CHAPTER 16

Configuring Trunking

This chapter describes the trunking feature provided in Cisco MDS 9000 switches. It includes the following sections:

About Trunking

Trunking, also known as VSAN trunking, is a feature specific to switches in the Cisco MDS 9000 Family. Trunking enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link. Trunking is supported on E ports and F ports. (See Figure 16-1 and Figure 16-2).

This section includes the following topics:

Trunking E Ports

Trunking the E ports enables interconnect ports to transmit and receive frames in more than one VSAN, over the same physical link, using enhanced ISL (EISL) frame format.

Figure 16-1 Trunking E Ports

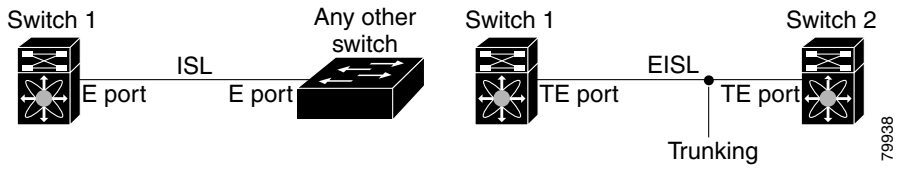

Note Trunking is not supported by internal ports on both the Cisco Fabric Switch for HP c_Class BladeSystem and the Cisco Fabric Switch for IBM BladeCenter.

Trunking F Ports

Trunking F ports allows interconnected ports to transmit and receive tagged frames in more than one VSAN, over the same physical link. Figure 16-2 represents the possible trunking scenarios in a SAN with MDS core switches, NPV switches, third-party core switches, and HBAs.

Figure 16-2 Trunking F Ports

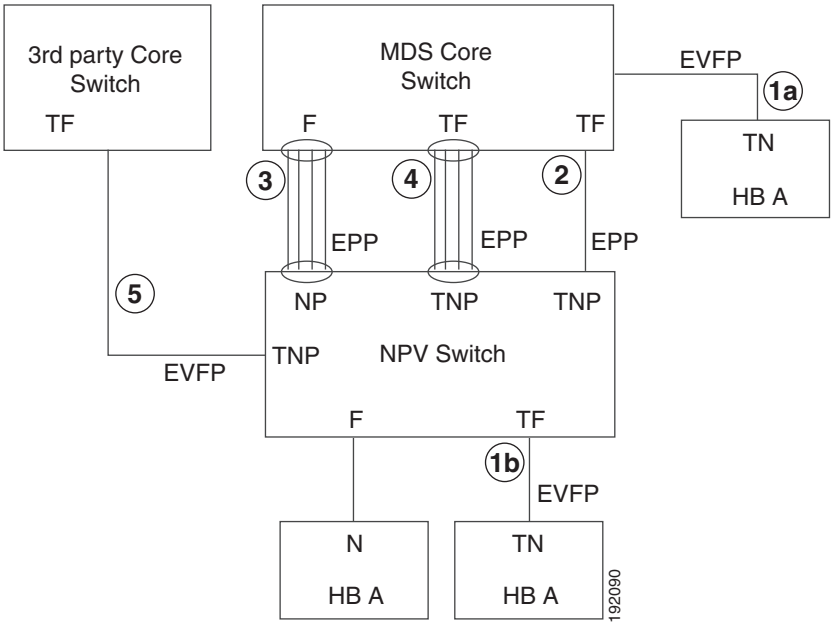

Link Number	Link Description
1a and 1b	F port trunk with N port.[1]
2	F port trunk with NP port.
3	F PortChannnel with NP port.
4	Trunked F PortChannel with NP port.
5	Trunking NP port with third-party core switch F port.[1]

1. These features are not supported currently.

Key Concepts

The trunking feature includes the following key concepts:

- TE port—If trunk mode is enabled in an E port and that port becomes operational as a trunking E port, it is referred to as a TE port.
- TF port—If trunk mode is enabled in an F port (see the link 2 in Figure 16-2) and that port becomes operational as a trunking F port, it is referred to as a TF port.
- TN port—If trunk mode is enabled (not currently supported) in an N port (see the link 1b in Figure 16-2) and that port becomes operational as a trunking N port, it is referred to as a TN port.
- TNP port—If trunk mode is enabled in an NP port (see the link 2 in Figure 16-2) and that port becomes operational as a trunking NP port, it is referred to as a TNP port.
- TF PortChannel—If trunk mode is enabled in an F PortChannel (see the link 4 in Figure 16-2) and that PortChannel becomes operational as a trunking F PortChannel, it is referred to as TF PortChannel. Cisco Port Trunking Protocol (PTP) is used to carry tagged frames.
- TF-TN port link—A single link can be established to connect an F port to an HBA to carry tagged frames (see the link 1a and 1b in Figure 16-2) using Exchange Virtual Fabrics Protocol (EVFP). A server can reach multiple VSANs through a TF port without inter-VSAN routing (IVR).
- TF-TNP port link—A single link can be established to connect an TF port to an TNP port using the PTP protocol to carry tagged frames (see the link 2 in Figure 16-2). PTP is used because PTP also supports trunking PortChannels.

  Note The TF-TNP port link between a third-party NPV core and a Cisco NPV switch is established using the EVFP protocol.

- A Fibre Channel VSAN is called Virtual Fabric and uses a VF_ID in place of the VSAN ID. By default, the VF_ID is 1 for all ports. When an N port supports trunking, a PWWN is defined for each VSAN and called as logical PWWN. In the case of MDS core switches, the PWWNs for which the N port requests additional FC_IDs are called virtual PWWNs.

Trunking Guidelines and Restrictions

The trunking feature includes the following guidelines and restrictions:

- F ports support trunking in Fx mode.

- The trunk-allowed VSANs configured for TE, TF, and TNP links are used by the trunking protocol to determine the allowed active VSANs in which frames can be received or transmitted.
- If a trunking enabled E port is connected to a third-party switch, the trunking protocol ensures seamless operation as an E port.
- Trunking F ports and trunking F PortChannels are not supported on the following hardware:
  - 91x4 switches, if NPIV is enabled and used as the NPIV core switch.
  - Generation 1 2-Gbps Fibre Channel switching modules.
- On core switches, the FC-SP authentication will be supported only for the physical FLOGI from the physical PWWN.
- No FC-SP authentication is supported by the NPV switch on the server F ports.
- MDS does not enforce the uniqueness of logical PWWNs across VSANs.
- DPVM is not supported on trunked F port logins.
- The DPVM feature is limited to the control of the port VSAN, since the EVFP protocol does not allow changing the VSAN on which a logical PWWN has done FLOGI.
- The port security configuration will be applied to both the first physical FLOGI and the per VSAN FLOGIs.
- Trunking is not supported on F ports that have FlexAttach enabled.
- On MDS 91x4 core switches, hard zoning can be done only on F ports that are doing either NPIV or trunking. However, in NPV mode, this restriction does not apply since zoning is enforced on the core F port.

Trunking Misconfiguration Examples

If you do not configure the VSANs correctly, issues with the connection may occur. For example, if you merge the traffic in two VSANs, both VSANs will be mismatched. The trunking protocol validates the VSAN interfaces at both ends of a link to avoid merging VSANs (see Figure 16-3).

Figure 16-3 VSAN Mismatch

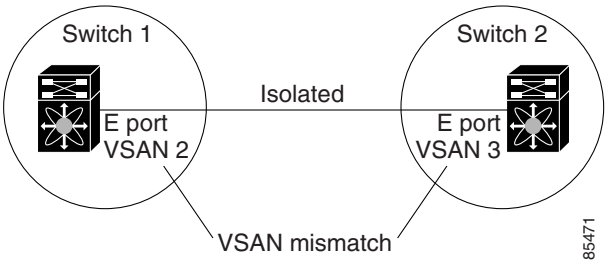

In this example, the trunking protocol detects potential VSAN merging and isolates the ports involved.

