

BB-6 modifications for Multiple FCFs on one VLAN

11-515v0

Bill Martin (Emulex)

Erik Smith (EMC)

Issue:

This is proposed modifications of FC-BB-6 based on the issue presented in 11-480v0. The remainder of this proposal is a PDF extraction from FC-BB-6 version revision 1.11

Revision History:

Revision 0 – 5 December 2011

Original proposal

The EOF field specifies the EOF delimiter for the encapsulated FC frame. The value of the EOF field shall be set as specified in table 24.

Table 24 – FCoE EOF field

Value	Description	Reference
41h	EOFn	FC-FS-3
42h	EOFt	FC-FS-3
49h	EOFni	FC-FS-3
50h	EOFa	FC-FS-3

7.9 FC-BB_E device initialization

7.9.1 FCoE Initialization Protocol (FIP) overview

The FCoE Initialization Protocol (FIP) is used to perform the functions of FC-BB_E device discovery, initialization, and maintenance. To perform these functions, encapsulated FIP operations (see 7.9.7.2) are specified.

The FIP Ethernet Type (see 7.9.7.1) is different than the FCoE Ethernet Type (see 7.8) to enable the distinction of discovery, initialization, and maintenance traffic from other FCoE traffic.

FIP frames are used to perform the following protocols:

- a) FIP VLAN discovery (see 7.9.2);
- b) FIP discovery (see 7.9.3);
- c) FCoE Virtual Link instantiation (see 7.9.4);
- d) FCoE Virtual Link maintenance (see 7.9.5); and
- e) Locally Unique N_Port_IDs (see 7.9.6).

All FIP protocols are performed on a per-VLAN basis. It is recommended to use the FIP VLAN Discovery protocol on the default VLAN (see IEEE 802.1Q-2005). All other FIP protocols shall be performed in each VLAN that provides FC-BB_E services.

In order to provide FC-BB_E services on a VLAN, FCoE and FIP protocols, other than FIP VLAN discovery, shall both be performed on that VLAN. ~~Support for multiple Fabrics per VLAN is outside the scope of this standard.~~

~~NOTE 16 – The security provisions of this standard are not sufficient to ensure that FCoE frames remain associated with the correct Fabric if multiple Fabrics are used on the same VLAN.~~

On ENodes, the ENode MAC address shall be used for all FIP frames, except the VN_Port FIP Keep Alive frame (see 7.9.8.5). On FCFs, the FCF-MAC address shall be used for all FIP frames.

ENode MACs shall listen to the All-ENode-MACs group address, FCF-MACs shall listen to the All-FCF-MACs group address, and both ENode MACs and FCF-MACs shall listen to the All-FCoE-MACs group address.

An ENode MAC shall discard a FIP message destined to an address other than its ENode MAC address or the All-ENode-MACs address.

7.9.2 FIP VLAN discovery protocol

When becoming operational, an ENode MAC or an FCF-MAC may invoke the FIP VLAN discovery protocol to discover the VLANs in the Lossless Ethernet network that provide FC-BB_E services. The FIP VLAN discovery protocol is not needed if these VLANs are already known or if VLANs are not used.

An ENode MAC may send a FIP VLAN Request frame to the All-FCF-MACs MAC address over an available VLAN (e.g., the port VLAN). VF_Port capable FCF-MACs that receive a FIP VLAN Request frame shall respond with a unicast FIP VLAN Notification frame over the same VLAN. The FIP VLAN Notification frame should provide the list of VLAN IDs over which the originating FCF offers FC-BB_E services. The ENode MAC that received a FIP VLAN Notification frame may enable one or more of these VLANs for subsequent operations. VF_Port capable FCF-MACs may limit the number of VLAN IDs listed in a FIP VLAN Notification frame on a per-requester basis.

A VF_Port capable FCF-MAC shall discard a multicast FIP VLAN Request frame that has a source address equal to its FCF-MAC address, and the FIP VLAN Request frame should be reported in a vendor specific way as an indication of a MAC address duplication.

If the configuration of VLANs on which a VF_Port capable FCF-MAC supports FC-BB_E services changes, that FCF-MAC should send a unicast FIP VLAN Notification frame to each ENode MAC address with which that FCF-MAC has established VN_Port to VF_Port Virtual Links. The unicast FIP VLAN Notification frame shall carry the revised list of VLAN IDs over which the originating VF_Port capable FCF-MAC offers FC-BB_E services.

A VE_Port capable FCF-MAC may send a FIP VLAN Request frame to the MAC address All-FCF-MACs over an available VLAN (e.g., the default VLAN). VE_Port capable FCF-MACs that receive a FIP VLAN Request frame shall respond with a unicast FIP VLAN Notification frame over the same VLAN. The FIP VLAN Notification frame carries the list of VLAN IDs over which the originating FCF offers FC-BB_E services. The VE_Port capable FCF-MAC that received a FIP VLAN Notification frame may enable one or more of these VLANs for subsequent operations.

A VE_Port capable FCF-MAC shall discard a multicast VLAN Request frame that has a source address equal to its FCF-MAC address. Such a VLAN Request frame should be reported in a vendor specific way as an indication of a MAC address duplication.

FCF-MACs shall listen to the All-FCF-MACs group address in the default VLAN and in other VLANs that ENodes or FCFs may use to invoke this protocol.

If the configuration of VLANs on which a VE_Port capable FCF-MAC supports FC-BB_E services changes, that FCF-MAC should send a unicast FIP VLAN Notification frame to each FCF-MAC address with which that FCF-MAC has established VE_Port to VE_Port Virtual Links. The unicast FIP VLAN Notification frame shall specify the revised list of VLAN IDs over which the originating VE_Port capable FCF-MAC offers FC-BB_E services.

7.9.3 FIP discovery protocol

7.9.3.1 Overview

On a network deploying multiple VLANs, the FIP discovery protocol is performed in the VLANs where FC-BB_E services are offered when these VLANs are known (e.g., upon performing the FIP VLAN discovery protocol (see 7.9.2)).

7.9.3.2 ENode/FCF discovery

The FCoE Controller of a VF_Port capable FCF-MAC shall periodically transmit multicast Discovery Advertisements (see 7.9.8.3) to the All-ENode-MACs group address every FKA_ADV_PERIOD. The FKA_ADV_PERIOD period shall be randomized by adding a random delay uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network. The FCoE Controller of a VF_Port capable FCF-MAC should begin transmitting unsolicited multicast Discovery Advertisements on completion of Fabric configuration (see FC-SW-5).

The FCoE Controller of an ENode MAC shall discard incompatible Discovery Advertisements and shall create an entry for each compatible FCF-MAC in an internal FCF list.

NOTE 17 – The internal data structures used to describe this protocol are a model to express the behavior, not an implementation requirement.

Each entry in the FCF list has the following flags:

- a) 'Max FCoE Size Verified' - set to zero for entries created from unsolicited multicast Discovery Advertisements, set to one when a solicited unicast Discovery Advertisement is received; and
- b) 'Available for Login' - reflects the value of the A bit provided by the most recently received Discovery Advertisement from that VF_Port capable FCF-MAC.

The FCoE Controller of an ENode MAC selects for login a subset of the FCF-MACs in the FCF list having the 'Available for Login' flag set to one (i.e., the FCF Login Set) on the basis of a local policy that should default to selecting the one(s) with higher priority (i.e., lower priority value) in the absence of explicit configuration of other selection criteria. A FIP FLOGI may be performed with an FCF-MAC in the FCF Login Set only if its 'Max FCoE Size Verified' flag is set to one. In order to perform a FIP FLOGI with an FCF-MAC in the FCF Login Set with the 'Max FCoE Size Verified' flag set to zero, the FCoE Controller of an ENode MAC shall transmit a unicast Discovery Solicitation (see 7.9.8.2) to that FCF-MAC address and receive a solicited unicast Discovery Advertisement in response.

The periodic reception of unsolicited multicast Discovery Advertisements allows the FCoE Controller of ENode MACs to continuously verify FCF-MAC connectivity. The Available for Login (A) bit in received Discovery Advertisements provides the information that the transmitting FCF-MAC is available for FIP FLOGI/FDISC, and this information is updated in the FCF list and FCF Login Set on reception of Discovery Advertisements. The A bit is informational and shall have no effect on existing logins.

When the FCoE Controller of an ENode MAC becomes operational it should discover VF_Port capable FCF-MACs with which it may perform FIP FLOGI by transmitting a multicast Discovery Solicitation to the All-FCF-MACs group address. In response to a Discovery Solicitation from an ENode MAC, a VF_Port capable FCF-MAC shall transmit a solicited unicast Discovery Advertisement to the soliciting ENode MAC if it is configured to allow a FIP FLOGI from that ENode. The solicited unicast Discovery Advertisement shall be transmitted to the MAC address specified in the MAC address descriptor in the received Discovery Solicitation. The solicited unicast Discovery Advertisement shall be transmitted within ADV_TOV (see table 54) upon reception of the Discovery Solicitation. Discovery Advertisements transmitted in response to a multicast Discovery Solicitation should be delayed by a random time uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network. This delay should not be applied to solicited unicast Discovery Advertisements sent in response to unicast Discovery Solicitations.

Solicited unicast Discovery Advertisements should not be transmitted until Fabric configuration (see FC-SW-5) is completed.

NOTE 18 – An ENode MAC may also wait to receive unsolicited multicast Discovery Advertisements and then send unicast Discovery Solicitations to the FCF-MACs selected for login from the FCF Login set.

A Discovery Solicitation shall carry in the Max FCoE Size descriptor the maximum FCoE PDU size the ENode MAC intends to use for FCoE traffic (see 7.9.7.3.7). The FIP PDU (see table 26) in a solicited unicast Discovery Advertisement shall be extended to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.9.8.3).

NOTE 19 – If an ENode transmits an FCoE frame with an FCoE PDU size that is greater than its maximum FCoE PDU size, the network may not deliver the FCoE frame.

An ENode MAC may generate multiple Discovery Solicitations.

NOTE 20 – As an example, an ENode MAC that does not receive a solicited unicast Discovery Advertisement in response to a Discovery Solicitation may transmit additional Discovery Solicitations, unicast or multicast.

Reception of a solicited unicast Discovery Advertisement from an FCF-MAC shall set the 'Max FCoE Size Verified' flag to one in the entry for that FCF-MAC in the FCF Login Set of an ENode MAC.

It is possible for an FCF to receive a multicast Discovery Solicitation from the same ENode MAC on multiple FCF-MACs. In this case, a separate solicited unicast Discovery Advertisement shall be transmitted by each of the FCF-MACs that received the Discovery Solicitation. The ENode MAC that transmitted the multicast Discovery Solicitation is able to determine that it received multiple solicited unicast Discovery Advertisements from the same FCF since the value of the Name_Identifier field in the Name_Identifier descriptor is the same in each of the solicited unicast Discovery Advertisements (see 7.9.8.3). In this case the ENode MAC should select the FCF-MAC for Fabric login with that FCF based on the value of the Priority descriptor in the Discovery Advertisements.

