

0.1 FC-BB_FCoE definitions

insert here.

0.2 List of commonly used acronyms and abbreviations

FIP FCoE Initialization Protocol
FPMA Fabric Provided MAC Address
SPMA Server Provided MAC Address

0.3 Defined Constants

This standard utilizes several defined constants. The value for each of these constants is specified in Table 1.

Table 1 – Defined Constants

Constant	Value	Description
FIP_TYPE	8914h	Value to be used in the Type field of the 802.3 frame
ALL_FCF_MACS	TBD	Group address for all FCFs
ALL_ENODE_MACS	TBD	Group address for all FCFs and ENodes
DEFAULT_FIP_PRIORITY	128	Default value to be used in the Priority Descriptor

1 FCoE Initialization Protocol

1.1 Introduction

The FCoE Initialization Protocol (FIP) provides the following services:

- A mechanism for devices on an FCoE fabric to discover other devices on the fabric
- May be used by ENodes to discover available FCFs that are willing to accept VN_Port connections from the ENode.
- May be used by FCFs to discover other FCFs to which VE_Port connections are possible.
- Discovery capability includes the peer devices' MAC address, the FC_MAP in use, the FCFs Switch_Name, the Fabric_Name, and the peer devices' Port_Name, and addressing modes (FPMA and/or SPMA).
- Fabric log in and log out
- Fabric Exchange Link Parameters

FIP frame formats differ from FCoE frame formats to facilitate processing by the mechanisms responsible for fabric discovery and initialization. Frame formats are defined for the discovery operations. For fabric log in and log out, FIP specifies an encapsulation of the standard Fibre Channel Link services ELS requests FLOGI, FDISC (for NPIV), fabric LOGO, ELP, and their associated replies. FIP supports NPIV operation utilizing an FDISC exchange similar to that performed in Fibre Channel.

In general, the discovery operation is performed to discover devices willing to accept connections from the discovering device and to determine certain parameters necessary to effect those connections. The discovering device then uses these data to establish connections (via FLOGI / FDISC) to the desired devices. Devices may continue to perform discovery operations to determine changes in peer device reachability.

1.2 Overview of the FIP discovery process

When an ENode joins an FCF fabric, it must discover the presence of FCFs into which it may potentially log. FIP provides a Discovery Solicitation and a Discovery Response for this purpose. In general, an ENode sends a Discovery Solicitation to the ALL_FCF_MACS multicast group. Among other things, the Discovery Solicitation contains the MAC address to which responding FCFs should address their Advertisements, and the addressing mode capability of the ENode (FPMA and/or SPMA).

When an FCF receives a Discovery Solicitation, if it is willing to allow the ENode to log into it, it responds with a unicast Discovery Advertisement. This Discovery Advertisement indicates the FCF's willingness to accept logins (FLOGI or NPIV FDISC) from the ENode.

An FCF may also generate Discovery Solicitations to discover the presence of other FCFs. In addition, an FCF transmits periodic Discovery Advertisements to the ALL_FCF_MACS and ALL_ENODE_MACS group address to make other devices on the FCoE fabric aware of new FCFs joining the fabric. ENodes may transmit additional Discovery Solicitations for the same purpose.

1.3 Overview of the FIP Login / Logout process

All FCoE devices use the FIP FLOGI, NPIV FDISC, fabric LOGO, and ELP procedures in lieu corresponding Fibre Channel services. In addition to the normal services FLOGI and NPIV FDISC per-

form, the FIP process provides a mechanism for the FCF to assign a MAC address to the VN_Port. Within the FLOGI request, the ENode indicates the addressing modes it supports (FPMA and/or SP-MA). If the ENode supports only SPMA, the ENode must propose a MAC address. If FPMA is supported, the ENode may propose a MAC address, but doing so is not required. Upon successful FLOGI or NPIV FDISC, the FCF responds with an assigned MAC address. If the ENode supports only SPMA, the assigned MAC address will be the proposed MAC address. If FPMA is supported, the FCF attempts to use the proposed MAC address if one is provided. If it is not able to do so (e.g. the proposed MAC address is improperly formed or duplicates an already assigned FPMA MAC Address), then the FCF may assign any valid FPMA address.

The FIP LOGO process enables the FCF to remove the assigned address.

