

FCoE Point To Point

VN_Port to VN_Port Network Configuration

Craig Carlson - Qlogic

Roger Hathorn – IBM

Fred Knight - NetApp

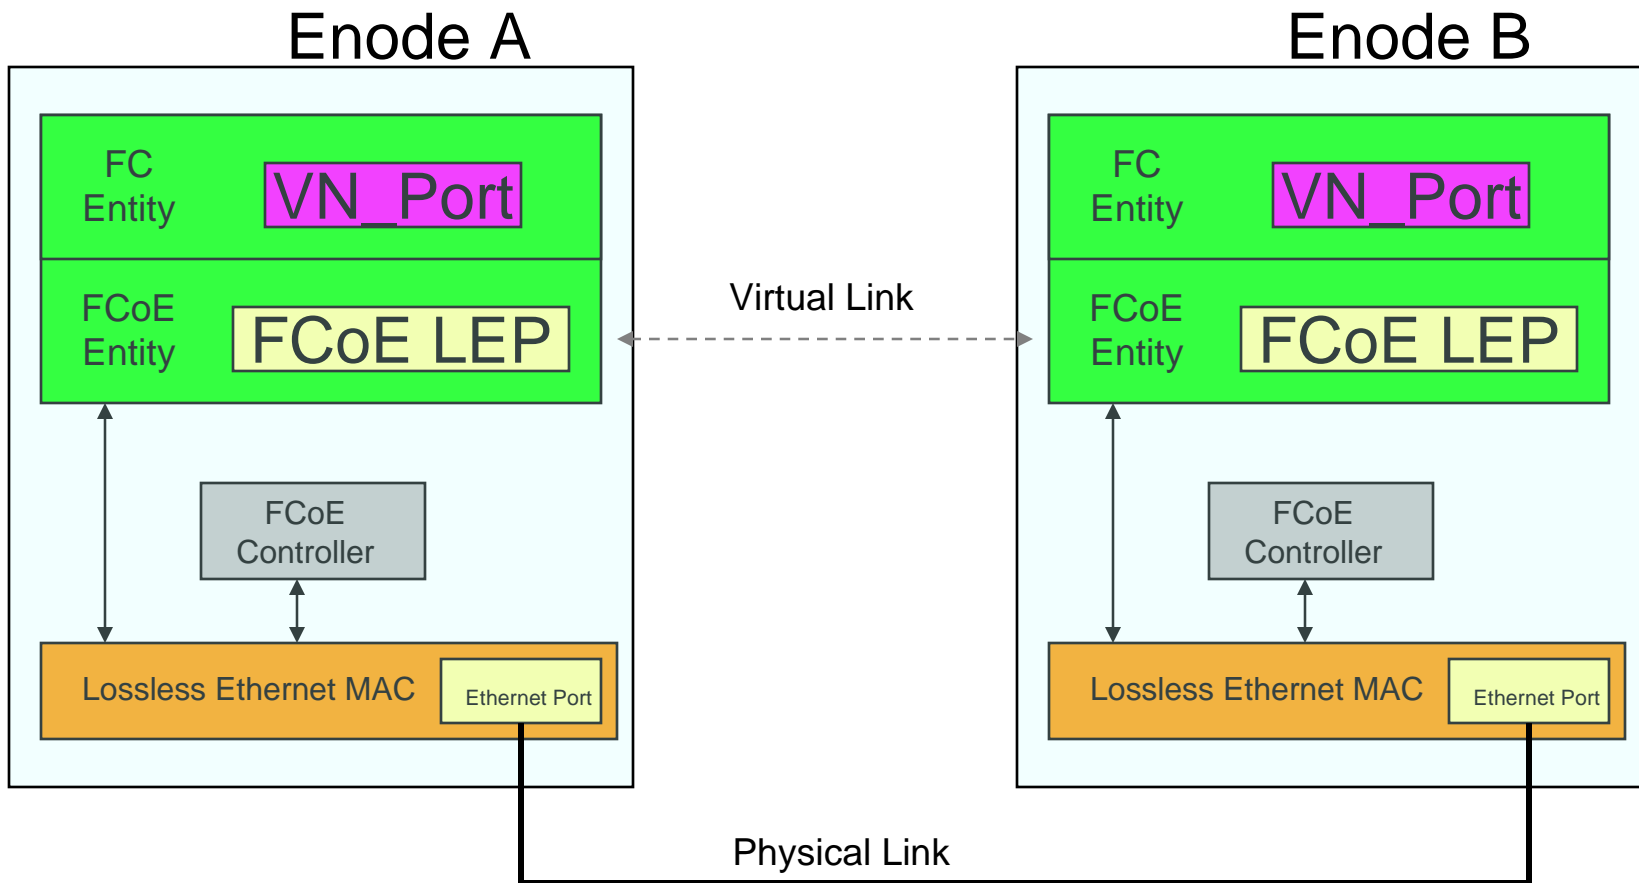
Lou Ricci - IBM

Erik Smith – EMC

Alan Spalding – QLogic

12/9/09

VN_Port to VN_Port Network Configuration



FCoE Point-to-Point - Motivation

- FC supports Point-to-Point
 - Very useful in testing and demo scenarios
 - Also useful in simple user configurations
 - Very popular in certain environments
- FCoE in BB-5 doesn't support Point-to-Point
 - Loss of this FC capability impacts FCoE promise of full FC capability
- Objective – Add Point-to-Point to FCoE

Point-to-Point Discovery Design Points and Requirements

- Must Co-exist with current ENode FCF Discovery
- Move to FCF Discovery and tear down point to point connections whenever:
 1. One or more FCFs are detected
 2. Multiple ENode(s) detected
- Point to point configuration not detected in properly configured fabric configuration
- Discovery completes in a timely manner

FC-LS Point to Point Login Process

- 6.2.2.4 Nx_Port response to FLOGI

