0.1 FC-BB_FCoE definitions

insert here.

0.2 List of commonly used acronyms and abbreviations

ACE Access Control Entry
ACL Access Control List
AFL Available FCF List
D_A_TOV Discovery Advertisement Time-Out Value
FCF Fibre Channel Forwarder
FIP FCoE Initialization Protocol
FPMA Fabric Provided MAC Address
MAC Media Access Control
VLAN Virtual Local Area Network

0.3 Approved references

IEEE Std 802.3™-2005, *Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.*

0.4 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
FCoE_Type	8906h	Value to be used in the type field of the 802.3 frame to indicate an FCoE payload
ALL_FCF_MACS	TBD	Group address for all FCFs
ALL_ENODE_MACS	TBD	Group address for all ENodes
DEFAULT_FIP_PRIORITY	128	Default value to be used in the Priority Descriptor
DEFAULT_FC-MAP	0EFC00h	Default value to be used for the FC-MAP
D_A_TOV	Default: 5 May be administratively configured to any integer value between 1 and 60, inclusive	Number of seconds between Discovery Advertisements sent by FCFs

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 FCF's Switch_Name, the Fabric_Name, and the peer devices' Node_Name, and addressing modes (FPMA and/or SPMA).
- Fabric log in and log out
- Fabric Exchange Link Parameters
- Link Reset

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), N_Port 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 FIP FLOGI and FIP NPIV_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, it responds with a unicast Discovery Advertisement. 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.

1.3 Overview of the FIP Login / Logout process

All FCoE devices use the FIP FLOGI, NPIV FDISC, N_Port LOGO, and ELP procedures in lieu corresponding Fibre Channel services. In addition to the normal services FLOGI and NPIV FDISC perform, 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 SPMA). 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 Link Reset Process

FIP provides an operation that enables an FCF to force an ENode to reset VN_Ports, including all or a subset of the N_Ports associated with a given ENode MAC. This operation is used by an FCF when a loss of connectivity is detected.

1.5 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 preformed (e.g. Discovery Solicitation, Discovery Advertisement, Fabric Login / Logout operation, or Link Reset). The flag bits convey operation specific data concerning addressing mode, advertisement type, device type, and ability or willingness to accept logins. 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.

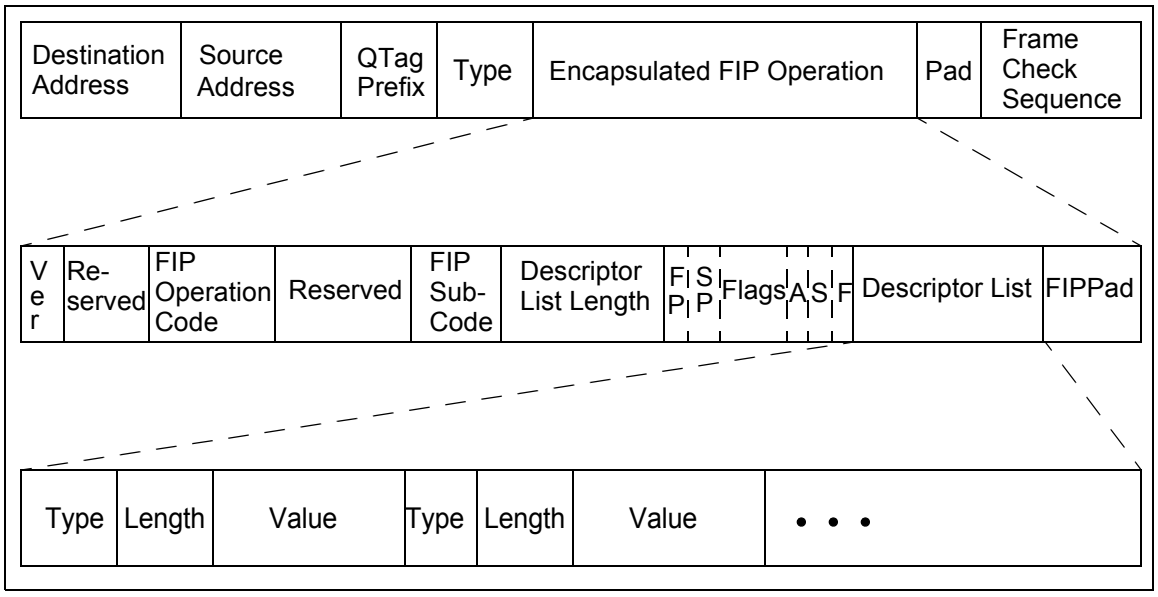

Figure 1 - FIP Frame Encapsulation Structure

1.5.1 Encapsulated FIP Operation Format

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

The Ver field shall be set to 1.

