FCoE: Common Addressing Structure

John L. Hufferd (Brocade)
Robert Snively (Brocade)
Steve Wilson (Brocade)
Suresh Vobbilisetty (Brocade)
Uri Elzur (Broadcom)
Pat Thaler (Broadcom)
Claudio DeSanti (Cisco)
Joe Pelissier (Cisco)
Landon Noll (Cisco)
David Black (EMC)
Parag Bhide (Emulex)
Bob Nixon (Emulex)
Joy Jiang (Finisar)
Michael Krause (HP)

Scott Carlson (IBM)
Dan Eisenhauer (IBM)
Roger Hathorn (IBM)
Renato Recio (IBM)
Gary Y Tsao (Intel)
Diego Crupnicoff (Mellanox)
Frederick Knight (NetApp)
Ed Bugnion (Nuova)
Silvano Gai (Nuova)
J.R. Rivers (Nuova)
Ed McGlaughlin (Qlogic)
Craig W. Carlson (Qlogic)

T11/08-038v2

Background

Two addressing approaches have been proposed

- SPMA (Server Provided MAC Addresses)
- FPMA (Fabric Provided MAC Addresses)
  - AKA Mapped MAC Addresses
- Extensive discussions in the industry have proven that:
  - Both have value
  - Needed a way to keep the value of both approaches
  - Both approaches “SHOULD” be supported

Therefore:

We created a new “FCoE Initialization Protocol” (FIP) which contained

- The Discovery Phase
- The Login Phase

FIP has its own Ethertype (Ethertype=FIP)

FIP has its own Frame format

- Independent from the FCoE frame format
- Easy to detect and direct to “control plane” of Ethernet Switches and FCFs

FIP and FCoE

Discovery Phase
- Provides visibility between ENodes and FCFs
- "Same" as today plus perhaps some extensions
  - Addressing Capabilities
  - Possible Duplicate MAC detection

Login Phase
- Creates association between ENode and FCF
- Two allowed alternatives
  - FPMA (i.e., Mapped)
  - SPMA
- Works with FLOGI, FLOGI ACC, LOGO, …

Data Transfer Phase
- PLOGI/PRLI
- All other FC protocol frames (ELS, FC4 ULPs. etc.)

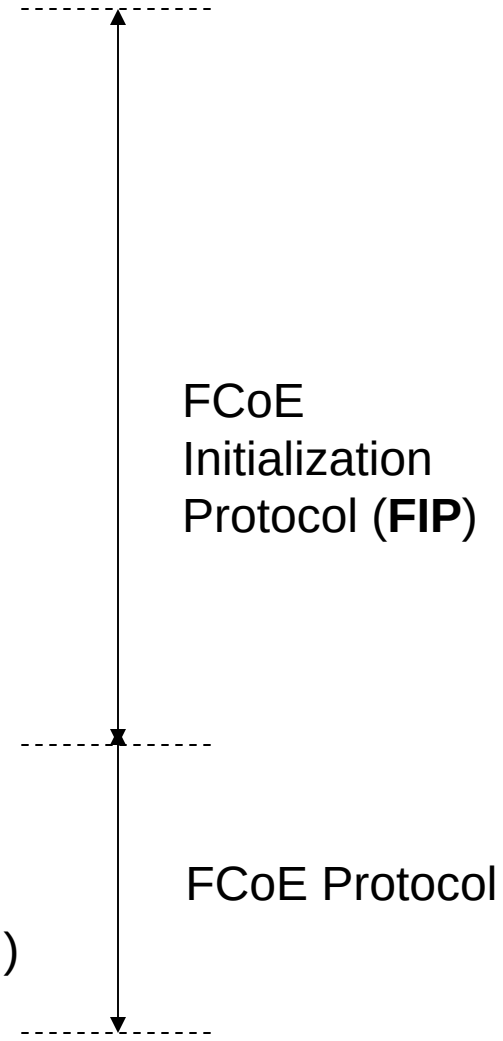

The FIP Protocol

Dedicated to control traffic, it includes:

- the Discovery Protocol
- FLOGI
- FDISC with S_ID = 0.0.0
- LOGO with D_ID = F_Port Controller
- Related Accepts

Separate Ethertype to simplify processing and to allow more flexibility in the frame format

Login Phase

ENode performs a Login with the FCF

- Sends FLOGI in the FIP Frame (Ethertype=FIP)

FCF accepts the Login

- Sends back an FLOGI ACC in the FIP Frame (Ethertype=FIP) – see details in “FCoE: Address Assignment Mechanism” T11-08-039v0
  - Includes the assigned MAC address
    - A Fabric Provided MAC Address
      (constructed by concatenating the 24-bit FC-MAP with the assigned N_Port_ID)

    or

    - The SA MAC Address used in the FLOGI (i.e., SPMA)

Separate Ethertype for FCoE Initialization Protocol (FIP)

- Simplifies the Snooping for FCoE link status changes
- Simplifies ACL deployment (including dynamic ACLs)
- Simplifies handling in the FCoE Controllers

ENode Login Requirements

Regardless of addressing method

- ENode simply uses the MAC address returned by the FCF in the FIP login response
  - After possibly checking MAC address consistency

NPIV functionality is preserved

Initial Login Flow ladder (2 Phases)

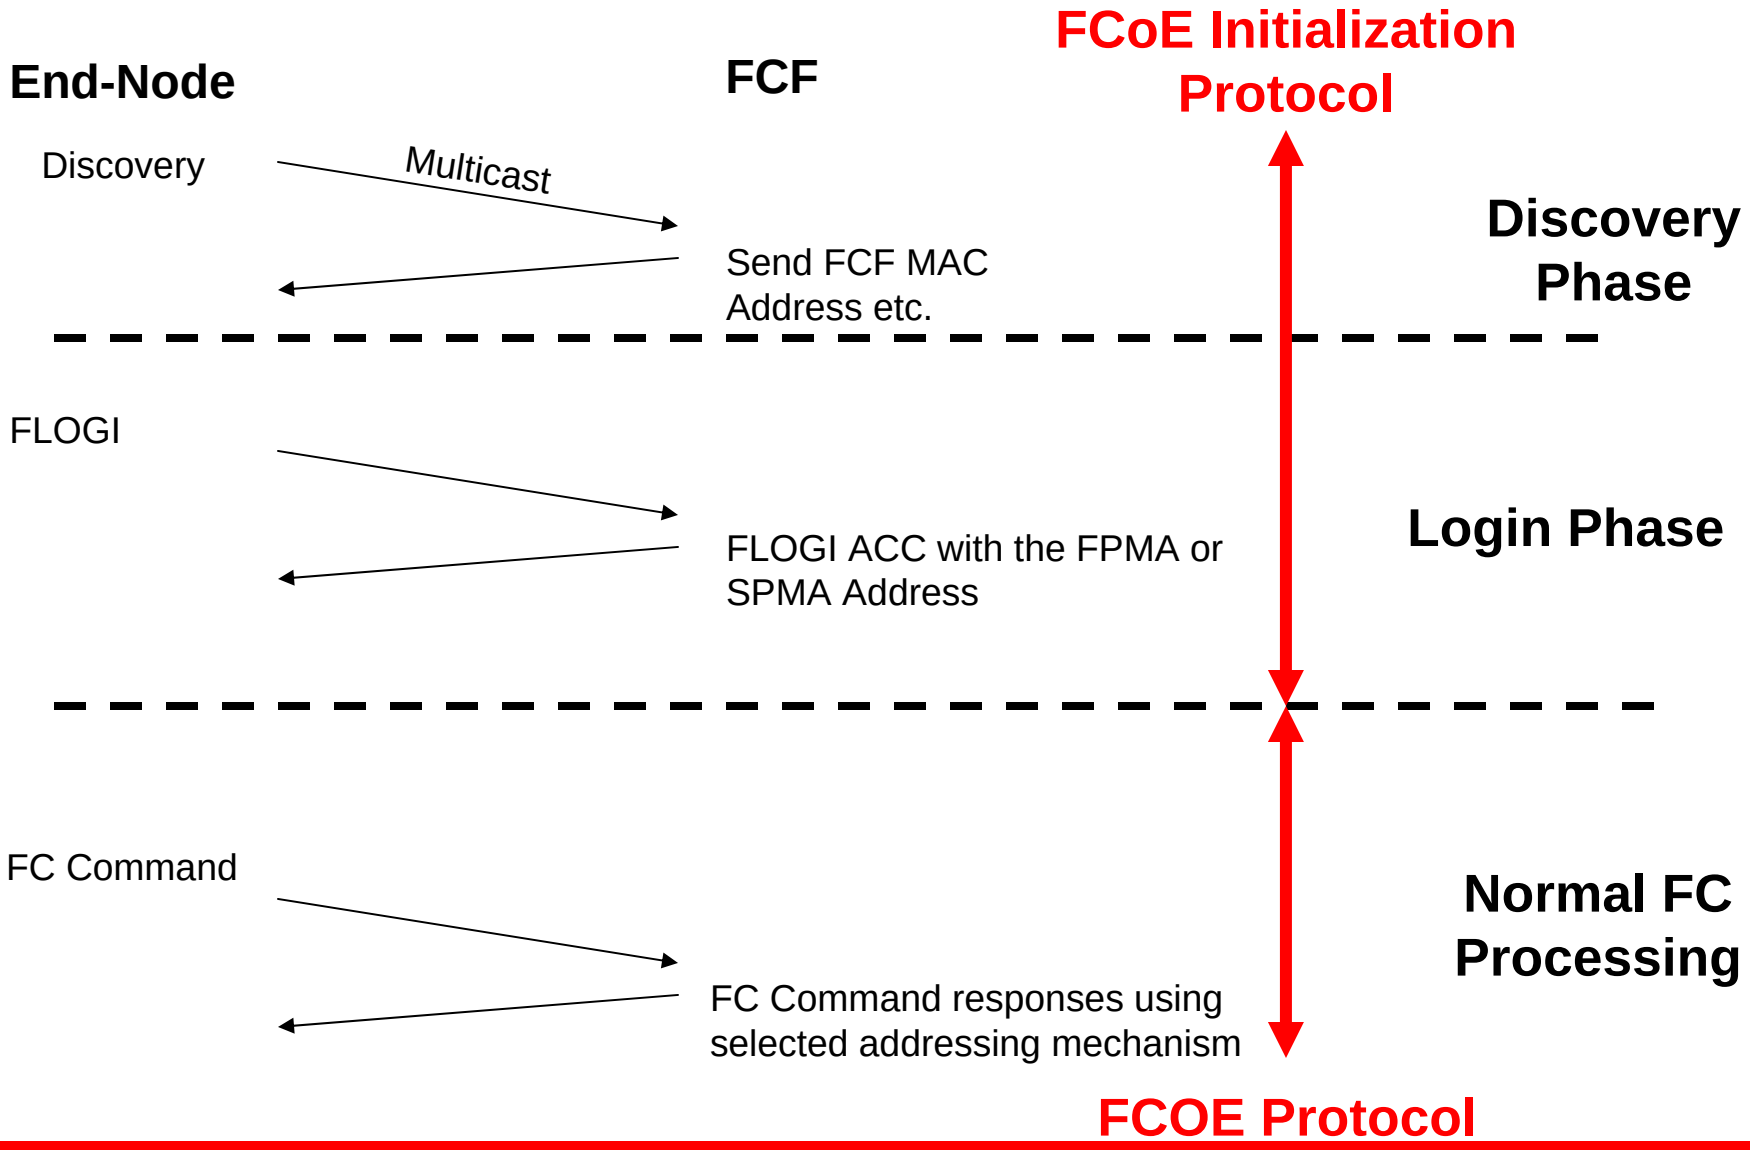

Subsequent Login Flow ladder (1 Phase)