1.4 FIP Frame Format

Figure 1 illustrates the structure of FIP frames. All FIP frames shall be formatted in accordance with IEEE Std 802.3™-2005 (802.3 Frame). All FIP frames shall carry the Type value FIP_TYPE. The source and destination MAC address settings and the frame payload are specified in the following subclauses. The MAC Client Data field within the 802.3 frame shall contain the Encapsulated FIP Operation.

Within the Encapsulated FIP Operation, the FIP Operation Code and FIP SubCode define the operation to be performed (e.g. Discovery Solicitation, Discovery Advertisement, Fabric Login / Logout operation). The flag bits convey operation specific data concerning addressing mode, advertisement type, and device type. A single Encapsulated FIP Operation may contain one or multiple Descriptors depending on the FIP operation. Descriptors are composed of groups of Type, Length, Value fields.

1.4.1 Encapsulated FIP Operation Format

An Encapsulated FIP operation shall be constructed as shown in Table 2.

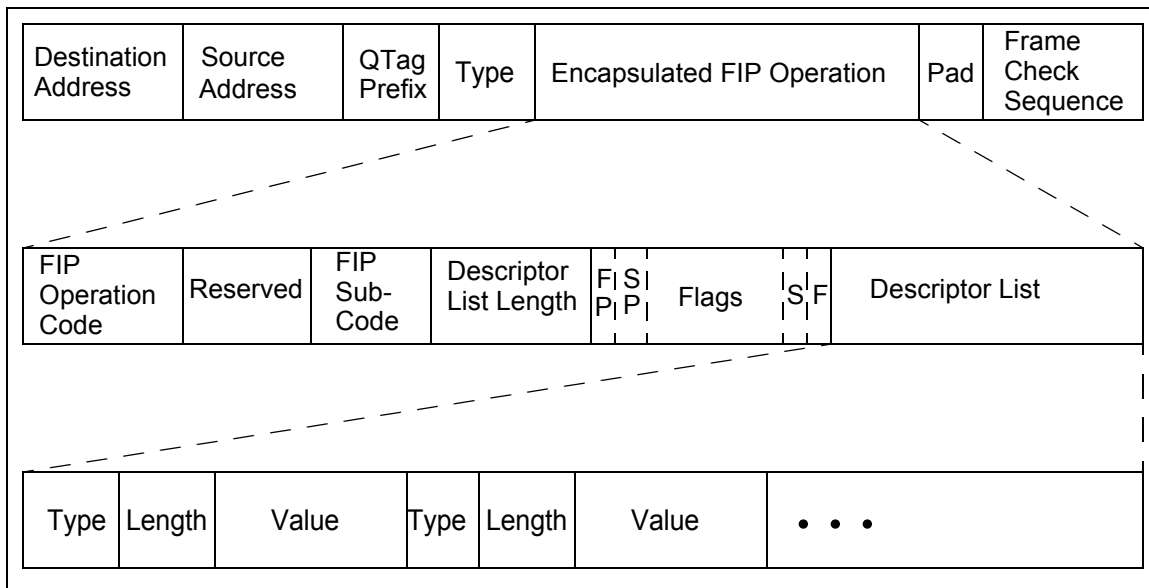


Figure 1 - FIP Frame Encapsulation Structure

The FIP Operation Code shall be set in accordance with Table 3. The Descriptor List Length field shall be set to the total length of all Descriptors within the Descriptor list, in 32 bit words.

The FP and SP Flags bits are the most and next most significant bits in the Flags field respectively. They shall be set in accordance with Table 4 for the specified operations and are reserved for all other operations.

The S Flags bit is the second least significant bit in the Flags field and shall be set to 1 for Discovery Advertisements that are initiated in response to a Discovery Solicitation. The S Flags bit shall be set to 0 for Discovery Advertisements that are not initiated in response to a Discovery Solicitation. The S Flags bit is reserved in all other cases.

The F Flags bit is the least significant bit in the Flags field and shall be set to 1 for Discovery FIP Operations if the originating device is an FCF. The F Flags bit shall be set to 0 for Discovery FIP Operations if the originating device is not an FCF. The F Flags bit is reserved in all other cases.

Table 2 – Encapsulated FIP Operation frame structure

Field	Size (Bytes)
FIP Operation Code	2
Reserved	1
FIP SubCode	1
Descriptor List Length	2
Flags	2
Descriptor List	Variable
Padding	Variable

All other Flags bits are reserved.