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 A Flags bit is the third least significant bit in the Flags field. In Discovery Advertisements originated from FCFs, it shall be set to 1 if the FCF is able to accept additional FIP FLOGIs or NPIV_FDISCs. It shall be set to zero if the FCF is not currently able to accept additional FIP FLOGIs or NPIV_FDISCs. In solicited Discovery Advertisements, the FCF may also set this bit to zero to indicate the FCF's unwillingness to accept FIP FLOGIs or NPIV_FDISCs from the soliciting device. This bit is reserved for all other FIP operations.

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 (Bits)
Ver	4
Reserved	12
FIP Operation Code	16
Reserved	8
FIP SubCode	8
Descriptor List Length	16
Flags	16
Descriptor List	Variable
FIPPad	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 1.5.3).

The last field is the FIPPad. This field enables solicited Discovery Advertisements to be extended to the maximum length supported by the solicitor to ensure verify that the fabric can transport frames of that size (see clause 1.5.3).

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
0003	01	FIP Link Reset
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
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.		

1.5.2 FIP Descriptors

A FIP Descriptor contains the fields listed in Table 5. The individual FIP Descriptors and their corresponding Type field value are listed in Table 6. All values for the Type field not listed in Table 6 are reserved.

The Length field in each descriptor shall be set to the length of the descriptor in 32-bit words inclusive of the Type, Length, and Value fields.

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 Type Values

Descriptor	Type Field Value
Priority	1
MAC_Address	2
FC-MAP	3
Switch_Node_Name	4
Fabric_Name	5
Max_Receive_Size	6
FLOGI	7
NPIV_FDISC	8
LOGO	9
ELP	10

1.5.2.1 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 contains the Priority field and shall contain the a priority value set by the originating FCF. Higher priorities are indicated by lower numerical values. Unless otherwise administratively configured, 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 1.5.3.3).

1.5.2.2 MAC_Address FIP Descriptor

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

1.5.2.3 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.

1.5.2.4 Switch_Node_Name FIP Descriptor

The Switch_Node_Name FIP Descriptor contains a 64 bit value field Switch_Node_Name. This field shall be set to the Switch_Name if originated from an FCF or the Node_Name if originated from an ENode.

1.5.2.5 Fabric_Name FIP Descriptor

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

1.5.2.6 Max_Receive_Size FIP Descriptor

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

1.5.2.7 FLOGI FIP Descriptor

The FLOGI FIP descriptor contains a variable length field that shall be set to one of the following three options (which may be encrypted in accordance with FC-SP):

1) A complete Fibre Channel FLOGI request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences and Login Extension Data 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 and Login Extension Data 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)

1.5.2.8 NPIV_FDISC FIP Descriptor

The NPIV_FDISC FIP descriptor contains a variable length field that shall be set to one of the following three options (which may be encrypted in accordance with FC-SP):

1) A complete Fibre Channel NPIV FDISC request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences and Login Extension Data 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 and Login Extension Data 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)

1.5.2.9 LOGO FIP Descriptor

The NPIV_FDISC FIP descriptor contains a variable length field that shall be set to one of the following three options (which may be encrypted in accordance with FC-SP):

1) A complete Fibre Channel N_Port LOGO request frame exclusive of the SOF, CRC, and EOF (multi-frame sequences are prohibited). It is prohibited to use this FIP Descriptor to 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 N_Port 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 N_Port LOGO request (multi-frame sequences are prohibited)

1.5.2.10 ELP FIP Descriptor

The ELP FIP descriptor contains a variable length field that shall be set to one of the following three options (which may be encrypted in accordance with FC-SP):

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 SW_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 SW_RJT exclusive of the SOF, CRC, and EOF corresponding to a previously received ELP request (multi-frame sequences are prohibited)

1.5.3 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.

