

7.12 Controlling FCF

A Controlling FCF may support one or more Virtual Domains (see 7.13).

The Controlling FCF model is specified in figure 2, with the items in brackets being optional.

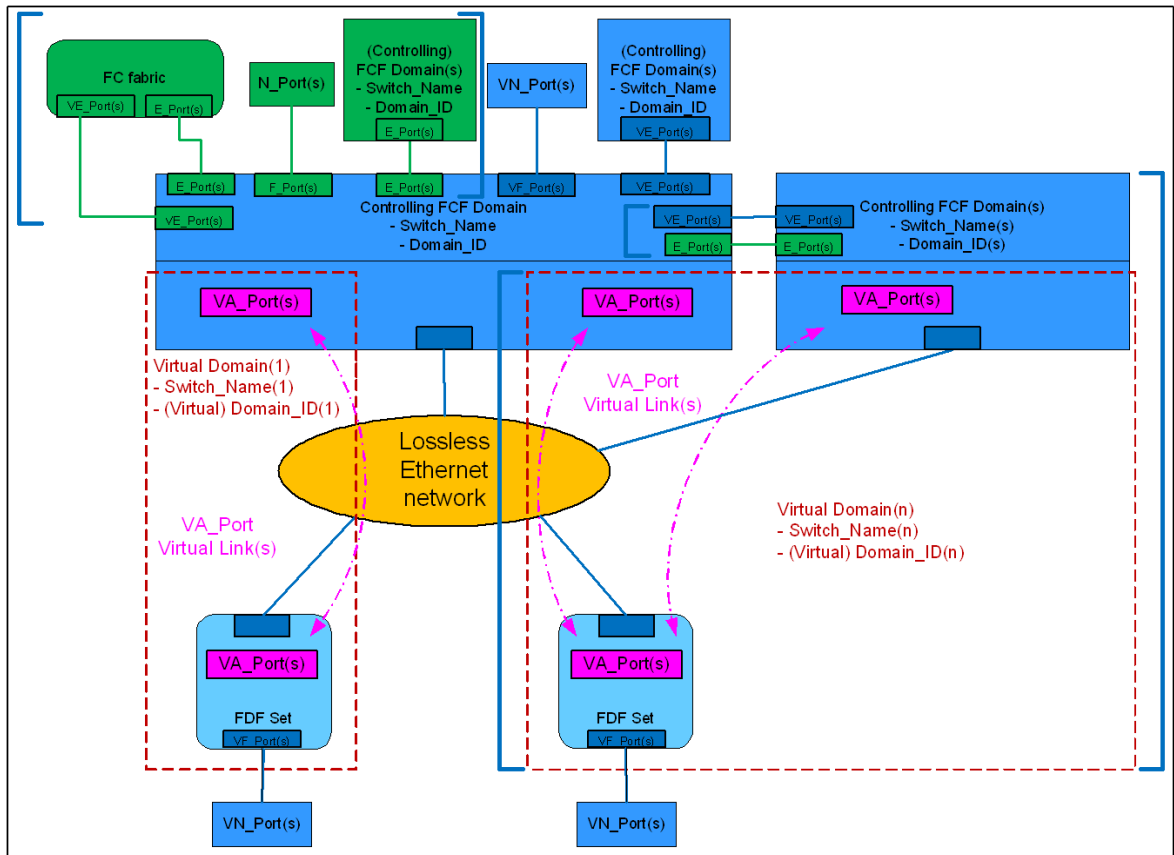


Figure 2 – Controlling FCF model

7.13 Virtual Domain model

7.13.1 Virtual Domain overview

A Virtual Domain consists of one or more Controlling FCFs (i.e., a Controlling FCF set) and one or more FDF(s) (i.e., an FDF set). A Virtual Domain shall have:

- a Virtual Domain Switch_Name (see 7.13.2);
- a Virtual Domain_ID (see 7.13.3); and
- an associated Principal Controlling FCF.

If a Virtual Domain consists of one Controlling FCF, then that Controlling FCF is the Principal Controlling FCF for the Virtual Domain.

If a Virtual Domain consists of two or more Controlling FCFs, then a Principal Controlling FCF shall be selected as specified in 7.13.6.2.

The Virtual Domain model is specified in figure 3, with the items in brackets being optional.