Following the Flags bits are one or more FIP Descriptors. The required FIP Descriptors are dependent upon the FIP Operation Code and task to be performed (See 1.4.3)

Two group addresses are provided for Discovery FIP Operations. The All_ENODE_MACS group address is intended to address all ENode devices. A typical use of this would be for periodic unsolicited Discovery Advertisements. The All_FCF_Devices group address is intended to address all FCFs. A typical use of this group address is to allow an ENode to discover reachable FCFs using a Discovery Solicitation.

Table 3 – FIP Operation Codes and SubCodes

FIP Operation Code Value (hex)	FIP SubCode Value (hex)	Operation
0001	01	Discovery Solicitation
0001	02	Discovery Advertisement
0002	01	Link Service Request
0002	02	Link Service Reply
All Others		Reserved

Table 4 – FP and SP bit usage (applies to discovery FIP operations only)

Bit	FIP Operation	Usage
FP	Discovery Solicitation or Advertisement	Set to 1 if device initiating frame supports FPMA Set to 0 if device initiating frame does not support FPMA
	FLOGI or NPIV FDISC Request	Set to 1 if FPMA Address is requested Set to 0 if FPMA Address is not supported (See Note 1)
	FLOGI or NPIV FDISC LS_ACC	Set to 1 if FPMA Address granted Set to 0 if SPMA Address granted
SP	Discovery Solicitation or Advertisement	Set to 1 if device initiating frame supports SPMA Set to 0 if device initiating frame does not support SPMA
	FLOGI or NPIV FDISC Request	Set to 1 if SPMA Address is requested Set to 0 if FPMA Address is not supported (See Note 1)
	FLOGI or NPIV FDISC LS_ACC	Set to 1 if SPMA Address granted Set to 0 if FPMA Address granted
<p>Note 1: It is valid to set both the FP and SP bits to 1 in FLOGI and NPIV FDISC requests to indicate that the device will accept either an FPMA or SPMA address; however, least one of these bits shall be set.</p>		

1.4.2 FIP Descriptors

A FIP Descriptor contains the fields listed in Table 5. The individual FIP Descriptors, Type field values, and Length field values are listed in Table 6. All values for the Type field not listed in Table 6 are reserved. Each Descriptor is described in detail in the following subclauses.

Table 5 – FIP Descriptor Structure

Field	Size (Bytes)
Type	1
Length	1
Value	Descriptor dependent, but always 2 plus a multiple of 4.

Table 6 – FIP Descriptors

Descriptor	Type Field Value	Length Field Value
Priority	1	1
MAC_Address	2	2
FC-MAP	3	2
Switch_Name	4	3
Fabric_Name	5	3
Max_Receive_Size	6	1
FLOGI	7	36 (request) 9(LS_ACC/ LS_RJT)
NPIV_FDISC	8	36 (request) 9(LS_ACC/ LS_RJT)
LOGO	9	1 (request) 9 (LS_ACC) 8 (LS_RJT)
ELP	10	33 (request) 9(LS_ACC/ LS_RJT)

1.4.2.1 Priority FIP Descriptor

The Priority FIP Descriptor contains a 16 bit Value field. The eight most significant bits are reserved. The eight least significant bits shall contain the value that the originating FCF inserts in a Priority. Higher priorities are indicated by higher numerical values. Unless otherwise configured (e.g. via management), this field shall be set to DEFAULT_FIP_PRIORITY (see table 1). This value is used by ENodes when multiple Discovery Advertisements are received (see 1.4.3.3)

1.4.2.2 MAC_Address FIP Descriptor

The MAC Address FIP Descriptor contains a 48 bit value field MAC_Address.

In FIP Discovery Solicitations, MAC_Address shall be set to the MAC address to which the device expects responses to be addressed. Discovery Advertisements sent in response to a Discovery Solicitation shall be addressed to the MAC address contained in this descriptor within the Discovery Solicitation.

In FIP Link Service Requests containing FLOGI or NPIV_FDISC FIP Descriptors originated from devices that support SPMA only, MAC_Address shall be set to the MAC Address to which the originating device expects FCoE frames addressed. In FIP Link Service Requests containing FLOGI or NPIV_FDISC Descriptors originating from devices that support FPMA only, MAC_Address shall be set to the MAC Address to which the originating device proposes FCoE frames sent to it be addressed (the FCF may or may not honor the proposal), or to all zeros (indicating the originating device is not proposing a MAC address). In FIP Link Service Requests containing a FLOGI or FDISC FIP Descriptor originated from devices that support both SPMA and FPMA, the MAC Address shall be set to one of the options previously specified in this paragraph.