It is possible for an ENode MAC to receive multiple unsolicited multicast Discovery Advertisements from multiple FCF-MACs of the same FCF. The ENode MAC is able to determine that those unsolicited multicast Discovery Advertisements are from the same FCF since the value of the Name_Identifier field in the Name_Identifier descriptor is the same in each of the unsolicited multicast Discovery Advertisements (see 7.9.8.3). In this case the ENode MAC should select the FCF-MAC for Fabric login with that FCF based on the value of the Priority descriptor in the Discovery Advertisements.

An ENode MAC shall discard any received Discovery Solicitation. A VF_Port capable FCF-MAC shall discard any Discovery Solicitation originated by a VE_Port capable FCF-MAC (i.e., having the F bit set to one (see 7.9.7.2)).

An ENode MAC shall discard an unsolicited multicast Discovery Advertisement that has a source address equal to its ENode MAC address. Such a Discovery Advertisement should be reported in a vendor specific way as an indication of a MAC address duplication.

A VF_Port capable FCF-MAC shall discard a multicast Discovery Solicitation that has a source address equal to its FCF-MAC address. Such a Discovery Solicitation should be reported in a vendor specific way as an indication of a MAC address duplication.

Reception of Discovery Advertisements for more than one Fabric on the same VLAN should be reported by an ENode MAC in a vendor specific manner and no subsequent VN_Port to VF_Port Virtual Links should be instantiated.

7.9.3.3 FCF/FCF discovery

The FCoE Controller of a VE_Port capable FCF-MAC shall periodically transmit multicast Discovery Advertisements (see 7.9.8.3) to the All-FCF-MACs group address every FKA_ADV_PERIOD. The FKA_ADV_PERIOD period shall be randomized by adding a random delay uniformly distributed between 0 and 100 ms to avoid synchronized bursts of multicast traffic within the Ethernet network.

On receiving Discovery Advertisements, the FCoE Controller of a VE_Port capable FCF-MAC shall create an entry per FCF-MAC in an internal FCF list.

NOTE 21 – The internal data structures used to describe this protocol are a model to express the behavior, not an implementation requirement.

Each entry in the FCF list has the following flags:

- a) 'Max FCoE Size Verified' - set to zero for entries created from unsolicited multicast Discovery Advertisements, set to one when a solicited unicast Discovery Advertisement is received; and
- b) 'Available for ELP' - reflects the value of the A bit provided by the most recently received Discovery Advertisement from that VE_Port capable FCF-MAC.

A FIP ELP may be performed with an FCF-MAC in the FCF list only if its 'Max FCoE Size Verified' flag is set to one. In order to perform a FIP ELP with an FCF-MAC in the FCF list with the 'Max FCoE Size Verified' flag set to zero, the FCoE Controller of a VE_Port capable FCF-MAC shall transmit a unicast Discovery Solicitation (see 7.9.8.2) to that FCF-MAC address and receive a solicited unicast Discovery Advertisement in response.

The periodic reception of unsolicited multicast Discovery Advertisements allow the FCoE Controller of VE_Port capable FCF-MACs to continuously verify the FCF-MACs connectivity. The 'Available for Login' (A) bit in received Discovery Advertisements provides the information that the transmitting FCF-MAC is available for FIP ELP, and this information is updated in the FCF list on reception of Advertisements. The A bit is informational and shall have no effect on existing VE_Port to VE_Port Virtual Links.

When the FCoE Controller for a VE_Port capable FCF-MAC becomes operational it **should** discover other VE_Port capable FCF-MACs by transmitting a multicast Discovery Solicitation to the All-FCF-MACs group address. In response to a Discovery Solicitation from an FCF-MAC, a VE_Port capable FCF-MAC shall transmit a solicited unicast Discovery Advertisement to the soliciting FCF-MAC if the FC-MAP value in the Discovery Solicitation is compatible with the FC-MAP configured on the FCF and if it is configured to allow a Virtual Link with that FCF.



The solicited unicast Discovery Advertisement shall be transmitted to the MAC address specified in the MAC address descriptor in the received Discovery Solicitation. The solicited unicast Discovery Advertisement shall be transmitted within ADV_TOV (see table 54) upon reception of the Discovery Solicitation. Discovery Advertisements transmitted in response to a multicast Discovery Solicitation should be delayed by a random time uniformly distributed between 0 and 100 ms to avoid

synchronized bursts of multicast traffic within the Ethernet network. This delay should not be applied to solicited unicast Discovery Advertisements sent in response to unicast Discovery Solicitations.

NOTE 22 – A VE_Port capable FCF-MAC may also wait to receive unsolicited multicast Discovery Advertisements and then send unicast Discovery Solicitations to the FCF-MACs in the FCF list.

A Discovery Solicitation shall specify in the Max FCoE Size descriptor (see 7.9.7.3.7) the maximum FCoE PDU size the VE_Port capable FCF-MAC intends to use for FCoE traffic. The FIP PDU (see table 26) in a solicited unicast Discovery Advertisement shall be extended to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.9.8.3).

NOTE 23 – If a VE_Port capable FCF-MAC transmits an FCoE frame with an FCoE PDU size that is greater than its maximum FCoE PDU size, the network may not deliver the FCoE frame.

A VE_Port capable FCF-MAC may generate multiple Discovery Solicitations.

NOTE 24 – As an example, a VE_Port capable FCF-MAC that does not receive a solicited unicast Discovery Advertisement in response to a Discovery Solicitation may transmit additional Discovery Solicitations.

Reception of a solicited unicast Discovery Advertisement from an FCF-MAC shall set the 'Max FCoE Size Verified' flag to one in the entry for that FCF-MAC in the FCF list of the receiving VE_Port capable FCF-MAC.

It is possible for an FCF to receive multicast Discovery Solicitations from the same VE_Port capable FCF-MAC on multiple FCF-MACs. In this case, a separate solicited unicast Discovery Advertisement shall be transmitted by each of the FCF-MACs that received the Discovery Solicitation. The VE_Port capable FCF-MAC that transmitted the multicast Discovery Solicitation is able to determine that it received multiple solicited unicast Discovery Advertisements from the same FCF since the value of the Name_Identifier field in the Name_Identifier descriptor is the same in each of the solicited unicast Discovery Advertisements (see 7.9.8.3).

It is possible for a VE_Port capable FCF-MAC to receive multiple unsolicited multicast Discovery Advertisements from multiple FCF-MACs of the same FCF. The VE_Port capable FCF-MAC is able to determine that those unsolicited multicast Discovery Advertisements are from the same FCF since the value of the Name_Identifier field in the Name_Identifier descriptor is the same in each of the unsolicited multicast Discovery Advertisements (see 7.9.8.3).

After receiving a Discovery Solicitation originated by an FCF (i.e., the F bit is set to one), an FCF-MAC shall perform the following verification checks:

- a) the Name_Identifier field value in the Discovery Solicitation is different than the Switch_Name of the recipient FCF (see 7.9.8.2.2); and
- b) the FP bit is set to one (see 7.9.7.2) and the FC-MAP value in the FC-MAP descriptor in the Discovery Solicitation is the same as the FC-MAP value of the recipient FCF.

If any verification check is false, then the Discovery Solicitation shall be discarded.

After receiving a Discovery Advertisement, an FCF-MAC shall perform the following verification checks:

- a) the Name_Identifier field value in the Discovery Advertisement is different than the Switch_Name of the recipient FCF (see 7.9.8.3); and

- b) the FP bit is set to one (see 7.9.7.2) and the FC-MAP value in the Fabric descriptor in the Discovery Advertisement is the same as the FC-MAP value of the recipient FCF.

If any verification check is false, then the Discovery Advertisement shall be discarded.

NOTE 25 – It is possible for an FCF to receive a multicast Discovery Solicitation or a multicast Discovery Advertisement that it originated because FIP frames sent to the All-FCF-MACs group address may be forwarded to other ports on the same FCF by intermediate Ethernet bridges.

A VE_Port capable FCF-MAC shall discard any Discovery Solicitation originated by an ENode (i.e., having the F bit set to zero (see 7.9.7.2)).

A VE_Port capable FCF-MAC shall discard a multicast Discovery Solicitation that has a source address equal to its FCF-MAC address. Such a Discovery Solicitation should be reported in a vendor specific way as an indication of a MAC address duplication.

Reception of Discovery Advertisements for more than one Fabric on the same VLAN should be reported by VE_Port capable FCF-MAC in a vendor specific manner and no subsequent VE_Port to VE_Port Virtual Links should be instantiated.

7.9.4 FCoE Virtual Link instantiation protocol

7.9.4.1 VN_Port to VF_Port Virtual Links

The FCoE Controller of an ENode MAC instantiates VN_Port to VF_Port Virtual Links on successful completion of a FIP Fabric login request. Fabric login (i.e., FLOGI, NPIV FDISC) shall be performed using FIP frames (see table 26) and the associated FIP descriptor type (see table 31). Fabric login (i.e., FLOGI, NPIV FDISC) shall not be performed using FCoE frames .

In addition to providing Fabric login, the FIP Fabric login provides a method to assign a MAC address for the VN_Port (see 7.9.8.4.2).

When the FCoE Controller of an ENode MAC transmits a FIP FLOGI Request or FIP NPIV FDISC Request it shall indicate that it intends to use FPMAs as VN_Port MAC addresses (i.e., set the FP bit to one and the SP bit to zero). The MAC address returned by the FCF in a FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC shall be used as the VN_Port MAC address (see 7.7).

The FCF shall return a properly formed FPMA MAC address in the FIP FLOGI LS_ACC or FIP NPIV FDISC LS_ACC (see 7.7).

Explicit VN_Port to VF_Port Virtual Link de-instantiation is performed by an ENode MAC by performing Fabric logout. Fabric logout (i.e., Fabric LOGO) shall be performed by an ENode using FIP frames (see table 26) and the associated FIP descriptor type (see table 31). Fabric logout shall not be performed using FCoE frames.

In addition to providing Fabric logout, the FIP Fabric logout provides a method to de-assign a MAC address for the VN_Port (see 7.9.8.4.3).

7.9.4.2 VE_Port to VE_Port Virtual Links

The FCoE Controller of a VE_Port capable FCF-MAC instantiates VE_Port to VE_Port Virtual Links on successful completion of a FIP ELP request. ELP shall be performed using FIP frames (see table 26) and the associated FIP descriptor type (see table 31). ELP shall not be performed using FCoE frames.

In addition to providing ELP, the FIP ELP provides a method to communicate the MAC address for the VE_Port (see 7.9.8.4.4).