The trunking protocol cannot detect merging of VSANs when a third-party switch is placed in between two Cisco MDS 9000 Family switches (see Figure 16-4).

Figure 16-4 Third-Party Switch VSAN Mismatch

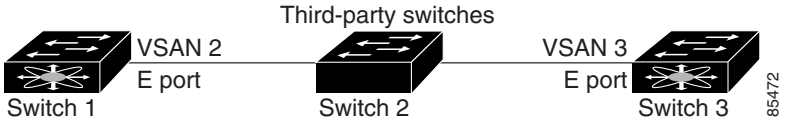

VSAN 2 and VSAN 3 are effectively merged with overlapping entries in the name server and the zone applications. The Cisco MDS 9000 Fabric Manager helps detect such topologies. Refer to the *Cisco MDS 9000 Family Fabric Manager Configuration Guide*.

Upgrade and Downgrade Restrictions

The trunking and channeling feature includes the following upgrade and downgrade restrictions:

- When F port trunking or channeling is configured on a link, the switch cannot be downgraded to Cisco MDS SAN-OS Release 3.x and NX-OS Release 4.1(1b), or earlier.
- Affect of an Upgrade on the EVFP Isolated VSAN—If you are upgrading from a SAN-OS Release 3.x to NX-OS Release 4.1(3a), and you have not created VSAN 4079, the NX-OS software will automatically create VSAN 4079 and reserve it for EVFP use.

  If VSAN 4079 is reserved for EVFP use, the switchport trunk allowed vsan command will filter out VSAN 4079 from the allowed list, as shown in the following example:

```
switch(config-if)# switchport trunk allowed vsan 1-4080
1-4078,4080
switch(config-if)#
```

  If you have created VSAN 4079, the upgrade to NX-OS Release 4.1(3a) will have no affect onVSAN 4079.

  If you downgrade after NX-OS Release 4.1(3a) creates VSAN 4079 and reserves it for EVFP use, the VSAN will no longer be reserved.

Difference Between TE Ports and TF-TNP Ports

In case of TE ports, the VSAN will in be init state when VSAN is coming up on that interface and when peers are in negotiating phase. Once the handshake is done, VSAN will be moved to up state in the successful case, and isolated state in the case of failure. Device Manager will show the port status as Amber during initializing state and it will be green once VSANs are up.

This example shows the Trunk VSAN states of a TE port:

```
Switch# show interface fc2/15
fc2/15 is trunking
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:4f:00:0d:ec:6d:2b:40
    Peer port WWN is 20:0a:00:0d:ec:3f:ab:80
    Admin port mode is auto, trunk mode is on
    snmp link state traps are enabled
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
    Rate mode is dedicated
    Transmit B2B Credit is 16
    Receive B2B Credit is 250
```

```
    B2B State Change Number is 14
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (isolated)                 (100-101)
    Trunk vsans (initializing)             ()
```

In case of TF ports, after the handshake, one of the allowed VSAN will be moved to Up state. And all other VSAN will be in init state even though the handshake with the peer is completed and successful. Each VSAN will be moved from initializing state to up state when a server or target logins through the trunked F or NP ports in the corresponding VSAN.

Note In case of TF or TNP ports, the Device Manager will show port status in Amber even after port is up and there is no failure. It will be changed to green once all the VSAN has successful logins.

This example shows a TF port information after the port is in up state:

```
sw7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2182-2183)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

This example shows the TF port information when a server logins on non-internal flogi vsan: vsan 2183 is moved to up state when server logins in vsan 2183.