If an Nx_Port receives a FLOGI, the Nx_Port shall respond to the received FLOGI with an LS_ACC reply Sequence with the OX_ID equal to the OX_ID of the received FLOGI and the Common Service Parameter Nx_Port/F_Port bit set to zero (i.e., an Nx_Port). This indicates a point-to-point connection with another Nx_Port. The D_ID of the LS_ACC shall be the S_ID of the received FLOGI. The Payload shall include the Service Parameters from the received FLOGI with all classes mark invalid, and the 64-bit N_Port_Name and 64-bit Node_Name of the connected Nx_Port. **If the received N_Port_Name is less than its N_Port_Name, the Nx_Port proceeds to N_Port Login.** If the received N_Port_Name is greater than its N_Port_Name, the Nx_Port waits for PLOGI from the attached N_Port.

VN_Port to VN_Port Virtual Link Establishment - overview

1. Determine with reasonable certainty that you are in a point to point configuration.
2. Perform point to point FIP discovery protocol very similar to FCF discovery.
3. Complete point to point FC login process using unchanged FC-LS protocol.

Point to Point Topology Discovery

To determine with reasonable certainty that you are in a point to point configuration...

- a) During VLAN discovery, receive a FIP VLAN Notification without the F bit set to one; or
- b) During E_node discovery, receive a FIP Discovery Advertisement without the F bit set to one.

