-46-233	up	up	<input type="checkbox"/>	00xeb(235)	0	0	0	1000	2000	30	60	2007/03/14-00:00:00
sw172-22-46-225	up	up	<input type="checkbox"/>	00xe8(232)	0	0	0	1000	2000	30	60	2007/03/29-11:00:00
sw172-22-46-174	up	up	<input type="checkbox"/>	00xed(237)	0	0	0	1000	2000	30	60	2007/03/14-00:00:00

- 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.

Send documentation comments to mdsfeedback-doc@cisco.com

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, page 25-6](#)
- [Configuring FSPF Link Cost, page 25-7](#)
- [About Hello Time Intervals, page 25-7](#)
- [Configuring Hello Time Intervals, page 25-8](#)
- [About Dead Time Intervals, page 25-8](#)
- [Configuring Dead Time Intervals, page 25-8](#)
- [About Retransmitting Intervals, page 25-8](#)
- [Configuring Retransmitting Intervals, page 25-9](#)
- [About Disabling FSPF for Specific Interfaces, page 25-9](#)
- [Disabling FSPF for Specific Interfaces, page 25-9](#)
- [Displaying the FSPF Database, page 25-10](#)
- [Viewing FSPF Statistics, page 25-11](#)

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.

Send documentation comments to mdsfeedback-doc@cisco.com

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

Switch	VSAN Id	Interface	Set To	Admin	Hello Interval	Dead Interval	ReTx Interval	Neighbor State	Neighbor Domain	Neighbor PortIndex	CreateTime
sw172-22-46-182	1	fc1/16	500	up	20	80	5	full	0x0a(218)	0xc1000f	2006.03/10-15:44:24
sw172-22-46-224	1	fc1/5	500	up	20	80	5	full	0xd7(215)	0xc10004	2006.03/12-20:24:38
sw172-22-46-220	1	fc1/1	250	up	20	80	5	full	0xd2(210)	0xc10300	2006.03/12-20:19:46
sw172-22-46-225	1	fc1/5	500	up	20	80	5	full	0xd9(217)	0xc10004	2006.03/12-20:24:42
sw172-22-46-224	1	fc1/9	500	up	20	80	5	full	0xd7(215)	0xc10008	2006.03/12-20:24:42
sw172-22-46-225	1	fc1/9	500	up	20	80	5	full	0xd9(217)	0xc10008	2006.03/12-20:24:42
sw172-22-46-220	1	fc1/12	500	up	20	80	5	full	0xd2(210)	0xc1030b	2006.03/12-20:19:46
sw172-22-46-224	1	fc1/13	500	up	20	80	5	full	0xd7(215)	0xc1000c	2006.03/12-20:24:42
sw172-22-46-225	1	fc1/13	500	up	20	80	5	full	0xd9(217)	0xc1000c	2006.03/12-20:24:42
sw172-22-46-220	1	fc1/13	250	up	20	80	5	full	0xd7(215)	0xc1000c	2006.03/12-21:06:00
sw172-22-46-224	1	fc1/21	500	up	20	80	5	full	0xd7(215)	0xc10008	2006.03/10-15:45:01
sw172-22-46-225	4001	fc1/5	500	up	20	80	5	full	0xeb(235)	0xc10004	2006.03/12-20:24:43
sw172-22-46-220	1	fc1/14	250	up	20	80	5	full	0xd7(215)	0xc1099d	2006.03/12-21:06:00
sw172-22-46-153	1	fc1/18	500	up	20	80	5	full	0xd9(217)	0xc10014	2006.03/10-15:45:01
sw172-22-46-224	4001	fc1/5	500	up	20	80	5	full	0xeb(235)	0xc10004	2006.03/12-20:24:38
sw172-22-46-225	4001	fc1/9	500	up	20	80	5	full	0xeb(235)	0xc10008	2006.03/12-20:24:42
sw172-22-46-220	1	fc2/5	500	up	20	80	5	full	0xd2(210)	0xc10104	2006.03/12-20:19:15
sw172-22-46-224	4001	fc1/9	500	up	20	80	5	full	0xeb(235)	0xc10008	2006.03/12-20:24:38
sw172-22-46-225	4001	fc1/13	500	up	20	80	5	full	0xeb(235)	0xc1000c	2006.03/12-20:24:43
sw172-22-46-220	1	fc2/9	500	up	20	80	5	full	0xd2(210)	0xc10108	2006.03/12-20:19:14
sw172-22-46-224	4001	fc1/13	500	up	20	80	5	full	0xeb(235)	0xc1000c	2006.03/12-20:24:38
sw172-22-46-225	4002	fc1/5	500	up	20	80	5	full	0xeb(235)	0xc10004	2006.03/12-20:24:42
sw172-22-46-224	4001	fc1/21	500	up	20	80	5	full	0xeb(235)	0xc10008	2006.03/12-18:38:27
sw172-22-46-220	1	fc2/10	500	up	20	80	5	full	0xd2(210)	0xc1010c	2006.03/12-20:19:15
sw172-22-46-225	4002	fc1/9	500	up	20	80	5	full	0xeb(235)	0xc10008	2006.03/12-20:24:42
sw172-22-46-224	4002	fc1/5	500	up	20	80	5	full	0xeb(235)	0xc10004	2006.03/12-20:24:48
sw172-22-46-220	1	fc2/15	500	up	20	80	5	full	0xd5(213)	0xc10000	2006.03/12-20:34:24
sw172-22-46-225	4002	fc1/13	500	up	20	80	5	full	0xeb(235)	0xc1000c	2006.03/12-20:24:42
sw172-22-46-224	4002	fc1/9	500	up	20	80	5	full	0xeb(235)	0xc10008	2006.03/12-20:24:38
sw172-22-46-220	1	fc2/16	500	up	20	80	5	full	0xd2(210)	0xc10118	2006.03/12-20:19:15
sw172-22-46-153	1	fc1/16	500	up	20	80	5	full	0xd8(216)	0xc1000f	2006.03/10-15:45:01
sw172-22-46-225	4005	fc1/17	1000	up	20	80	5	full	0x75(117)	0xc3	2006.03/12-20:24:44
sw172-22-46-224	4002	fc1/13	500	up	20	80	5	full	0xeb(235)	0xc1000c	2006.03/12-20:24:48
sw172-22-46-220	1	fc3/2	100	up	20	80	5	full	0xd7(231)	0xc10201	2006.03/12-21:05:42

