Claudio DeSanti (Cisco), Erik Smith (EMC),
Craig Carlson (QLogic), Ed McGlaughlin (QLogic)
Pat Thaler (Broadcom), Fred Knight (NetApp),
Roger Hathorn (IBM)

1 VA_Port Protocols

NOTE 1 – This document is written in terms of native Fibre Channel operations. The same operations apply also to FCoE, with physical FC links replaced by FCoE Virtual Links. Specific references to a protocol are provided when needed. See T11/11-223v1 for the terminology used in this document.

1.1 VA_Port SW_ILSs

1.1.1 Overview

The VA_Port SW_ILSs have the same high-order byte in their command code, denoted here as XXh. Table 1 shows the VA_Port SW_ILSs command codes.

Table 1 – VA_Port SW_ILSs Command Codes

Encoded Value	Description	Abbreviation
XX00 0001h	VN_Port Reachability Notification	VNRN
XX00 0002h	VN_Port Unreachability Notification	VNUN
XX00 0003h	FCDF Reachability Notification	FDRN
XX00 0004h	FCDF Unreachability Notification	FDUN
XX00 0005h	N_Port_ID Route Distribution	NPRD
XX00 0006h	N_Port_ID and Zoning ACL Distribution	NPZD
XX00 0007h	Active Zoning ACL Distribution	AZAD
XX00 0008h	Distributed Switch Membership Distribution	DFMD

The VA_Port SW_ILSs are used to exchange information between Controlling Switches and FCDFs (i.e., they are not used to exchange information between FCDFs). When a Distributed Switch includes cascaded FCDFs, the intermediate FCDFs relay the SW_ILSs as shown in figure 1. To facili-

tate this, all VA_Port SW_ILSs include the originating and destination FCDF or Controlling Switch Switch_Names in the first two fields of their payload.

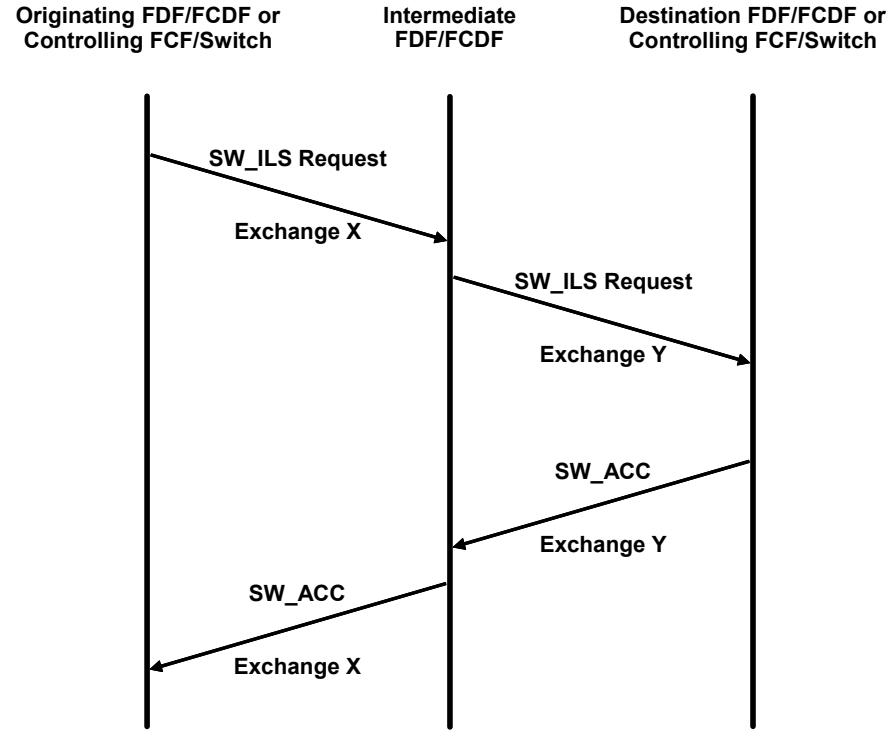

Figure 1 – VA_Port SW_ILS Relay

1.1.2 VN_Port Reachability Notification (VNRN)

The VN_Port Reachability Notification SW_ILS is used by an FCDF to communicate to the Primary Controlling Switch that a VN_Port is attempting Fabric login through an FLOGI Request or a NPIV FDISC Request. If the FCDF does not have an ASL with the Primary Controlling Switch, the VNRN SW_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs.

VNRN Request Sequence

Addressing: the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port.

Payload: the format of the VNRN Request Sequence Payload is shown in table 2.

Table 2 – VNRN Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0001h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8
F_Port_Name	8
Physical Port Number	4
FLOGI/NPIV FDISC Parameters	116

Originating FCDF Switch_Name: contains the Switch_Name of the requesting FCDF.

Destination Controlling Switch Switch_Name: contains the Switch_Name of the destination Controlling Switch.

F_Port_Name: contains the F_Port_Name of the VF_Port to which the requesting VN_Port is being associated.

Physical Port Number: contains the physical port number where the FLOGI or NPIV FDISC Request has been received.

FLOGI/NPIV FDISC Parameters: contains the payload of the received FLOGI or NPIV FDISC (see FC-LS-2).