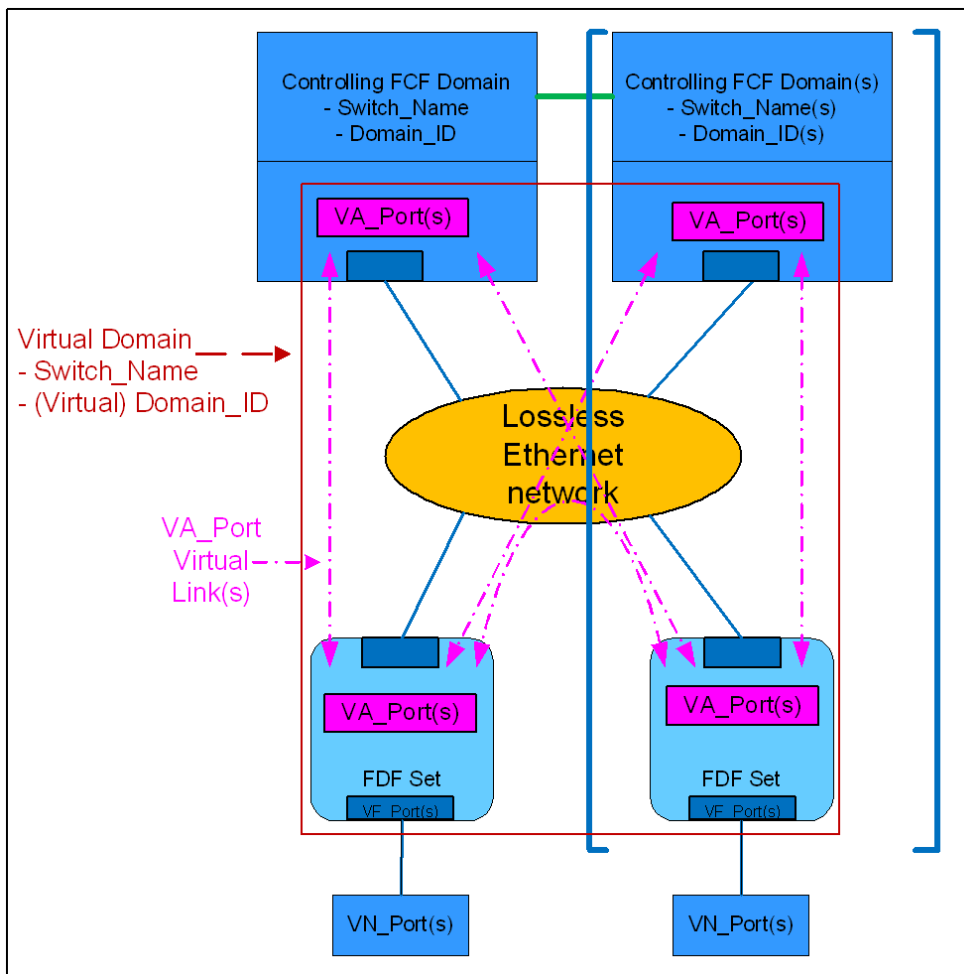


Figure 3 – Virtual Domain model

7.13.2 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.

7.13.3 Virtual Domain_ID (VD_ID)

7.13.3.1 Overview

Each Virtual Domain shall have one VD_ID. The VD_ID shall be assigned statically (see FC-SW-5 and 7.13.3.3.2) or dynamically (see 7.13.3.3).

The default value for the VD_ID is 00h.

7.13.3.2 Statically assigned VD_ID

If Domain_IDs are assigned statically (see FC-SW-5), then the administrator shall set the VD_ID to a value other than the default value (i.e., 00h), and the Principal Controlling FCF shall:

- 1) transmit a VDD SW_ILS request to each peer Controlling FCF, if any; and
- 2) instantiate the Virtual Domain (see 7.13.4).

7.13.3.3 Dynamically assigned VD_ID

7.13.3.3.1 Overview

If Domain_IDs are assigned dynamically (see FC-SW-5), then the VD_ID shall be set to an administratively configured value (see 7.13.3.3.2) or a non-administratively configured value (see 7.13.3.3.3).

If the VD_ID is set to an administratively configured value, then the Insistent Domain_ID feature may be enabled to allow the Virtual Domain to be instantiated only if the requested Domain_ID is granted.