7.9.4.3 VN_Port to VN_Port Virtual Links

A VN2VN ENode MAC, operating in either multipoint or point-to-point mode, instantiates VN_Port to VN_Port Virtual Links on successful completion of a point-to-point FIP FLOGI, as defined in FC-LS-2. Both FIP FLOGI Request and LS_ACC shall have the Locally Unique N_Port_ID of the originating VN2VN_Port as S_ID, the Locally Unique N_Port_ID of the destination VN2VN_Port as D_ID for the point-to-point FLOGI protocol, and the originating VN_Port FPMA in the MAC Address descriptor. The MAC addresses of the FIP FLOGI Request and LS_ACC shall be the ENode MAC addresses of the involved VN2VN_Ports. As specified in FC-LS-2, the VN2VN_Port with the greater N_Port_Name proceeds to N_Port Login, with the PLOGI ELSs encapsulated in FCoE. Both FCoE PLOGI Request and LS_ACC shall have the Locally Unique N_Port_ID of the originating VN2VN_Port as S_ID and the Locally Unique N_Port_ID of the destination VN2VN_Port as D_ID. Upon completion of FCoE PLOGI the involved VN_Ports operate in point-to-point mode (see FC-LS-2).

A point-to-point FIP FLOGI Request coming from a VN2VN_Port not listed in the VN2VN Neighbor Set shall be rejected with reason code 'FIP Error' (i.e., 20h) and reason code explanation 'VN2VN_Port Not in Neighbor Set' (i.e., 62h, see FC-LS-2).

A VN_Port to VN_Port Virtual Link is explicitly deinstantiated by performing a FIP LOGO, that deinstantiates the FCoE_LEPs and performs a N_Port logout. The S_ID and D_ID on the encapsulated LOGO ELS shall be set to the Locally Unique N_Port_IDs of the involved VN2VN_Ports.

7.9.5 FCoE Virtual Link maintenance protocol

7.9.5.1 Virtual Link maintenance protocol overview

VN_Port to VF_Port Virtual Links (see figure 30) and VE_Port to VE_Port Virtual Links (see figure 31) overlay over a Lossless Ethernet network. The Virtual Link maintenance protocol specifies how to deal with faults that may occur in a Lossless Ethernet network.

7.9.5.2 VN_Port to VF_Port Virtual Link maintenance protocol

To deal with local physical layer faults, an ENode MAC shall de-instantiate all its VN_Ports to VF_Port Virtual Links upon detecting that its physical layer is not operational. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports. A VF_Port capable FCF-MAC shall de-instantiate all its VF_Ports upon detecting that its physical layer is not operational.

To deal with non-local faults, the FCoE Controllers of an ENode MAC and of a VF_Port capable FCF-MAC shall continuously verify the state of the VN_Port to VF_Port Virtual Link by transmitting appropriate FIP frames and by verifying received FIP frames. This behavior may be disabled by VF_Port capable FCF-MACs under administrative control by setting to one the D bit in the FKA_ADV_Period descriptor in Discovery Advertisements (see 7.9.7.3.13). The D bit in the FKA_ADV_Period descriptor may be set to one only in a direct-attach topology (i.e., when an ENode is directly connected to an FCF without any intermediate Ethernet bridges).

The FCoE Controller of an ENode MAC shall transmit a unicast FIP Keep Alive frame on behalf of the ENode MAC (i.e., with the ENode MAC address as source MAC address and without a Vx_Port Identification descriptor in the FIP Descriptor list (see table 52)) to each VF_Port capable FCF-MAC with which it has VN_Ports logged in. This ENode FIP Keep Alive frame shall be transmitted every

FKA_ADV_PERIOD. The FKA_ADV_PERIOD is obtained from the Discovery Advertisements received from the VF_Port capable FCF-MACs with which the ENode MAC has VN_Ports logged in.

In addition, the FCoE Controller of an ENode MAC shall transmit a unicast FIP Keep Alive frame on behalf of each VN_Port (i.e., with the VN_Port MAC address as source MAC address and containing a Vx_Port Identification descriptor for that VN_Port in the FIP Descriptor list, see table 52) to the VF_Port capable FCF-MAC with which that VN_Port is logged in. This VN_Port FIP Keep Alive frame shall be transmitted every FKA_VN_PERIOD.

The FCoE Controller of an ENode MAC shall monitor the status of a VF_Port with which it has VN_Ports logged in by verifying reception of unsolicited multicast Discovery Advertisements from that VF_Port capable FCF-MAC. Unsolicited multicast Discovery Advertisements are expected to be received every FKA_ADV_PERIOD. If unsolicited multicast Discovery Advertisements are not received within $2.5 * FKA_ADV_PERIOD$, all the VN_Port to VF_Port Virtual Links with that VF_Port shall be implicitly de-instantiated. This condition should be counted as a Virtual Link failure and shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports. That FCF-MAC shall be removed from the FCF Login Set (see 7.9.3.2). A subsequent FIP Fabric Login may be performed with an FCF-MAC in the current FCF Login Set as specified in see 7.9.3.2.

The FCoE Controller of a VF_Port capable FCF-MAC shall transmit an unsolicited multicast Discovery Advertisement to the All-ENode-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller of a VF_Port capable FCF-MAC shall monitor the status of an ENode MAC with which it has active VN_Port to VF_Port Virtual Links by verifying the reception of FIP Keep Alive frames from that ENode MAC and its VN_Ports. VN_Port FIP Keep Alive frames (i.e., those containing a Vx_Port Identification descriptor) are expected to be received every FKA_VN_PERIOD and ENode FIP Keep Alive frames (i.e., those not containing a Vx_Port Identification descriptor) are expected to be received every FKA_ADV_PERIOD.

If VN_Port FIP Keep Alive frames are not received within $2.5 * FKA_VN_PERIOD$, the associated VN_Port to VF_Port Virtual Link shall be explicitly de-instantiated (i.e., a FIP Clear Virtual Links frame listing the unreachable VN_Port shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Port. If ENode FIP Keep Alive frames are not received within $2.5 * FKA_ADV_PERIOD$, all associated VN_Port to VF_Port Virtual Links shall be explicitly de-instantiated (i.e., a FIP Clear Virtual Links frame listing all the unreachable VN_Ports shall be generated). This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

NOTE 26 – The use of a faster FKA_ADV_PERIOD rate for ENode FIP Keep Alive frames allows fast response times against loss of connectivity in the Ethernet realm with limited overhead. The use of VN_Port FIP Keep Alive frames transmitted at a slower FKA_VN_PERIOD rate allows the clearing of state associated with each individual VN_Port when that VN_Port becomes not operational.

Explicit VN_Port to VF_Port Virtual Link de-instantiation is invoked by a VF_Port capable FCF-MAC by transmitting a FIP Clear Virtual Links frame (see 7.9.8.6). A FIP Clear Virtual Links frame transmitted to an ENode MAC with logged in VN_Ports provides the list of VN_Ports to be removed. An ENode MAC shall de-instantiate the VN_Ports listed in a FIP Clear Virtual Links frame upon reception of the FIP frame. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

The size of a FIP Clear Virtual Links frame shall not exceed the standard Ethernet MAC Client Data size (i.e., 1 500 bytes for basic frames and 1 504 bytes for Q-tagged frames, see IEEE 802.3-2008). If the list of VN_Ports to be removed does not fit in a single FIP frame, multiple FIP frames should be transmitted to convey the entire list.

On receiving a VN_Port FIP Keep Alive frame coming from a VN_Port that is not logged, the FCoE Controller of a VF_Port capable FCF-MAC shall transmit a FIP Clear Virtual Links frame listing that VN_Port.

On receiving an ENode FIP Keep Alive frame coming from an ENode MAC that is not logged in, the FCoE Controller of a VF_Port capable FCF-MAC shall transmit a FIP Clear Virtual Links frame listing no VN_Ports. A FIP Clear Virtual Links frame listing no VN_Ports shall be handled by an ENode MAC by de-instantiating all VN_Port to VF_Port Virtual Links with that VF_Port capable FCF-MAC. This condition shall be handled as an implicit Fabric logout (see FC-LS-2) for the involved VN_Ports.

FIP Clear Virtual Links frames may be generated by FCFs whenever appropriate to speed-up fault recovery.

NOTE 27 – As an example, in certain topologies an FCF may generate a FIP Clear Virtual Links frame to de-instantiate the VN_Port to VF_Port Virtual Links affected by a local physical layer fault on other ports upon detection of that fault.

The FKA_ADV_PERIOD value (see table 54) may be changed on a FCF under administrative control. When this happens, each VF_Port capable FCF-MAC of the FCF shall advertise the updated FKA_ADV_PERIOD in subsequent unsolicited multicast Discovery Advertisements.

When the FKA_ADV_PERIOD value is decreased, a VF_Port capable FCF-MAC shall transmit unsolicited multicast Discovery Advertisements at the interval specified by the updated value, but shall not use the updated value for detection of missing ENode FIP Keep Alives until five times the old value has elapsed since the transmission of the first updated unsolicited multicast Discovery Advertisement.

When the FKA_ADV_PERIOD value is increased, a VF_Port capable FCF-MAC shall transmit unsolicited multicast Discovery Advertisements at the interval specified by the old value until five times the updated value has elapsed since the transmission of the first updated unsolicited multicast Discovery Advertisement, but shall use the updated value for detection of missing ENode FIP Keep Alives.

On detecting the updated value, an ENode having VN_Port to VF_Port Virtual Links instantiated with that FCF shall transmit ENode FIP Keep Alive frames at the interval specified by the updated FKA_ADV_PERIOD value and shall use the updated value for detection of missing unsolicited multicast Discovery Advertisements.

7.9.5.3 VE_Port to VE_Port Virtual Link maintenance protocol

To deal with local physical layer faults, a VE_Port capable FCF-MAC shall de-instantiate all its VE_Port to VE_Port Virtual Links upon detecting that its physical layer is not operational.

To deal with non-local faults, the FCoE Controllers for VE_Port capable FCF-MACs shall continuously verify the state of a VE_Port to VE_Port Virtual Link by transmitting unsolicited multicast Discovery Advertisements and by verifying received unsolicited multicast Discovery Advertisements.

The FCoE Controller for a VE_Port capable FCF-MAC shall transmit a Discovery Advertisement to the All-FCF-MACs group address every FKA_ADV_PERIOD.

The FCoE Controller for a VE_Port capable FCF-MAC shall monitor the status of a VE_Port to VE_Port Virtual Link by verifying the reception of unsolicited multicast Discovery Advertisements. Unsolicited multicast Discovery Advertisements are expected to be received every FKA_ADV_PERIOD. If unsolicited multicast Discovery Advertisements are not received within 2.5 *

FKA_ADV_PERIOD, the VE_Port to VE_Port Virtual Link associated with that FCF-MAC shall be explicitly de-instantiated along with the associated VE_Port. That FCF-MAC shall be removed from the FCF List (see 7.9.3.3).