VNRN Reply Sequence

SW_RJT: SW_RJT indicates the rejection of the VNRN Request Sequence. As a result, a (FIP) FLOGI LS_RJT or a (FIP) NPIV FDISC LS_RJT is sent as response to the (FIP) FLOGI Request or (FIP) NPIV FDISC Request that caused the issuance of the VNRN Request.

SW_ACC: SW_ACC indicates the acceptance of the VNRN Request Sequence for processing. The format of the VNRN SW_ACC Payload is shown in table 3.

Table 3 – VNRN SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8

1.1.3 VN_Port Unreachability Notification (VNUN)

The VN_Port Unreachability Notification SW_ILS is used by an FCDF to communicate to the Primary Controlling Switch that one or more of its VN_Ports have been logged out. If the FCDF does not have an ASL with the Primary Controlling Switch, the VNUN SW_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs.

VNUN Request Sequence

Addressing: the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port.

Payload: the format of the VNUN Request Sequence Payload is shown in table 4.

Table 4 – VNUN Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0002h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8
Flags	1
Unreachable N_Port_ID	3
Unreachable N_Port_Name	8
F_Port_Name	8

Originating FCDF Switch_Name: contains the Switch_Name of the requesting FCDF.

Destination Controlling Switch Switch_Name: contains the Switch_Name of the destination Controlling Switch.

Flags: 8 flag bits. The following flag bits are defined:

Bit 8 .. 1: reserved.

Bit 0: indicates if only one VN_Port is unreachable or if all the VN_Ports associated to a VF_Port are unreachable. This flag is set to zero to indicate the only one VN_Port in unreachable and to one to indicate that all the VN_Ports associated to a VF_Port are unreachable.

Unreachable N_Port_ID: when bit 0 of the flag field is set to zero contains the N_Port_ID of the unreachable VN_Port. When bit 0 of the flag field is set to one contains 000000h.

Unreachable N_Port_Name: when bit 0 of the flag field is set to zero contains the N_Port_Name of the unreachable VN_Port. When bit 0 of the flag field is set to one contains 00000000 00000000h.

F_Port_Name: contains the F_Port_Name of the involved VF_Port.

VNUN Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the VNUN Request Sequence for processing. The format of the VNUN SW_ACC Payload is shown in table 5.

Table 5 – VNUN SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8

1.1.4 FCDF Reachability Notification (FDRN)

The FCDF Reachability Notification SW_ILS is used by an FCDF to communicate to the Primary Controlling Switch that it has instantiated an ASL with another FCDF or with the Secondary Controlling Switch. If the FCDF does not have an ASL with the Primary Controlling Switch, the FDRN SW_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs.

The FDRN SW_ILS is also used between Primary and Secondary Controlling Switch to keep their state synchronized.

FDRN Request Sequence

Addressing: when used between a FCDF and the Primary Controlling Switch the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port. When used between the two Controlling Switches the S_ID field shall be set to FFFFFDh, indicating the originating VE_Port, and the D_ID field shall be set to FFFFFDh, indicating the destination VE_Port.

Payload: the format of the FDRN Request Sequence Payload is shown in table 6.

Table 6 – FDRN Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0003h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8
Reachable FCDF or Controlling Switch Switch_Name	8
Reserved	2
Link Cost	2

Originating FCDF Switch_Name: contains the Switch_Name of the requesting FCDF.

Destination Controlling Switch Switch_Name: contains the Switch_Name of the destination Controlling Switch.

Reachable FCDF or Controlling Switch Switch_Name: contains the Switch_Name of the adjacent entity with which the ASL has been instantiated.

Link Cost: contains the cost of the instantiated ASL.

FDRN Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the FDRN Request Sequence for processing. The format of the FDRN SW_ACC Payload is shown in table 7.

Table 7 – FDRN SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8

1.1.5 FCDF Unreachability Notification (FDUN)

The FCDF Unreachability Notification SW_ILS is used by an FCDF to communicate to the Primary Controlling Switch that it has deinstantiated an ASL with another FCDF or with the Secondary Controlling Switch. If the FCDF does not have an ASL with the Primary Controlling Switch, the FDUN SW_ILS is relayed to the Primary Controlling Switch by the intermediate FCDFs.

The FDUN SW_ILS is also used between Primary and Secondary Controlling Switch to keep their state synchronized.

FDUN Request Sequence

Addressing: when used between a FCDF and the Primary Controlling Switch the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port. When used between the two Controlling Switches the S_ID field shall be set to FFFFFDh, indicating the originating VE_Port, and the D_ID field shall be set to FFFFFDh, indicating the destination VE_Port.

Payload: the format of the FDUN Request Sequence Payload is shown in table 8.

Table 8 – FDUN Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0004h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8
Unreachable FCDF or Controlling Switch Switch_Name	8

Originating FCDF Switch_Name: contains the Switch_Name of the requesting FCDF.

Destination Controlling Switch Switch_Name: contains the Switch_Name of the destination Controlling Switch.

Unreachable FCDF or Controlling Switch Switch_Name: contains the Switch_Name of the adjacent entity with which the ASL has been deinstantiated.

FDUN Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the FDUN Request Sequence for processing. The format of the FDUN SW_ACC Payload is shown in table 9.