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.

Send documentation comments to mdsfeedback-doc@cisco.com

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.

Send documentation comments to mdsfeedback-doc@cisco.com

**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.

Send documentation comments to mdsfeedback-doc@cisco.com

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

The screenshot shows the 'FSPF' dialog box for 'sw172-22-46-220'. The 'LSDB LSRs' tab is selected, displaying a table with the following columns: VSAN Id, Admin Status, Oper Status, SetTo Default?, RegionId, DomainId, SPF HoldTime, SPF Delay, LSR Min Arrival (ms), LSR Min Interval (ms), LSR Refresh Time (min), LSR Max Age (min), CreateTime, and CheckSum. The table contains 10 rows of data, all with 'up' status and a 'SetTo Default?' checkbox.

VSAN Id	Admin Status	Oper Status	SetTo Default?	RegionId	DomainId	SPF HoldTime	SPF Delay	LSR Min Arrival (ms)	LSR Min Interval (ms)	LSR Refresh Time (min)	LSR Max Age (min)	CreateTime	CheckSum
1	up	up	<input type="checkbox"/>		0x67(103)	0	0	1000	2000	30	60	2007/04/09-19:14:47	331654
2	up	up	<input type="checkbox"/>		0xef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	328940
3	up	up	<input type="checkbox"/>		0x2(2)	0	0	1000	2000	30	60	2007/04/09-19:14:47	192896
444	up	up	<input type="checkbox"/>		0x11(17)	0	0	1000	2000	30	60	2007/04/09-19:14:47	413687
501	up	up	<input type="checkbox"/>		0xe3(227)	0	0	1000	2000	30	60	2007/04/09-19:14:47	266935
666	up	up	<input type="checkbox"/>		0x1b(27)	0	0	1000	2000	30	60	2007/04/09-19:14:47	363053
999	up	up	<input type="checkbox"/>		0xe7(231)	0	0	1000	2000	30	60	2007/04/09-19:14:47	421291
4001	up	up	<input type="checkbox"/>		0xef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	229951
4002	up	up	<input type="checkbox"/>		0xef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	297089
4003	up	up	<input type="checkbox"/>		0xef(239)	0	0	1000	2000	30	60	2007/04/09-19:14:47	310734

10 row(s)

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

You see the FSPF database information shown in [Figure 25-7](#).