Explicit VE_Port to VE_Port Virtual Link de-instantiation is performed by transmitting a FIP Clear Virtual Links frame from a VE_Port capable FCF-MAC. A FIP Clear Virtual Links frame transmitted to a VE_Port capable FCF-MAC provides the destination FCF-MAC address in the Vx_Port Identification descriptor. A VE_Port capable FCF-MAC shall de-instantiate the VE_Port listed in a FIP Clear Virtual Links frame upon reception of the FIP frame.

The FKA_ADV_PERIOD value (see table 54) may be changed on a FCF under administrative control. When this happens, each VE_Port capable FCF-MAC of the FCF shall advertise the updated FKA_ADV_PERIOD in subsequent unsolicited multicast Discovery Advertisements and shall transmit them at the interval specified by the updated value.

On detecting the updated value in received unsolicited multicast Discovery Advertisements, a VE_Port capable FCF-MAC shall use the updated value for detection of missing unsolicited multicast Discovery Advertisements.

7.9.5.4 VN_Port to VN_Port Virtual Link maintenance protocol

To deal with local physical layer faults, a VN2VN ENode MAC shall de-instantiate all its VN_Port to VN_Port Virtual Links upon detecting that its physical layer is not operational.

To deal with non-local faults, the FCoE Controller of a VN2VN ENode MAC operating in multipoint mode shall continuously verify the state of the VN_Port to VN_Port Virtual Links by verifying received N_Port_ID Beacons per each VN2VN_Port in the VN2VN Login Set. The FCoE Controller of a VN2VN ENode MAC operating in point-to-point mode shall continuously verify the state of the VN_Port to VN_Port Virtual Link by verifying received N_Port_ID P2P Beacons per the VN2VN_Port in the VN2VN Login Set.

N_Port_ID Beacons and N_Port_ID P2P Beacons are expected to be received every BEACON_PERIOD. If N_Port_ID Beacons or N_Port_ID P2P Beacons from a VN2VN_Port are not received within $2.5 * BEACON_PERIOD$, that VN2VN_Port shall be removed from the VN2VN Neighbor Set and from the VN2VN Login Set and VN_Port to VN_Port Virtual Links with that remote VN2VN_Port, if any, shall be implicitly de-instantiated.

7.9.6 Locally Unique N_Port_IDs

7.9.6.1 Overview

The Locally Unique N_Port_IDs protocol enables the establishment of VN_Port to VN_Port Virtual Links across a Lossless Ethernet network. The Locally Unique N_Port_ID protocol supports both multipoint (see 7.9.6.2) and point-to-point operations (see 7.9.6.3).

Two multicast MAC addresses are used by the protocol: All-VN2VN-ENode-MACs, used when a VN2VN ENode MAC operates in multipoint mode, and All-PT2PT-ENode-MACs, used when a VN2VN ENode MAC operates in point-to-point mode.

When becoming operational, a VN2VN ENode MAC shall enable reception of frames sent to both MAC addresses, All-VN2VN-ENode-MACs and All-PT2PT-ENode-MACs, and shall select a tentative Locally Unique N_Port_ID in the range 000001h to 00FFFEh. A tentative Locally Unique N_Port_ID should be a recorded Locally Unique N_Port_ID if available (see 7.9.6.4), otherwise a tentative Locally Unique N_Port_ID should be generated using a pseudo-random number generator with a

uniform distribution in the range 000001h to 00FFFEh. The N_Port_Name of the VN2VN_Port for which a Locally Unique N_Port_ID is being selected should be used as one of the parameters to seed the pseudo-random number generator at its first invocation. In this way, different ENode MACs generate different numbers and a VN2VN_Port is usually associated with the same N_Port_ID each time it is enabled. The method to generate tentative Locally Unique N_Port_IDs shall be chosen so that different VN2VN ENode MACs are not likely to generate the same sequence of numbers (see annex F).

VN ENode MACs are not likely to generate the same sequence of numbers (see annex F).

7.9.6.2 Multipoint operation

7.9.6.2.1 Probing a Locally Unique N_Port_ID

After selecting a tentative Locally Unique N_Port_ID, a VN2VN ENode MAC operating in multipoint mode shall probe the Lossless Ethernet network to which it is connected to verify if the Locally Unique N_Port_ID is already in use. The VN2VN ENode MAC shall wait for a random time interval selected uniformly in the range zero to PROBE_WAIT milliseconds, and shall then transmit two multicast N_Port_ID Probe Requests (see 7.9.8.10) to All-VN2VN-ENode-MACs, spaced PROBE_WAIT milliseconds apart. If the selected tentative Locally Unique N_Port_ID is a recorded one, then the REC/P2P bit of the N_Port_ID Probe Requests shall be set to one, otherwise it shall be set to zero.

If during the period from the beginning of the probing process until ANNOUNCE_WAIT milliseconds after the last N_Port_ID Probe Request is sent the VN2VN ENode MAC receives any N_Port_ID Probe Reply (see 7.9.8.11), N_Port_ID Claim Notification (see 7.9.8.12) claiming the selected tentative Locally Unique N_Port_ID, or N_Port_ID Beacon (see 7.9.8.14) announcing the selected tentative Locally Unique N_Port_ID, then the VN2VN ENode MAC shall select a new random tentative Locally Unique N_Port_ID and repeat the process (i.e., wait for a random time interval and then transmit two N_Port_ID Probe Requests).

If during the same period the VN2VN ENode MAC receives any N_Port_ID Probe Request probing the selected tentative Locally Unique N_Port_ID, then the VN2VN ENode MAC shall check the N_Port_Name carried in the Vx_Port_Identification descriptor of the received N_Port_ID Probe Request and the value of the REC/P2P bit (see 7.9.7.2). If the REC/P2P bit is set to zero and the selected tentative Locally Unique N_Port_ID is a recorded one, then the VN2VN ENode MAC shall keep its selection and reply to the received N_Port_ID Probe Request with a N_Port_ID Probe Reply. If the REC/P2P bit is set to one and the selected tentative Locally Unique N_Port_ID is not a recorded one, then the VN2VN ENode MAC shall select a new random tentative Locally Unique N_Port_ID and repeat the process (i.e., wait for a random time interval and then transmit two N_Port_ID Probe Requests) without replying to the received N_Port_ID Probe Request. If the REC/P2P bit is set to one and the selected tentative Locally Unique N_Port_ID is a recorded one, or the REC/P2P bit is set to zero and the selected tentative Locally Unique N_Port_ID is not a recorded one, then the processing is based on the value of the received N_Port_Name. If the received N_Port_Name is greater than the N_Port_Name of its VN2VN_Port then the VN2VN ENode MAC shall select a new random tentative Locally Unique N_Port_ID and repeat the process (i.e., wait for a random time interval and then transmit two N_Port_ID Probe Requests) without replying to the received N_Port_ID Probe Request, otherwise it shall keep its selection and reply to the received N_Port_ID Probe Request with a N_Port_ID Probe Reply.

N_Port_ID Probe Requests probing a different Locally Unique N_Port_ID, N_Port_ID Claim Notification claiming a different Locally Unique N_Port_ID, and N_Port_ID Beacon announcing a different Locally Unique N_Port_ID received during the same period shall be ignored.

A VN2VN ENode MAC should maintain a counter of the number of Locally Unique N_Port_ID conflicts it has experienced while trying to acquire a Locally Unique N_Port_ID. If the number of conflicts exceeds ten, then the VN2VN ENode MAC should report the situation in a vendor specific way as an indication of potential network failure and shall limit the rate at which it probes for new Locally Unique N_Port_IDs to no more than one new N_Port_ID per RATE_LIMIT_INTERVAL. This is to prevent multicast storms in pathological failure cases, such as a rogue VN2VN ENode MAC that answers all N_Port_ID Probe Requests, causing legitimate VN2VN ENode MACs to go into an infinite loop while attempting to select a usable Locally Unique N_Port_ID.

NOTE 28 – An ENode MAC connecting to a network with 1 300 ENodes has a 98% chance of selecting an unused Locally Unique N_Port_ID on the first try and a 99.96% chance of selecting an unused Locally Unique N_Port_ID within two tries. The probability that it will have to try more than ten times is about 10^{-17} .

At any time, upon receiving a N_Port_ID P2P Claim Notification (see 7.9.8.12) or a N_Port_ID P2P Beacon (see 7.9.8.14), a VN2VN ENode MAC operating in multipoint mode shall transmit a N_Port_ID Probe Request probing its N_Port_ID to All-VN2VN-ENode-MACs to indicate to the sending VN2VN ENode MAC operating in point-to-point mode that it is not in its expected network configuration.

7.9.6.2.2 Claiming a Locally Unique N_Port_ID

If by ANNOUNCE_WAIT milliseconds after the transmission of the last N_Port_ID Probe Request no conflicting N_Port_ID Probe Reply, N_Port_ID Probe Request, N_Port_ID Claim Notification, or N_Port_ID Beacon has been received, then the VN2VN ENode MAC shall claim the selected tentative Locally Unique N_Port_ID by transmitting a multicast N_Port_ID Claim Notification to All-VN2VN-ENode-MACs. A N_Port_ID Claim Notification notifies the other VN2VN ENode MACs on the Lossless Ethernet network of the selected Locally Unique N_Port_ID.

A VN2VN ENode MAC maintains the list of the reachable VN2VN_Ports in a VN2VN Neighbor Set. The VN2VN Neighbor Set is updated upon receiving N_Port_ID Claim Notifications and N_Port_ID Claim Responses (see 7.9.8.13) and checked upon receiving N_Port_ID Beacons.

Upon receiving a N_Port_ID Claim Notification claiming a different Locally Unique N_Port_ID, a VN2VN ENode MAC shall add the notifying VN2VN_Port to its VN2VN Neighbor Set and shall reply with a unicast N_Port_ID Claim Response. The N_Port_ID Claim Response shall be delayed by a random time uniformly distributed between 0 and 100 ms. The N_Port_ID Claim Notification provides the maximum FCoE size the VN2VN ENode MAC intends to use for VN_Port to VN_Port Virtual Links, and the N_Port_ID Claim Response FIP PDU shall have a length equal to the minimum of the FCoE sizes supported by the two involved VN2VN ENode MACs.

Upon receiving a non conflicting N_Port_ID Claim Response a VN2VN ENode MAC shall add the responding VN2VN_Port to its VN2VN Neighbor Set. A VN2VN ENode MAC determines from its VN2VN Neighbor Set the VN2VN ENode MACs with which to establish VN_Port to VN_Port Virtual Links (i.e., the VN2VN Login Set), also on the basis of their FC-4 support. A VN2VN ENode may become ready to instantiate VN_Port to VN_Port Virtual Links upon receiving a N_Port_ID Claim Response. A VN2VN ENode wishing to discover all its VN2VN neighbors in order to select the VN2VN Login Set should wait for ANNOUNCE_WAIT milliseconds (i.e., the N_Port_ID Claim Responses timeout) upon transmitting a N_Port_ID Claim Notification to collect N_Port_ID Claim Responses.