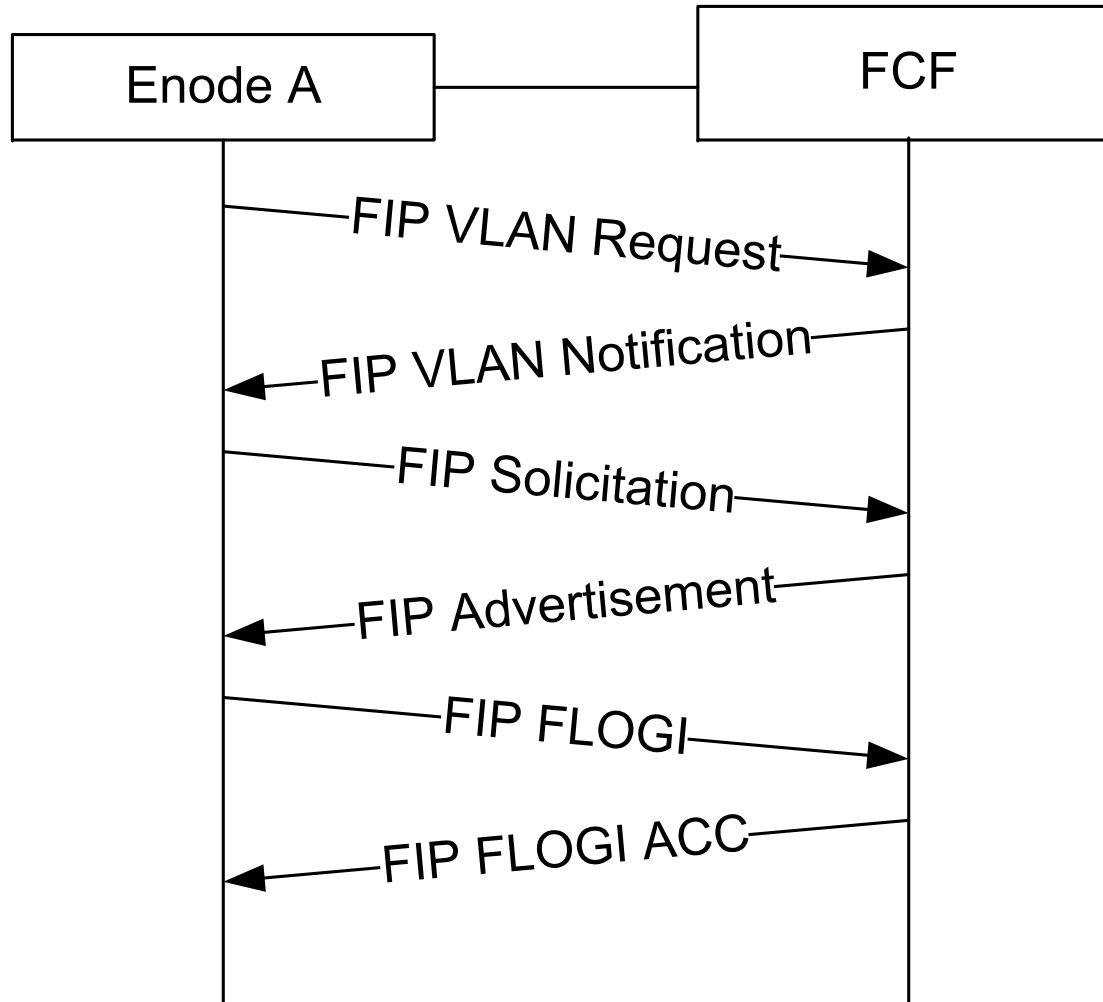
Note: Could define new P (=point to point) bit in discovery requests/responses to allow modified protocols rather than using F=0 which could be an error in BB-5 protocol

ENode FIP Discovery

Today

- FIP VLAN Discovery (optional)
 - VLAN Request (SA = ENode MAC, DA = ALL-FCF-MACs)
<MAC=Enode MAC>
 - Wait for VLAN Notification (SA = FCF MAC, DA = ENode MAC)
<MAC=FCF MAC, VLAN VID=x, VLAN VID=y>
 - May retry a number of times
- FIP FCF Discovery
 - FIP Discovery Solicitation (SA = ENode MAC, DA = ALL-FCF-MACs)
<MAC=Enode MAC, Name=Node Name, Max FCoE Size>
 - FIP Discovery Advertisement(s) (SA = FCF MAC, DA = ENode MAC)
<F=1, MAC=FCF MAC, FKA_ADV_PERIOD, Priority, Name, Fabric Descriptor, FIP Pad>
- Table of FCFs created.

ENode FIP Discovery Today



ENode FIP Discovery

With Point-to-Point Support

- Transmit two FIP VLAN Discovery Requests
 - I. DA = ALL-FCF-MACs
 - II. DA = ALL-ENode-MACs, VLAN Descriptor indicating suggested FCoE VID (shall only be one).
- Wait for FIP VLAN Notification(s)
 - I. SA = FCF MAC, DA = ENode MAC, F bit = 1; or
 - II. SA = ENode-B MAC, DA = ENode-A MAC, F bit = 0
- May retry a number of times if no response
- If properly configured, should not receive both FCF (F=1) and ENode (F=0) VLAN Notifications.

