FCoE Point To Point

VN_Port to VN_Port Network Configuration

Roger Hathorn, Lou Ricci – IBM

Fred Knight - NetApp

Erik Smith – EMC

Alan Spalding, Craig Carlson – QLogic

02/03/2010

VN_Port to VN_Port Network Configuration

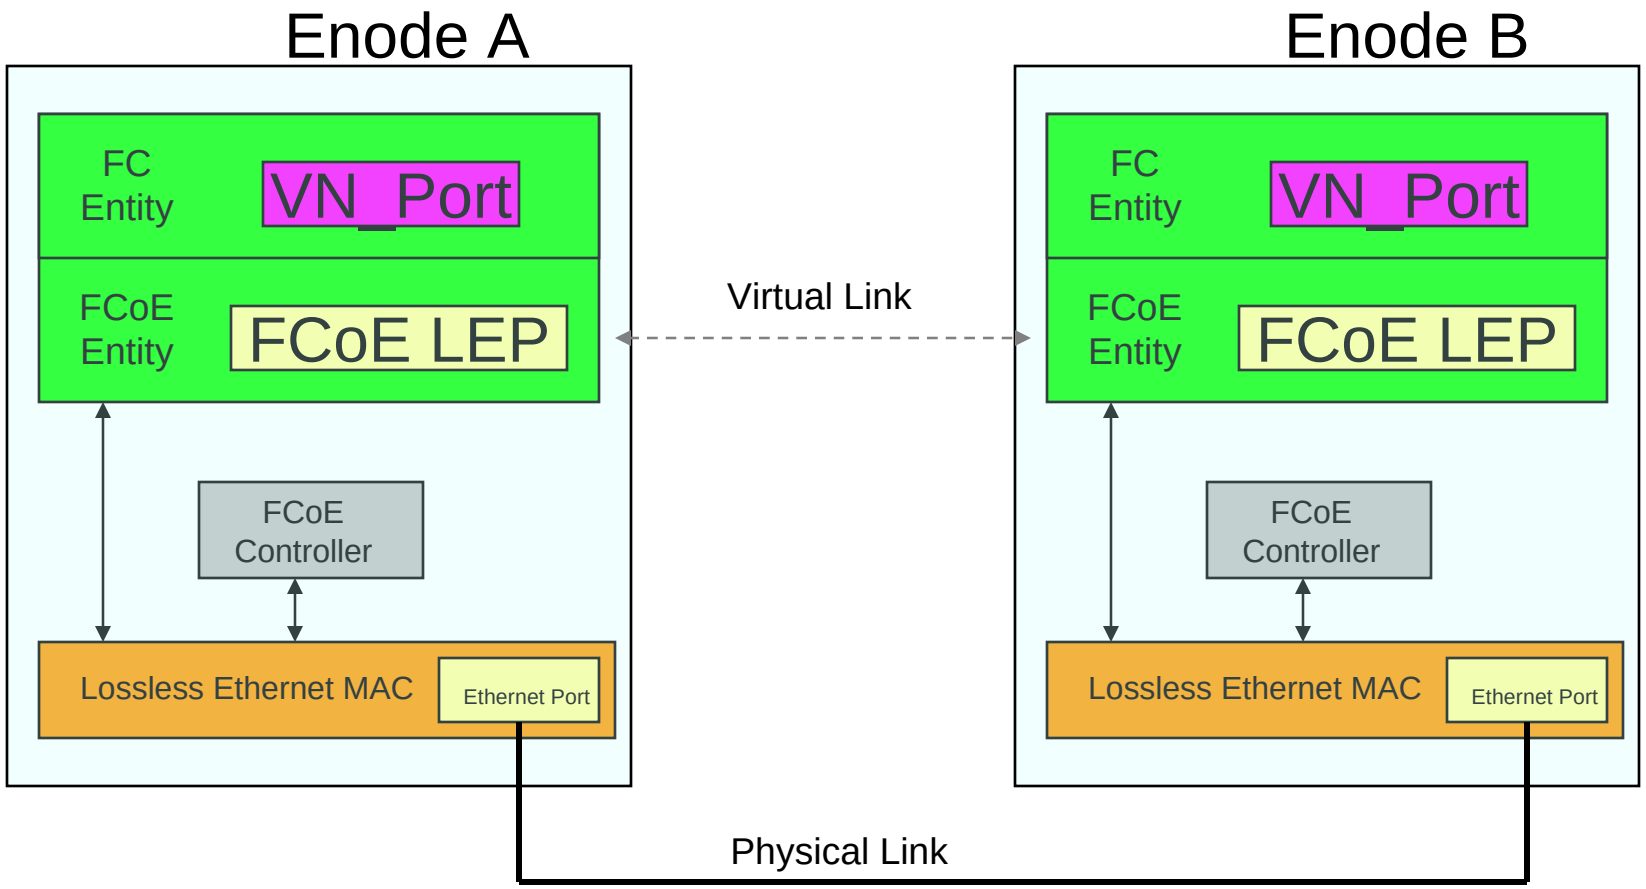

FCoE Point-to-Point - Motivation

- 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

- Topology is discovered, not configured.
- Discovery process must be the same process and work in EITHER a point to point or fabric topology.
- De-instantiate point to point virtual link and re-start discovery whenever:
  1. One or more FCFs are detected
  2. Multiple ENode(s) detected
- A point to point configuration is not detected in properly configured fabric topology