7.13.3.3.2 Administratively configured VD_ID

If the VD_ID is administratively configured (i.e., the 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 Principal Controlling FCF shall transmit a VDD SW_ILS request (see 7.13.5.4) to each peer Controlling FCF, if any, and instantiate the Virtual Domain (see 7.13.4).

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 Principal Controlling FCF shall transmit a VDD SW_ILS request (see 7.13.5.4) to each peer Controlling FCF, if any, and instantiate the Virtual Domain.

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

7.13.3.3.3 Non-administratively configured VD_ID

If the VD_ID is not administratively configured (i.e., the 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 Principal Controlling FCF shall transmit a VDD SW_ILS request to each peer Controlling FCF, if any, and instantiate the Virtual Domain.

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

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

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

Table 58 – Virtual Domain SW_ILSs and descriptor order

Virtual Domain SW_ILS	Command code/ subcommand code ^a	Expected Descriptors and Order
Build Virtual Domain (BVD)	XXh/01h	1) Header
Virtual Domain Distribution (VDD)	XXh/02h	1) Header 2) Virtual Domain_ID 3) Name_Identifier(s)
Virtual Domain Heartbeat (VDHB)	XXh/03h	1) Header
Virtual Domain Build (VDB)	XXh/04h	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)
a)		

Table 58 – Virtual Domain SW_ILSs and descriptor order (Continued)

Virtual Domain SW_ILS	Command code/ subcommand code ^a	Expected Descriptors and Order
Distributed Switch/FCF Membership Distribution	XXh/08h	1) Header 2) Name_Identifier(s) 3) Integrity Check
a)		

7.13.5.2 Virtual Domain SW_ILS descriptors

7.13.5.2.1 Overview

Virtual Domain SW_ILS descriptors are specified in table 59.

Table 59 – Virtual Domain SW_ILS descriptors

Type	Descriptor	Reference
0	Reserved	
1	Header	7.13.5.2.2
2	Domain_ID	7.13.5.2.3
3	Name_Identifier	7.13.5.2.4
4	Address identifier	7.13.5.2.5
5	FCF priority	7.13.5.2.5
10	Reject	7.13.5.2.7

7.13.5.2.2 Header descriptor

The Header descriptor is used in operations as specified in table 58.

The Header descriptor format shall be as specified in table 60.

Table 60 – 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.

7.13.5.2.3 Domain_ID descriptor

The Domain_ID descriptor is used as specified in table 58.

The Domain_ID descriptor format shall be as specified in table 61.

Table 61 – Domain_ID descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	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).

7.13.5.2.4 Name_Identifier descriptor

The Name_Identifier descriptor is used in operations as specified in table 58.

The Name_Identifier descriptor format shall be as specified in table 62.

Table 62 – Name_Identifier descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	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)																															
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 63 – 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.

Protocol:

Virtual Domain Distribution (VDD) request Sequence
 Virtual Domain Distribution (VDD) response Sequence

Addressing: The S_ID shall be set to FFFCxxh designating the Domain Controller ID of the Controlling FCF that generates the VDD SW_ILS. The D_ID shall be set to FFFCyh to designate the Domain Controller ID of the recipient Controlling FCF.

Request payload: The format of the VDD request payload is shown in table 69.

Table 69 – VDD request payload

Word	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0-1	Common																															
2-6	Header descriptor																															
7	Virtual Domain_ID descriptor																															
8-10	Virtual Domain Switch_Name Name_Identifier descriptor																															
11-13	Principal Controlling FCF Switch_Name Name_Identifier descriptor																															
14-n	Controlling FCF Switch_Name Name_Identifier descriptor(s)																															

Response payload: The format of the VDD response payload is shown in table 70.

Table 70 – VDD response payload

Word	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0-1	Common																															
2	Reject descriptor (if R=1 and A=1)																															

7.13.5.5 Virtual Domain Heartbeat (VDHB) SW_ILS

The Virtual Domain Heartbeat SW_ILS is used by Controlling FCFs to monitor the state of peer Controlling FCF(s) in the Virtual Domain.

Use of the VDHB SW_ILS for redundancy purposes is described in 7.13.6.