Send documentation comments to mdsfeedback-doc@cisco.com

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

VSAN Id	DomainId	AdvDomainId	Age	IncarnationNumber	CheckSum	Links	External
1, 0x42 (66)	0x67(103)	0x67(103)	230	0x80000177	0x1d5f	5	true
1, 0x61 (97)	0x61(97)	0x61(97)	1253	0x800000d3	0xd50d	4	false
1, 0x62 (98)	0x62(98)	0x62(98)	1262	0x800000d8	0x2a97	4	false
1, 0x63 (99)	0x63(99)	0x63(99)	237	0x800000d8	0xcf4	9	false
1, 0x64 (100)	0x64(100)	0x64(100)	836	0x800000d9	0xa8ed	10	false
1, 0x65 (101)	0x65(101)	0x65(101)	831	0x800000da	0x17ac	9	false
1, 0x66 (102)	0x66(102)	0x66(102)	831	0x800000d0	0xa391	3	false
1, 0x67 (103)	0x67(103)	0x67(103)	830	0x800000e6	0x36d	15	false
1, 0x68 (104)	0x68(104)	0x68(104)	1181	0x800000dd	0x9ee4	6	false
1, 0xd5 (213)	0xd5(213)	0xd5(213)	1013	0x80000901	0xe6f3	2	false
1, 0xd6 (214)	0xd6(214)	0xd6(214)	1447	0x8000090c	0xf821	3	false
2, 0x1 (1)	0x1(1)	0x1(1)	1257	0x80000936	0x45bb	4	false
2, 0x4 (4)	0x4(4)	0x4(4)	1191	0x80000a1c	0x615a	2	false

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

Switch	Spt	Computations	Error Rx	Errors	Checksum	LSU Rx	LSU Tx	LSU Retx	LSA Rx	LSA Tx	Hello Rx	Hello Tx	Max Age
sw172-22-46-220		143	17	0	616	2138	6	2129	606	37223	37240	12	

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

Send documentation comments to mdsfeedback-doc@cisco.com

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, page 25-12](#)
- [Configuring Fibre Channel Routes, page 25-12](#)
- [About Broadcast and Multicast Routing, page 25-14](#)
- [About Multicast Root Switch, page 25-14](#)
- [Setting the Multicast Root Switch, page 25-14](#)

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.

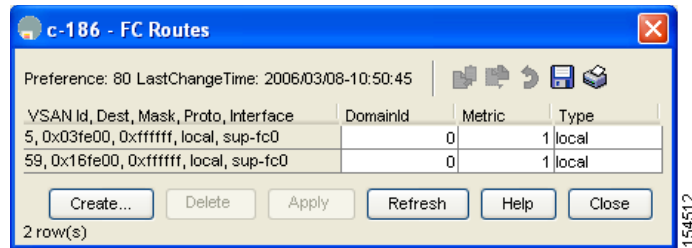
To configure a Fibre Channel route using Device Manager, follow these steps:

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

You see the FC Static Route Configuration dialog box shown in [Figure 25-10](#).

Send documentation comments to mdsfeedback-doc@cisco.com


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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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.