Within solicited Discovery Advertisements, the FIPPad Field shall be of the length required to create an 802.3 frame with an overall length that matches the MAX_Receive_Size Descriptor received in the solicitation. The value of the field is reserved. In all other cases, the FIPPad Field shall be of zero length (i.e. not present).

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:
Discovery Solicitation	ENode	MAC_Address, Switch_Node_Name, Max_Receive_Size
Discovery Solicitation	FCF	MAC_Address, FC-MAP, Switch_Node_Name, Max_Receive_Size
Discovery Advertise-ment	FCF	MAC_Address, FC-MAP, Switch_Node_Name, Fabric_Name
FLOGI Request	ENode	FLOGI, MAC_Address
FLOGI LS_ACC	FCF	FLOGI, MAC_Address
FLOGI LS_RJT	FCF	FLOGI
NPIV FDISC	ENode	NPIV_FDISC, MAC_Address
NPIV FDISC LS_ACC	FCF	NPIV_FDISC, MAC_Address
NPIV FDISC LS_RJT	FCF	NPIV_FDISC
N_Port LOGO	ENode or FCF	LOGO, MAC_Address
N_Port LOGO LS_ACC	ENode or FCF	LOGO, MAC_Address
N_Port LOGO LS_RJT	ENode or FCF	LOGO
ELP	FCF	ELP, MAC_Address
ELP SW_ACC	FCF	ELP, MAC_Address
ELP SW_RJT	FCF	ELP
Link Reset	FCF	MAC_Address, Switch_Node_Name, one or more additional MAC_Address

1.5.3.1 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 ALL_FCF_MACS group address.

The MAC_Address field within the MAC_Address FIP Descriptor shall be set to the MAC address to which the ENode expects responses to the Discovery Solicitation to be addressed.

The Max_Receive_Size field within the Max_Receive_Size Descriptor shall be set to the maximum size 802.3 Frame that the ENode is able to receive, in octets (counted from the first octet of the Destination Address through the last octet of the FCS, inclusive).

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

If the above checks are both true, the receiving FCF shall construct a Discovery Advertisement and transmit it to the soliciting device as specified in clause 1.5.3.3.

1.5.3.2 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.

The MAC_Address field within the MAC_Address FIP Descriptor shall be set to the MAC address to which the FCF expects responses to the Discovery Solicitation to be addressed.

In Discovery Solicitations 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 originated from FCFs that do not support FPMA, FC-MAP is reserved.

Unless otherwise administratively configured, FCFs that support FPMA shall use DEFAULT_FC-MAP (see table 1) as the FC-MAP value.

In Discovery Solicitations originated from FCFs, the Fabric_Name field within the Fabric_Name descriptor shall be set to the Fabric_Name assigned to the fabric.

The Max_Receive_Size field within the Max_Receive_Size Descriptor shall be set to the maximum size 802.3 Frame that the FCF is able to receive, in octets (counted from the first octet of the Destination Address through the last octet of the FCS, inclusive).

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

– The Switch_Node_Name field value in the Discovery Solicitation *differs from* 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 as specified in clause 1.5.3.3.

Note: it is possible for an FCF to receive a Discovery Solicitation from itself due to the fact that Discovery Solicitations sent to the ALL_FCF_MACS group address may be forwarded to other ports on the same FCF by intervening Ethernet bridges. These frames are easily detected and discarded by the fact that the Switch_Node_Name field will match the Switch_Name of the receiving FCF.

1.5.3.3 Discovery Advertisements

FCFs transmit Discovery Advertisements in response to Discovery Solicitations as specified in subclause 1.5.3.1 and subclause 1.5.3.2. Discovery Advertisements sent in response to a Discovery Solicitation shall be addressed to the MAC address contained in MAC_Address descriptor within the Discovery Solicitation. The Discovery Advertisement shall be sent within D_A_TOV seconds of receipt of the Discovery Solicitation (see table 1).

Note that an FCF may receive Discovery Solicitations from the same remote device on multiple MACs. In this case, a separate solicited Discovery Advertisement shall be constructed and transmitted on each of the MACs from which a Discovery Solicitation was received. The soliciting device may

determine that it received multiple solicited Discovery Advertisements from the same FCF by the fact that the value of the Switch_Node_Name field will be identical in each of the solicited Discovery Advertisements.