```
w7# show interface fc1/13
fc1/13 is trunking (Not all VSANs UP on the trunk)
    Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
    Port WWN is 20:0d:00:0d:ec:6d:2b:40
    Admin port mode is FX, trunk mode is on
    snmp link state traps are enabled
    Port mode is TF
    Port vsan is 1
    Speed is 4 Gbps
    Rate mode is shared
    Transmit B2B Credit is 16
    Receive B2B Credit is 32
    Receive data field Size is 2112
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1,100-101,1101,1163-1166,1216,2172,2
182-2183)
    Trunk vsans (up)                       (1,2183)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             (1101,1163-1166,1216,2172,2182)
```

Enabling the Trunking Protocols

This section explains how to enable or disable the required trunking and channeling protocols represented in Figure 16-2 and includes the following topics:

About Trunking Protocols

The trunking protocol is important for trunking operations on the ports. The protocols enable the following activities:

- Dynamic negotiation of operational trunk mode.
- Selection of a common set of trunk-allowed VSANs.
- Detection of a VSAN mismatch across an ISL.

Table 16-1 specifies the protocols used for trunking and channeling.

Table 16-1 Supported Trunking Protocols

Trunk Link	Default
TE-TE port link	Cisco EPP (PTP)
TF-TN port link[1]	FC-LS Rev 1.62 EVFP
TF-TNP port link	Cisco EPP (PTP)
E or F PortChannel	Cisco EPP (PCP)
TF Port Channel	Cisco EPP (PTP and PCP)
Third-party TF-TNP port link[1]	FC-LS Rev 1.62 EVFP

1. These features are not currently supported.

By default, the trunking protocol is enabled on E ports and disabled on F ports. If the trunking protocol is disabled on a switch, no port on that switch can apply new trunk configurations. Existing trunk configurations are not affected. The TE port continues to function in trunk mode, but only supports traffic in VSANs that it negotiated with previously (when the trunking protocol was enabled). Also, other switches that are directly connected to this switch are similarly affected on the connected interfaces. In some cases, you may need to merge traffic from different port VSANs across a non-trunking ISL. If so, disable the trunking protocol.

Note We recommend that both ends of a trunking link belong to the same port VSAN. On certain switches or fabric switches where the port VSANs are different, one end returns an error and the other end is not connected.

Tip To avoid inconsistent configurations, disable all ports with a **shutdown** command before enabling or disabling the trunking protocols.

Enabling the Cisco Trunking and Channeling Protocols

To enable or disable the Cisco trunking and channeling protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# trunk protocol enable` `switch(config)#`	Enables the Cisco PTP trunking protocol (default).
	`switch(config)# no trunk protocol enable` `switch(config)#`	Disables the Cisco PTP trunking protocol.

Enabling the F Port Trunking and Channeling Protocol

To enable or disable the F port trunking and channeling protocol, follow these steps:

	Command	Purpose
Step 1	`switch# config tasf`	Enters configuration mode.
