



Configuring Fibre Channel Routing Services and Protocols

This chapter contains the following sections:

- [Information About Fibre Channel Routing Services and Protocols, on page 1](#)

Information About 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 the E mode and TE mode Fibre Channel interfaces on Cisco SAN 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. FSPF provides the following capabilities:

- Dynamically computes routes throughout a fabric by establishing the shortest and quickest path between any two switches.
- Selects 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.
- 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 algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.



Note The FSPF feature can be used on any topology.

Information 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 algorithm, but there is a static dynamic option for a more robust, efficient, and incremental Dijkstra algorithm. The reconvergence time is fast and efficient as the route computation is done on a per VSAN basis.



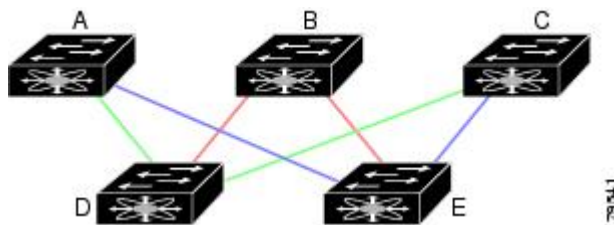
Note The FSPF feature can be used on any topology.

FSPF Examples

Fault Tolerant Fabric Example

The following figure 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 1: Fault Tolerant Fabric



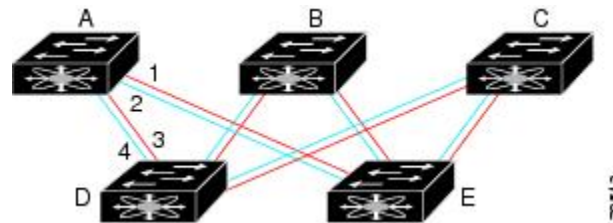
For example, if all links are of equal speed, the FSPF calculates two equal paths from A to C: A-D-C (green) and A-E-C (blue).

Redundant Link Example

To improve on the topology, each connection between any pair of switches can be replicated; two or more links can be present between a pair of switches. The following figure shows this arrangement. Because Cisco SAN switches support SAN port channels, 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 SAN port channel. This configuration also improves the resiliency of the network. The failure of a link in a SAN port channel does not trigger a route change, which reduces the risks of routing loops, traffic loss, or fabric downtime for route reconfiguration.

Figure 2: Fault Tolerant Fabric with Redundant Links



For example, if all links are of equal speed and no SAN port channels 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 SAN port channels exist, these paths are reduced to two.

FSPF Global Configuration

By default, FSPF is enabled on Cisco SAN switches.

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.

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.

Link State Records

Each time a new switch enters the fabric, a link state record (LSR) is sent to the neighboring switches and is then flooded throughout the fabric.

The following table displays the default settings for switch responses.

Table 1: 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

You can configure an FSPF feature for the entire VSAN.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fspf config vsan vsan-id Example: switch(config)# fspf config vsan 14	Enters FSPF global configuration mode for the specified VSAN. Note User needs to configure the VSAN on which FSPF is being configured.
Step 3	spf static Example: switch-config-(fspf-config)# spf static	Forces static SPF computation for the dynamic (default) incremental VSAN.
Step 4	spf hold-time value Example: switch-config-(fspf-config)# spf hold-time 10	Configures the hold time between two route computations in milliseconds (msec) for the entire VSAN. The default value is 0.

	Command or Action	Purpose
		Note If the specified time is shorter, the routing is faster. However, the processor consumption increases accordingly.
Step 5	region <i>region-id</i> Example: switch-config-(fspf-config)# region 1	Configures the autonomous region for this VSAN and specifies the region ID.

Resetting FSPF to the Default Configuration

You can return the FSPF VSAN global configuration to its factory default.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	no fspf config vsan <i>vsan-id</i> Example: switch(config)# no fspf config vsan 24	Deletes the FSPF configuration for the specified VSAN.

Enabling or Disabling FSPF

You can enable or disable FSPF routing protocols.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fspf enable vsan <i>vsan-id</i> Example: switch(config)# fspf enable vsan 567	Enables the FSPF routing protocol in the specified VSAN.
Step 3	no fspf enable vsan <i>vsan-id</i> Example: switch(config)# no fspf enable vsan 567	Disables the FSPF routing protocol in the specified VSAN.