Table 9 – FDUN SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8

1.1.6 N_Port_ID Route Distribution (NPRD)

The N_Port_ID Route Distribution SW_ILS is used by the Primary Controlling Switch to communicate to an FCDF the N_Port_ID routing information for the Distributed Switch. If the Primary Controlling Switch does not have an ASL with the destination FCDF, the NPRD SW_ILS is relayed to the destination FCDF by the intermediate FCDFs.

NPRD Request Sequence

Addressing: the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port.

Payload: the format of the NPRD Request Sequence Payload is shown in table 10.

Table 10 – NPRD Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0005h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8
Primary Controlling Switch Switch_Name	8
Number of Paths to the Primary Controlling Switch (j)	4
Next-hop Switch_Name #1	8
Path #1 cost	4
Next-hop Switch_Name #2	8
Path #2 cost	4
...	
Next-hop Switch_Name #j	8
Path #j cost	4
Secondary Controlling FCF Switch_Name	8
Number of Paths to the Secondary Controlling Switch (g)	4
Next-hop Switch_Name #1	8
Path #1 cost	4
Next-hop Switch_Name #2	8
Path #2 cost	4
...	
Next-hop Switch_Name #g	8
Path #g cost	4
Number of N_Port_ID Range Entries (p)	4
N_Port_ID Range Entry #1	see table 11
N_Port_ID Range Entry #2	see table 11
…	
N_Port_ID Range Entry #p	see table 11
Number of Reachable Domain_ID Entries (r)	4
Reachable Domain_ID Entry #1	see table 12
Reachable Domain_ID Entry #2	see table 12
…	
Reachable Domain_ID Entry #r	see table 12

Originating Controlling Switch Switch_Name: contains the Switch_Name of the requesting Controlling Switch.

Destination FCDF Switch_Name: contains the Switch_Name of the destination FCDF.

Primary Controlling Switch Switch_Name: contains the Switch_Name of the Primary Controlling Switch.

Number of Paths to the Primary Controlling Switch: contains the number of paths toward the Primary Controlling Switch. Each path that follows is expressed as the Switch_Name of the next-hop FCDF or Controlling Switch followed by its cost.

NOTE 2 – Paths toward the Primary Controlling Switch are fundamental for the operation of an FCDF. Specifying higher cost paths enables more redundancy, because if the lowest cost path toward the Primary Controlling Switch fails, a higher cost path may be used.

Secondary Controlling Switch Switch_Name: contains the Switch_Name of the Secondary Controlling Switch.

Number of Paths to the Secondary Controlling Switch: contains the number of paths toward the Secondary Controlling Switch. Each path that follows is expressed as the Switch_Name of the next-hop FCDF or Controlling Switch followed by its cost.

Number of N_Port_ID Range Entries: contains the number of N_Port_ID Range Entries that follow. The N_Port_ID Range Entry format is shown in table 11.

Table 11 – N_Port_ID Range Entry Format

Item	Size (bytes)
Destination FCDF Switch_Name	8
Number of Equal Cost Paths to the Destination FCDF (w)	4
Next-hop Switch_Name #1	8
Next-hop Switch_Name #2	8
...	
Next-hop Switch_Name #w	8
Number of N_Port_ID Ranges (q)	4
N_Port_ID Range #1	4
N_Port_ID Range #2	4
…	
N_Port_ID Range #q	4

Destination FCDF Switch_Name: contains the Switch_Name of the FCDF to which the subsequent next-hops and N_Port_ID Ranges refer.

Number of Equal Cost Paths to the Destination FCDF: contains the number of equal cost paths having the lowest cost toward the destination FCDF. Each path that follows is expressed as the Switch_Name of the next-hop FCDF or Controlling Switch.

Number of N_Port_ID Ranges: contains the number of N_Port_ID Range Entries that follow. The N_Port_ID Range is defined by an N_Port_ID in the least significant three bytes, and by the number of bits defining the range in the most significant byte (e.g., the range 020200h .. 02027Fh is expressed as ‘7 || 020200h’).

Number of Reachable Domain_ID Entries: contains the number of Reachable Domain_ID Entries that follow. The Reachable Domain_ID Entry format is shown in table 12.

Table 12 – Reachable Domain_ID Entry Format

Item	Size (bytes)
Destination Domain_ID	4
Number of Equal Cost Paths to the Destination Domain_ID (y)	4
Next-hop Switch_Name #1	8
Next-hop Switch_Name #2	8
...	
Next-hop Switch_Name #y	8

Destination Domain_ID: contains the reachable destination Domain_ID. The three most significant bytes of this field are reserved.

Number of Equal Cost Paths to the Destination Domain_ID: contains the number of equal cost paths having the lowest cost toward the destination Domain_ID. Each path that follows is expressed as the Switch_Name of the next-hop FCDF or Controlling Switch.

NPRD Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the NPRD Request Sequence for processing. The format of the NPRD SW_ACC Payload is shown in table 13.