In addition, an FCF shall transmit a Discovery Advertisement to the ALL_FCF_MACS (for VE_Ports) or ALL_ENODE_MACS (for VF_Ports) multicast group every D_A_TOV * 10 seconds +/- 50% (see table 1). FCFs should randomize the time that Discovery Advertisements are sent within this window to avoid large bursts of multicast traffic within the network.

Editor's Note: Table 1 currently defines D_A_TOV with a default of five and allows administrative assignment of between one and 60. No doubt this will generate much discussion, but its a starting point.

Note that within solicited Discovery Advertisements, the FIPPad Field is of the length required to create an 802.3 frame with an overall length that matches the MAX_Receive_Size Descriptor received in the solicitation. In non-solicited Discovery Advertisements, the FIPPad Field is of zero length (i.e. not present). See clause 1.5.3 for the normative requirements of this padding.

Unless otherwise administratively configured, 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 within at least 2 x D_A_TOV seconds after transmitting the Discovery Solicitation. 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. Discovery Advertisements received after this period may but are not required to be considered.

Editor's Note: Note that the requirement is worded as a shall, which is the editor's preference. However, there currently exists a wide variety of opinions on how strong this should be stated. The opinions seem to range from "shall" to here is a suggestion, do what you want. This will need to be worked out.

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

Unless otherwise administratively configured, FCFs that support FPMA shall use DEFAULT_FC-MAP (see table 1) as the FC-MAP value.

In Discovery Advertisements originated from FCFs, the Fabric_Name field within the Fabric_Name descriptor shall be set to the Fabric_Name assigned to the fabric.

1.5.3.4 FLOGI and NPIV FDISC Processing

FCoE devices shall use the FIP FLOGI and NPIV FDISC protocol to perform FLOGI and NPIV related FDICS 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 FIP FLOGI or NPIV_FDISC Operations originated from devices that support SPMA, the MAC_Address field within the MAC_Address Descriptor shall be set to the proposed MAC Address to which the originating device expects FCoE frames addressed. In FIP FLOGI or NPIV_FDISC Opera-

tions 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).

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. In FIP FLOGI or NPIV_FDISC Operation Replies containing an LS_ACC, the MAC_Address field within the MAC_Address descriptor 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 FLOGI or NPIV_FDISC 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. If the originator of the FLOGI or FDISC supports both FPMA and SPMA, the assigned MAC address shall be either the MAC address proposed in the FIP Link Service Request, or a properly formed FPMA MAC address assigned by the FCF. 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.

Unless otherwise administratively configured, an FCF that supports FPMA shall use DEFAULT_FC-MAP as the value for FC-MAP (see table 1).

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.

The 802.3 Frame Source Address in a FIP NPIV_FDISC shall be the same as the 802.3 frame Source Address of the prior FIP FLOGI associated with the FIP NPIV_FDISC. A successful FIP FLOGI creates a VF_Port. Subsequent FIP NPIV_FDISCs with the same 802.3 Frame Source Address as the FIP FLOGI associate additional VN_Ports to the single VF_Port.

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

1.5.3.5 N_Port LOGO Processing

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

In FIP LOGO FIP Operations, the MAC_Address field within the MAC_Address Descriptor shall be set to the MAC address previously assigned to the VN_Port (by an FCF) that is being logged out.

1.5.3.6 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.

1.5.3.7 Link Reset

The FIP Link Reset operation may be used by an FCF to remove all VN_Ports, VF_Ports, and Virtual Links between and FCF-MAC and multiple ENode MACs. This may be used, for example, to remove state when a loss of connectivity is detected.

A FIP Link Reset operation shall be transmitted from an FCF MAC to a ENode MAC after a period of 10 * D_A_TOV seconds have elapsed without receiving any LKA ELS from one or more of the VN_Ports associated with the ENode MAC. The FIP Link Reset operation shall contain all of the MAC addresses of each VN_Port that from which no traffic has been received for this time period.

Editor's Note: This requires that VN_Ports transmit LKAs at a minimal rate. Such a normative statement needs to be made elsewhere in the specification (this is not exactly a FIP requirement).

The FIP Link Reset operation shall be generated only by FCFs. Enodes shall not generate a FIP Link Reset operation. The originating FCF shall set the first MAC_Address descriptor to the MAC address of the FCF MAC generating the request. It shall set the Switch_Node_Name to the Switch_Name of the originating FCF. It shall set the remaining MAC_Address descriptors to the VN_Port MAC addresses that are to be reset. The frame shall be unicast and addressed to MAC address used by the ENode for FIP operations.