Clearing FSPF Counters for the VSAN

You can clear the FSPF statistics counters for the entire VSAN.

Procedure

	Command or Action	Purpose
Step 1	clear fspf counters vsan <i>vsan-id</i> Example: switch# clear fspf counters vsan 345	Clears the FSPF statistics counters for the specified VSAN. If an interface reference is not specified, all counters are cleared.

FSPF Interface Configuration

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

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

You can configure FSPF link cost.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fspf cost value vsan <i>vsan-id</i> Example: switch(config-if)# fspf cost 500 vsan 38	Configures the cost for the selected interface in the specified VSAN.

Hello Time Intervals

You can set the FSPF Hello time interval to specify the interval between the periodic hello messages that are 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

You can configure the FSPF Hello time interval.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fspf hello-interval value vsan vsan-id Example: <pre>switch(config-if)# fspf hello-interval 25 vsan 10</pre>	Specifies the hello message interval to verify the health of the link in the VSAN. The default is 20 seconds.

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

You can configure the FSPF dead time interval.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fspf dead-interval value vsan vsan-id Example: <pre>switch(config-if)# fspf dead-interval 60 vsan 101</pre>	Specifies the maximum interval for the specified VSAN before which a hello message must be received on the selected interface before the neighbor is considered lost. The default is 80 seconds.

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

You can configure the FSPF retransmit time interval.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fspf retransmit-interval value vsan vsan-id Example: <pre>switch(config-if)# fspf retransmit-interval 10 vsan 25</pre>	Specifies the retransmit time interval for unacknowledged link state updates in the specified VSAN. The default is 5 seconds.

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.


Note

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

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: switch# configure terminal switch(config)#	Enters global configuration mode.
Step 2	fspf passive vsan <i>vsan-id</i> Example: switch(config-if)# fspf passive vsan 24	Disables FSPF for the specified interface in the specified VSAN.
Step 3	no fspf passive vsan <i>vsan-id</i> Example: switch(config-if)# no fspf passive vsan 23	Reenables FSPF for the specified interface in the specified VSAN.

Clearing FSPF Counters for an Interface

You can clear the FSPF statistics counters for an interface.

Procedure

	Command or Action	Purpose
Step 1	switch# clear fspf counters vsan_ <i>vsan-id</i> interface <i>fc_slot/port</i> .	Clears the FSPF statistics counters for the specified interface in the specified VSAN.
Step 2	clear fspf counters vsan_ <i>vsan-id</i> interface <i>type_if-number</i> .	switch# clear fspf counters vsan 12 interface pwnn 9

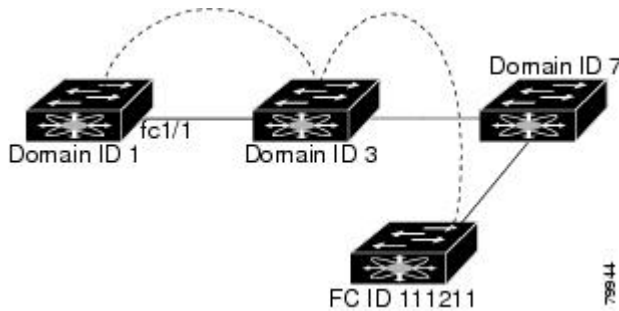
FSPF Routes

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

Fibre Channel Routes

Each port implements a 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 the following figure).

Figure 3: Fibre Channel Routes



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, Cisco SAN switches 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 a 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.

