SNIA

Education

Fibre Channel over Ethernet (FCoE)

John L Hufferd, Consultant
Hufferd Enterprises

SNIA Legal Notice

- The material contained in this tutorial is copyrighted by the SNIA.
- Member companies and individual members may use this material in presentations and literature under the following conditions:
  - Any slide or slides used must be reproduced in their entirety without modification
  - The SNIA must be acknowledged as the source of any material used in the body of any document containing material from these presentations.
- This presentation is a project of the SNIA Education Committee.
- Neither the author nor the presenter is an attorney and nothing in this presentation is intended to be, or should be construed as legal advice or an opinion of counsel. If you need legal advice or a legal opinion please contact your attorney.
- The information presented herein represents the author's personal opinion and current understanding of the relevant issues involved. The author, the presenter, and the SNIA do not assume any responsibility or liability for damages arising out of any reliance on or use of this information.

  NO WARRANTIES, EXPRESS OR IMPLIED. USE AT YOUR OWN RISK.

Abstract

- A new concept has just been standardized in the Fibre Channel (T11) standards committee; it is called Fibre Channel over Ethernet (FCoE)
- The FCoE standard specifies the encapsulation of Fibre Channel frames into Ethernet Frames and the amalgamation of these technologies into a network fabric that can support Fibre Channel protocols and other protocols such as TCP/IP, UDP/IP etc.
- The tutorial will show the Fundamentals of the FCoE concept and describe how it might be exploited in a Data Center environment

Agenda

- Introduction
- FCoE Fabrics & Convergence
- Architecture
- Discovery & Link Instantiation
- Topologies
- Scenarios
- Summary

Introduction

- This presentation provides an overview of Fibre Channel over Ethernet (FCoE)
- One should think about FCoE as placing the FC protocol on a new physical link
  - New Lossless Ethernet links instead of physical FC links
  - But it is still Fibre Channel
- The protocol has been defined in the INCITS Fibre Channel (T11) technical committee
- The new Lossless Ethernet links is being defined in the IEEE 802.1

FCoE Fabrics & Convergence

FCoE Fabrics (part 1)
(This is NOT Traditional Ethernet)

- **FCoE requires specific Ethernet extensions to be implemented**
  - Lossless switches and fabrics (e.g., supporting IEEE 802.3 PAUSE) configurations are required
  - Jumbo frame support is strongly recommended (not a standard, but widely available)

- **Deployments of FCoE should utilize the advances in Ethernet currently being discussed in IEEE 802.1, specifically:**
  - Priority-based Flow Control (PFC) →802.1Qbb
  - Enhanced Transmission Selection (ETS) → 802.1Qaz
  - DCB (capability) eXchange (DCBX) Protocol → 802.1Qaz
  - Possible future→Congestion Notification (802.1Qau), & Multi-pathing (IETF– TRILL)

- **These 802.1 advances are important for Converged Flows (Messaging, Clustering and Storage)**

  This set of functions is called CEE – Converged Enhanced Ethernet (intended for a Data Center Environment) or (in the IEEE) DCB -- Data Center Bridging

- FCoE Fabrics must be built with FCoE – CEE/DCB Switches that:
  - Are called FC Forwarder → FCF
  - Are part of a lossless Ethernet Fabric and have CEE/DCB Lossless Ethernet ports
  - Also provide functions of traditional FC switches (capabilities and services)

Converged Enhanced Ethernet (CEE) & FCoE Switch (FCF) with FC connections

Implementations are combining the features and capabilities of a CEE/DCB Switch with the features and capabilities of a FC switch which will:

- Support Ethernet and IP standards for switching, pathing and routing
- Support FC standards for switching, pathing and routing
- Support current and Converged Enhanced Ethernet Standards
- Adapt between FCoE and FC