When ready to instantiate VN_Port to VN_Port Virtual Links, a VN2VN ENode MAC shall transmit a multicast N_Port_ID Beacon to All-VN2VN-ENode-MACs and shall continue to transmit multicast N_Port_ID Beacons periodically every BEACON_PERIOD milliseconds. The BEACON_PERIOD period shall be randomized by adding a random delay uniformly distributed between 0 and 100 ms to

avoid synchronized bursts of multicast traffic within the Ethernet network. Upon transmitting the first N_Port_ID Beacon the VN2VN ENode MAC may instantiate VN_Port to VN_Port Virtual Links (see 7.9.4.3).

If after having claimed a Locally Unique N_Port_ID a VN2VN ENode MAC receives a N_Port_ID Claim Notification claiming its Locally Unique N_Port_ID or a N_Port_ID Beacon announcing its Locally Unique N_Port_ID, then the VN2VN ENode MAC shall check the N_Port_Name carried in the Vx_Port_Identification descriptor of the received N_Port_ID Claim Notification or N_Port_ID Beacon.

NOTE 29 – This situation may happen only when joining two previously separated networks. In this case there is a chance of N_Port_ID conflict, resulting in one or more VN2VN_Ports needing to change their Locally Unique N_Port_IDs.

If the received N_Port_Name is greater than the N_Port_Name of its VN2VN_Port then the VN2VN ENode MAC shall implicitly deinstantiate all its VN_Port to VN_Port Virtual Links and select a new Locally Unique N_Port_ID (see 7.9.6.2.1). If the received N_Port_Name is lower than the N_Port_Name of its VN2VN_Port then the VN2VN ENode MAC shall transmit a N_Port_ID Claim Notification to All-VN2VN-ENode-MACs, to notify all other VN2VN ENode MACs on the network that it has that N_Port_ID. A VN2VN ENode MAC with a VN2VN_Port that has a Virtual Link with a VN2VN_Port that is changing its N_Port_ID shall implicitly deinstantiate that Virtual Link upon receiving a N_Port_ID Claim Notification associating that N_Port_ID to a different N_Port_Name.

Reception of a N_Port_ID Beacon from a VN2VN ENode MAC not listed in the VN2VN Neighbor Set is an indication of a network join. If a N_Port_ID Claim Notification has not been sent in the past ANNOUNCE_WAIT milliseconds, a VN2VN ENode MAC shall transmit a N_Port_ID Claim Notification to All-VN2VN-ENode-MACs upon receiving such a N_Port_ID Beacon. This N_Port_ID Claim Notification shall be delayed by a random time uniformly distributed between 0 and 100 ms.

If after having claimed a Locally Unique N_Port_ID a VN2VN ENode MAC receives a N_Port_ID Probe Request probing its Locally Unique N_Port_ID, then the VN2VN ENode MAC shall reply to the N_Port_ID Probe Request with a N_Port_ID Probe Reply.

N_Port_ID Probe Requests probing a different Locally Unique N_Port_ID shall be ignored.

At any time, upon receiving a N_Port_ID P2P Claim Notification (see 7.9.8.12) or a N_Port_ID P2P Beacon (see 7.9.8.14), a VN2VN ENode MAC operating in multipoint mode shall transmit a N_Port_ID Probe Request probing its N_Port_ID to All-VN2VN-ENode-MACs to indicate to the sending VN2VN ENode MAC operating in point-to-point mode that it is not in its expected network configuration.

7.9.6.3 Point-to-point operation

7.9.6.3.1 Claiming a Locally Unique N_Port_ID

After selecting a tentative Locally Unique N_Port_ID, a VN2VN ENode MAC operating in point-to-point mode shall claim it by sending a N_Port_ID P2P Claim Notification to All-PT2PT-ENode-MACs.

NOTE 30 – A VN2VN ENode MAC operating in point-to-point mode does not send N_Port_ID Probe Requests.

The N_Port_ID P2P Claim Notification (see 7.9.8.12) provides the maximum FCoE size the VN2VN ENode MAC intends to use for VN_Port to VN_Port Virtual Links. A responding N_Port_ID P2P Claim Response FIP PDU (see 7.9.8.13) shall have a length equal to the minimum of the FCoE sizes supported by the two involved VN2VN ENode MACs.

If no N_Port_ID P2P Claim Responses are received within ANNOUNCE_WAIT milliseconds, the VN2VN ENode MAC shall continue sending a N_Port_ID P2P Claim Notification every ANNOUNCE_WAIT milliseconds until a N_Port_ID Claim Response is received.

If more than one N_Port_ID P2P Claim Responses are received within ANNOUNCE_WAIT milliseconds, the VN2VN ENode MAC shall cease the point-to-point operations.

If one N_Port_ID P2P Claim Response is received within ANNOUNCE_WAIT milliseconds, the VN2VN ENode MAC shall add the responding VN2VN_Port to its Neighbor Set, shall transmit a multicast N_Port_ID P2P Beacon (see 7.9.8.14) to All-PT2PT-ENode-MACs and shall continue to transmit multicast N_Port_ID P2P Beacons periodically every BEACON_PERIOD milliseconds. Upon transmitting the first N_Port_ID P2P Beacon the VN2VN ENode MAC may instantiate a VN_Port to VN_Port Virtual Link with the VN2VN neighbor (see 7.9.4.3). A VN2VN ENode MAC may perform this processing as soon as a N_Port_ID Claim Response is received.

If a N_Port_ID P2P Claim Notification claiming its Locally Unique N_Port_ID is received, then the VN2VN ENode MAC shall check the N_Port_Name carried in the Vx_Port_Identification descriptor of the received N_Port_ID P2P Claim Notification. If the received N_Port_Name is greater than the N_Port_Name of its VN2VN_Port then the VN2VN ENode MAC shall select a new random Locally Unique N_Port_ID and shall reply with a unicast N_Port_ID P2P Claim Response indicating the new Locally Unique N_Port_ID. If the received N_Port_Name is lower than the N_Port_Name of its VN2VN_Port then the VN2VN ENode MAC shall keep its selected Locally Unique N_Port_ID and shall reply with a unicast N_Port_ID P2P Claim Response indicating the originally selected Locally Unique N_Port_ID.

If a N_Port_ID P2P Claim Notification claiming a different Locally Unique N_Port_ID is received, then the VN2VN ENode MAC shall reply with a unicast N_Port_ID P2P Claim Response.

If after transmitting the first N_Port_ID P2P Beacon the VN2VN ENode MAC receives a N_Port_ID P2P Claim Notification or a N_Port_ID P2P Claim Response or a N_Port_ID P2P Beacon from a VN2VN_Port different than the one present in its Neighbor Set, the VN2VN ENode MAC shall cease the point-to-point operations.

At any time, upon receiving a N_Port_ID Probe Request, a N_Port_ID Claim Notification, a N_Port_ID Beacon, or a FIP Advertisement, a VN2VN ENode MAC operating in point-to-point mode shall cease the point-to-point operations.

7.9.6.4 Persistence of Locally Unique N_Port_IDs

VN2VN ENode MACs that are equipped with persistent storage may record the Locally Unique N_Port_ID they have selected for themselves. The recorded Locally Unique N_Port_ID may be also administratively configured. On booting, VN2VN ENode MACs with a recorded Locally Unique N_Port_ID should use that Locally Unique N_Port_ID as their first tentative Locally Unique N_Port_ID. This increases the stability of N_Port_IDs (e.g., if some VN2VN ENode MACs are powered off, then when they are powered on they resume using the previous Locally Unique N_Port_IDs, instead of picking different N_Port_IDs and potentially having to resolve conflicts).

Depending on local configuration, a VN2VN ENode MAC may stop the processing of the Locally Unique N_Port_ID selection protocol and report the situation in a vendor specific way if it detects that its recorded Locally Unique N_Port_ID is already in use on the network.

7.9.6.5 Locally Unique N_Port_ID timers and constants

Table 25 shows the timers and constants of the Locally Unique N_Port_IDs selection protocol.

Table 25 – Locally Unique N_Port_IDs timers and constants

Parameter	Value
PROBE_WAIT	100 ms
ANNOUNCE_WAIT	400 ms
RATE_LIMIT_INTERVAL	10 000 ms
BEACON_PERIOD	8 000 ms
All-VN2VN-ENode-MACs	01-10-18-01-00-04
All-PT2PT-ENode-MACs	01-10-18-01-00-05

7.9.7 FIP frames

7.9.7.1 FIP frame format

A FIP frame is an Ethernet frame (see IEEE 802.3-2008) containing a FIP PDU. FIP frames shall be formatted in accordance with IEEE 802.3-2008 and the MAC Client Data field within the Ethernet frame shall contain a FIP PDU (see table 26). The use of an 802.1Q tag header is optional and additional IEEE 802.1 defined tags may be present in a FIP frame. See Annex B for examples of FIP frames.

The format of a FIP PDU is specified in table 26.

Table 26 – FIP PDU format

Word	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0	Bit 3	Bit 2	Bit 1	Bit 0
0	Type = FIP_TYPE												Version				Reserved											
1	(MSB)																											
n	Encapsulated FIP operation (n words)																(LSB)											

The Type field in the Ethernet header shall be set to FIP_TYPE (see table 54). The Type field is not part of the FIP PDU.

The Version field shall be set to FIP_FRAME_VER (see table 54).

The Encapsulated FIP operation field shall be as specified in 7.9.7.2.

A received FIP frame containing no encapsulated FIP operation shall be discarded.

7.9.7.2 Encapsulated FIP operation

The format of the encapsulated FIP operation is specified in table 27.

Table 27 – Encapsulated FIP operation format

Word	Bit	31	30	29	28	27	26	25	24	23	22	21	20	19	18	17	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1	0
0	FIP Protocol Code														Reserved						FIP Subcode												
1	FIP Descriptor List Length														F	S	Reserved						R	A	S	F	P	P	P	P	P	P	
2	FIP Descriptor List (m words)																																
n-1	FIP_Pad (n-m-2 words)																																

The FIP Protocol Code and FIP Subcode field values and operations are specified in table 28.