- Discovery, an more importantly, recovery completes in a timely manner (no waiting)

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.

ENode - FCF 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 (to each discovered VLAN)
  - 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 - FCF Discovery
Today

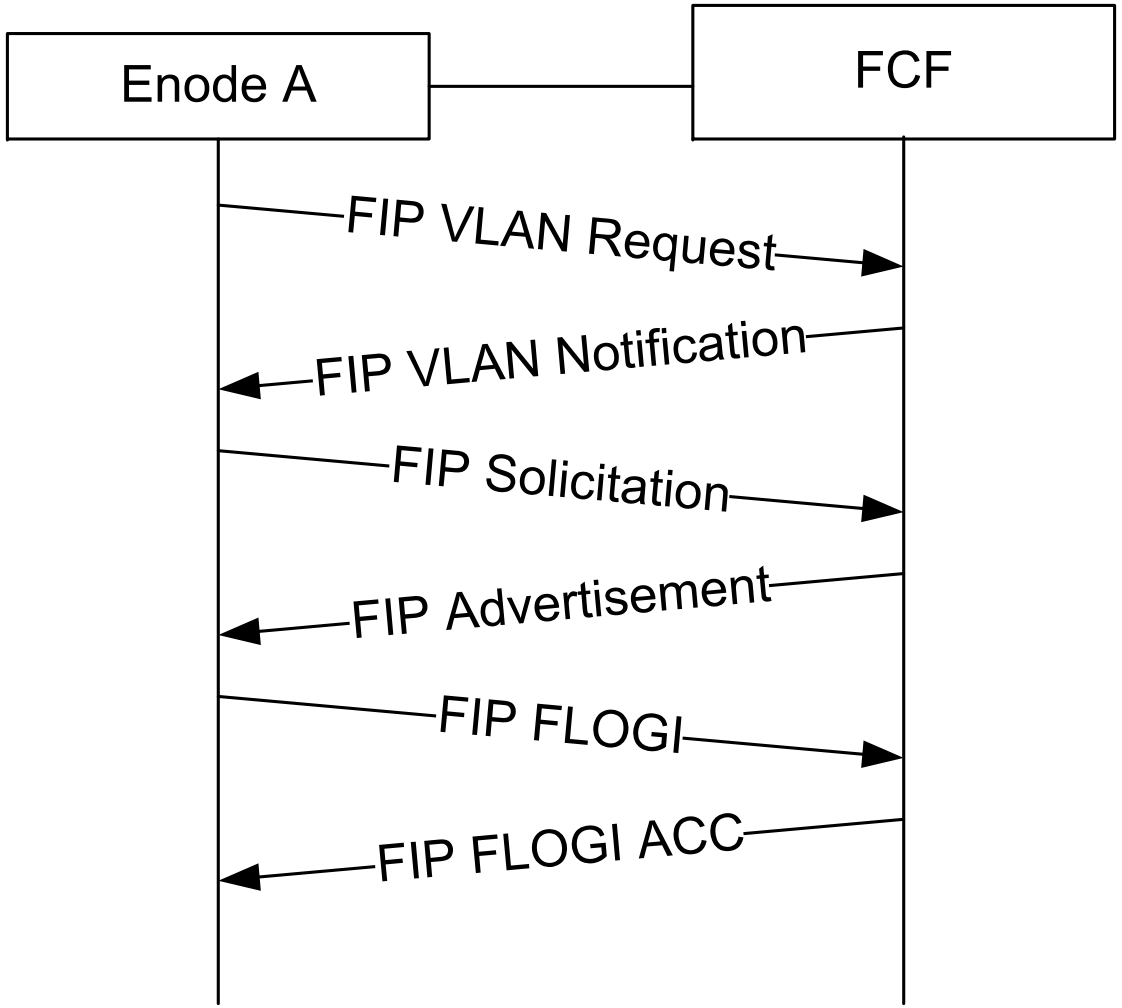

New ENode – Enode Discovery and Login

Point-to-Point

- ENode Discovery
  - Each ENode transmits Multicast FIP Discovery Solicitation to ALL-ENODES-MACs.
  - Each ENode responds with unicast Advertisement
    - Padded to Max-FCoE-Size
- Login
  - FIP FLOGI (SA = ENode-A MAC, DA = ENode-B MAC)
  - FIP FLOGI_ACC (SA = ENode-B MAC, DA = ENode-A MAC) <MAC Address = Enode-A MAC>
  - FCoE PLOGI from ENode with Higher WWPN