In FIP Link Service Requests containing the LOGO FIP Descriptor, MAC_Address shall be set to the MAC address previously assigned to the VN_Port (by an FCF) that is being logged out.

In FIP Link Service Replies containing an LS_ACC for an FLOGI or FDISC, MAC_Address shall be set to the MAC address assigned by the FCF that the originator of the FLOGI or FDISC is to use as the destination MAC address for FCoE frames addressed to it. If the originator of the FLOGI or FDISC supports SPMA only, the assigned MAC address shall be the MAC address proposed in the FIP Link Service Request. If the originator of the FLOGI or FDISC supports FPMA only, the assigned MAC address shall be a properly formed FPMA MAC address. In this case, the assigned MAC address should be the MAC address proposed in the FIP Link Service Request if that MAC address is a properly formed FPMA MAC address. For the purpose of this paragraph, a properly formed FPMA MAC address is one in which the 24 most significant bits equal the fabric's FC-MAP value and the least significant 24 bits are unique within the fabric for all FPMA addresses assigned with the same FC-MAP.

1.4.2.3 FC-MAP FIP Descriptor

The FC-MAP FIP Descriptor contains a 48 bit value field. The 24 most significant bits are reserved. The 24 least significant bits contain a FC-MAP.

In Discovery Solicitations and Advertisements originated from FCFs that support FPMA, the FC-MAP value shall be set to the FC-MAP that the FCF is using. In Discovery Solicitations and Advertisements originated from FCFs that do not support FPMA, FC-MAP is reserved.

1.4.2.4 Switch_Name FIP Descriptor

The Switch_Name FIP Descriptor contains a 64 bit value field Swich_Name.

In Discovery Solicitations and Advertisements originated from FCFs, Switch_Name shall be set to the Switch_Name assigned to the FCF.

1.4.2.5 Fabric_Name FIP Descriptor

The Fabric_Name FIP Descriptor contains a 64 bit value field Fabric_Name.

In Discovery Solicitations and Advertisements originated from FCFs, Fabric_Name shall be set to the Fabric_Name assigned to the FCF.

1.4.2.6 Max_Receive_Size FIP Descriptor

The Max_Receive_Size FIP Descriptor contains a 16 bit value field Max_Receive_Size.

In Discovery Solicitations, Max_Receive_Size shall be set to the maximum size 802.3 Frame that the originating device is able to receive (from the first octet of the Destination Address through the last octet of the FCS, inclusive).

1.4.2.7 FLOGI FIP Descriptor

The FLOGI FIP descriptor contains either a 144 byte value field or a 36 byte value field. It shall be set to one of the following three options:

- 1) A complete Fibre Channel FLOGI request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences are prohibited)
- 2) A complete Fibre Channel FLOGI LS_ACC exclusive of the SOF, CRC, and EOF, corresponding to a previously received FLOGI request (multi-frame sequences are prohibited)
- 3) A complete Fibre Channel FLOGI LS_RJT exclusive of the SOF, CRC, and EOF corresponding to a previously received FLOGI request (multi-frame sequences are prohibited)

1.4.2.8 NPIV_FDISC FIP Descriptor

The NPIV_FDISC FIP descriptor contains either a 144 byte value field or a 36 byte value field. It shall be set to one of the following three options:

- 1) A complete Fibre Channel NPIV FDISC request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences are prohibited). It is prohibited to use this FIP Descriptor for FDISC requests not related to NPIV (use an FDISC frame embedded in an FCoE frame instead for this purpose).
- 2) A complete Fibre Channel FDISC LS_ACC exclusive of the SOF, CRC, and EOF, corresponding to a previously received NPIV FDISC request (multi-frame sequences are prohibited)
- 3) A complete Fibre Channel FDISC LS_RJT exclusive of the SOF, CRC, and EOF corresponding to a previously received NPIV FDISC request (multi-frame sequences are prohibited)

1.4.2.9 LOGO FIP Descriptor

The NPIV_FDISC FIP descriptor contains either a 44 byte value field or a 36 byte value field. It shall be set to one of the following three options:

- 1) A complete Fibre Channel fabric LOGO request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences are prohibited). It is prohibited to use this FIP Descriptor for process LOGO requests (use a LOGO frame embedded in an FCoE frame instead for this purpose).
- 2) A complete Fibre Channel LOGO LS_ACC exclusive of the SOF, CRC, and EOF, corresponding to a previously received fabric LOGO request (multi-frame sequences are prohibited)
- 3) A complete Fibre Channel LOGO LS_RJT exclusive of the SOF, CRC, and EOF corresponding to a previously received fabric LOGO request (multi-frame sequences are prohibited)

1.4.2.10 ELP FIP Descriptor

The ELP FIP descriptor contains either a 132 byte value field or a 36 byte value field. It shall be set to one of the following three options:

- 1) A complete Fibre Channel ELP request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences are prohibited)
- 2) A complete Fibre Channel ELP LS_ACC exclusive of the SOF, CRC, and EOF, corresponding to a previously received ELP request (multi-frame sequences are prohibited)
- 3) A complete Fibre Channel ELP LS_RJT exclusive of the SOF, CRC, and EOF corresponding to a previously received ELP request (multi-frame sequences are prohibited)

1.4.3 FIP Frame Construction and Reception

The construction of FIP Frames for various purposes are explicitly defined including which FIP descriptors are to be included and their order. Originators of FIP frames shall include the FIP descriptors listed in Table 7 in the order listed. No other FIP Descriptors shall be included. Furthermore, devices shall emit only those FIP Frames listed as valid in Table 7 for the device type. The Pad Field shall be of the length specified in Table 7 and may contain any value.

On reception, the receiving device shall verify that the FIP frame includes the FIP Descriptors required in Table 7.

Note: unless otherwise specified, the FIP Descriptors may be received in any order. Additional FIP Descriptors may be present. If additional FIP Descriptors are present, they shall be ignored. This is to provide flexibility for future standardization.

The following subclauses describe the purpose of each FIP operation and provide additional operation specific requirements.

Table 7 – FIP Frame Construction

Frame Type	Valid if Initiated From:	FIP Descriptors and order of inclusion:	Pad Size (Bytes)
Discovery Solicitation	ENode	MAC_Address, Max_Receive_Size	24
Discovery Solicitation	FCF	MAC_Address, FC-MAP, Swich_Name, Max_Receive_Size	4
Discovery Advertisement	FCF	MAC_Address, FC-MAP, Switch_Name, Fabric_Name	Note 1
FLOGI Request	ENode	FLOGI, MAC_Address	0
FLOGI LS_ACC	FCF	FLOGI, MAC_Address	0
FLOGI LS_RJT	FCF	FLOGI	0
NPIV FDISC	ENode	NPIV_FDISC, MAC_Address	0
NPIV FDISC LS_ACC	FCF	NPIV_FDISC, MAC_Address	0
NPIV FDISC LS_RJT	FCF	NPIV_FDISC	0
Fabric LOGO	ENode or FCF	LOGO, MAC_Address	0
Fabric LOGO LS_ACC	ENode or FCF	LOGO, MAC_Address	0
Fabric LOGO LS_RJT	ENode or FCF	LOGO	0
ELP	FCF	ELP, MAC_Address	0
ELP LS_ACC	FCF	ELP, MAC_Address	0
ELP LS_RJT	FCF	ELP	0
Note 1: If the Discovery Advertisement is in response to a Discovery Solicitation, the Pad Size shall be such that the entire 802.3 frame is of the length specified in the Max_Receive_Size FIP Descriptor within the Discovery Solicitation; otherwise, the Pad Size shall be zero (i.e. no padding).			

1.4.3.1 Discovery Solicitation from an ENode

ENodes may transmit Discovery Solicitations to FCFs to request the FCF to reply with an Discovery Advertisement. The Discovery Solicitation shall be addressed to an individual FCF or to the

Non-FCF devices receiving a FIP Frame identified as a Discovery Solicitation, regardless of how it is formatted, shall ignore the frame.

FCF devices receiving a Discovery Solicitation from an ENode (i.e. Flags F bit set to zero) shall construct a Discovery Advertisement and transmit it to the soliciting device using the MAC_Address in the Discovery Solicitation MAC_Address FIP Descriptor as the Destination Address if the FCF is willing to accept logins (via the FIP FLOGI or NPIV FDISC processes) from the ENode. If the FCF is not willing to accept such logins, then it shall ignore the Discovery Solicitation.