Table 28 – FIP Protocol Code and FIP Subcode field values

FIP Protocol Code	FIP Subcode	FIP Operation	Reference
0001h	01h	Discovery Solicitation	7.9.8.2
	02h	Discovery Advertisement	7.9.8.3
0002h	01h	Virtual Link Instantiation Request	7.9.8.4
	02h	Virtual Link Instantiation Reply	7.9.8.4
0003h	01h	FIP Keep Alive	7.9.8.5
	02h	FIP Clear Virtual Links	7.9.8.6
0004h	01h	FIP VLAN Request	7.9.8.7
	02h	FIP VLAN Notification	7.9.8.8
0005h	01h	N_Port_ID Probe Request	7.9.8.10
	02h	N_Port_ID Probe Reply	7.9.8.11
	03h	N_Port_ID Claim Notification	7.9.8.12
	04h	N_Port_ID Claim Response	7.9.8.13
	05h	N_Port_ID Beacon	7.9.8.14
FFF8h to FFFEh	00h to FFh	Vendor Specific	7.9.8.9
All others	All others	Reserved	

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

The Fabric Provided (FP) bit and Server Provided (SP) bit settings are dependent on the FIP operation and shall be set as specified in table 29

Table 29 – FP bit and SP bit setting

Bit	FIP Protocol	FIP Operation (see table 52)	Setting
FP	Discovery	Discovery Solicitation	Shall be set to one.
		Discovery Advertisement	Shall be set to one.
	Virtual Link Instantiation	FIP FLOGI Request	Shall be set to one.
		FIP NPIV FDISC Request	Shall be set to one.
		FIP FLOGI LS_ACC	Shall be set to one.
		FIP NPIV FDISC LS_ACC	Shall be set to one.
	All others	All others	Reserved
SP	Discovery	Discovery Solicitation	Shall be set to zero.
		Discovery Advertisement	Shall be set to zero.
	Virtual Link Instantiation	FIP FLOGI Request	Shall be set to zero.
		FIP NPIV FDISC Request	Shall be set to zero.
		FIP FLOGI LS_ACC	Shall be set to zero.
		FIP NPIV FDISC LS_ACC	Shall be set to zero.
	All others	All others	Reserved

The Solicited (S) bit shall be set to one in solicited unicast Discovery Advertisements (i.e., Discovery Advertisements transmitted in response to a Discovery Solicitation). The S bit shall be set to zero in unsolicited multicast Discovery Advertisements (i.e., Discovery Advertisements not transmitted in response to a Discovery Solicitation). The S bit is reserved for all other FIP operations.

The FCF (F) bit shall be set to one in a Discovery Solicitation, Discovery Advertisement, or FIP VLAN Request if the originating device is an FCF. The F bit shall be set to zero in a Discovery Solicitation, Discovery Advertisement, or in a FIP VLAN Request if the originating device is not an FCF. The F bit is reserved for all other FIP operations.

The Available for Login (A) bit shall be set to one in a Discovery Advertisement if the originating FCF is available to process FIP FLOGI, FIP NPIV FDISC, or FIP ELP Requests (see 7.9.3). The A bit shall be set to zero in a Discovery Advertisement if the originating FCF is not available to process FIP FLOGI, FIP NPIV FDISC, or FIP ELP Requests. The A bit is reserved for all other FIP operations.

Bit 3 of word 1 of the encapsulated FIP operation is called the REC/P2P flag.

The FIP Descriptor List field shall contain one or more FIP descriptors (see 7.9.7.3).

The FIP_Pad field shall be used in solicited unicast Discovery Advertisements to extend the FIP PDU (see table 26) to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation that the Discovery Advertisement is responding to (see 7.9.8.3). The FIP_Pad field shall be used in N_Port_ID Claim Responses to extend the FIP PDU to

have a length that matches the minimum of the FCoE sizes supported by the two involved VN2VN ENode MACs. The FIP_Pad field shall be of zero length (i.e., not present) for all other FIP operations.

Received FIP frames shall be checked for correct formatting before any FIP descriptor processing occurs. A malformed FIP frame shall be discarded and should be reported in a vendor specific way. The checks for correct formatting include:

- a) the FIP Descriptor List Length value matches the sum of the descriptors' lengths in the FIP Descriptor List;
- b) the FIP Protocol Code field and FIP Subcode field are valid (see table 28);
- c) the critical descriptors (see 7.9.7.3.1) required by the FIP Protocol Code and FIP Subcode are present;
- d) no critical descriptors other than the ones required by the FIP Protocol Code and FIP Subcode are present; and
- e) descriptors use valid values for MAC addresses (see 7.9.7.3.3), FKA_ADV_PERIOD (see table 54) and VLAN IDs (see 7.9.7.3.15).

7.9.7.3 FIP descriptors

7.9.7.3.1 FIP descriptor overview

FIP descriptors are specified using a TLV format (i.e., Type, Length, Value). The length field value shall be specified as the number of words in the FIP descriptor including the TLV format. FIP descriptor type values are split into two ranges, critical and non-critical, as specified in table 30.

Table 30 – FIP descriptor type value ranges

Range	Value	Description
Critical	0 to 127	An FCoE Controller that receives a FIP frame with an unknown critical descriptor shall discard the FIP frame.
Non-critical	128 to 255	An FCoE Controller that receives a FIP frame with one or more unknown non-critical descriptors shall ignore the unknown descriptors and continue to process the FIP frame.

A descriptor with an invalid length value shall be considered invalid. If that descriptor is critical, the entire FIP operation shall be discarded. If that descriptor is not critical, the descriptor shall be ignored.

The FIP descriptor types are specified in table 31. See table 52 for how FIP descriptors are used.

Table 31 – FIP descriptor types

Range	Type	FIP Descriptor	Reference
Critical	0	Reserved	
	1	Priority	7.9.7.3.2
	2	MAC address	7.9.7.3.3
	3	FC-MAP	7.9.7.3.4
	4	Name_Identifier	7.9.7.3.5
	5	Fabric	7.9.7.3.6
	6	Max FCoE Size	7.9.7.3.7
	7	FLOGI ^a	7.9.7.3.8
	8	NPIV FDISC ^a	7.9.7.3.9
	9	LOGO ^a	7.9.7.3.10
	10	ELP ^a	7.9.7.3.11
	11	Vx_Port Identification	7.9.7.3.12
	12	FKA_ADV_Period	7.9.7.3.13
	13	Vendor_ID	7.9.7.3.14
	14	VLAN	7.9.7.3.15
	15	FC-4 Attributes	7.9.7.3.17
16 to 127	Reserved		
Non-critical	128	Clear Virtual Links Reason Code	7.9.7.3.18
	129 to 240	Reserved	
	241 to 254	Vendor Specific	7.9.7.3.16
	255	Reserved	
a The FC CRC, SOF, and EOF shall not be included in the FIP descriptor.			

7.9.7.3.2 FIP Priority descriptor

The FIP Priority descriptor is used in FIP operations as specified in table 52. An ENode may use the value provided in the Priority descriptor of received Discovery Advertisements to select the FCF-MAC to which to perform FIP FLOGI. The default value for the Priority field is DEFAULT_FIP_PRIORITY (see table 54). The highest priority value is 0 and the lowest priority value is 255 (i.e., lower numerical values indicate higher priorities).

The FIP Name_Identifier descriptor format shall be as specified in table 35.

Table 35 – FIP Name_Identifier descriptor format

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

Name_Identifier: the Name_Identifier (see FC-FS-3) contained in the descriptor.

7.9.7.3.6 FIP Fabric descriptor

The FIP Fabric descriptor is used in FIP operations as specified in table 52.

The FIP Fabric descriptor format shall be as specified in table 36.

Table 36 – FIP Fabric descriptor format

Word	Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
0		Type = 05h								Length = 04h								Reserved				VF_ID											
1		Reserved								(MSB)								FC-MAP								(LSB)							
2		(MSB)																															
3		Fabric_Name																(LSB)															

VF_ID: the VF_ID (see see FC-FS-3) associated with the Fabric, if any.

FC-MAP: the value to be used as the most significant 24 bits in FPMAs (see 7.7).

Fabric_Name: the Fabric_Name (see FC-FS-3) identifying the Fabric.

7.9.7.3.7 FIP Max FCoE Size descriptor

The FIP Max FCoE Size descriptor is used in FIP operations as specified in table 52.

The FIP Max FCoE Size descriptor format shall be as specified in table 37.

Table 37 – FIP Max FCoE Size descriptor format

Word	Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0		Type = 06h								Length = 01h								Max_FCoE_Size														

Max_FCoE_Size: the size in bytes that the FIP PDU (see table 26) in a solicited unicast Discovery Advertisement is requested to be extended to.

7.9.7.3.8 FIP FLOGI descriptor

The FIP FLOGI descriptor is used in FIP operations as specified in table 52.

The FIP FLOGI descriptor format shall be as specified in table 38

Table 38 – FIP FLOGI descriptor format

Word	Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0		Type = 07h								Length								Reserved														
1		(MSB)								FLOGI Request or FLOGI																						
n										LS_ACC/LS_RJT								(LSB)														

Length: shall be set to 36 for a FLOGI Request and FLOGI LS_ACC, or to 9 for a FLOGI LS_RJT.

FLOGI Request or FLOGI LS_ACC/LS_RJT: an encapsulated FLOGI Request, FLOGI LS_ACC, or FLOGI LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field. In an FLOGI Request or FLOGI LS_ACC, the Payload bit shall be set to zero (see FC-LS-2).

7.9.7.3.9 FIP NPIV FDISC descriptor

The FIP NPIV FDISC descriptor is used in FIP operations as specified in table 52.

The FIP NPIV FDISC descriptor format shall be as specified in table 39

Table 39 – FIP NPIV FDISC descriptor format

Word	Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0		Type = 08h								Length								Reserved														
1		(MSB)								NPIV FDISC Request or NPIV FDISC																						
n										LS_ACC/LS_RJT								(LSB)														

Length: shall be set to 36 for an FDISC Request and FDISC LS_ACC, or to 9 for an FDISC LS_RJT.

NPIV FDISC Request or NPIV FDISC LS_ACC/LS_RJT: an encapsulated FDISC Request, FDISC LS_ACC, or FLOGI LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field. In an FLOGI Request or FLOGI LS_ACC, the Payload bit shall be set to zero (see FC-LS-2).

7.9.7.3.10 FIP LOGO descriptor

The FIP LOGO descriptor is used in FIP operations as specified in table 52.

The FIP LOGO descriptor format shall be as specified in table 40

Table 40 – FIP LOGO descriptor format

Word	Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
0		Type = 09h								Length								Reserved														
1		(MSB)								LOGO Request or LOGO																						
n										LS_ACC/LS_RJT								(LSB)														

Length: shall be set to 11 for a LOGO Request, 8 for a LOGO LS_ACC, or to 9 for a LOGO LS_RJT.

LOGO Request or LOGO LS_ACC/LS_RJT: an encapsulated LOGO Request, LOGO LS_ACC, or LOGO LS_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an ELS payload but without the CRC field (see FC-LS-2).

7.9.7.3.11 FIP ELP descriptor

The FIP ELP descriptor is used in FIP operations as specified in table 52.

The FIP ELP descriptor format shall be as specified in table 41.

Table 41 – FIP ELP descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
0	Type = 0Ah				Length								Reserved																			
1	(MSB)																															
n	ELP Request or ELP SW_ACC/SW_RJT																(LSB)															

Length: shall be set to 33 for an ELP Request and ELP SW_ACC, or to 9 for an ELP SW_RJT.