ENode Topology Discovery

With Point-to-Point Support

- Enode-FCF and Enode-Enode discovery execute simultaneously: Transmit two FIP Discovery Requests
  - I. FIP FCF Discovery Solicitation (DA = ALL-FCF-MACs)
  - II. FIP ENode Discovery Solicitation (DA = ALL-ENode-MACs)
- Wait for FIP Response(s)
  - I. FIP FCF Advertisment (SA = FCF MAC, DA = ENode MAC, **F bit = 1**; OR
  - II. FIP Enode Advertisement (SA = ENode-B MAC, DA = ENode-A MAC, **F bit = 0**)
- Response indicates topology.
  - a) FCF Discovery Advertisement (F bit = 1) = Fabric Topology. Continue with FLOGI; OR
  - b) Enode Discovery Advertisement (F bit = 0) = Point to Point Topology. Continue with FLOGI.
- If properly configured, will not receive both responses.

Enode Topology Discovery

Point-to-Point Topology

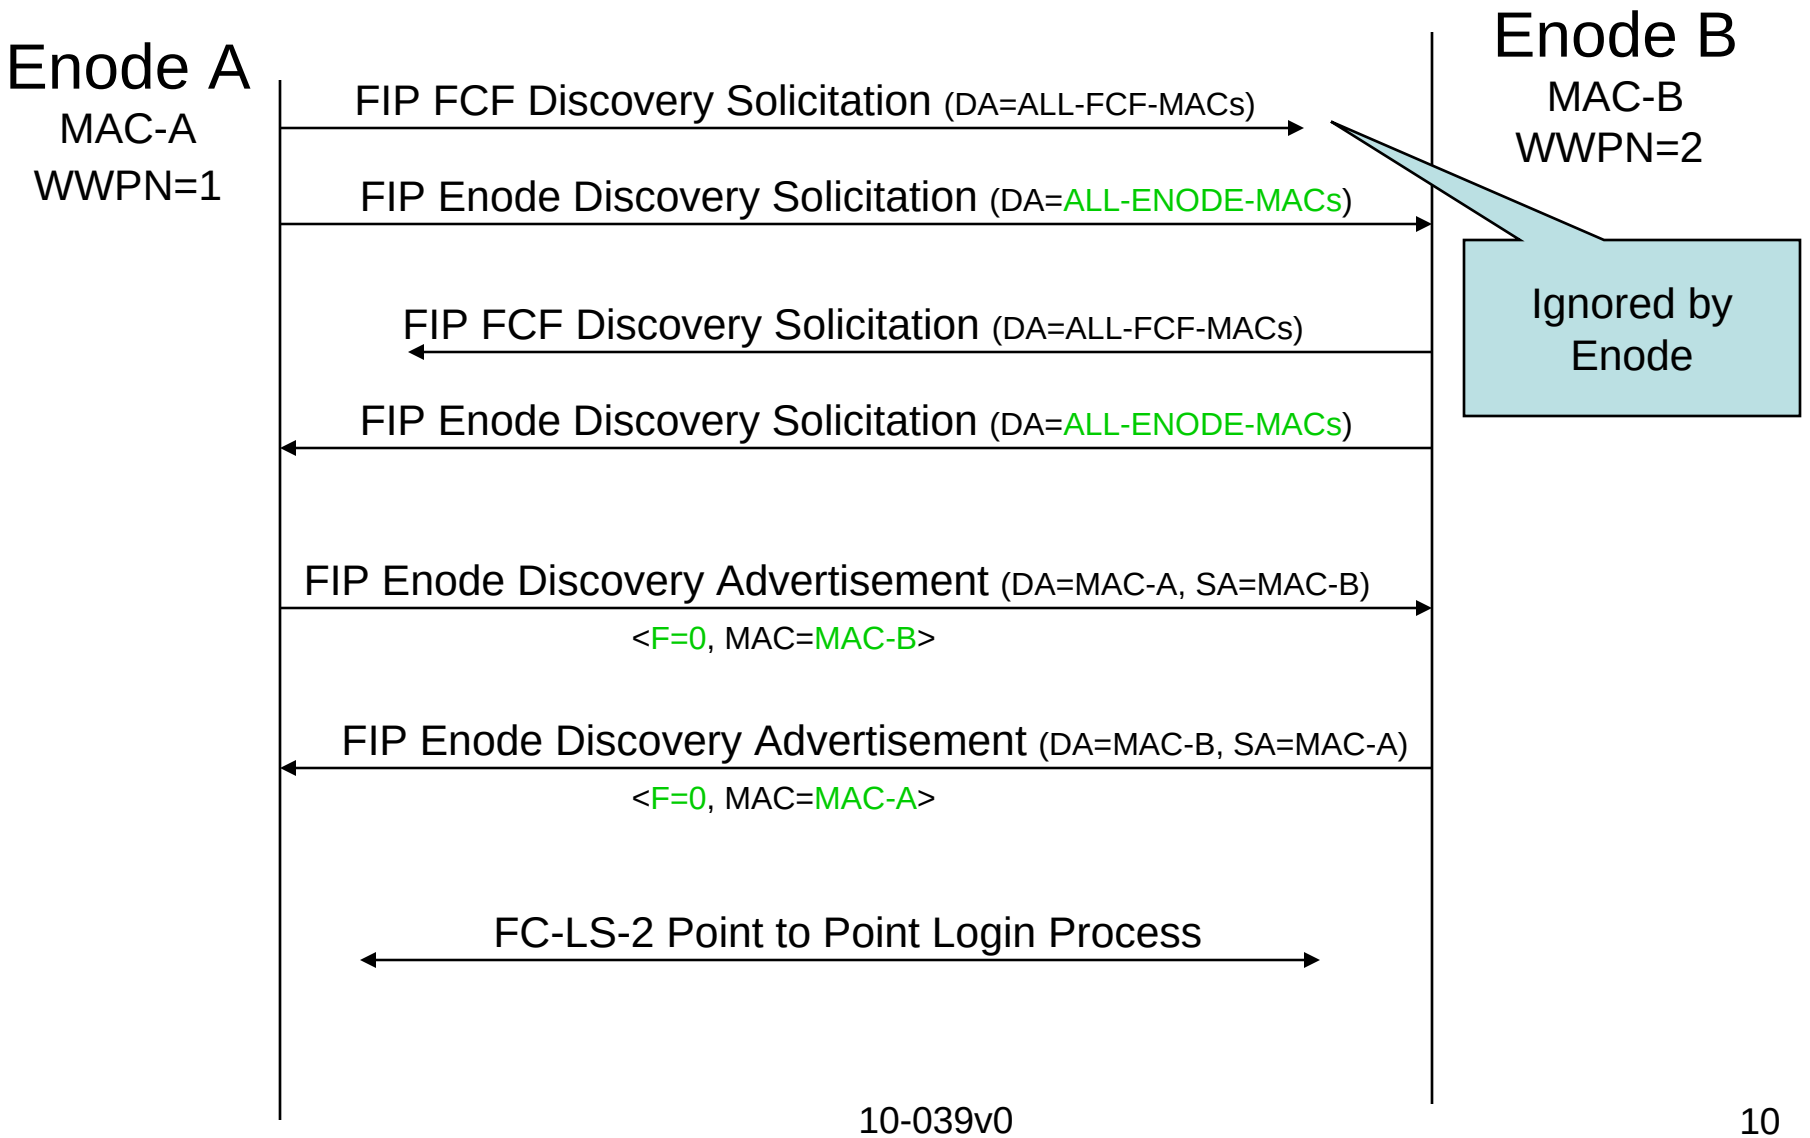

VN_Port to VN_Port Login Process

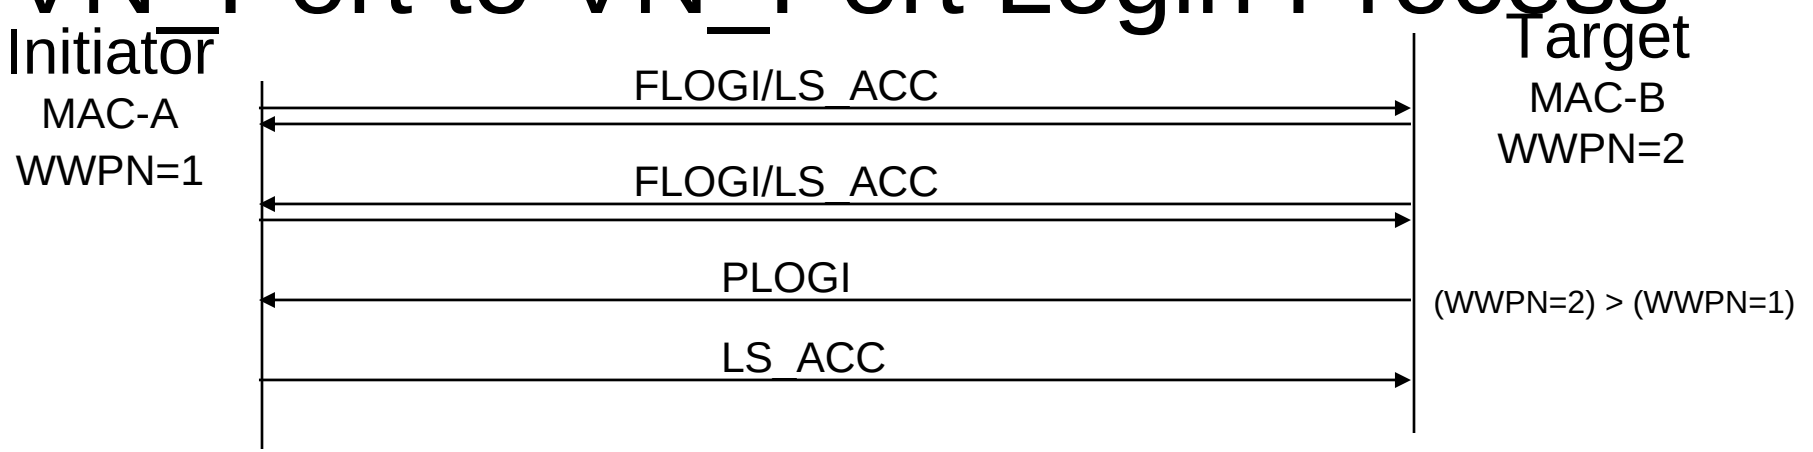

- Performed after receiving Enode Discovery Advertisement (a solicited discovery advertisement with the “F” bit set to 0.)
- Follows the FC-LS-2 point to point login process.
- Uses FIP FLOGI and LS_ACC with the requester’s ENode MAC address in the MAC address Descriptor (like SPMA)
- The Enode MAC addressed are used for FCoE traffic.

ENode Topology Discovery

With Point-to-Point Support
(and optional VLAN Discovery)

- Enode-FCF and Enode-Enode discovery execute simultaneously: Transmit two FIP Discovery Requests
  I. FIP VLAN Discovery DA = ALL-FCF-MACs
  II. FIP ENode Solicitation DA = ALL-ENode-MACs
- Wait for FIP Response(s)
  I. FIP VLAN Notification (SA = FCF MAC, DA = ENode MAC, F bit = 1; OR
  II. FIP Enode Advertisement (SA = ENode-B MAC, DA = ENode-A MAC, F bit = 0)
- Response indicates topology.
  a) VLAN Notification = Fabric Topology. Continue with FCF Solicitation; OR
  b) Enode Advertisement = Point to Point Topology. Continue with FLOGI.
- If properly configured, will not receive both responses.

Enode Topology Discovery

Point-to-Point Topology
(with optional VLAN Discovery)