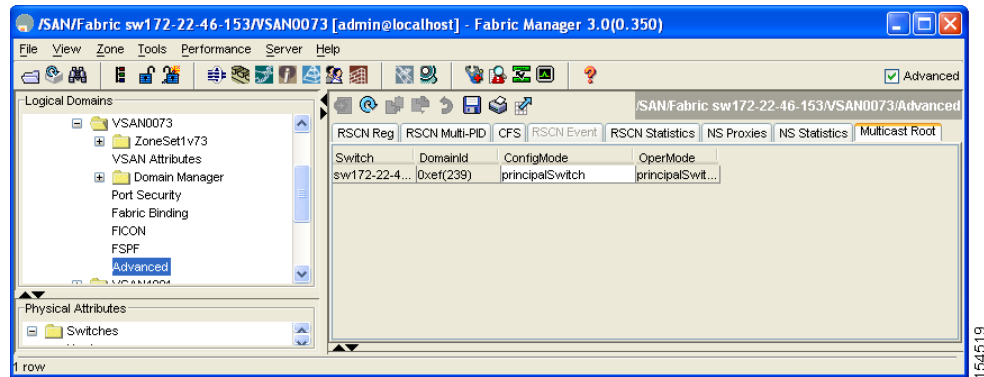
Setting the Multicast Root Switch

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](#).

Send documentation comments to mdsfeedback-doc@cisco.com

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:

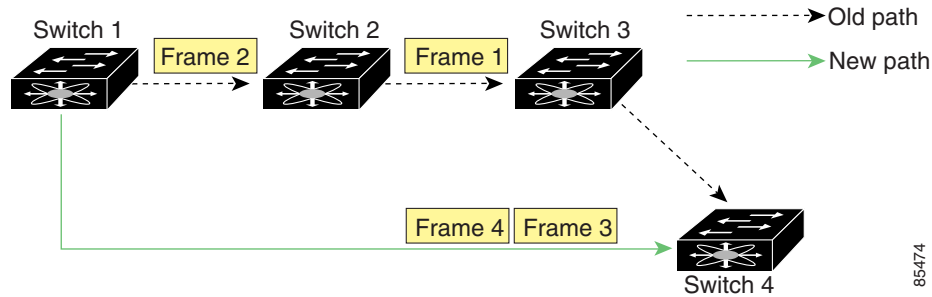
- [About Reordering Network Frames, page 25-15](#)
- [About Reordering PortChannel Frames, page 25-17](#)
- [About Enabling In-Order Delivery, page 25-17](#)
- [Enabling In-Order Delivery Globally, page 25-18](#)
- [Enabling In-Order Delivery for a VSAN, page 25-18](#)
- [Configuring the Drop Latency Time, page 25-18](#)

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.

Send documentation comments to mdsfeedback-doc@cisco.com

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.

[Send documentation comments to mdsfeedback-doc@cisco.com](mailto:mdsfeedback-doc@cisco.com)

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 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.

Send documentation comments to mdsfeedback-doc@cisco.com

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 Release 1.3(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

Switch	Id	Name	Mtu	LoadBalancing	InterOp	Admin	Oper	FICON	InOrder Delivery	Network Latency
sw172-22-46-225	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-223	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-222	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-220	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-233	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-221	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-174	1	VSAN0001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-225	4001	VSAN4001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-222	4001	VSAN4001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-223	73	VSAN0073	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-220	73	VSAN0073	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000
sw172-22-46-233	4001	VSAN4001	2112	srcld/Destld/Oxld	default	active	up	false	<input checked="" type="checkbox"/>	2000
sw172-22-46-233	4001	VSAN4001	2112	srcld/Destld/Oxld	default	active	up	false	<input type="checkbox"/>	2000

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:

Send documentation comments to mdsfeedback-doc@cisco.com

- 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

Switch	Id	Name	Mtu	LoadBalancing	InterOp	Admin	Oper	FICON	InOrder	Network	Latency
sw172-22-46-225	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-223	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-222	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-220	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-233	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-221	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-174	1	VSAN0001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-225	4001	VSAN4001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-222	4001	VSAN4001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-223	73	VSAN0073	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	
sw172-22-46-220	73	VSAN0073	2112	srcId/DestId/OxId	default	active	up	false	<input checked="" type="checkbox"/>	2000	
sw172-22-46-233	4001	VSAN4001	2112	srcId/DestId/OxId	default	active	up	false	<input type="checkbox"/>	2000	

- 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).

Send documentation comments to mdsfeedback-doc@cisco.com

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.