ELP Request or ELP SW_ACC/SW_RJT: an encapsulated ELP Request, ELP SW_ACC, or ELP SW_RJT shall be a complete Fibre Channel frame content (see FC-FS-3) with a Fibre Channel Frame_Header and an SW_ILS payload but without the CRC field (see FC-SW-5).

7.9.7.3.12 FIP Vx_Port Identification descriptor

The FIP Vx_Port Identification descriptor is used in FIP operations as specified in table 52.

The FIP Vx_Port Identification descriptor format shall be as specified in table 42.

Table 42 – FIP Vx_Port Identification descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
0	Type = 0Bh				Length = 05h								(MSB)																			
1	MAC address (LSB)																															
2	Reserved				(MSB)								Address Identifier (LSB)																			
3	(MSB)																															
4	Port_Name																(LSB)															

MAC address: the MAC address associated with the referred VN_Port or VE_Port. Valid MAC addresses are unicast addresses.

Address Identifier: the address identifier associated with the referred VN_Port or the value FFFFDh for a VE_Port.

Port_Name: the N_Port_Name of the referred VN_Port or the E_Port_Name of the referred VE_Port.

7.9.7.3.13 FIP FKA_ADV_Period descriptor

The FIP FKA_ADV_Period descriptor is used in FIP operations as specified in table 52.

The FIP FKA_ADV_Period descriptor format shall be as specified in table 43.

Table 43 – FIP FKA_ADV_Period descriptor format

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

D: the value of the D bit shall be administratively configurable on FCFs. The D bit shall be set to zero unless administratively set to one. VE_Port capable FCF-MACs shall ignore the value of the D bit provided in received Discovery Advertisements. ENode MACs shall process the value of the D bit provided in received Discovery Advertisements.

When the D bit is set to zero, the receiving ENode MAC shall verify periodic reception of Discovery Advertisements and shall transmit periodic ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.9.5.2. The VF_Port capable FCF-MAC shall verify periodic reception of ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.9.5.2.

When the D bit is set to one, the receiving ENode MAC may verify periodic reception of Discovery Advertisements and should not transmit periodic ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.9.5.2. The VF_Port capable FCF-MAC shall not verify periodic reception of ENode FIP Keep Alive and VN_Port FIP Keep Alive frames as specified in 7.9.5.2 and should discard possible received ENode FIP Keep Alive and VN_Port FIP Keep Alive frames.

FKA_ADV_PERIOD: the value of the advertised FKA_ADV_PERIOD (see table 54). See table 54 for the range of valid values for FKA_ADV_PERIOD.

7.9.7.3.14 FIP Vendor_ID descriptor

The FIP Vendor_ID descriptor is used in FIP operations as specified in table 52.

The FIP Vendor_ID descriptor format shall be as specified in table 44.

Table 44 – FIP Vendor_ID descriptor format

Word	Bit 3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
0	Type = 0Dh				Length = 03h								Reserved																			
1	(MSB)																															
2	Vendor_ID																(LSB)															

Vendor_ID: the vendor's Vendor_ID value.

7.9.7.3.15 FIP VLAN descriptor

The FIP VLAN descriptor is used in in FIP operations as specified in table 52.

The FIP VLAN descriptor format shall be as specified in table 45

Table 45 – FIP VLAN descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type = 0Eh								Length = 01h								Reserved				FCoE VID											

FCoE VID: the VLAN ID of a VLAN where FCoE services may be available. The range of valid values for VLAN IDs is 001h to FFEh.

7.9.7.3.16 FIP Vendor Specific descriptors

FIP Vendor Specific descriptors are non-critical and may be used in any FIP frame. An FC-BB_E device shall not require the use of any FIP Vendor Specific descriptor in order to operate in accordance with this standard.

The FIP Vendor Specific descriptor format is specified in table 46.

Table 46 – FIP Vendor Specific descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type								Length								Reserved															
1	(MSB)																															
2	Vendor_ID														(LSB)																	
3	(MSB)																															
n	Vendor Specific information														(LSB)																	

Type: the FIP Vendor Specific descriptors are identified by a type value in the range 241 to 254, inclusive.

Length: shall be set to the length in words of the descriptor.

Vendor_ID: the vendor’s Vendor_ID value.

Vendor Specific Information: defined by the vendor.

7.9.7.3.17 FIP FC-4 Attributes descriptor

The FIP FC-4 Attributes descriptor format shall be as specified in table 47.

Table 47 – FIP FC-4 Attributes descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0			
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0		
0	Type = 0Fh								Length = 29h								Reserved																	
1 to 8	MSB																FC-4 Types																LSB	
9 to 40	MSB																FC-4 Features																LSB	

FC-4 Types: the FC-4 Types object, as specified in FC-GS-6.

FC-4 Features: the FC-4 Features object, as specified in FC-GS-6.

7.9.7.3.18 FIP Clear Virtual Links Reason Code descriptor

The FIP Clear Virtual Links Reason Code descriptor format shall be as specified in table 48.

Table 48 – FIP Clear Virtual Links Reason Code descriptor format

Bit	3	3	2	2	2	2	2	2	2	2	1	1	1	1	1	1	1	1	1	1	0	9	8	7	6	5	4	3	2	1	0	
Word	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0	9	8	7	6	5	4	3	2	1	0
0	Type = 80h								Length = 02h								Reserved															
1	Reserved								Reason Code								Reason Code Explanation				Vendor Specific											

Reason Code: shall be set to one of the Clear Virtual Links Reason Codes specified in table 49.

Table 49 – FIP Clear Virtual Links Reason Codes

Encoded value (bits 23 to 16)	Description
00h	Timeout
01h	Vx_Port state change
02h to FFh	Reserved

Reason Code Explanation: shall be set to one of the Clear Virtual Links Reason Code Explanations specified in table 50.

Table 50 – FIP Clear Virtual Links Reason Code Explanations

Encoded value (bits 15 to 8)	Description
00h	No additional explanation
01h	ENode FIP Keep Alive timeout
02h	VN_Port FIP Keep Alive timeout
03h	Configuration change
04h	Fabric offline
05h to FFh	Reserved

The Reason Code field and Reason Code Explanation field may be used to indicate a particular event has occurred. FIP Clear Virtual Links Reason Code descriptor examples are specified in table 51.

Table 51 – FIP Clear Virtual Links Reason Code descriptor examples

Reason Code	Reason Code Explanation	Description
00h	01h	2.5 * FKA_ADV_PERIOD has passed since the last ENode FIP Keep Alive frame was received by the FCF.
00h	02h	2.5 * FKA_VN_PERIOD has passed since the last VN_Port FIP Keep Alive frame was received by the FCF.
01h	03h	A configuration change has been performed on the FCF that requires the Virtual Link to be de-instantiated.
01h	04h	The Fibre Channel functionality within the FCF has been disabled or is no longer accessible.

Vendor Specific: defined by the vendor.

7.9.8 FIP operations

7.9.8.1 FIP operations overview

Table 52 specifies the FIP descriptors required in each FIP operation and the recommended order in which they should be encapsulated by a transmitting FCoE Controller. In certain cases, as indicated, strict ordering is required. A receiving FCoE Controller shall process unknown descriptors according to the criticality of the FIP descriptor (see 7.9.7.3.1). Unless otherwise specified (e.g., for a FIP

FLOGI Request), a receiving FCoE Controller shall be able to process the FIP descriptors in any order.

NOTE 31 – The ability to process FIP descriptors in any order is to provide flexibility for future protocol extensions,

A FIP operation shall contain the expected critical descriptors (see table 52) and may contain additional non-critical descriptors. If some critical descriptors are missing or unexpected, the FIP operation shall be discarded and it should be reported in a vendor specific manner.

Table 52 – FIP operation descriptors and order

FIP Operation	FIP Protocol Code/Subcode ^a	Originator	Expected Descriptors and Order
Discovery Solicitation (see 7.9.8.2.1)	0001h/01h	ENode	1) MAC address 2) Name_Identifier 3) Max FCoE Size
Discovery Solicitation (see 7.9.8.2.2)	0001h/01h	FCF	1) MAC address 2) FC-MAP 3) Name_Identifier 4) Max FCoE Size
Discovery Advertisement (see 7.9.8.3)	0001h/02h	FCF	1) Priority 2) MAC address 3) Name_Identifier 4) Fabric 5) FKA_ADV_Period
FIP FLOGI Request ^b (see 7.9.8.4.2)	0002h/01h	ENode	1) FLOGI 2) MAC address
FIP FLOGI LS_ACC ^b (see 7.9.8.4.2)	0002h/02h	FCF	1) FLOGI 2) MAC address
FIP FLOGI LS_RJT ^b (see 7.9.8.4.2)	0002h/02h	FCF	1) FLOGI

- a) Virtual Link instantiation requests are differentiated by the contained descriptors.
- b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here.
- c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Links frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP).
- d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames.
- e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor.
- f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present).
- g) The FIP Clear Virtual Links Reason Code descriptor is optional in FIP Clear Virtual Links frames (i.e., the FIP Clear Virtual Links Reason Code descriptor may not be present).

Table 52 – FIP operation descriptors and order (Continued)

FIP Operation	FIP Protocol Code/Subcode ^a	Originator	Expected Descriptors and Order
FIP NPIV FDISC Request ^b (see 7.9.8.4.2)	0002h/01h	ENode	1) NPIV FDISC 2) MAC address
FIP NPIV FDISC LS_ACC ^b (see 7.9.8.4.2)	0002h/02h	FCF	1) NPIV FDISC 2) MAC address
FIP NPIV FDISC LS_RJT ^b (see 7.9.8.4.2)	0002h/02h	FCF	1) NPIV FDISC
FIP Fabric LOGO ^b (see 7.9.8.4.3)	0002h/01h	ENode ^c	1) LOGO 2) MAC address
FIP Fabric LOGO LS_ACC ^b (see 7.9.8.4.3)	0002h/02h	FCF ^c	1) LOGO 2) MAC address
FIP Fabric LOGO LS_RJT ^b (see 7.9.8.4.3)	0002h/02h	FCF ^c	1) LOGO
FIP ELP Request ^b (see 7.9.8.4.4)	0002h/01h	FCF	1) ELP 2) MAC address
FIP ELP SW_ACC ^b (see 7.9.8.4.4)	0002h/02h	FCF	1) ELP 2) MAC address
FIP ELP SW_RJT ^b (see 7.9.8.4.4)	0002h/02h	FCF	1) ELP
FIP Keep Alive (see 7.9.8.5)	0003h/01h	ENode	1) MAC address 2) Vx_Port Identification ^d
<p>a) Virtual Link instantiation requests are differentiated by the contained descriptors.</p> <p>b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here.</p> <p>c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Links frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP).</p> <p>d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames.</p> <p>e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor.</p> <p>f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present).</p> <p>g) The FIP Clear Virtual Links Reason Code descriptor is optional in FIP Clear Virtual Links frames (i.e., the FIP Clear Virtual Links Reason Code descriptor may not be present).</p>			

