

7.12 Virtual Domain (equivalent to Distributed Switch/FCF)

7.12.1 Virtual Domain Switch_Name

A unique Switch_Name shall be associated with each Virtual Domain and this Name_Identifier shall persist (e.g., across failover) as the Switch_Name for the Virtual Domain.

7.12.2 Virtual Domain_ID (VD_ID) acquisition

7.12.2.1 Overview

The Domain_ID for the Virtual Domain is an administratively configured value or requested from the fabric by the Principal Controlling FCF. The default value for the Virtual Domain_ID is 00h.

7.12.2.2 Administratively configured VD_ID

If the VD_ID is administratively configured (i.e., VD_ID is not equal to 00h), then after obtaining a Domain_ID from the fabric, the Principal Controlling FCF shall transmit an RDI SW_ILS request (see FC-SW-6) with the Requested Switch_Name field set to the Virtual Domain Switch_Name and the Requested Domain_ID field set to the administratively configured VD_ID value.

If Insistent Domain_ID is enabled and the response to the RDI SW_ILS is an SW_ACC and the Granted Domain_ID field is set to the administratively configured VD_ID value, then the Virtual Domain is instantiated and the Principal Controlling FCF shall transmit a VDD SW_ILS request (see 7.12.3.3) to each peer Controlling FCF.

If Insistent Domain_ID is enabled and the response to the RDI SW_ILS is an SW_ACC and the Granted Domain_ID field is not set to the administratively configured VD_ID value, then the Virtual Domain is not instantiated.

If Insistent Domain_ID is not enabled and the response to the RDI SW_ILS is an SW_ACC, then the Virtual Domain is instantiated and the Principal Controlling FCF shall transmit a VDD SW_ILS request (see 7.12.3.3) to each peer Controlling FCF.

If the response to the RDI SW_ILS is an SW_RJT, then the Virtual Domain is not instantiated.

7.12.2.3 Non-administratively configured VD_ID

If the VD_ID is not administratively configured (i.e., VD_ID is equal to 00h), then after obtaining a Domain_ID from the fabric, the Principal Controlling FCF shall transmit an RDI SW_ILS request with the Requested Switch_Name field set to the Virtual Domain Switch_Name and the Requested Domain_ID field set to zero.

If the response to the RDI SW_ILS request is an SW_ACC, then the Virtual Domain is instantiated and the Principal Controlling FCF shall transmit a VDD SW_ILS request to each peer Controlling FCF.

If the response to the RDI SW_ILS is an SW_RJT, then the Virtual Domain is not instantiated.

Editor's Note: How is/does allocated Domain_ID return to the pool of available Domain_IDs if Virtual Domain is not instantiated?

7.12.2.4 Virtual Domain instantiation

Upon successful completion of the VDD SW_ILS requests as specified in 7.12.2.2 and 7.12.2.3, the Principal Controlling FCF shall transmit an LSU SW_ILS with a link state record describing the Virtual Domain.

In addition, each Controlling FCF shall list the Virtual Domain_ID as a directly attached Domain_ID in their link state record. An example FSPF topology is specified in figure 2, where C is the Virtual Domain_ID and A and B are the Domain_IDs of the two Controlling FCFs.

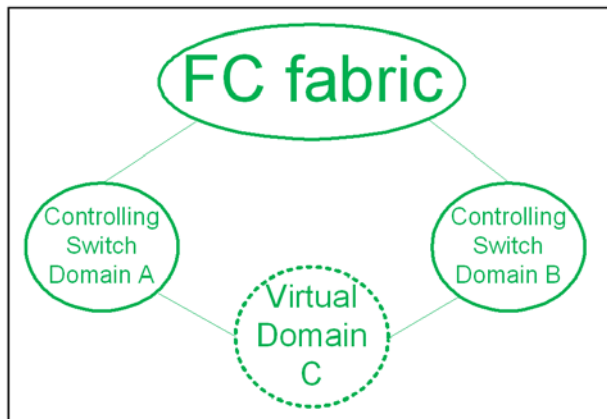


Figure 2 – Virtual Domain FSPF topology example

7.12.3 Virtual Domain SW_ILSs

7.12.3.1 Overview

All Virtual Domain SW_ILS requests and responses have a common format in the first two words of the payload as specified in table 53.

Table 53 – Virtual Domain SW_ILS common payload format

Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	Command code										R	A	Reserved										Subcommand code									
1	Descriptor list length																															

Command code: shall be set to the command code for the Virtual Domain SW_ILS as specified in table 54.