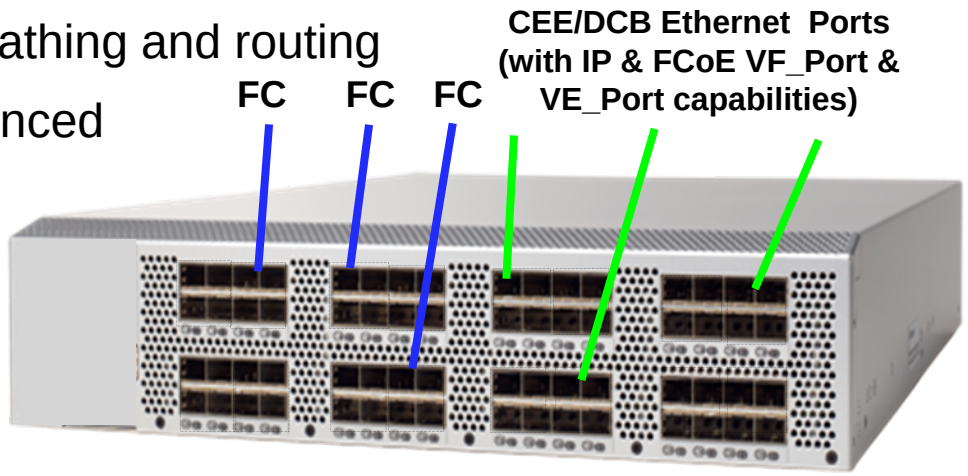

The FCoE Ports have F_Port or E_Port functions

Called VF_Ports and VE_Ports

(Because many logical (virtual) ports can share one physical port)