Table 13 – NPRD SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8

1.1.7 N_Port_ID and Zoning ACL Distribution (NPZD)

The N_Port_ID and Zoning ACL Distribution SW_ILS is used by the Primary Controlling Switch to communicate to an FCDF and to the Secondary Controlling Switch the allocation of an N_Port_ID and its associated Zoning ACL information and/or the deallocation of one or more N_Port_IDs and their associated Zoning ACL information. Upon receiving an NPZD Request, an FCDF shall update its Zoning enforcement according to the received Zoning ACLs only for the listed Principal N_Port_IDs. If the Primary Controlling Switch does not have an ASL with the destination FCDF, the NPZD SW_ILS is relayed to the destination FCDF by the intermediate FCDFs.

NPZD Request Sequence

Addressing: when used between a FCDF and the Primary Controlling Switch the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port. When used between the two Controlling Switches the S_ID field shall be set to FFFFFDh, indicating the originating VE_Port, and the D_ID field shall be set to FFFFFDh, indicating the destination VE_Port.

Payload: the format of the NPZD Request Sequence Payload is shown in table 14.

Table 14 – NPZD Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0006h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF or Controlling Switch Switch_Name	8
Number of Allocation / Deallocation Entries (z)	4
Allocation / Deallocation Entry #1	see table 15
Deallocation Entry #2	see table 15
...	
Deallocation Entry #z	see table 15
Number of Peering Entries (h)	4
Peering Entry #1	see table 16
Peering Entry #2	see table 16
…	
Peering Entry #h	see table 16

Originating Controlling Switch Switch_Name: contains the Switch_Name of the requesting Controlling Switch.

Destination FCDF or Controlling Switch Switch_Name: contains the Switch_Name of the destination FCDF or Controlling Switch.

Number of Allocation / Deallocation Entries: contains the number of Allocation / Deallocation Entries that follow. Only one Allocation Entry may be present, multiple Deallocation Entries may be present. The Allocation / Deallocation Entry format is shown in table 15.

Table 15 – Allocation / Deallocation Entry Format

Item	Size (bytes)
Flags	4
Allocated / Deallocated N_Port_ID	4
N_Port_Name associated with the Allocated/Deallocated N_Port_ID	8
Switch_Name of the FCDF associated with the Allocated/Deallocated N_Port_ID	8
FLOGI / NPIV FDISC LS_ACC Parameters	116

Flags: 32 flag bits. The following flag bits are defined:

Bit 32 .. 2: reserved.

Bit 1: indicates if the FLOGI / NPIV FDISC LS_ACC Parameters field is present in the payload. The field is present when this flag is set to one and not present when this flag is set to zero. This

flag shall not be set to one when bit 0 indicates deallocation (i.e., the FLOGI / NPIV FDISC LS_ACC Parameters field may be present only when an N_Port_ID allocation is performed).

Bit 0: indicates if the operation is an allocation or a deallocation. This flag is set to zero to indicate allocation and to one to indicate deallocation.

Allocated / Deallocated N_Port_ID: contains the N_Port_ID that the Primary Controlling FCF allocated or deallocated in the least significant three bytes. The most significant byte is reserved.

N_Port_Name associated with the Allocated/Deallocated N_Port_ID: contains the N_Port_Name of the VN_Port for which an N_Port_ID is allocated or deallocated.

Switch_Name of the FCDF associated with the Allocated/Deallocated N_Port_ID: contains the Switch_Name of the FCDF associated with the VN_Port for which an N_Port_ID is allocated or deallocated.

FLOGI / NPIV FDISC LS_ACC Parameters: this field is present when bit 1 of the flags field is set to one. It contains the payload of the LS_ACC generated by the Primary Controlling FCF in response to the FLOGI or NPIV FDISC payload provided in the VNRN Request Sequence.

Number of Peering Entries: contains the number of Peering Entries that follow. The Peering Entry format is shown in table 16.

Table 16 – Peering Entry Format

Item	Size (bytes)
Principal N_Port_ID	4
Number of Allowed Peers (k)	4
Peer N_Port_ID #1	4
Peer N_Port_ID #2	4
…	
Peer N_Port_ID #q	4

Principal N_Port_ID: contains the N_Port_ID to which the subsequent Peer N_Port_IDs refer.

Number of Allowed Peers: contains the number of N_Port_IDs to which the Principal N_Port_ID is allowed to communicate.

Peer N_Port_ID: contains an N_Port_ID in the least significant three bytes and the most significant byte is reserved.

NPZD Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the NPZD Request Sequence for processing. The format of the NPZD SW_ACC Payload is shown in table 17.

Table 17 – NPZD SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8

1.1.8 Active Zoning ACL Distribution (AZAD)

The Active Zoning ACL Distribution SW_ILS is used by the Primary Controlling Switch to communicate to an FCDF new Zoning ACL information when a new Zone Set is activated in the fabric. Upon receiving an NPZD Request, an FCDF shall completely replace its Zoning enforcement according to the received Zoning ACLs. If the Primary Controlling Switch does not have an ASL with the destination FCDF, the AZAD SW_ILS is relayed to the destination FCDF by the intermediate FCDFs.

AZAD Request Sequence

Addressing: the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port.

Payload: the format of the AZAD Request Sequence Payload is shown in table 18.

Table 18 – AZAD Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0007h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8
Number of Peering Entries (m)	4
Peering Entry #1	see table 16
Peering Entry #2	see table 16
…	
Peering Entry #m	see table 16

Originating Controlling Switch Switch_Name: contains the Switch_Name of the requesting Controlling Switch.