Step 2	`switch(config)# feature fport-channel-trunk` `switch(config)#`	Enables the F port trunking and channeling protocol (default).
	`switch(config)# no feature fport-channel-trunk` `switch(config)#`	Disables the F port trunking and Channeling protocol.

Note The trunking protocols must be enabled to support trunking, and NPIV must be enabled on the core switch to activate a TF-TNP link. To enable NPIV, use the **feature npiv** command.

Configuring Trunk Mode and VSAN List

This section includes the following topics:

About Trunk Modes

By default, trunk mode is enabled on all Fibre Channel interfaces (Mode: E, F, FL, Fx, ST, and SD) on non-NPV switches. On NPV switches, by default, trunk mode is disabled. You can configure trunk mode as on (enabled), off (disabled), or auto (automatic). The trunk mode configuration at the two ends of an ISL, between two switches, determine the trunking state of the link and the port modes at both ends (see Table 16-2).

Table 16-2 Trunk Mode Status Between Switches

Your Trunk Mode Configuration			Resulting State and Port Mode	
Port Type	**Switch 1**	**Switch 2**	**Trunking State**	**Port Mode**
E ports	On	Auto or on	Trunking (EISL)	TE port
	Off	Auto, on, or off	No trunking (ISL)	E port
	Auto	Auto	No trunking (ISL)	E port
Port Type	**Core Switch**	**NPV Switch**	**Trunking State**	**Link Mode**
F and NP ports	On	Auto or on	Trunking	TF-TNP link
	Auto	On	Trunking	TF-TNP link
	Off	Auto, on, or off	No trunking	F-NP link

Tip The preferred configuration on the Cisco MDS 9000 Family switches is one side of the trunk set to auto and the other side set to on.

Note When connected to a third-party switch, the trunk mode configuration on E ports has no effect. The ISL is always in a trunking disabled state. In the case of F ports, if the third-party core switch ACC's physical FLOGI with the EVFP bit is configured, then EVFP protocol enables trunking on the link.

Configuring Trunk Mode

To configure trunk mode, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# switchport trunk mode on`	Enables (default) the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode off`	Disables the trunk mode for the specified interface.
	`switch(config-if)# switchport trunk mode auto`	Configures the trunk mode to **auto** mode, which provides automatic sensing for the interface.

About Trunk-Allowed VSAN Lists and VF_IDs

Each Fibre Channel interface has an associated trunk-allowed VSAN list. In TE-port mode, frames are transmitted and received in one or more VSANs specified in this list. By default, the VSAN range (1 through 4093) is included in the trunk-allowed list.

The common set of VSANs that are configured and active in the switch are included in the trunk-allowed VSAN list for an interface, and they are called *allowed-active* VSANs. The trunking protocol uses the list of allowed-active VSANs at the two ends of an ISL to determine the list of operational VSANs in which traffic is allowed.