ENode FIP Discovery

Point-to-Point VLAN Discovery Rules

- An ENode receiving a VLAN Discovery Request (with P=1?) shall compare the received MAC address to it's own ENode MAC address.
 - If the received MAC address is higher, then the VLAN Notification shall contain the VLAN descriptor suggested in the received VLAN Discovery Request
 - If the received MAC address is lower, then the ENode shall assign it's own VLAN descriptor in the VLAN notification.
- An ENode may move on to FIP FCF/ENode Discovery at any point
 - if VLAN Discovery fails or is not performed, use “default” VLAN ID
 - User configurable default; or
 - Priority tagged (null VID)

ENode FIP Discovery

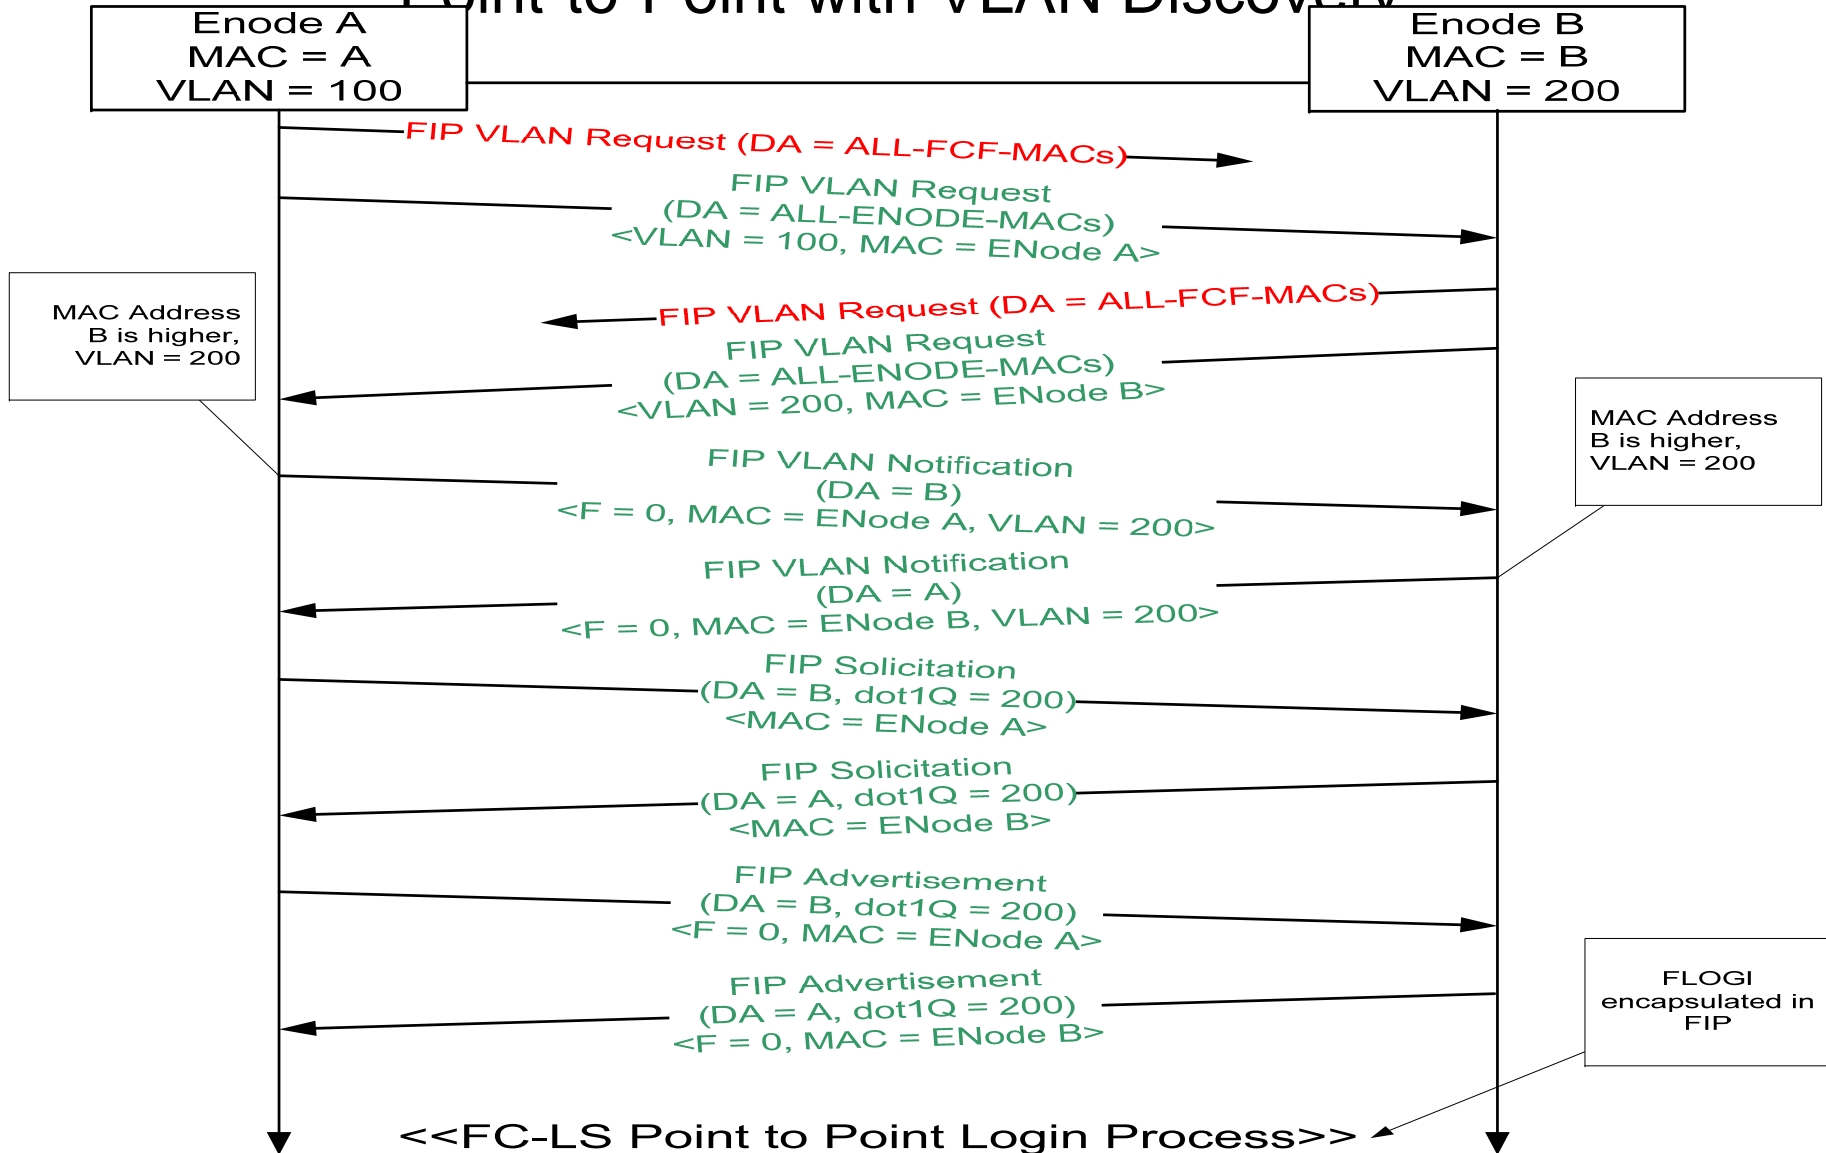
Point-to-Point (cont.)