R: shall be set to zero for a Virtual Domain SW_ILS request and shall be set to one for a Virtual Domain SW_ILS response.

A: if the R bit is set to one, then the A bit shall be set to zero for a Virtual Domain SW_ILS Accept or set to one for a Virtual Domain SW_ILS Reject. If the R bit is set to zero, then the A bit shall be ignored.

Subcommand code: shall be set to the subcommand code for the Virtual Domain SW_ILS as specified in table 54.

Descriptor list length: shall be set to the length in words of all descriptor(s) that follow. A receiver shall ignore any data beyond the length specified in this field.

Table 54 specifies the Virtual Domain SW_ILS command codes, subcommand codes, expected descriptors and order.

Table 54 – Virtual Domain SW_ILSs and descriptor order

Virtual Domain SW_ILS	Command code/ subcommand code ^a	Expected Descriptors and Order
Virtual Domain Distribution (VDD)	XXh/01h	1) Header 2) Virtual Domain_ID 3) Name_Identifier(s)
Virtual Domain Heartbeat (VDHB)	XXh/02h	1) Header
Virtual Domain Build (VDB)	XXh/03h	1) Header 2) Domain_ID(s)
VN_Port Reachability Notification	XXh/01h	1) Header 2) Name_Identifier 3) FLOGI/NPIV FDISC
VN_Port Unreachability Notification	XXh/02h	1) Header 2) Name_Identifier 3) Address Identifier
FDCF/FCF Reachability Notification	XXh/03h	1) Header 2) Name_Identifier 3) Link Cost
FDCF/FCF Unreachability Notification	XXh/04h	1) Header 2) Name_Identifier
Address Identifier Route Distribution	XXh/05h	1) Header 2) Controlling Switch/FCF Path(s) 3) Address Identifier Range(s) 4) Reachable Domain_ID(s)
Address Identifier and Zoning ACL Distribution	XXh/06h	1) Header 2) Flags 3) Address Identifier 4) Name_Identifier 5) Name_Identifier 6) FLOGI/NPIV FDISC 7) Peering Entry(s)
Active Zoning ACL Distribution	XXh/07h	1) Header 2) Peering Entry(s)
Distributed Switch/FCF Membership Distribution	XXh/08h	1) Header 2) Name_Identifier(s) 3) Integrity Check
a)		

7.12.3.2 Virtual Domain SW_ILS descriptors

7.12.3.2.1 Overview

Virtual Domain SW_ILS descriptors are specified in table 55.

Table 55 – Virtual Domain SW_ILS descriptors

Type	Descriptor	Reference
0	Reserved	
1	Header	7.12.3.2.2
2	Domain_ID	7.12.3.2.3
3	Name_Identifier	7.12.3.2.4
4	Address identifier	7.12.3.2.5
10	Reject	7.12.3.2.6

7.12.3.2.2 Header descriptor

The Header descriptor is used in operations as specified in table 54.

The Header descriptor format shall be as specified in table 56.

Table 56 – Header descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
0	Type = 01h																Length = 05h															
1	(MSB)																															
2	Originating Switch_Name																(LSB)															
3	(MSB)																															
4	Destination Switch_Name																(LSB)															

NOTE 37 – Changed Type to 16 bits per WG discussion. Also changed Length to 16 bits.

NOTE 38 – Add Reserved word 1?

Originating Switch_Name: the Switch_Name of the device that originated the request.

Destination Switch_Name: the Switch_Name of the destination device for the request.

7.12.3.2.3 Domain_ID descriptor

The Domain_ID descriptor is used as specified in table 54.

The Domain_ID descriptor format shall be as specified in table 57.

Table 57 – Domain_ID descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0		
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type = 02h																Length = 02h															
1	Reserved																						Domain_ID									

Domain_ID: shall be set to a valid Domain_ID value (see FC-SW-6).

7.12.3.2.4 Name_Identifier descriptor

The Name_Identifier descriptor is used in operations as specified in table 54.

The Name_Identifier descriptor format shall be as specified in table 58.

Table 58 – Name_Identifier descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0		
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type = 03h																Length = 04h															
1	Reserved																						Name ID Type									
2	(MSB)																Name_Identifier															
3	Name_Identifier																(LSB)															

Name ID Type: the type of device the Name_Identifier is associated with (i.e., Controlling FCF, N_Port, FCDF/FDF, Virtual Domain).