Destination FCDF Switch_Name: contains the Switch_Name of the destination FCDF.

Number of Peering Entries: contains the number of Peering Entries that follow. The Peering Entry format is shown in table 16.

AZAD Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the AZAD Request Sequence for processing. The format of the AZAD SW_ACC Payload is shown in table 19.

Table 19 – AZAD SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8

1.1.9 Distributed Switch Membership Distribution (DFMD)

The Distributed Switch Membership Distribution SW_ILS is used by the Primary Controlling Switch to communicate to an FCDF the identities of the Primary and Secondary Controlling Switches and of all the FCDFs that compose the Distributed Switch. The DFMD payload may be integrity protected by a cryptographic hash; in this case the involved entities shall be provided with a shared key. If the Primary Controlling Switch does not have an ASL with the destination FCDF, the DFMD SW_ILS is relayed to the destination FCDF by the intermediate FCDFs.

DFMD Request Sequence

Addressing: the S_ID field shall be set to FFFFF9h, indicating the originating VA_Port, and the D_ID field shall be set to FFFFF9h, indicating the destination VA_Port.

Payload: the format of the DFMD Request Sequence Payload is shown in table 20.

Table 20 – DFMD Request Payload

Item	Size (bytes)
SW_ILS Code = XX00 0008h	4
Originating Controlling Switch Switch_Name	8
Destination FCDF Switch_Name	8
Primary Controlling Switch Switch_Name	8
Secondary Controlling Switch Switch_Name	8
Number of FCDFs (n)	4
FCDF Switch_Name #1	8
FCDF Switch_Name #2	8
…	
FCDF Switch_Name #n	8
Integrity Type	4
Integrity Check Value Length	4
Integrity Check Value	variable

Originating Controlling Switch Switch_Name: contains the Switch_Name of the requesting Controlling Switch.

Destination FCDF Switch_Name: contains the Switch_Name of the destination FCDF.

Primary Controlling Switch Switch_Name: contains the Switch_Name of the Primary Controlling Switch.

Secondary Controlling Switch Switch_Name: contains the Switch_Name of the Secondary Controlling Switch.

Number of FCDFs: contains the number of FCDF Switch_Names that follow. This list of FCDF Switch_Names is the FCDF Set of the Distributed Switch. If the number of FCDF Switch_Names is zero, then any FCDF is allowed in the Distributed Switch.

Integrity Type: indicates, in the least significant byte, the type of cryptographic integrity that protects the DFMD Payload. The defined values are:

00h: No integrity

01h: HMAC-SHA-256-128 integrity

02h .. FFh: Reserved

Integrity Check Value Length: contains the length expressed in bytes of the Integrity Check Value.

Integrity Check Value: contains the cryptographic hash of the DFMD payload computed using the shared key according to the specified Integrity Type.

DFMD Reply Sequence

SW_ACC: SW_ACC indicates the acceptance of the DFMD Request Sequence for processing. The format of the DFMD SW_ACC Payload is shown in table 21.

Table 21 – DFMD SW_ACC Payload

Item	Size (bytes)
SW_ILS Code = 0200 0000h	4
Originating FCDF Switch_Name	8
Destination Controlling Switch Switch_Name	8
Number of Physical Ports	4
RNID Specific Node-Identification Data	see FC-SB-4

Number of Physical Ports: contains the number of physical ports the FCDF has.

1.2 FCDF Handling of Well Known Addresses

N_Ports use Well Known Addresses (WKAs) and Domain Controller address identifiers to exchange information with the Fabric, either through ELSs or through the Common Transport protocol.

An FCDF supports VF_Ports, therefore it shall terminate FC frames destined to the F_Port Controller WKA. This implies local processing by the FCDF of the FLOGI, FDISC, LOGO, and RLS ELSs.

The handling of other WKAs and Domain Controllers address identifiers is performed by the Primary Controlling Switch, therefore an FCDF shall forward all FC frames having as D_ID the address identifiers listed in table 22 to the Primary Controlling Switch through a VA_Port. The NPRD SW_ILS provides to FCDFs the routing information needed to reach the Primary Controlling Switch.

Table 22 – Forwarded Domain Controller and Well Known Address Identifiers

Address Value	Description
FFFC01h .. FFFCFEh	Domain Controller Address Identifiers
FFFFF4h	Event Service WKA
FFFFF6h	Clock Synchronization Service WKA
FFFFF7h	Security Key Distribution Service WKA
FFFFFAh	Management Service WKA
FFFFFBh	Time Service WKA
FFFFFCh	Directory Service WKA
FFFFFDh	Fabric Controller WKA

The AISLs used for the redundancy protocol between the Primary and Secondary Controlling Switch are used as paths to reach the Primary Controlling Switch when an FCDF is connected to the Secondary Controlling Switch but not anymore to the Primary one. In order to do so, the Secondary Controlling Switch shall forward to the Primary Controlling Switch over the AISLs:

a) any FC frame destined to the address identifier FFFFF9h (i.e., the VA_Port Controller); and

b) any FC frame destined to the address identifiers shown in table 22 when they are received from a VA_Port.