In Figure 16-5, switch 1 has VSANs 1 through 5, switch 2 has VSANs 1 through 3, and switch 3 has VSANs 1, 2, 4, and 5 with a default configuration of trunk-allowed VSANs. All VSANs configured in all three switches are allowed-active. However, only the common set of allowed-active VSANs at the ends of the ISL become operational as shown in Figure 16-5.

For all F, N, and NP ports, the default VF_ID is 1 when there is no VF_ID configured. The trunk-allowed VF_ID list on a port is same as the list of trunk-allowed VSANs. VF_ID 4094 is called the control VF_ID and it is used to define the list of trunk-allowed VF-IDs when trunking is enabled on the link.

If F port trunking and channeling is enabled, or if **switchport trunk mode on** is configured in npv mode for any interface, or if NP PortChannel is configured, the VSAN and VF-ID ranges available for configuration are as follows:

Table 16-3 VSAN and VF-ID Reservations

VSAN or VF-ID	Description
000h	Cannot be used as Virtual Fabric Identifier
001h(1) to EFFh(3839)	This VSAN range is available for user configuration
F00h(3840) to FEEh(4078)	Reserved VSANs and they are not available for user configuration.
FEFh(4079)	EVFP isolated VSAN
FF0h(4080) to FFEh(4094)	Used for vendor-specific VSANs
FFFh	Cannot be used as Virtual Fabric Identifier

Note If the VF_ID of the F port and the N port do not match, then no tagged frames can be exchanged.

Figure 16-5 Default Allowed-Active VSAN Configuration

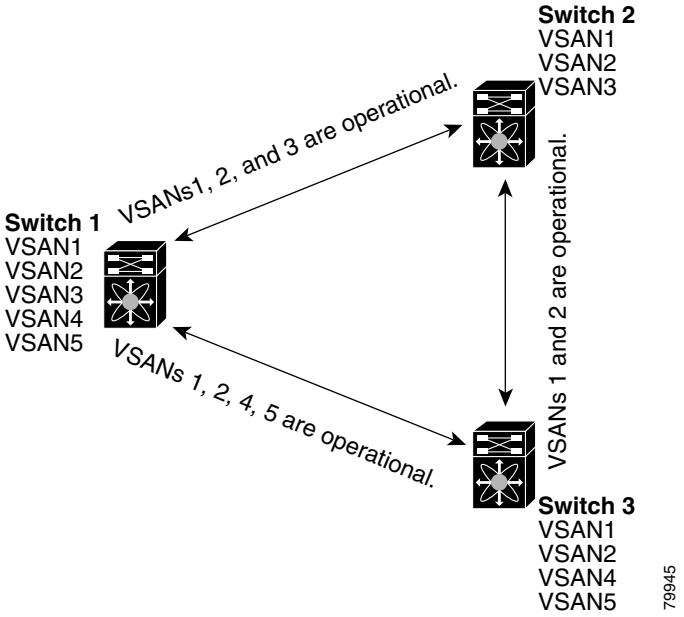

You can configure a select set of VSANs (from the allowed-active list) to control access to the VSANs specified in a trunking ISL.

Using Figure 16-5 as an example, you can configure the list of allowed VSANs on a per-interface basis (see Figure 16-6). For example, if VSANs 2 and 4 are removed from the allowed VSAN list of ISLs connecting to switch 1, the operational allowed list of VSANs for each ISL would be as follows:

- The ISL between switch 1 and switch 2 includes VSAN 1 and VSAN 3.
- The ISL between switch 2 and switch 3 includes VSAN 1 and VSAN 2.
- The ISL between switch 3 and switch 1 includes VSAN 1, 2, and 5.

Consequently, VSAN 2 can only be routed from switch 1 through switch 3 to switch 2.

Figure 16-6 Operational and Allowed VSAN Configuration

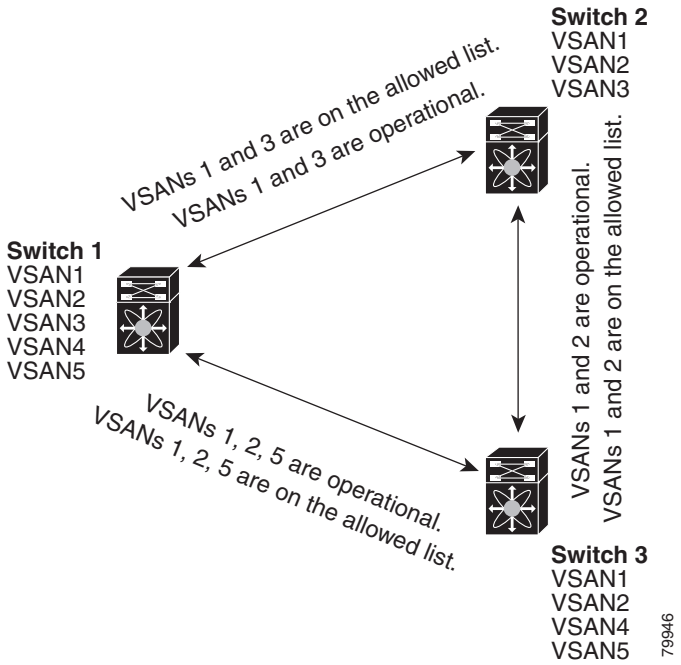

Configuring an Allowed-Active List of VSANs

To configure an allowed-active list of VSANs for an interface, follow these steps:

	Command	Purpose
Step 1	`switch# config t`	Enters configuration mode.
Step 2	`switch(config)# interface fc1/1` `switch(config-if)#`	Configures the specified interface.
Step 3	`switch(config-if)# switchport trunk allowed vsan 2-4`	Changes the allowed list for the specified VSANs.
	`switch(config-if)# switchport trunk allowed vsan add 5` `updated trunking membership`	Expands the specified VSAN (5) to the new allowed list.
	`switch(config-if)# no switchport trunk allowed vsan 2-4`	Deletes VSANs 2, 3, and 4.
	`switch(config-if)# no switchport trunk allowed vsan add 5`	Deletes the expanded allowed list.

Example F Port Trunking Configuration

This example shows how to configure trunking and bring up the TF-TNP link between anF port in the NPIV core switch, and an NP port in the NPV switch:

Step 1 Enable the F port trunking and channeling protocol on the MDS core switch:

```
switch(config)# feature fport-channel-trunk
```

Step 2 Enable NPIV on the MDS core switch:

```
switch(config)# feature npiv
```

Step 3 Step 3Create the PortChannel on the MDS core switch:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport mode F
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel mode active
switch(config-if)# exit
```

Step 4 Configure the PortChannel member interfaces on the MDS core switch in dedicated mode::

```
switch(config)# interface fc1/4-6
switch(config-if)# shut
switch(config-if)# switchport mode F
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Step 5 Create the PortChannel in dedicated mode on the NPV switch:

```
switch(config)# interface port-channel 2
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# no shut
switch(config-if)# exit
```

Step 6 Configure the PortChannel member interfaces on the NPV switch in dedicated mode:

```
switch(config)# interface fc3/1-3
switch(config-if)# shut
switch(config-if)# switchport mode NP
switch(config-if)# switchport speed 4000
switch(config-if)# switchport rate-mode dedicated
switch(config-if)# switchport trunk mode on
switch(config-if)# channel-group 2
switch(config-if)# no shut
switch(config-if)# exit
```

Displaying Trunking Information

The **show interface** command is invoked from the EXEC mode and displays trunking configurations for a TE port. Without any arguments, this command displays the information for all of the configured interfaces in the switch. See Examples 16-1 to 16-3.

Example 16-1 Displays a Trunked Fibre Channel Interface

```
switch# show interface fc1/13
fc1/13 is trunking
    Hardware is Fibre Channel
    Port WWN is 20:0d:00:05:30:00:58:1e
    Peer port WWN is 20:0d:00:05:30:00:59:1e
    Admin port mode is auto, trunk mode is on
    Port mode is TE
    Port vsan is 1
    Speed is 2 Gbps
```

```
    Receive B2B Credit is 255
    Beacon is turned off
    Trunk vsans (admin allowed and active) (1)
    Trunk vsans (up)                       (1)
    Trunk vsans (isolated)                 ()
    Trunk vsans (initializing)             ()
    5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
    5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
      233996 frames input, 14154208 bytes, 0 discards
        0 CRC, 0 unknown class
        0 too long, 0 too short
      236 frames output, 13818044 bytes, 0 discards
      11 input OLS, 12 LRR, 10 NOS, 28 loop inits
      34 output OLS, 19 LRR, 17 NOS, 12 loop inits
```

Example 16-2 Displays the Trunking Protocol

```
switch# show trunk protocol
Trunk protocol is enabled
```

Example 16-3 Displays Per VSAN Information on Trunk Ports

```
switch# show interface trunk vsan 1-1000
fc3/1 is not trunking
...
fc3/7 is trunking
    Vsan 1000 is down (Isolation due to vsan not configured on peer)
...
fc3/10 is trunking
    Vsan 1 is up, FCID is 0x760001
    Vsan 2 is up, FCID is 0x6f0001

fc3/11 is trunking
    Belongs to port-channel 6
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
...
port-channel 6 is trunking
    Vsan 1 is up, FCID is 0xef0000
    Vsan 2 is up, FCID is 0xef0000
```

Default Settings

Table 16-4 lists the default settings for trunking parameters.

Table 16-4 Default Trunk Configuration Parameters

Parameters	Default
Switch port trunk mode	ON on non-NPV and MDS core switches. OFF on NPV switches.
Allowed VSAN list	1 to 4093 user-defined VSAN IDs.
Allowed VF-ID list	1 to 4093 user-defined VF-IDs.
Trunking protocol on E ports	Enabled.
Trunking protocol on F ports	Disabled.