(VLAN Discovery was successful)

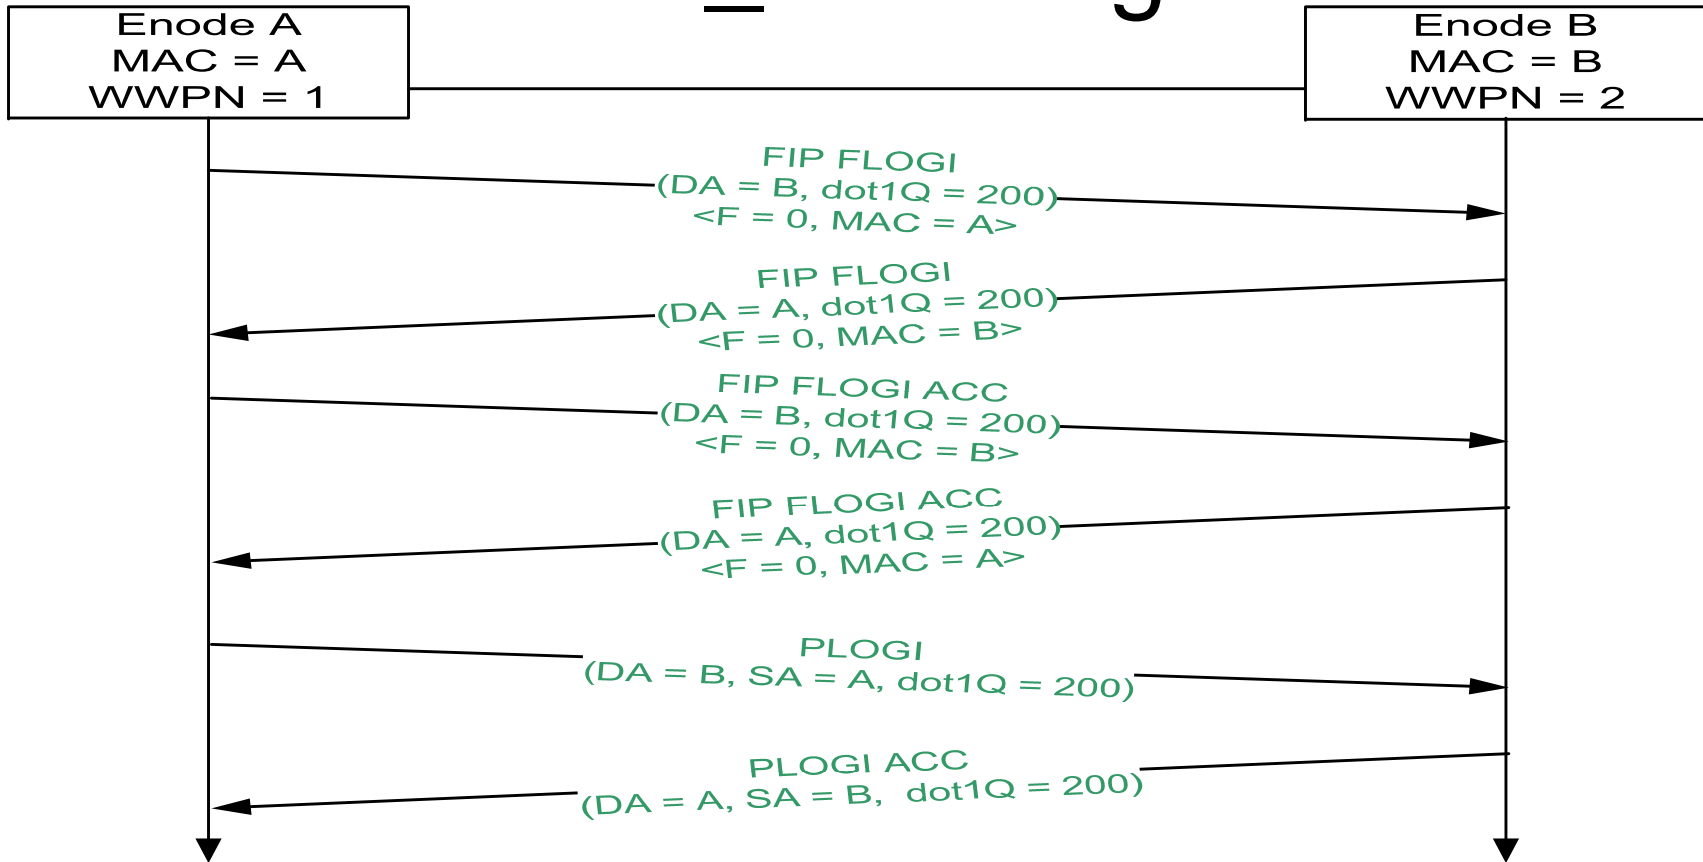
- **ENode Discovery** (performed using the negotiated FCoE VID)
 - Each ENode transmits one FIP Discovery Solicitation to the DA of the other ENode (Other ENode DA MAC obtained in the VLAN Discovery phase).
 - Each ENode responds with unicast Advertisement
 - Padded to Max-FCoE-Size
- **Login**
 - FIP FLOGI (SA = ENode MAC, DA = ENode MAC from ENode Advertisement)
 - FIP FLOGI_ACC
 - FCoE PLOGI from ENode with Higher WWPN only

ENode FIP Discovery

Point-to-Point with VLAN Discovery



VN Port to VN_Port Login Process



- Performed after receiving Point to Point Advertisement (a solicited advertisement with the “F” bit set to 0.)
- Follows the FC-LS-2 point to point login process.

ENode FIP Discovery

Point-to-Point (cont.)