Table 59 – Name identifier types

Value	Description
0	Reserved
1	Controlling FCF
2	FDF
3	Virtual Domain
4	N_Port

Name_Identifier: the Name_Identifier (see FC-FS-3) associated with the device.

7.12.3.2.5 Address identifier descriptor

The Address identifier descriptor is used as specified in table 54.

Request payload: The format of the VDHB request payload is shown in table 64.

Table 64 – VDHB request payload

Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0-1	Common																															
2-6	Header descriptor																															

7.12.3.5 Virtual Domain Build (VDB) SW_ILS

The Virtual Domain Build SW_ILS is used by Controlling FCFs to select a Principal Controlling FCF for the Virtual Domain.

Use of the VDB SW_ILS for redundancy purposes is described in 7.12.4.

Protocol:

Virtual Domain Build (VDB) request Sequence

Addressing: The S_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the originating Controlling FCF. The D_ID field shall be set to FFFFFDh, indicating the Fabric Controller of the destination Controlling FCF.

Request payload: The format of the VDB request payload is shown in table 64.

Table 65 – VDB request payload

Word	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0-1	Common																															
2-6	Header descriptor																															
7-9	Principal Controlling FCF Switch_Name Name_Identifier descriptor																															
	Controlling FCF list																															
10-n	Controlling FCF Switch_Name Name_Identifier descriptor(s)																															

7.12.4 Virtual Domain redundancy protocol

7.12.4.1 Overview

To provide enhanced reliability, two or more Controlling FCFs may be connected together via point-to-points links (i.e., the Controlling FCF set) with the requirement that each Controlling FCF in the set is able to communicate with the same set of FDFs. An example configuration is shown in figure 3.

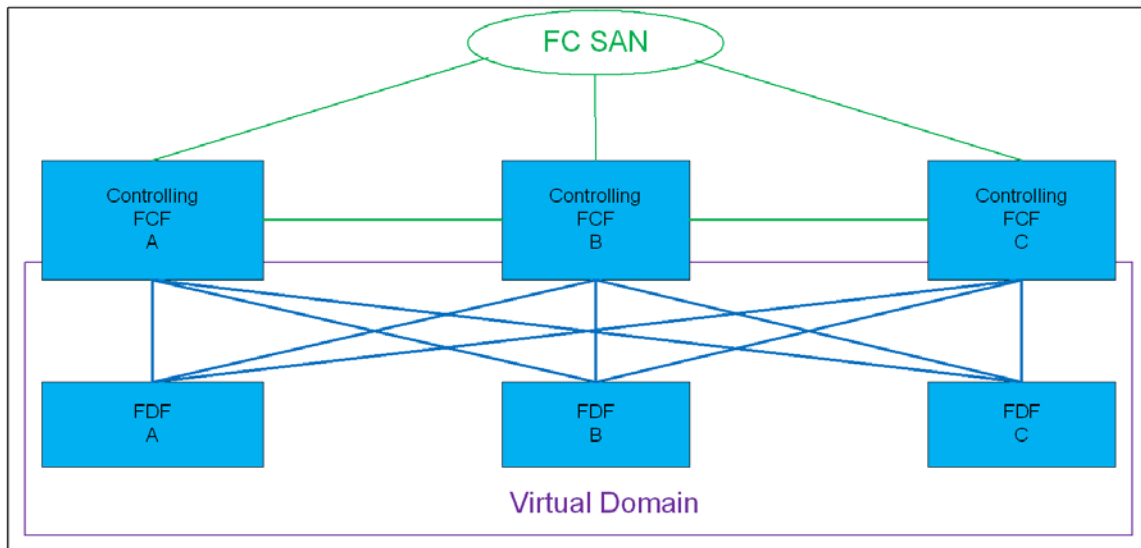


Figure 3 – Example redundant Virtual Domain configuration

In a redundant Virtual Domain configuration the Principal Controlling FCF selection process (see 7.12.4.2) to determine the Principal Controlling FCF for the Virtual Domain.

The Principal Controlling FCF for the Virtual Domain is responsible for:

- a) transmitting an RDI SW_ILS to obtain the Virtual Domain_ID for the Virtual Domain; and
- b) transmitting a VDD SW_ILS to each Controlling FCF in the Controlling FCF set.

7.12.4.2 Principal Controlling FCF selection

The Principal Controlling FCF selection state machine uses the following variable and timer:

Current_Principal: set to the highest known Principal Controlling FCF Switch_Name.