Connections to an Integrated CEE/DCB - FCoE Switch

- Fibre Channel is carried over lossless Ethernet as a L3 protocol

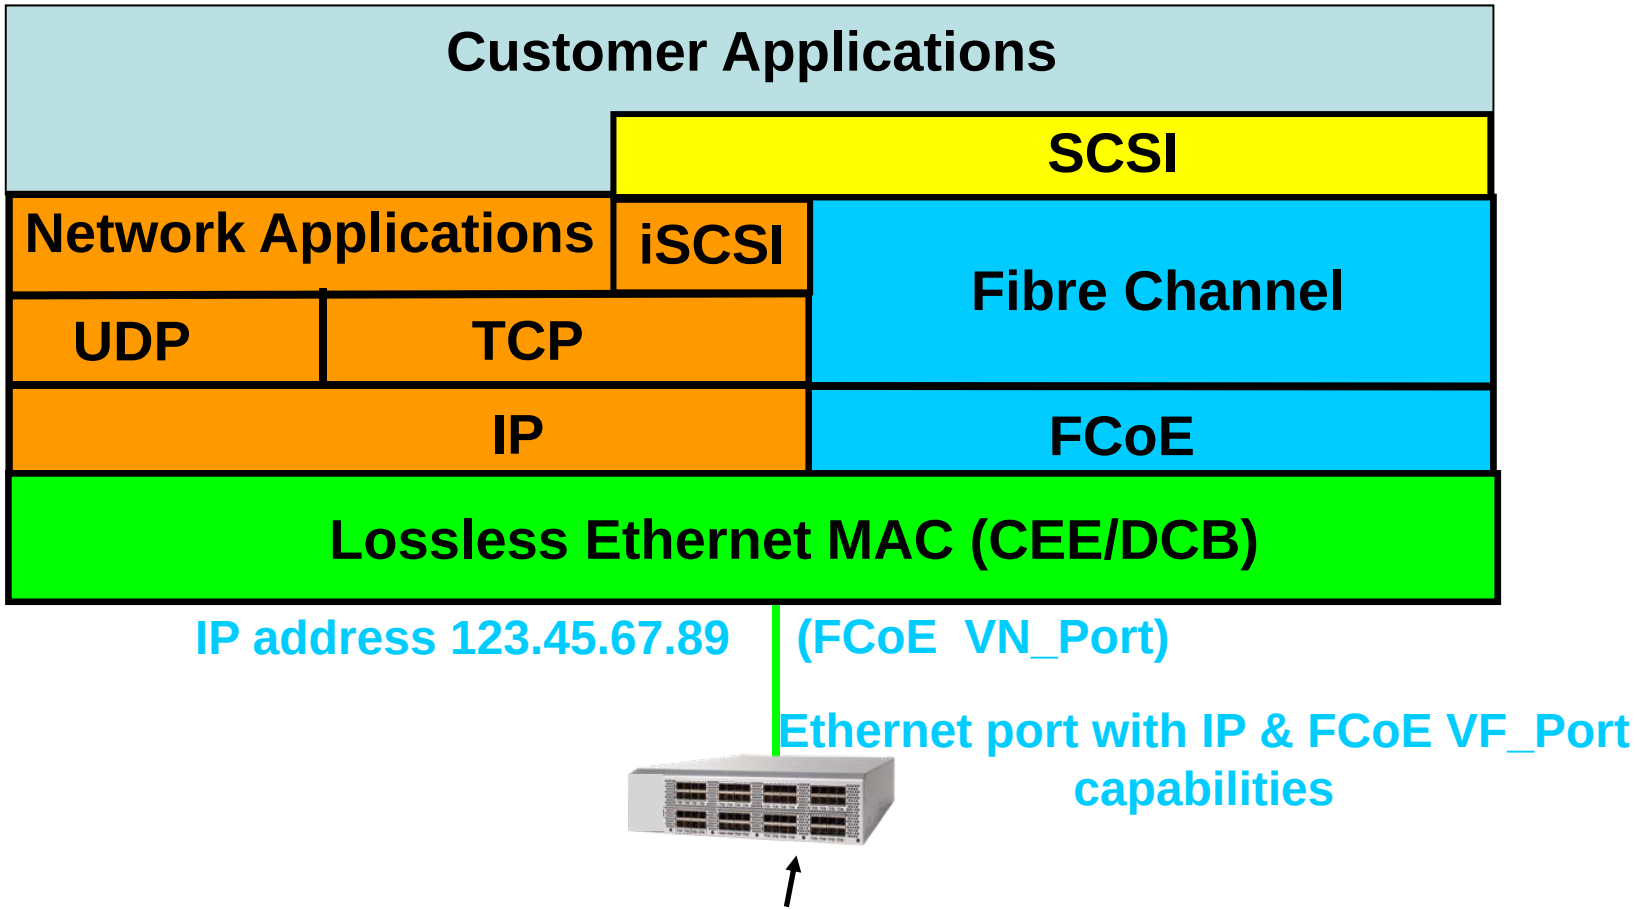

FCoE Fabrics (part 2)
(This is NOT Traditional Ethernet)

- CEE/DCB only Ethernet switches may also exist in an FCoE capable Fabric but one or more switches with FCoE capabilities must also exist
- FCoE fabrics must inter-operate seamlessly with real FC Fabrics
- FC services must operate identically on FCoE fabrics and Fibre Channel fabrics
- FCoE must support all Fibre Channel advanced features (e.g. virtual fabrics, IFR, security, etc.) transparently
- **FCoE will** not **require** changes to FC software (Apps, Drivers, etc.)
  - However, vendors will enhance Drivers & Mgnt to exploit new capabilities
- FCoE is NOT a replacement for FCIP
  - FCIP is for inter-switch links **beyond the Data Center**
  - FCIP uses TCP/IP

