

## **BB-6 modifications for Multiple FCFs on one VLAN**

13-077v0

Bill Martin (Emulex)

Erik Smith (EMC)

Claudio DeSanti (Cisco)

### **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 (11-515v0)

Original proposal

Revision 1 – 6 February 2012 (12-019v0)

Modifications from work group and additional input from Claudio

Revision 2 – 8 February 2012 (12-019v1)

Modifications from work group

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