VD_S_Timer: Virtual Domain stability timer that expires upon VD_S_TOV.

The Principal Controlling FCF selection process state machine is specified in figure 4.

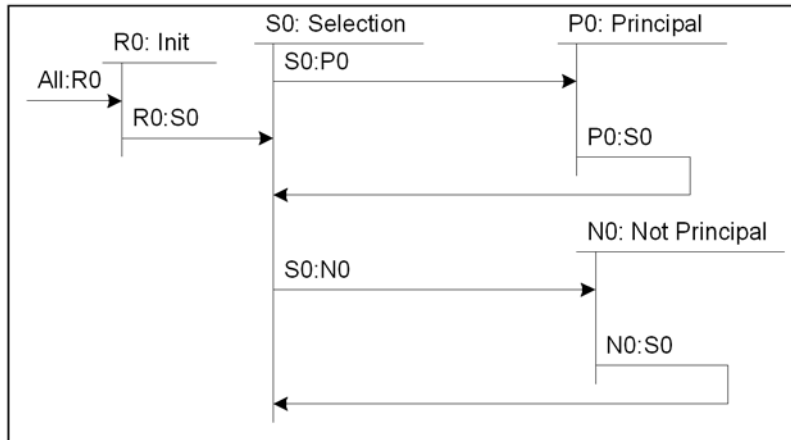


Figure 4 – Principal Controlling FCF selection state machine

Transition All:R0. This transition occurs when the redundancy protocol is disabled.

State R0:Init. In this state the Controlling FCF waits to start the selection process.

Transition R0:S0. This transition occurs when:

- a) the Controlling FCF has completed fabric configuration;
- b) the Controlling FCF set and FDF set has been configured; and
- c) the redundancy protocol is enabled.

State S0: Selection. In this state the Controlling FCF determines the Principal Controlling FCF. Upon entering this state the Controlling FCF shall:

- a) start the VD_S_Timer;
- b) set Current_Principal to the highest known Principal Controlling FCF Switch_Name; and
- c) transmit a VDB SW_ILS to each neighbor Controlling FCF with a higher Switch_Name.

While in this state:

- a) if the Controlling FCF receives a VDB SW_ILS with the Principal Controlling FCF Switch_Name less than or equal to the Controlling FCF Switch_Name, and Current_Principal is equal to the Controlling FCF Switch_Name, then it shall update the Controlling FCF list and transmit a VDB SW_ILS to each downstream Controlling FCF;
- b) if the Controlling FCF receives a VDB SW_ILS with the Principal Controlling FCF Switch_Name less than Current_Principal, but greater than or equal to the Controlling FCF Switch_Name, then the VDB SW_ILS shall be forwarded to each upstream Controlling FCF;
- c) if the Controlling FCF receives a VDB SW_ILS with the Principal Controlling FCF Switch_Name less than Current_Principal and less than the Controlling FCF Switch_Name, then the VDB SW_ILS shall be discarded; and
- d) if the Controlling FCF receives a VDB SW_ILS with the Principal Controlling FCF Switch_Name greater than or equal to Current_Principal, then it shall update the Controlling FCF list and transmit the VDB SW_ILS to each downstream Controlling FCF.

Editor's Note: in step b) if there are no new Controlling FCF(s) in the Controlling FCF list, then the VDB does not need to be forwarded.

Transition S0:P0. This transition occurs when VD_S_TOV expires and Current_Principal is equal to the Controlling Switch_Name.

Transition S0:N0. This transition occurs when VD_S_TOV expires and Current_Principal is not equal to the Controlling Switch_Name.

State P0: Principal. In this state the Controlling FCF shall operate as the Principal Controlling FCF.

Transition P0:S0. This transition occurs when the Controlling FCF receives a VDB SW_ILS.

State N0: Not Principal. In this state the Controlling FCF shall not operate as the Principal Controlling FCF.

Transition N0:S0. This transition occurs when the Controlling FCF receives a VDB SW_ILS.

An example Principal Controlling FCF selection process is shown in figure 5.




Figure 5 – Example Principal Controlling FCF selection process

Editor’s Note: Add text to describe figure 5.

7.12.5 Timers and constants

Virtual Domain timers and constants are specified in 66.

Table 66 – Virtual Domain timers and constants

Timer/Constant	Value	Description
VD_S_TOV	TBD	Virtual Domain stability timeout value
VD_ID	00h	Default Virtual Domain_ID value