1.3 Use of VA_Port Protocols

1.3.1 Distributed Switch Operations

In a Distributed Switch, the Primary Controlling Switch defines the routes for the FCDF topology and performs N_Port_ID allocations and deallocations for all its controlled FCDFs. The two Controlling Switches keep their state synchronized.

When becoming operational (i.e., when it has defined if it behaves as Primary or Secondary), a Controlling Switch instantiates ASLs with the FCDFs that are directly reachable and are part of its FCDF Set, as described in T11/11-223v1.

Upon instantiating an ASL with an FCDF, the Primary Controlling Switch shall initiate an FDRN Exchange (see 1.1.4) describing that link with the Secondary Controlling Switch to keep the state synchronized. Upon completion of this FDRN Exchange, the Primary Controlling Switch shall provide to that FCDF the Distributed Switch Membership information through a DFMD Exchange (see 1.1.9). At this point the instantiated ASL becomes part of the Distributed Switch internal topology (i.e., the set of ASLs internal to the Distributed Switch). The Primary Controlling Switch shall recompute the N_Port_ID routes and distribute them to each FCDF belonging to the Distributed Switch through NPRD Exchanges (see 1.1.6).

Upon deinstantiating an ASL with an FCDF, the Primary Controlling Switch shall initiate an FDUN Exchange (see 1.1.5) describing that disappeared link with the Secondary Controlling Switch to keep the state synchronized. Upon completion of this FDUN Exchange, the Primary Controlling Switch shall recompute the N_Port_ID routes and distribute them to each FCDF belonging to the Distributed Switch through NPRD Exchanges.

When becoming operational, an FCDF waits for a Controlling Switch or another FCDF to initiate an ELP Exchange with it, in order to set up a ASL. Upon completing the DFMD Exchange with the Primary Controlling Switch, the FCDF becomes able to initiate ELP Requests to instantiate other ASLs with other FCDFs. Upon completing the NPRD Exchange with the Primary Controlling Switch, an FCDF becomes able to set up proper forwarding tables to forward FC frames inside and outside the Distributed Switch. At this point the FCDF enables its ports for logins from Nodes; any FLOGI received on a FCDF port before this point is responded by the FCDF with a LS_RJT having reason code ‘Logical busy’ and reason code explanation ‘No additional explanation’.

Upon instantiating a ASL with another FCDF or with the Secondary Controlling Switch, an FCDF shall perform a FDRN Exchange with the Primary Controlling Switch to inform it of the new link. Upon completing a FDRN Exchange with an FCDF, the Primary Controlling Switch shall initiate another FDRN Exchange with the same parameters with the Secondary Controlling Switch to keep the state synchronized. After completing this FDRN Exchange the primary Controlling Switch shall provide to the newly reported FCDF the Distributed FCF Membership information through a DFMD Exchange (see 1.1.9), if that FCDF did not not already receive a DFMD Exchange in a previous step. At this point the instantiated ASL becomes part of the Distributed FCF internal topology (i.e., the set of ASLs internal to the Distributed Switch). Upon completion of this DFMD Exchange, the Primary Controlling Switch shall recompute the N_Port_ID routes and distribute them to each FCDF belonging to the Distributed Switch through NPRD Exchanges.

NOTE 3 – An ASL with the Secondary Controlling Switch may be instantiated before the ASL with the Primary Controlling Switch. The FCDF recognizes the Primary Controlling Switch because it is the one from which it receives the DFMD Request. In this case, the FCDF initiates with the Primary Controlling Switch the FDRN Exchange describing the link with the Secondary Controlling Switch upon completing the DFMD Exchange.

Upon deinstantiating an ASL with another FCDF or with the Secondary Controlling Switch, an FCDF shall perform a FDUN Exchange with the Primary Controlling Switch to inform it of the disappeared link. Upon completing a FDUN Exchange with an FCDF, the Primary Controlling Switch shall initiate another FDUN Exchange with the same parameters with the Secondary Controlling Switch to keep the state synchronized. Upon completion of this FDUN Exchange, the Primary Controlling Switch shall recompute the N_Port_ID routes and distribute them to each FCDF belonging to the Distributed Switch through NPRD Exchanges.

Upon receiving on a port a FLOGI Request or a NPIV FDISC Request from a Node, an FCDF shall send a VNRN Request (see 1.1.2) to the Primary Controlling Switch to inform it of the newly reachable VN_Port. If the Primary Controlling Switch rejects the VNRN Request, the FCDF shall also reject the FLOGI Request or NPIV FDISC Request. If the Primary Controlling Switch accepts the VNRN Request, it performs the following processing:

a) if the VNRN Request carried a FLOGI Request and that VN_Port was not already logged in or if the VNRN Request carried a NPIV FDISC Request, then the Primary Controlling Switch shall allocate to the newly reachable VN_Port an N_Port_ID from the Virtual Domain_ID; or

b) if the VNRN Request carried a FLOGI Request and that VN_Port was already logged in, then the Primary Controlling Switch shall implicitly log out that VN_Port and all the VN_Ports associated to the VF_Port that VN_Port was associated with and then allocate to that VN_Port an N_Port_ID from the Virtual Domain_ID.