1.4.3.2 Discovery Solicitation from an FCF

FCFs may transmit Discovery Solicitations to FCFs to request the FCF to reply with a Discovery Advertisement. The Discovery Solicitation shall be addressed to an individual FCF or to the ALL_FCF_MACS group address.

Non-FCF devices receiving a FIP Frame identified as a Discovery Solicitation, regardless of how it is formatted, shall ignore the frame.

FCF devices receiving a Discovery Solicitation from an FCF (i.e. Flags F bit set to one) shall make the following checks:

- The Switch_Name value in the Discovery Solicitation matches the switch name of the receiving FCF
- The FC-MAP value in the Discovery Solicitation is zero or it matches the FC-MAP of the receiving FCF

If the above checks are both true, the receiving FCF shall construct a Discovery Advertisement and transmit it to the soliciting device using the MAC_Address in the Discovery Solicitation MAC_Address FIP Descriptor as the Destination Address, if the FCF is willing to accept VE_Port connections from the soliciting FCF. If the FCF is not willing to accept VE_Port connections from the soliciting FCF, then it shall ignore the Discovery Solicitation.

1.4.3.3 Discovery Advertisements

FCFs transmit Discovery Advertisements in response to Discover Solicitations as specified in subclause 1.4.3.1 and subclause 1.4.3.2. In addition, an FCF shall transmit a Discovery Advertisement to the ALL_FCOE_MACS (for VE_Ports) or ALL_ENODE_MACS (for VF_Ports) multicast group periodically.

FCFs and ENodes shall not form VN_Ports, VF_Ports or VE_Ports (as appropriate) with peer devices unless the Discovery Advertisement contains the highest Priority received from all received Discovery Solicitations. If the same highest Discovery Advertisement Priority is received from multiple peer devices, then such Ports may be established with any set of devices from which these Discovery Advertisements were received.

Editor's Note: How long do we wait to ensure we have heard all the advertisements?

Editor's Note: What happens if we have connectivity, and a higher priority device is discovered?

Editor's Note: How often do we send periodic advertisements?

1.4.3.4 FLOGI and NPIV FDISC Processing

FCoE devices shall use the FIP FLOGI and NPIV FDISC protocol to perform FLOGI and NPIV related FDISC processes. It is prohibited to transmit an FLOGI, an FDISC related to NPIV, or their associated replies encapsulated in an FCoE frame.

The FIP FLOGI and NPIV FDISC operations provide a mechanism for assigning MAC addresses to the VN_Port. Each of these requests enable the ENode to indicate to the FCF whether it supports FPMA, SPMA, or both. In addition, the ENode proposes a specific address if it supports SPMA only, and may propose a specific address if it supports FPMA or both SPMA and FPMA. An ENode that supports FPMA may elect not to propose a specific address by filling the MAC Address field with zeros.

An FCF accepting an FIP FLOGI or NPIV FDISC from an ENode assigns the ENode a MAC address to be used with the VN_Port. If the ENode supports only SPMA, then the assigned MAC address will be the MAC address from the request. If the ENode supports FPMA or both SPMA and FPMA, the

FCF should attempt to honor the MAC address proposal as long as the proposed MAC address is valid and conforms with the addressing mode(s) supported by the FCF.

FCFs shall reject FIP FLOGI and NPIV FDISC requests from ENodes that support only SPMA and propose a MAC Address that is not a unicast address. In addition, FCFs shall reject FIP FLOGI and NPIV FDISC requests for an addressing mode (SPMA or FPMA) not supported by the FCF. Finally, if the FCF supports both FPMA and SPMA and the ENode supports only SPMA, the FCF shall reject FIP FLOGI and NPIV FDISC requests that contain a proposed MAC address in which the 24 most significant bits match the FC-MAP in use by the FCF.

If both the FCF and ENode support both SPMA and FPMA, the FCF may assign an address of either form.

1.4.3.5 Fabric LOGO Processing

FCoE devices shall use the FIP LOGO protocol to perform fabric LOGO processing. It is prohibited to use a LOGO embedded in an FCoE frame to perform fabric LOGO processing (however, process LOGO is performed by embedding a LOGO in an FCoE frame rather than a FIP frame).

1.4.3.6 ELP Processing

FCoE devices shall use the FIP ELP protocol to perform ELP processing. It is prohibited to use an ELP embedded in an FCoE frame to perform fabric ELP processing.