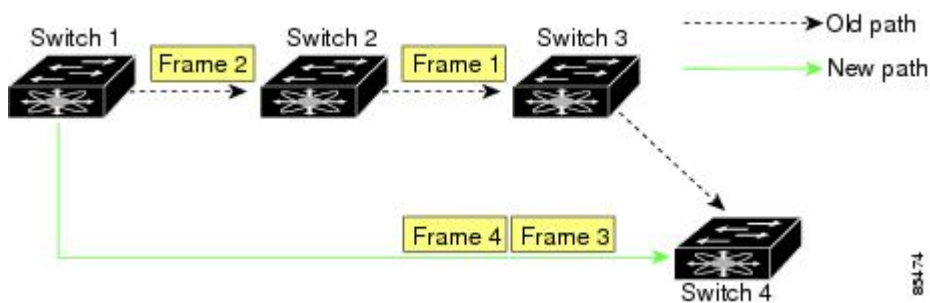
Use IOD only if your environment cannot support out-of-order frame delivery.

If you enable IOD, the graceful shutdown feature is not implemented.

Reordering Network Frames

When you experience a route change in the network, the new selected path might be faster or less congested than the old route (See the following figure).

Figure 4: Route Change Delivery



In the figure above, the new path from Switch 1 to Switch 4 is faster. In this scenario, Frame 3 and Frame 4 might be delivered before Frame 1 and Frame 2.

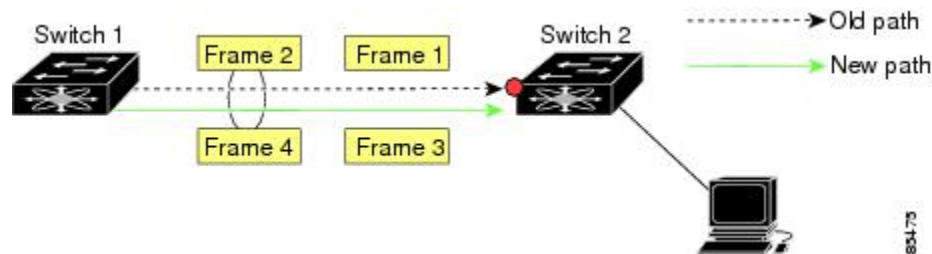
If the in-order guarantee feature is enabled, the frames within the network are delivered as follows:

- 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.

Reordering SAN Port Channel Frames

When a link change occurs in a SAN port channel, the frames for the same exchange or the same flow can switch from one path to another faster path (See the following figure).

Figure 5: Link Congestion Delivery



In the figure above, 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.

When the in-order delivery feature is enabled and a port channel link change occurs, the frames crossing the SAN port channel are delivered 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 network latency drop period has elapsed and all old frames are flushed.

Frames that cannot be delivered in order through the old path within the network latency drop period are dropped.

Related Topics

[Configuring the Drop Latency Time](#), on page 13

About Enabling In-Order Delivery

You can enable IOD for a specific VSAN or for the entire switch. By default, IOD is disabled on Cisco SAN switches.

We recommend that you enable this feature only when devices that cannot handle any out-of-order frames are present in the switch. Load-balancing algorithms within the switch 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 the hardware without any performance degradation. However, if the fabric encounters a failure and the in-order delivery 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

You can enable in-order delivery for the switch.

Procedure

	Command or Action	Purpose
Step 1	configuration terminal Example: switch# configuration terminal switch(config)#	Enters global configuration mode.
Step 2	in-order-guarantee Example: switch(config)# in-order-guarantee	Enables in-order delivery in the switch.
Step 3	no in-order-guarantee Example: switch(config)# no in-order-guarantee	Reverts the switch to the factory defaults and disables the in-order delivery feature.

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.

Procedure

	Command or Action	Purpose
Step 1	configuration terminal Example: switch# configuration terminal switch(config)#	Enters configuration mode.
Step 2	in-order-guarantee vsan <i>vsan-id</i> Example: switch(config)# in-order-guarantee vsan 30	Enables in-order delivery in the specified VSAN.
Step 3	no in-order-guarantee vsan <i>vsan-id</i> Example: switch(config)# no in-order-guarantee vsan 30	Reverts the switch to the factory defaults and disables the in-order delivery feature in the specified VSAN.