(VLAN Discovery was NOT successful or not performed)

- ENode Discovery

Using “default” VLAN ID, each ENode transmits:

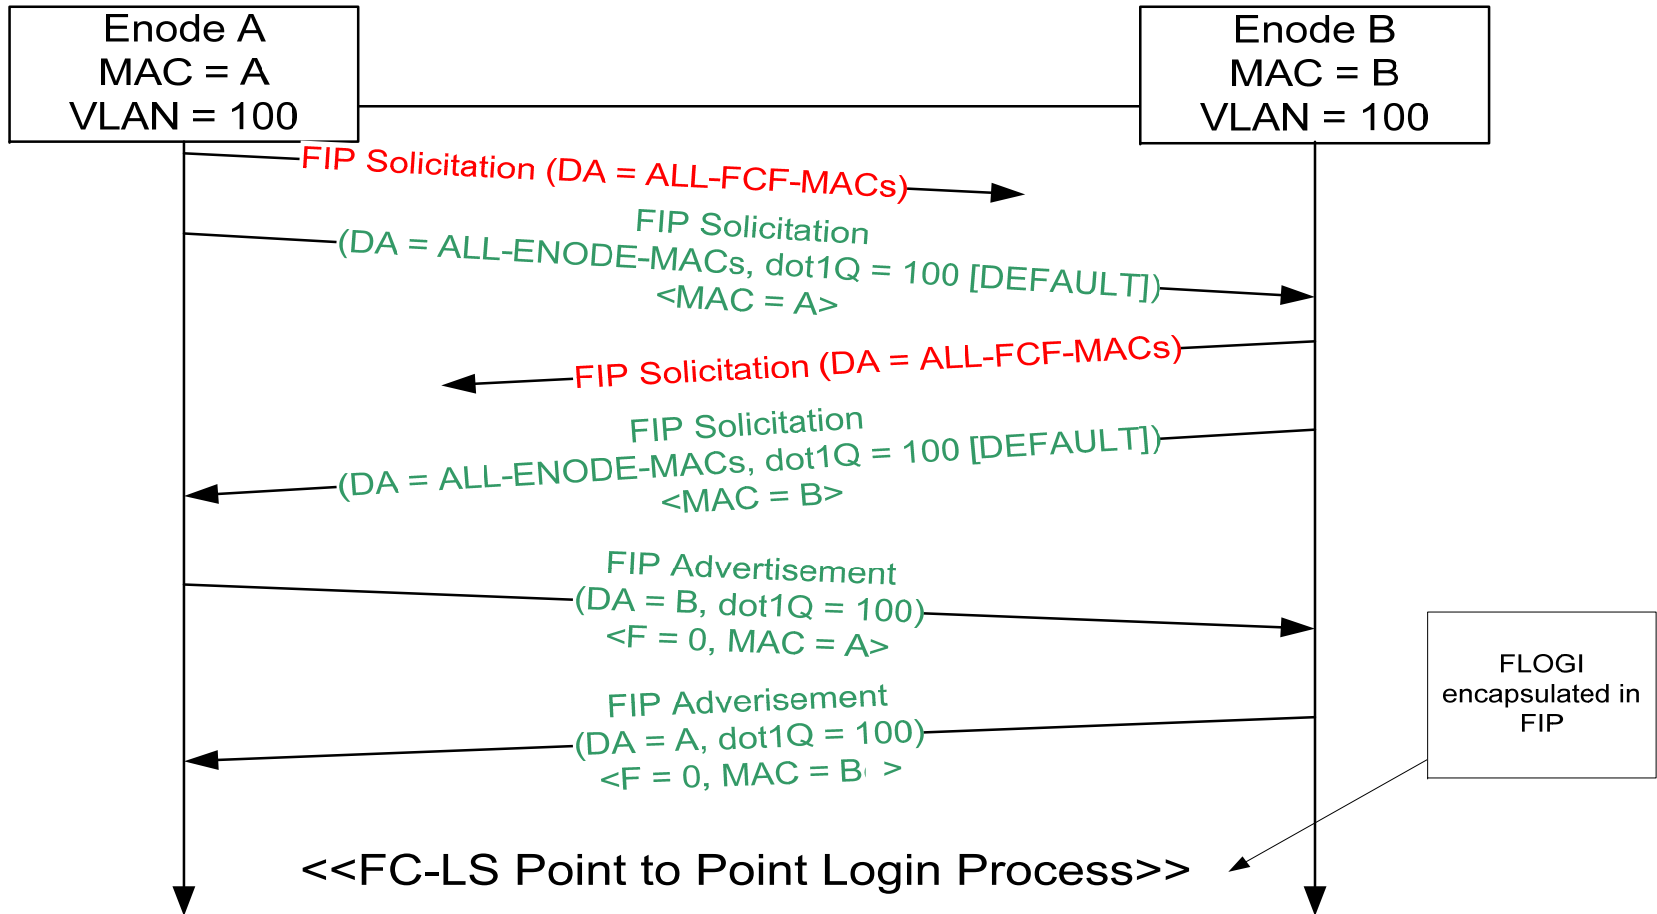
- I. FIP Solicitation (DA = ALL-FCF-MACs)
 - II. FIP Solicitation (DA = ALL-ENode-MACs)
- Each ENode responds with unicast Advertisement
 - Padded to Max-FCoE-Size