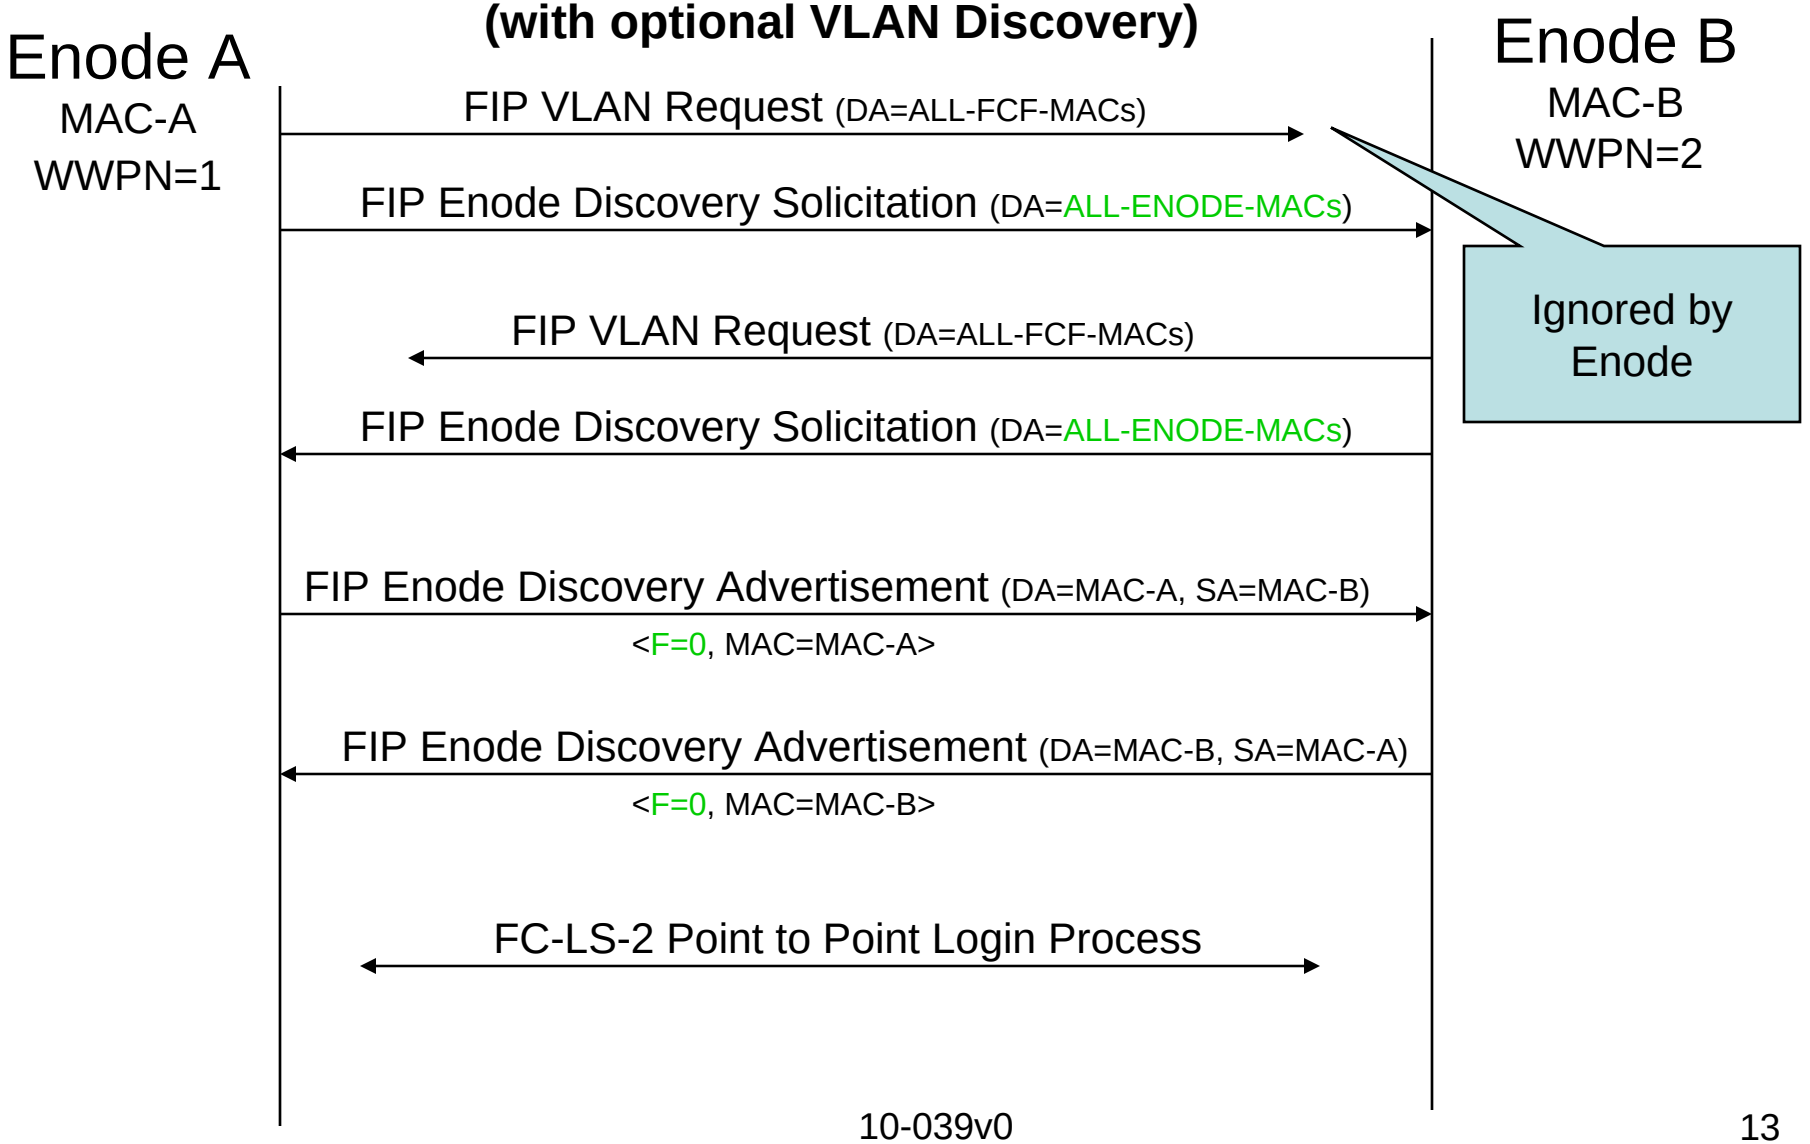

Enode Topology Discovery

Fabric Topology
(with optional VLAN Discovery)