The Compelling Value of Convergence is at the (Server Edge) Interface

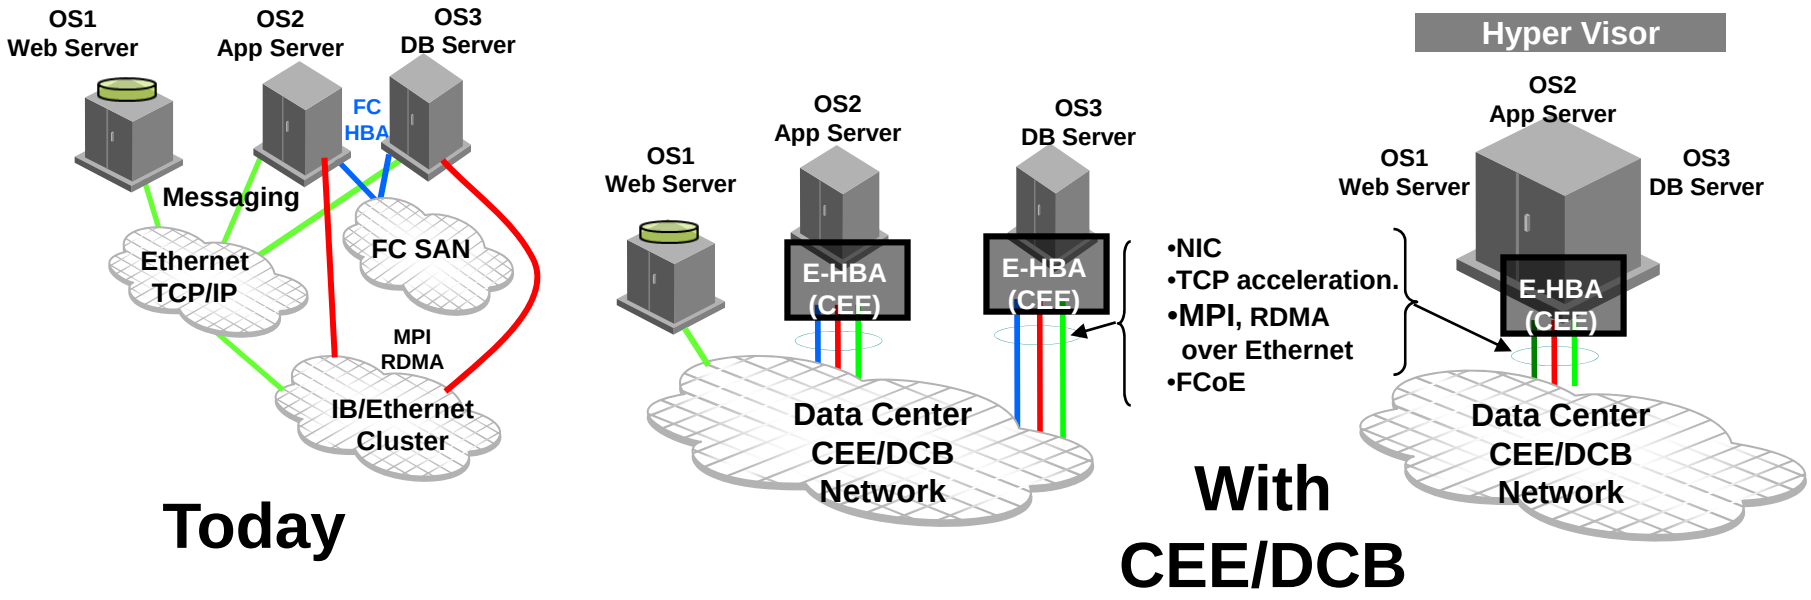

Converged Enhanced Ethernet (CEE/DCB)

(A "Lossless" Ethernet with Priority based Flow Control and Scheduling)

- **Dramatic Interface reduction in adapters, switch ports, cabling, power, & cooling**
  - 4-6 cables can be reduced to 2 Interfaces/cables per server
- **Seamless connection to the installed base of existing SANs and LANs**
- **Effective sharing of high bandwidth links**

But Even a Total Data Center (CEE/DCB) Fabric Requires Phase in (starting at the Server Edge Interface)