- Login

- FIP FLOGI (SA = ENode MAC, DA = ENode MAC from ENode Advertisement)
- FIP FLOGI_ACC FCoE PLOGI from ENode with Higher WWPN

ENode FIP Discovery

Point-to-Point without VLAN Discovery



ENode FIP Discovery

Point-to-Point

Important notes

- Both ENodes perform FCF and ENode Discovery in parallel
- Highest MAC Address decides “winning” ENode for VLAN ID assignment
- If FCF frame is received during Discovery process then cease ENode Discovery and complete FIP FCF Discovery
- If any frame from more than one Source MAC Address is received during the Discovery process then restart ENode/FCF Discovery or fail initialization

ENode FIP Discovery

Point-to-Point

Termination conditions

Terminate Point to Point communication

When?

- Receipt of any FIP frame with the F bit set to 1
- Receipt of any frame with Source MAC \neq “attached” ENode MAC address

How?

- FCoE LOGO to “attached” ENode
- Restart FCF/ENode Discovery process

FIP Frames for Point-To-Point

- No new FIP frames, just some content changes
- VLAN REQUEST - ALL-ENode-MACs
 - Same descriptors as used for Requests to ALL-FCF-MACs with one addition
 - MAC Address=Enode MAC
 - Name_Identifier (optional)
 - VLAN (NEW) (and maybe a “P” bit to qualify that this is OK)
- VLAN Notification (from one ENode to another)
 - No changes except that MAC address descriptor specifies the ENode MAC rather than an FCF MAC

FIP Frames for Point-To-Point (cont.)

- Discovery Solicitation – ALL-ENode-MACs
 - Flags field = 0 – no MAC Addressing specified
 - Same descriptors as ALL-FCF-MACs – MAC Address, Name_Identifier, Max FCoE_Size
- Discovery Advertisement (from ENode)
 - Flags field: S = 1, A = 1, F = 0
 - Descriptor List – Priority, MAC Address=Enode MAC, Name Identifier, (Fabric – None), FKA_ADV_Period
 - Pad to received Max_FCoE_Size

FIP Frames for Point-To-Point (cont.)

- FLOGI
 - Flags field = 0 – no MAC Addressing specified
 - Same descriptors as ALL-FCF-MACs –
 - MAC Address (requested MAC address for FCoE) **set to Enode MAC address**
- FLOGI_ACC (from ENode)
 - Flags field: F = 0 – no MAC Addressing specified
 - Same descriptors as FCF FLOGI_ACC
 - **MAC Address is set to requested MAC address from FLOGI (the Enode MAC address)**

Security Concerns?

- If attached to a properly configured bridge, initial ingress ACL will block All-Enode-MACS solicitation or VLAN request

**DA = All-FCF-MACs, Type = FIP_TYPE,
permit;Type = FIP_TYPE, deny;**

VN_Port to VN_Port Link Maintenance ?

To Maintain Compatibility with current standard, here are some rules to apply to Virtual Link Maintenance.

- The FCoE Controller for both ENode MACs that have discovered a point to point configuration (FIP FLOGI LS_ACC received with Nx_Port/F_Port bit = 0) generate FIP Keep Alive messages;
 - every FKA_ADV_PERIOD using the ENode's MAC address
- FKA_ADV_PERIOD is the smaller of that sent and received in Point to Point ENode Advertisements.
- The FCoE Controller for both ENode MACs that have discovered a point to point configuration verify that at least one FKA is received within $2.5 * FKA_ADV_PERIOD$
- Using Ethernet physical link failures can eliminate this requirement to send and monitor FKA.

Ensuring we are point to point

- Suggest an annex or informative note that states that LLDP could be used to ensure we are in a point to point configuration before attempting Enode discovery.
 - Break Layers and use LLDP – a Systems Capabilities TLV is received indicating “Station Only” (bit 7) (See IEEE-802.1AB-2005)
- This is not recommended for normative text

Thank You!!

(if we actually made it this far)