Table 52 – FIP operation descriptors and order (Continued)

FIP Operation	FIP Protocol Code/Subcode ^a	Originator	Expected Descriptors and Order
FIP Clear Virtual Links (see 7.9.8.6)	0003h/02h	FCF	1) MAC address 2) Name_Identifier 3) Vx_Port Identification(s) ^e 4) Clear Virtual Links Reason Code ^g
FIP VLAN Request (see 7.9.8.7)	0004h/01h	ENode or FCF	1) MAC address 2) Name_Identifier ^f
FIP VLAN Notification (see 7.9.8.8)	0004h/02h	FCF	1) MAC address 2) VLAN(s)
FIP Vendor Specific (see 7.9.8.9)	FFF8h to FFFEh / 00h to FFh	ENode or FCF	1) Vendor_ID 2) Descriptor(s)
N_Port_ID Probe Request (see 7.9.8.10)	0005h / 01h	VN2VN ENode	1) MAC address 2) Name_Identifier 3) Vx_Port Identification
N_Port_ID Probe Reply (see 7.9.8.11)	0005h / 02h	VN2VN ENode	1) MAC address 2) Name_Identifier 3) Vx_Port Identification
N_Port_ID Claim Notification (see 7.9.8.12)	0005h / 03h	VN2VN ENode	1) MAC address 2) Name_Identifier 3) Vx_Port Identification 4) FC-4 Attributes 5) Max FCoE Size

- a) Virtual Link instantiation requests are differentiated by the contained descriptors.
- b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here.
- c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Links frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP).
- d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames.
- e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor.
- f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present).
- g) The FIP Clear Virtual Links Reason Code descriptor is optional in FIP Clear Virtual Links frames (i.e., the FIP Clear Virtual Links Reason Code descriptor may not be present).

Table 52 – FIP operation descriptors and order (Continued)

FIP Operation	FIP Protocol Code/Subcode ^a	Originator	Expected Descriptors and Order
N_Port_ID Claim Response (see 7.9.8.13)	0005h / 04h	VN2VN ENode	1) MAC address 2) Name_Identifier 3) Vx_Port Identification 4) FC-4 Attributes 5) Max FCoE Size
N_Port_ID Beacon (see 7.9.8.14)	0005h / 05h	VN2VN ENode	1) MAC address 2) Name_Identifier 3) Vx_Port Identification
a) Virtual Link instantiation requests are differentiated by the contained descriptors. b) Strict ordering of the FIP descriptors is required in transmission. In this way the encapsulated ELS or SW_ILS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in an FCoE PDU. A receiving FCoE Controller is not required to be able to process these FIP operations in any order other than that specified here. c) FCFs are allowed to generate FIP LOGO, however a FIP Clear Virtual Links frame is the recommended method for an FCF to de-instantiate a Virtual Link, except for the specific cases where LOGO is required (see FC-SP). d) The Vx_Port Identification descriptor is present only in VN_Port FIP Keep Alive frames, it is not present in ENode FIP Keep Alive frames. e) A FIP Clear Virtual Links frame intended to de-instantiate VN_Port to VF_Port Virtual Links contains zero or more Vx_Port Identification descriptors. A FIP Clear Virtual Links frame intended to de-instantiate a VE_Port to VE_Port Virtual Link contains one Vx_Port Identification descriptor. f) The Name_Identifier descriptor is optional in FIP VLAN Request frames (i.e., the Name_Identifier descriptor may be not present). g) The FIP Clear Virtual Links Reason Code descriptor is optional in FIP Clear Virtual Links frames (i.e., the FIP Clear Virtual Links Reason Code descriptor may not be present).			

7.9.8.2 FIP Discovery Solicitation

7.9.8.2.1 ENode FIP Discovery Solicitation

As specified in table 52, a Discovery Solicitation operation originated by the FCoE Controller of an ENode MAC contains a MAC address descriptor (see 7.9.7.3.3), a Name_Identifier descriptor (see 7.9.7.3.5), and a Max FCoE Size descriptor (see 7.9.7.3.7).

A Discovery Solicitation frame may be unicast (i.e., addressed to a specific FCF-MAC) or multicast (i.e., addressed to the All-FCF-MACs group address).

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent solicited Discovery Advertisements from VF_Port capable FCF-MACs.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Node_Name of the ENode or to zero.

NOTE 32 – The Name_Identifier field may be set to zero if the Node_Name is ambiguous or not yet available when the Discovery Solicitation is transmitted.

The Max_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE PDU size the ENode MAC intends to use for FCoE traffic. The Max_FCoE_Size value shall be specified as the number of octets starting with and including Version field, up to and including the Reserved field following the EOF field (see table 22).

7.9.8.2.2 FCF FIP Discovery Solicitation

As specified in table 52, a Discovery Solicitation operation originated by the FCoE Controller of a VE_Port capable FCF-MAC contains a MAC address descriptor (see 7.9.7.3.3), an FC-MAP descriptor (see 7.9.7.3.4), a Name_Identifier descriptor (see 7.9.7.3.5), and a Max FCoE Size descriptor (see 7.9.7.3.7).

A Discovery Solicitation frame may be unicast (i.e., addressed to a specific FCF-MAC) or multicast (i.e., addressed to the All-FCF-MACs group address).

The MAC address field in the MAC address descriptor shall be set to the MAC address to use for subsequent solicited Discovery Advertisements from VE_Port capable FCF-MACs.

The FC-MAP field in the FC-MAP descriptor shall be set to the FC-MAP value the FCF-MAC is using. If the FC-MAP value is not administratively configured, then the FC-MAP value shall be set to DEFAULT_FC-MAP (see table 54).

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the FCF.

The Max_FCoE_Size field in the Max FCoE Size descriptor shall be set to the maximum FCoE frame size the VE_Port capable FCF-MAC intends to use for FCoE traffic. The Max_FCoE_Size value shall be specified as the number of octets starting with and including the Version field, up to and including the Reserved field following the EOF field (see table 22).

7.9.8.3 FIP Discovery Advertisements

As specified in table 52, a Discovery Advertisement operation contains a Priority descriptor (see 7.9.7.3.2), a MAC address descriptor (see 7.9.7.3.3), a Name_Identifier descriptor (see 7.9.7.3.5), a Fabric descriptor (see 7.9.7.3.6), and a FKA_ADV_Period descriptor (see 7.9.7.3.13).

When a Discovery Advertisement frame is solicited, it shall be unicast (i.e., addressed to a specific ENode MAC or FCF-MAC address). When a Discovery Advertisement frame is unsolicited, it shall be multicast (i.e., addressed to the All-ENode-MACs or to the All-FCF-MACs group addresses).

The Priority field in the Priority descriptor shall be set to the value the originating FCF-MAC is using. If the priority value is not administratively configured, then the priority value shall be set to DEFAULT_FIP_PRIORITY (see table 54).

The MAC address field in the MAC address descriptor shall be set to the originating FCF-MAC address.

The Name_Identifier field in the Name_Identifier descriptor shall be set to the Switch_Name of the originating FCF.

Discovery Advertisements shall only contain a single Fabric descriptor. All Discovery Advertisements from an FCF in a VLAN shall contain the same single Fabric descriptor. The VF_ID field in the Fabric descriptor shall be set to the VF_ID identifying the advertised FC Fabric. If a VF_ID is not defined for the advertised FC Fabric, the VF_ID field shall be set to zero. The FC-MAP field in the Fabric

descriptor shall be set to the FC-MAP value the FCF is using. If the FC-MAP value is not administratively configured, then the FC-MAP value shall be set to DEFAULT_FC-MAP (see table 54). The Fabric_Name field in the Fabric descriptor shall be set to the Fabric_Name for the originating FCF.

The FKA_ADV_PERIOD field in the FKA_ADV_Period descriptor shall be set to the FKA_ADV_PERIOD value the FCF is advertising (see table 54).

The FIP_Pad field shall be used to extend the FIP PDU (see table 26) to have a length that matches the Max_FCoE_Size field value in the Max FCoE Size descriptor in the Discovery Solicitation to which the Discovery Advertisement is responding. The FIP_Pad field value shall be set to zero and not checked in reception. For an unsolicited Discovery Advertisement, the FIP_Pad field shall be of zero length (i.e., not present).

7.9.8.4 FIP Virtual Link Instantiation Requests and Replies

7.9.8.4.1 FIP Virtual Link Instantiation Requests and Replies overview

FIP Virtual Link Instantiation Requests and Replies encapsulates an ELS or an SW_ILS. The encapsulated ELS or SW_ILS shall be a single-frame Sequence. FIP Virtual Link Instantiation Requests and Replies are used to perform:

- a) Fabric login between ENode MACs and VF_Port capable FCF-MACs (see 7.9.8.4.2);
- b) Fabric logout between ENode MACs and VF_Port capable FCF-MACs (see 7.9.8.4.3); and
- c) Exchange Link Parameters between VE_Port capable FCF-MACs (see 7.9.8.4.4).

7.9.8.4.2 Fabric login

As specified in table 52:

- a) a FIP FLOGI Request operation contains a FLOGI descriptor (see 7.9.7.3.8) and a MAC address descriptor (see 7.9.7.3.3);
- b) a FIP FLOGI LS_ACC operation contains a FLOGI descriptor (see 7.9.7.3.8) and a MAC address descriptor (see 7.9.7.3.3);
- c) a FIP FLOGI LS_RJT operation contains a FLOGI descriptor (see 7.9.7.3.8).
- d) a FIP NPIV FDISC Request operation contains a NPIV FDISC descriptor (see 7.9.7.3.9) and a MAC address descriptor (see 7.9.7.3.3);
- e) a FIP NPIV FDISC LS_ACC operation contains a NPIV FDISC descriptor (see 7.9.7.3.9) and a MAC address descriptor (see 7.9.7.3.3); and
- f) a FIP NPIV FDISC LS_RJT operation contains a NPIV FDISC descriptor (see 7.9.7.3.9).

The FLOGI or NPIV FDISC descriptor shall be the first descriptor in the operation.

NOTE 33 – In this way the encapsulated Fibre Channel ELS results at an offset in the FIP PDU equal to the offset it would have if it was encapsulated in a FCoE PDU.

The FLOGI or NPIV FDISC descriptor shall contain:

- a) a complete Fibre Channel frame content with a Fibre Channel Frame_Header (see FC-FS-3) and an FLOGI Request payload or an NPIV FDISC Request payload (see FC-LS-2) but without the CRC field (see FC-FS-3) for FIP FLOGI Request or FIP NPIV FDISC Request operations. The Payload bit in the FLOGI Request payload or NPIV FDISC Request payload shall be set to zero (see FC-LS-2);