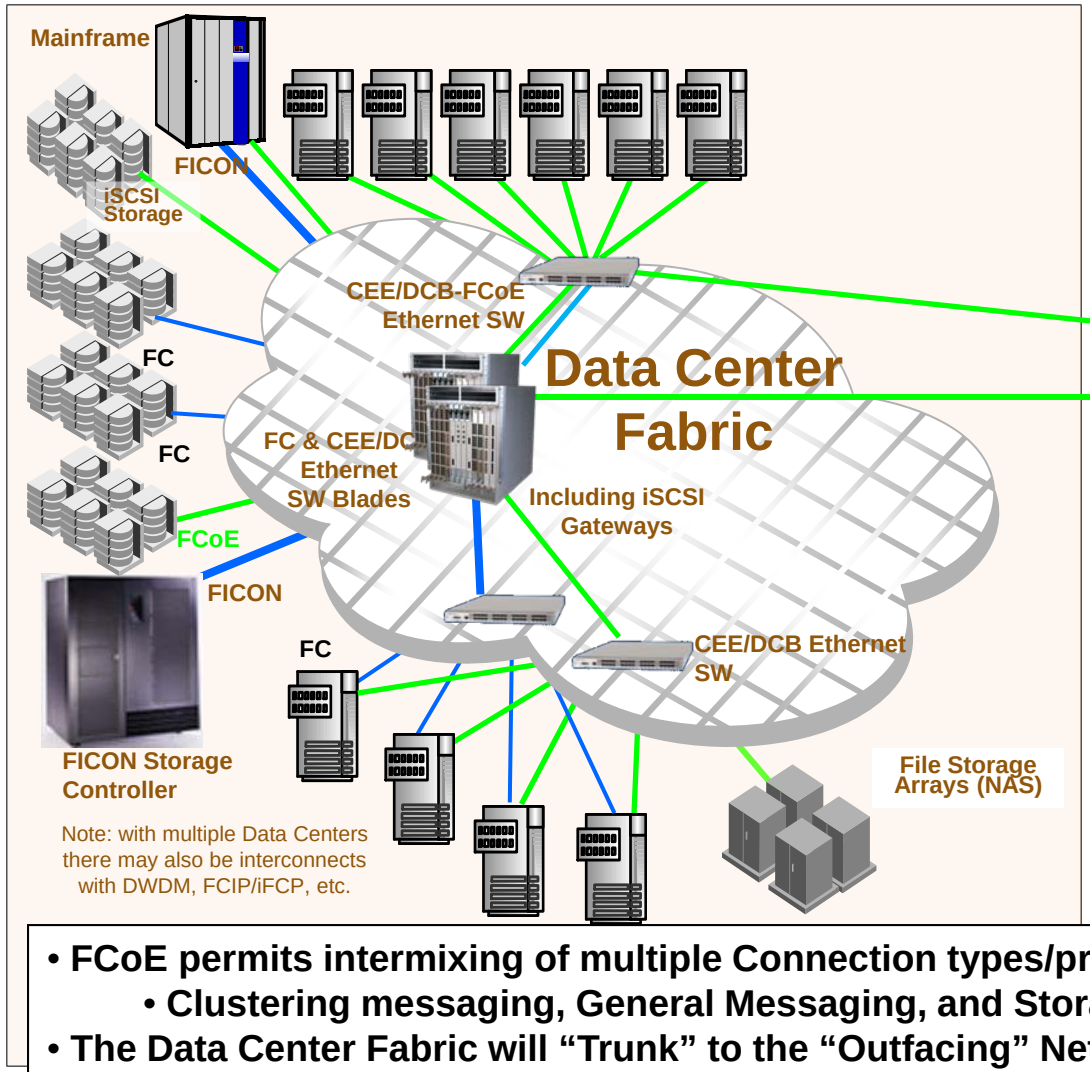

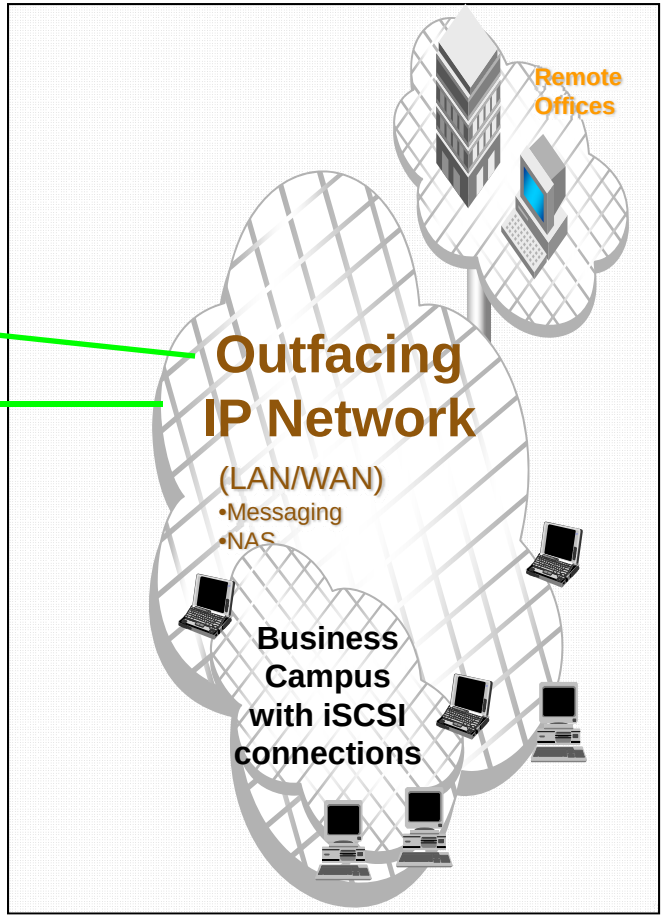

- FCoE permits intermixing of multiple Connection types/protocols
  - Clustering messaging, General Messaging, and Storage
- The Data Center Fabric will "Trunk" to the "Outfacing" Network