Displaying the In-Order Delivery Status

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

```
switch# show in-order-guarantee
global inorder delivery configuration:guaranteed
VSAN specific settings
vsan 1 inorder delivery:guaranteed
```

```

vsan 101 inorder delivery:not guaranteed
vsan 1000 inorder delivery:guaranteed
vsan 1001 inorder delivery:guaranteed
vsan 1682 inorder delivery:guaranteed
vsan 2001 inorder delivery:guaranteed
vsan 2009 inorder delivery:guaranteed
vsan 2456 inorder delivery:guaranteed
vsan 3277 inorder delivery:guaranteed
vsan 3451 inorder delivery:guaranteed
vsan 3452 inorder delivery:guaranteed

```

Configuring the Drop Latency Time

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

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fcdroplateny network <i>value</i> Example: <pre>switch(config)# fcdroplateny network 1000</pre>	Configures network drop latency time for the network. The valid range is from 0 to 60000 msec. The default is 2000 msec. Note The network drop latency must be computed as the sum of all switch latencies of the longest path in the network.
Step 3	fcdroplateny network <i>value</i> vsan <i>vsan-id</i> Example: <pre>switch(config)# fcdroplateny network 1000 vsan 12</pre>	Configures network drop latency time for the specified VSAN.
Step 4	no fcdroplateny network <i>value</i> Example: <pre>switch(config)# no fcdroplateny network 1000</pre>	Removes the current fcdroplateny network configuration and reverts the switch to the factory defaults.

Displaying Latency Information

You can view the configured latency parameters by using the **show fcdroplateny** command:

```

switch# show fcdroplateny
switch latency value:500 milliseconds

```

```

global network latency value:2000 milliseconds
VSAN specific network latency settings
vsan 1 network latency:5000 milliseconds
vsan 2 network latency:2000 milliseconds
vsan 103 network latency:2000 milliseconds
vsan 460 network latency:500 milliseconds

```

Flow Statistics Configuration

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.

Flow Statistics

If you enable flow counters, you can enable a maximum of 1000 entries for aggregate flow and flow statistics. Be sure to assign an unused flow index for each new flow. The number space for flow index is shared between the aggregate flow statistics and the flow statistics.

Counting Aggregated Flow Statistics

You can count the aggregated flow statistics for a VSAN.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fcflow stats aggregated index <i>value</i> vsan <i>vsan-id</i> Example: <pre>switch(config)# fcflow stats aggregated index 20 vsan 12</pre>	Enables the aggregated flow counter.
Step 3	no fcflow stats aggregated index <i>value</i> vsan <i>vsan-id</i> Example: <pre>switch(config)# no fcflow stats aggregated index 20 vsan 12</pre>	Disables the aggregated flow counter.

Counting Individual Flow Statistics

You can count the flow statistics for a source and destination FC ID in a VSAN.

Procedure

	Command or Action	Purpose
Step 1	configure terminal Example: <pre>switch# configure terminal switch(config)#</pre>	Enters global configuration mode.
Step 2	fcflow stats index value dest-fcid source-fcid netmask vsan vsan-id Example: <pre>switch(config)# fcflow stats index 10 0x123aff 0x070128 0xffffffff vsan 15</pre>	Enables the flow counter. Note The source ID and the destination ID are specified in FC ID hex format (for example, 0x123aff). The mask can be one of 0xff0000 or 0xffff.
Step 3	no fcflow stats aggregated index value vsan vsan-id Example: <pre>switch(config)# no fcflow stats aggregated index 11 vsan 200</pre>	Disables the flow counter.

Clearing FIB Statistics

Use the **clear fcflow stats** command to clear the aggregated flow counter:

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

The following example shows how to clear the flow counters for source and destination FC IDs:

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

Displaying Flow Statistics

Use the **show fcflow stats** commands to view flow statistics:

```
switch# show fcflow stats aggregated
```

```
Idx      VSAN      frames
-----
        6          1    42871
```

The following example shows how to display flow statistics:

```
switch# show fcflow stats
```

The following example shows how to display flow index usage:

```
switch# show fcflow stats usage
```

```
2 flows configured
Configured flows : 3,7
```

The following example shows how to display global FSPF information for a specific VSAN:

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

The following example shows how to display a summary of the FSPF database for a specified VSAN. If no additional parameters are specified, all LSRs in the database are displayed:

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

The following example shows how to display FSPF interface information:

Default Settings for FSPF

The following table lists the default settings for FSPF features.

Table 2: 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)
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