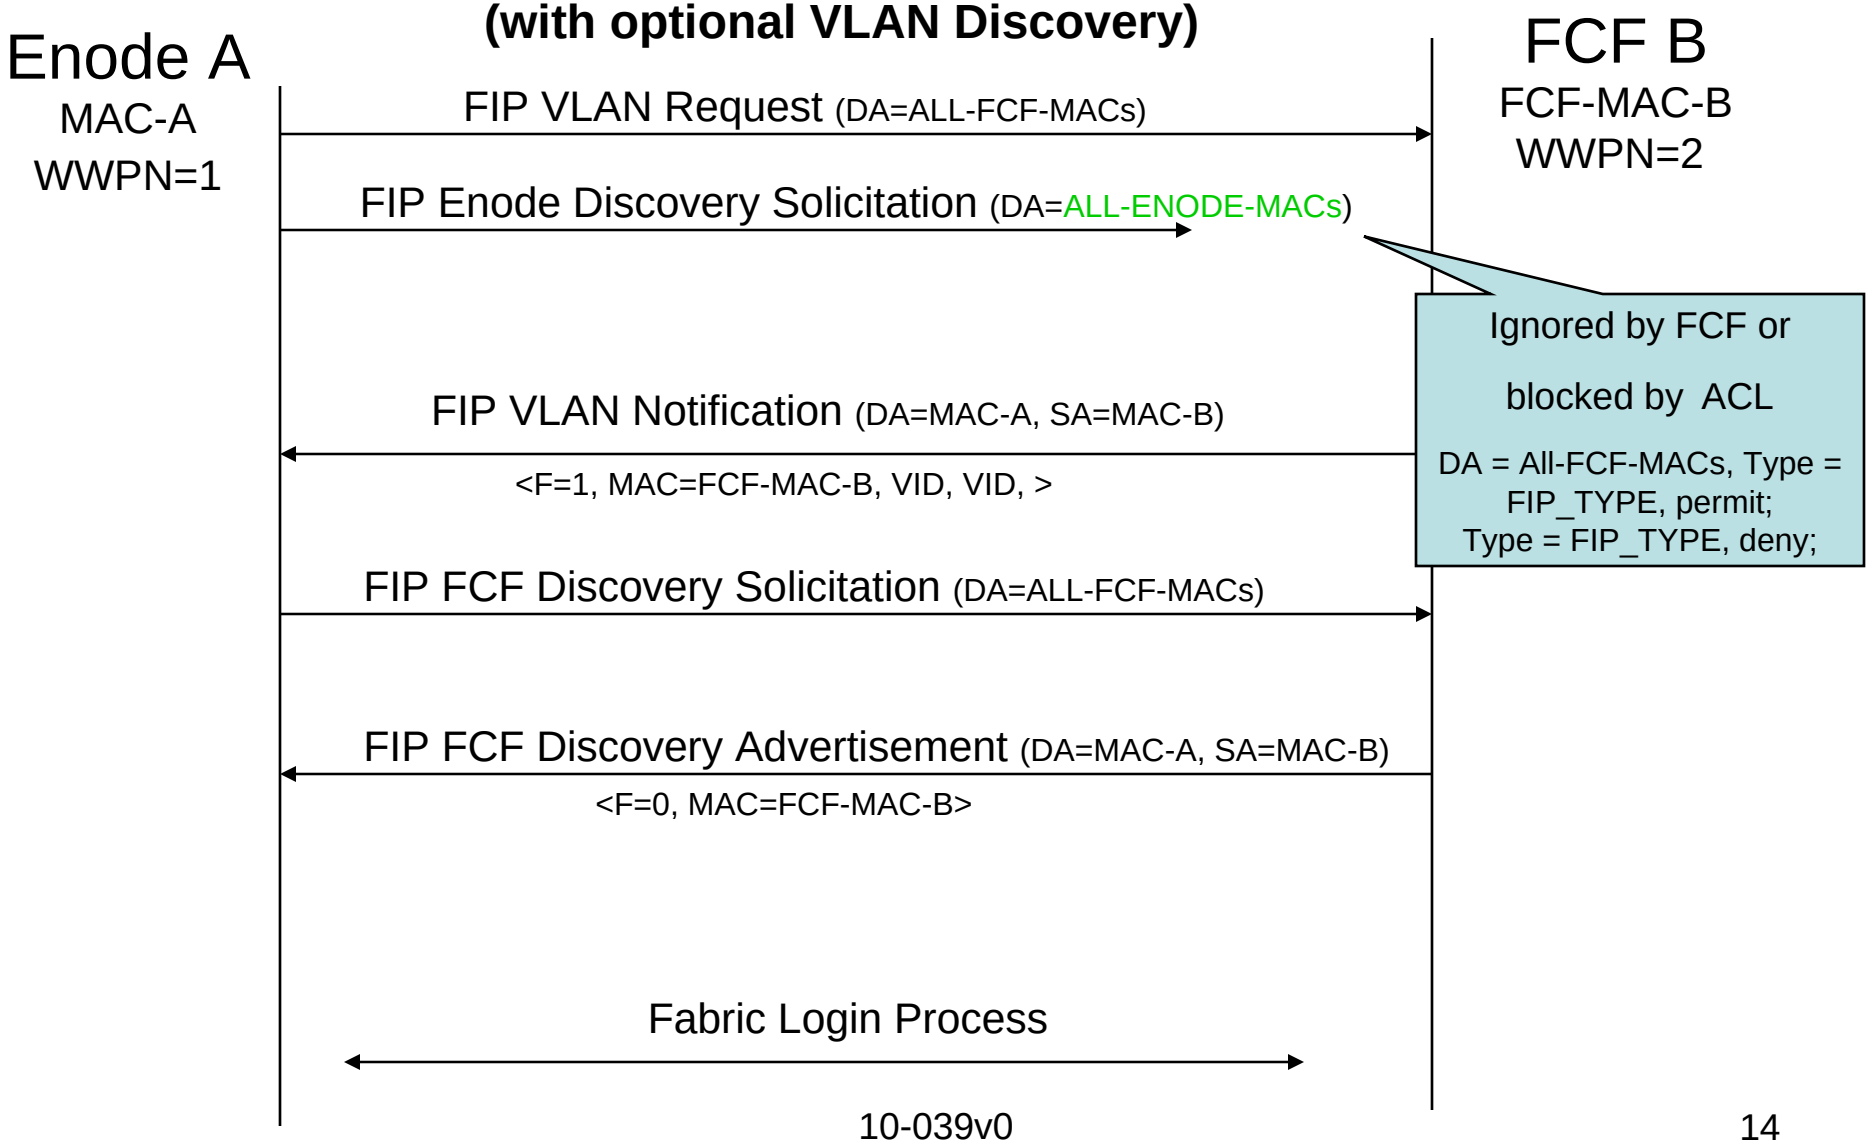

ENode FIP Discovery

Point-to-Point

Important notes

- Both ENodes perform FCF and ENode Discovery in parallel
- If FCF frame (F=1) is received during Discovery process then cease Enode Discovery and complete FIP FCF Discovery
- If any frame from more than one ENode SA is received during the Discovery process then restart ENode/FCF Discovery or fail initialization (This is an indication of misconfiguration)
- There are No delays or waiting

ENode FIP Discovery

Point-to-Point

Termination conditions

Terminate Point to Point communication

When? – Detect a misconfiguration

- Receipt of any FIP frame with the F bit set to 1
- Receipt of any frame with Source MAC <> “attached” Enode MAC address

How?

- FCoE LOGO to “attached” Enode
- De-instantiate VN_Port-VN_Port virtual Link
- Restart FCF/Enode Discovery process

FIP Frames for Point-To-Point

- No new FIP frames, just some content changes
- 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 = 0
  - 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.

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

- In case the ACL is not in place, if any FIP frames are received from more than one Enode, or both an Enode and and FCF, a configuration error is detected and Fabric topology is assumed.

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 each ENode MAC that has discovered a point to point configuration and successfully instantiates a VN_Port to VN_Port virtual link (FIP FLOGI LS_ACC received with Nx_Port/F_Port bit = 0) generates FIP Keep Alive messages;
  - every FKA_ADV_PERIOD using the ENode's MAC address
- The FKA_ADV_PERIOD to be used is the smaller of that sent and received in Point to Point Solicited ENode Advertisements (The higher of the desired frequencies).
- The FCoE Controller of an ENode MAC that has instantiated a VN_Port to VN_Port virtual link verifies that at least one FKA is received within 2.5*FKA_ADV_PERIOD on that virtual link.

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!!

(again)