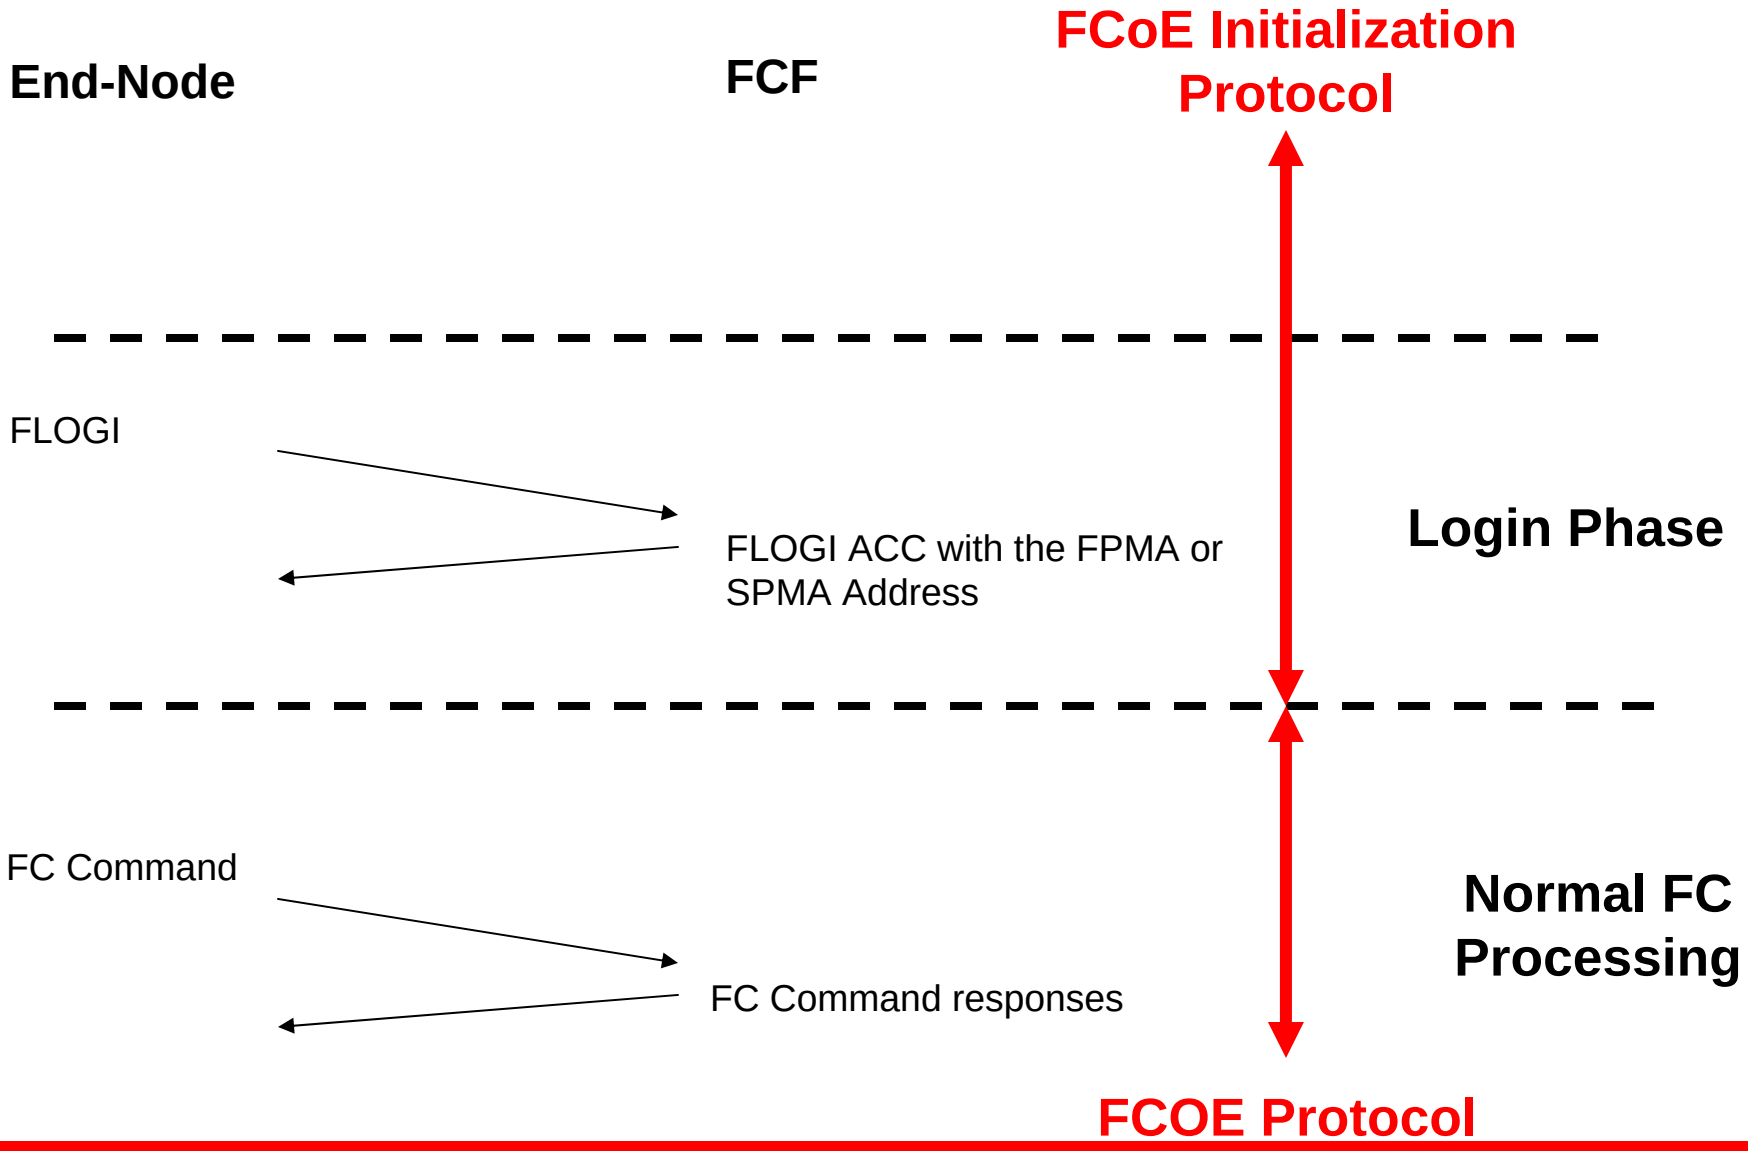

Summary

One unified approach to ENode MAC addressing

- Either FPMA or SPMA Methods may be used
  - Capability is indicated by the ENode
  - Method is Selected by the FCF
  - ENode can almost not care about the method
- FCF may support its favorite Method
  - If FCF can support both: then Admin may set the preference
- Permits a common method of creating Dynamic ACLs
- Does not require any frame format change for the FCoE protocol
- Both methods “SHOULD” be supported

Motion

To accept that FC-BB-5 define both Fabric Provided and Server Provided MAC Addresses, without mandating support for either one of them, through the unified FCoE Initialization Protocol (FIP) described at high level in this document

To accept that the FIP protocol use a different Ethertype than FCoE frames

To accept that FCoE frames have the FCoE frame format specified in T11/07-479v0 regardless of the addressing scheme used