Protocol:

Virtual Domain Heartbeat (VDHB) 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 VDHB request payload is shown in table 71.

Table 71 – 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																															

7.13.5.6 Exchange Virtual Domain Parameters (EVDP) SW_ILS

The Exchange Virtual Domain Parameters SW_ILS is used by Controlling FCFs to exchange Virtual Domain parameters and select a Principal Controlling FCF for the Virtual Domain.

Use of the EVDP SW_ILS for redundancy purposes is described in 7.13.6.

Protocol:

Exchange Virtual Domain Parameters (EVDP) 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 EVDP request payload is shown in table 72.

Table 72 – EVDP 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																															
7-8	FCF_Priority descriptor																															
9-11	Principal Controlling FCF Switch_Name Name_Identifier descriptor																															
Controlling FCF list																																
12-n	Controlling FCF Switch_Name Name_Identifier descriptor(s)																															

7.13.6 Controlling FCF redundancy protocol

7.13.6.1 Overview

To provide enhanced reliability, two or more Controlling FCFs may be connected together via point-to-points links with a requirement that each Controlling FCF in the FCF set is able to communicate with the same set of FDFs. An example redundant Controlling FCF configuration is shown in figure 5.

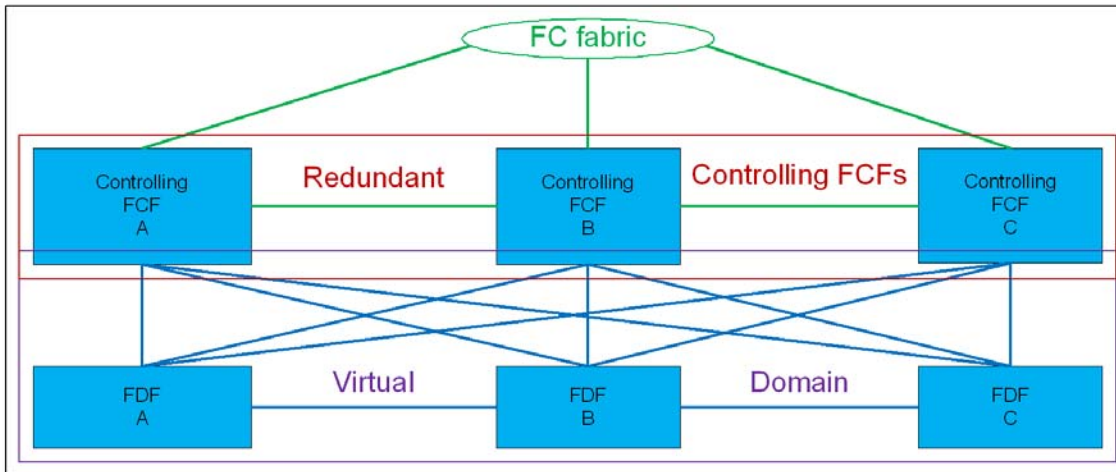


Figure 5 – Example redundant Controlling FCF configuration

In a redundant Controlling FCF configuration the Principal Controlling FCF selection process (see 7.13.6.2) is used 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.

7.13.6.2 Principal Controlling FCF selection

7.13.6.2.1 Overview

A Principal Controlling FCF is determined by configuration at a specific Controlling FCF (see 7.13.6.2.2) or by the Principal Controlling FCF selection process state machine (see 7.13.6.2.3).

7.13.6.2.2 Principal Controlling FCF by configuration

If the Principal Controlling FCF is determined by configuration, then the selected Principal Controlling FCF shall transmit a VDD SW_ILS declaring itself the Principal Controlling FCF to each Controlling FCF in the FCF set.

7.13.6.2.3 Principal Controlling FCF selection process state machine

The Principal Controlling FCF selection process state machine is initiated by a Controlling FCF transmitting a Build Virtual Domain (BVD) SW_ILS on all Controlling FCF virtual links. The BVD SW_ILS is then flooded to all Controlling FCFs to signal the transmission of the Exchange Virtual Domain Parameters (EVDP) SW_ILS.

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

Current_Principal: set to the lowest known FCF_Priority||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 6.