The Primary Controlling Switch shall also recompute the Zoning ACLs for the affected N_Port_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling Switch shall distribute the Zoning ACLs and N_Port_ID allocation/deallocation information to the Secondary Controlling Switch and to each FCDF belonging to the Distributed Switch except the one from which the VNRN has been received through an appropriate NPZD Exchange (see 1.1.7). The NPZD Request sent to the Secondary Controlling Switch shall include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry and shall carry no Peering Entries. The NPZD Requests sent to the other FCDFs shall not include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry. Upon receiving all SW_ACC from the Secondary Controlling Switch and the other FCDFs or upon the expiration of a one second timeout, whichever comes first, the Primary Controlling Switch shall send an NPZD Request also to the FCDF from which the VNRN has been received. This NPZD Request shall include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry. Upon receiving this NPZD Request the FCDF that sent the VNRN Request shall accept the FLOGI Request or NPIV FDISC Request and complete the N_Port login.

When a VN_Port is logged out or when a VF_Port is deintantiated, an FCDF shall perform a VNUN Exchange (see 1.1.3) with the Primary Controlling Switch to inform it that the VN_Port is now unreachable or that all the VN_Ports associated with that VF_Port are unreachable. Upon completing a VNUN Exchange, the Primary Controlling Switch shall deallocate the N_Port_ID(s) assigned to the affected VN_Port(s), recompute the Zoning ACLs for the affected N_Port_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling Switch shall then distribute this information to the Secondary Controlling Switch and to each FCDF belonging to the Distributed Switch through NPZD Requests indicating N_Port_ID(s) deallocation (see 1.1.7).

When a new Zone Set is activated in the Fabric, the Primary Controlling Switch shall recompute the Zoning ACLs for all N_Port_IDs allocated in the Virtual Domain_ID and distribute them to the FCDFs of the Distributed Switch through AZAD Exchanges (see 1.1.8).

Upon receiving on a port a FLOGI Request or a NPIV FDISC Request from a Node, a Controlling Switch shall allocate to the newly reachable VN_Port an N_Port_ID from its Domain_ID (i.e., not from the Virtual Domain_ID) if it accepts the received FLOGI or NPIV FDISC Request.

1.3.2 Distributed FCF Operations

In a Distributed FCF, the Primary Controlling FCF defines the routes for the FDF topology and performs N_Port_ID allocations and deallocations for all its controlled FDFs. The two Controlling FCFs keep their state synchronized.

When becoming operational (i.e., when it has defined if it behaves as Primary or Secondary), a Controlling FCF performs FIP discovery as described in T11/11-026v3. A Controlling FCF instantiates VA_Port to VA_Port Virtual Links with the discovered FDFs that are directly reachable and are part of its FDF Set by initiating FIP ELP Exchanges, as described in T11/11-026v3.

Upon instantiating a VA_Port to VA_Port Virtual Link with an FDF, the Primary Controlling FCF shall initiate an FDRN Exchange (see 1.1.4) describing that Virtual Link with the Secondary Controlling FCF to keep the state synchronized. Upon completion of this FDRN Exchange, the Primary Controlling FCF shall provide to that FDF the Distributed FCF Membership information through a DFMD Exchange (see 1.1.9). At this point the instantiated VA_Port to VA_Port Virtual Link becomes part of the Distributed FCF internal topology (i.e., the set of VA_Port to VA_Port Virtual Links internal to the Distributed FCF). The Primary Controlling FCF shall recompute the N_Port_ID routes and distribute them to each FDF belonging to the Distributed FCF through NPRD Exchanges (see 1.1.6).

Upon deinstantiating a VA_Port to VA_Port Virtual Link with an FDF, the Primary Controlling FCF shall initiate an FDUN Exchange (see 1.1.5) describing that Virtual Link with the Secondary Controlling FCF to keep the state synchronized. Upon completion of this FDUN Exchange, the Primary Controlling FCF shall recompute the N_Port_ID routes and distribute them to each FDF belonging to the Distributed FCF through NPRD Exchanges.

When becoming operational, an FDF participates in FIP discovery as described in T11/11-026v3. It then waits for a Controlling FCF or another FDF to initiate a FIP ELP Exchange with it, in order to set up a VA_Port to VA_Port Virtual Link. Upon completing the DFMD Exchange with the Primary Controlling FCF, the FDF becomes able to initiate FIP ELP Requests to instantiate other VA_Port to VA_Port Virtual Links with other FDFs. Upon completing the NPRD Exchange with the Primary Controlling FCF, an FDF becomes able to set up proper forwarding tables to forward FCoE frames inside and outside the Distributed FCF. At this point the FDF enables its VF_Port capable FCF-MACs for FIP discovery and FIP logins from ENodes.