Editor's Note: Part of the intent of this process is to enable snooping of this frame by intervening Ethernet bridges to clean up ACLs.

Upon reception of a FIP Link Reset operation, an FCF shall verify that the frame was received from a peer FCF by verifying that the Switch_Node_Name and the MAC_Address Descriptors are valid. If so, the FCF shall remove all of the VN_Port state for each VN_Port indicated by a MAC_Address descriptor. If the Switch_Node_Name or MAC_Address Descriptor of the peer FCF is invalid, then the frame shall be discarded.

FIP Link Reset operations shall be ignored by any device other than an FCF.

1.6 FCF Discovery Process *(informative)*

The precise process by which an ENode or FCF discovers available FCFs and decides to which of the available FCFs it wishes to login or establish VE_Ports with is beyond the scope of this standard. This clause describes a process using FIP operations that may be used for this purpose; other methods are also permissible.

As noted in clause 1.5.3.3, FCFs periodically send out Discovery Advertisements to all ENodes and FCFs using the ALL_ENODE_MACS and ALL_FCF_MACS group addresses. Monitoring these advertisements, an ENode or FCF may produce an Available FCF List (AFL). The longest time an FCF may wait before sending another advertisement is 15 * D_A_TOV seconds. Therefore, waiting at least 30 * D_A_TOV seconds will provide reasonable assurance that all advertisements are received.

Each entry in the AFL contains the data provided by the advertisement and a flag to indicate "Max Receive Size Verified." This flag is set FALSE for entries that are created from the receipt of unsolicited Discovery Advertisements.

From this list, ENodes select a subset of the AFL for Login and FCFs select a subset of the AFL for establishing VE_Ports. Unless otherwise administratively configured, the set contains the entries with the highest priority (indicated by the lowest value of Priority). For each member of this set, a Discov-

ery Solicitation is sent to the corresponding FCF. When the solicited Discovery Advertisement is received, the data in the AFL is updated and the “Max. Receive Size Verified” flag is set TRUE. (Solicited Discovery Advertisements are padded to the maximum receive size; therefore, reception of a Solicited Discovery Advertisement verifies that the path supports the frame size). The solicited Discovery Advertisement is required to be sent within D_A_TOV seconds of receipt of the Discovery Solicitation. Thus if no solicited Discovery Advertisement is received within 2 * D_A_TOV seconds, it may be reasonably assumed that either the path does not support the maximum frame size or that either the Discovery Solicitation or the solicited Discovery Advertisement was lost by the fabric for some reason (e.g. FCS error). A few repeated attempts without receiving the solicited Discovery Advertisement provides reasonable assurance that the fabric will not support the maximum length frame between the ENode and that particular FCF.

The ENode may then log into the FCFs in the AFL for which the A Flags bit is set to 1 (see 1.5.1) and the Max. Receive Size Verified flag is set TRUE. Likewise, an FCF may commence FIP ELP processing with FCFs in the AFL for which the A Flags bit is set to 1 and the MAx. Receive Size Verified flag is set TRUE.

Note that this process may take some time (on the order of 30 * D_A_TOV seconds). The time required for this process may be significantly reduced. To accomplish this, the ENode or FCF may send a Discovery Solicitation to the ALL_FCF_MACS group address. This will cause all FCFs to respond with a solicited Discovery Advertisement. Waiting for approximately 2 * D_A_TOV seconds will provide reasonable assurance that all solicited Discovery Advertisements are received. It may be desirable to repeat the process to discover FCFs that were not discovered in the first attempt due to the failure of the fabric to deliver the frames (e.g. FCF errors).

From the received solicited Discovery Advertisements, the AFL is created. In this case, the Max. Receive Size Verified flag may be set to TRUE, thus eliminating the need to perform extra Discovery Solicitations. After the AFL is created, an ENode may log into the FCFs within the AFL that have the highest priority (indicated by the lowest value of Priority), or other FCFs within the AFL based on administrative configuration. Similarly, FCFs may commence FIP ELP processing with the FCFs within the AFL that have the highest priority (indicated by the lowest value of Priority), or other FCFs within the AFL based on administrative configuration.