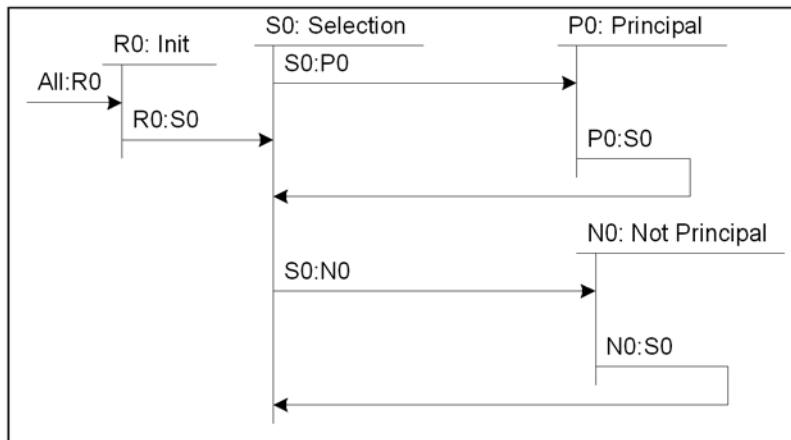


Figure 6 – Principal Controlling FCF selection state machine

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:

- the Controlling FCF has completed fabric configuration;
- the Controlling FCF set and FDF set have been configured; and
- the redundancy protocol is enabled (e.g., the Controlling FCF set contains more than one Controlling FCF) or a BVD SW_ILS is received.

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

- start the VD_S_Timer;
- set Current_Principal to the lowest known FCF_Priority||Principal Controlling FCF Switch_Name; and
- transmit a EVDP SW_ILS to the neighbor Controlling FCF with the lowest FCF_Priority||Switch_Name, if any.

Editor's Note: Can FCF_Priority be provided at ELP time?

While in this state:

- if the Controlling FCF receives a EVDP SW_ILS with the FCF_Priority||Principal Controlling FCF Switch_Name less than or equal to Current_Principal, then it shall update the Controlling FCF list and transmit the EVDP SW_ILS to each downstream Controlling FCF;
- if the Controlling FCF receives a EVDP SW_ILS with the FCF_Priority||Principal Controlling FCF Switch_Name greater than Current_Principal and Current_Principal does not equal

FCF_Priority||Controlling FCF Switch_Name, then the EVDP SW_ILS shall be forwarded to each upstream Controlling FCF; and

- c) If the Controlling FCF receives a EVDP SW_ILS with the FCF_Priority||Principal Controlling FCF Switch_Name greater than Current_Principal and Current_Principal is equal to FCF_Priority||Controlling FCF Switch_Name, then the Controlling FCF shall update the Controlling FCF list and transmit a EVDP SW_ILS to each downstream Controlling FCF.

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

Transition S0:N0. This transition occurs when VD_S_TOV expires and Current_Principal is not equal to FCF_Priority||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 EVDP 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 EVDP SW_ILS.

An example Principal Controlling FCF selection process is shown in figure 7.

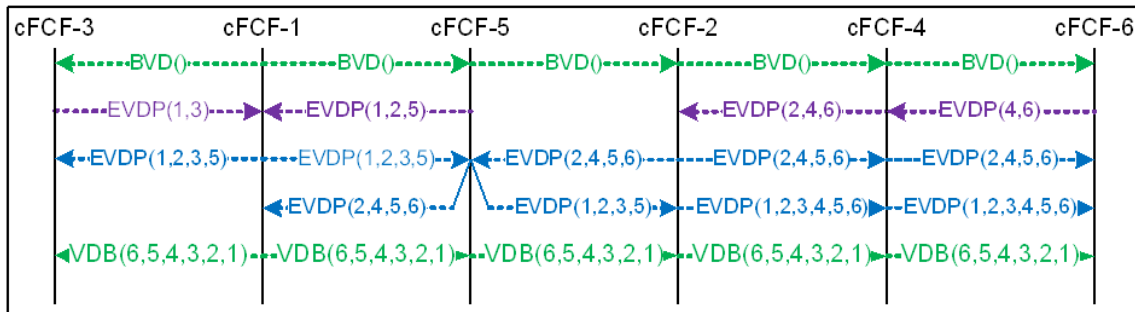


Figure 7 – Example Principal Controlling FCF selection process

7.13.7 Timers and constants

Virtual Domain timers and constants are specified in table 73.

Table 73 – 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