Upon instantiating a VA_Port to VA_Port Virtual Link with another FDF or with the Secondary Controlling FCF, an FDF shall perform a FDRN Exchange with the Primary Controlling FCF to inform it of the new Virtual Link. Upon completing a FDRN Exchange with an FDF, the Primary Controlling FCF shall initiate another FDRN Exchange with the same parameters with the Secondary Controlling FCF to keep the state synchronized. After completing this FDRN Exchange the primary Controlling FCF shall provide to the newly reported FDF the Distributed FCF Membership information through a DFMD Exchange (see 1.1.9), if that FDF did not not already receive a DFMD Exchange in a previous step. At this point the instantiated VA_Port to VA_Port Virtual Link becomes part of the Distributed FCF internal topology (i.e., the set of VA_Port to VA_Port Virtual Links internal to the Distributed FCF). Upon completion of this DFMD Exchange, the Primary Controlling FCF shall recompute the N_Port_ID routes and distribute them to each FDF belonging to the Distributed FCF through NPRD Exchanges.

NOTE 4 – A VA_Port to VA_Port Virtual Link with the Secondary Controlling FCF may be instantiated before the VA_Port to VA_Port Virtual Link with the Primary Controlling FCF. The FDF recognizes the Primary Controlling FCF because it is the one from which it receives the DFMD Request. In this case, the FDF initiates with the Primary Controlling FCF the FDRN Exchange describing the Virtual Link with the Secondary Controlling FCF upon completing the DFMD Exchange.

Upon deinstantiating a VA_Port to VA_Port Virtual Link with another FDF or with the Secondary Controlling FCF, an FDF shall perform a FDUN Exchange with the Primary Controlling FCF to inform it of the disappeared Virtual Link. Upon completing a FDUN Exchange with an FDF, the Primary Controlling FCF shall initiate another FDUN Exchange with the same parameters with the Secondary Controlling FCF to keep the state synchronized. Upon completion of this FDUN Exchange, the Primary Controlling FCF shall recompute the N_Port_ID routes and distribute them to each FDF belonging to the Distributed FCF through NPRD Exchanges.

Upon receiving on a VF_Port capable FDF-MAC a FIP FLOGI Request or a FIP NPIV FDISC Request from an ENode MAC, an FDF shall send a VNRN Request (see 1.1.2) to the Primary Controlling FCF to inform it of the newly reachable VN_Port. If the Primary Controlling FCF rejects the VNRN Request, the FDF shall also reject the FIP FLOGI Request or FIP NPIV FDISC Request. If the Primary Controlling FCF accepts the VNRN Request, it performs the following processing:

a) if the VNRN Request carried a FLOGI Request and that VN_Port was not already logged in or if the VNRN Request carried a NPIV FDISC Request, then the Primary Controlling Switch shall allocate to the newly reachable VN_Port an N_Port_ID from the Virtual Domain_ID; or

b) if the VNRN Request carried a FLOGI Request and that VN_Port was already logged in, then the Primary Controlling Switch shall implicitly log out that VN_Port and all the VN_Ports associated to the VF_Port that VN_Port was associated with and then allocate to that VN_Port an N_Port_ID from the Virtual Domain_ID.

The Primary Controlling FCF shall also recompute the Zoning ACLs for the affected N_Port_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling FCF shall distribute the Zoning ACLs and N_Port_ID allocation/deallocation information to the Secondary Controlling FCF and to each FDF belonging to the Distributed FCF except the one from which the VNRN has been received through an appropriate NPZD Exchange (see 1.1.7). The NPZD Request sent to the Secondary Controlling FCF shall include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry and shall carry no Peering Entries. The NPZD Requests sent to the other FDFs shall not include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry. Upon receiving all SW_ACC from the Secondary Controlling FCF and the other FDFs or upon the expiration of a one second timeout, whichever comes first, the Primary Controlling FCF shall send an NPZD Request also to the FDF from which the VNRN has been received. This NPZD Request shall include the FLOGI / NPIV FDISC LS_ACC Parameters in the allocation entry. Upon receiving this NPZD Request the FDF that sent the VNRN Request shall accept the FIP FLOGI Request or FIP NPIV FDISC Request and complete the instantiation of the VN_Port to VF_Port Virtual Link.

When a VN_Port to VF_Port Virtual Link is deinstantiated or when a VF_Port is deintantiated, an FDF shall perform a VNUN Exchange (see 1.1.3) with the Primary Controlling FCF to inform it that the VN_Port is now unreachable or that all the VN_Ports associated with that VF_Port are unreachable. Upon completing a VNUN Exchange, the Primary Controlling FCF shall deallocate the N_Port_ID previously assigned to the affected VN_Port(s), recompute the Zoning ACLs for the affected N_Port_IDs, generate appropriate RSCN(s), and update the Fibre Channel Name Server. The Primary Controlling FCF shall then distribute this information to the Secondary Controlling FCF and to each FDF belonging to the Distributed FCF through NPZD Exchanges indicating N_Port_ID(s) deallocation (see 1.1.7).

When a new Zone Set is activated in the Fabric, the Primary Controlling FCF shall recompute the Zoning ACLs for all N_Port_IDs allocated in the Virtual Domain_ID and distribute them to the FDFs of the Distributed FCF through AZAD Exchanges (see 1.1.8).

Upon receiving on a VF_Port capable FCF-MAC a FIP FLOGI Request or a FIP NPIV FDISC Request from an ENode MAC, a Controlling FCF shall allocate to the newly reachable VN_Port an

N_Port_ID from its Domain_ID (i.e., not from the Virtual Domain_ID) if it accepts the received FIP FLOGI or FIP NPIV FDISC Request.