- But many Customers may want keep a mixed environment on-going

FC Link
FICON Link
Ethernet Link

Architecture

FC Encapsulation Into Ethernet Frames
(2 FCoE Related Packet types)

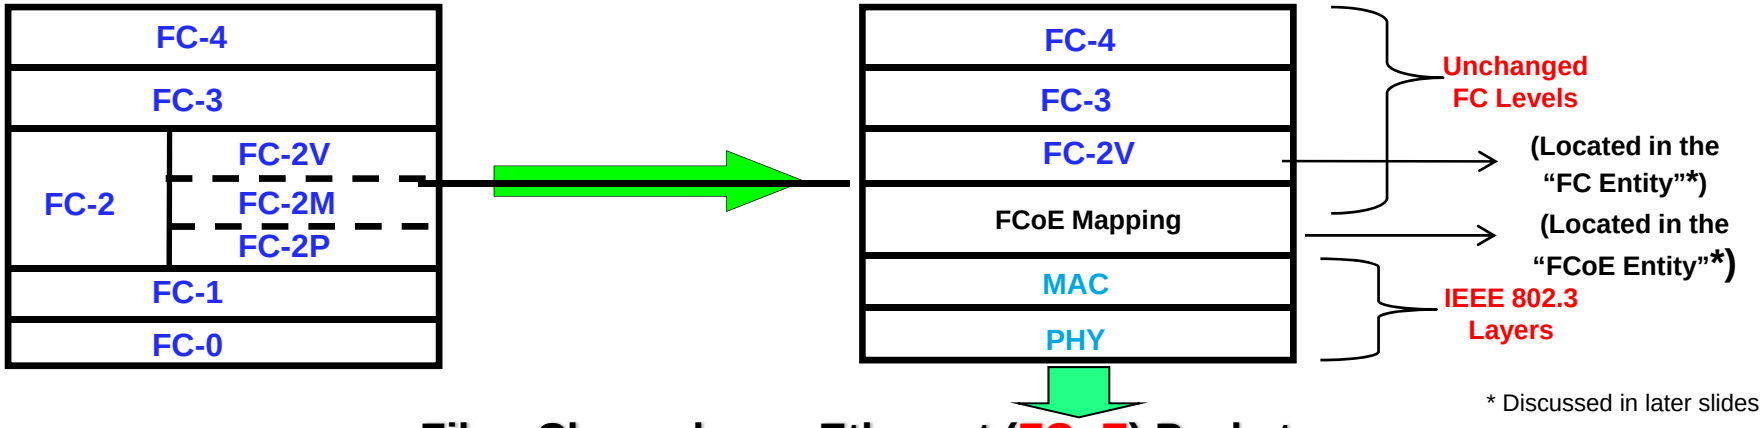

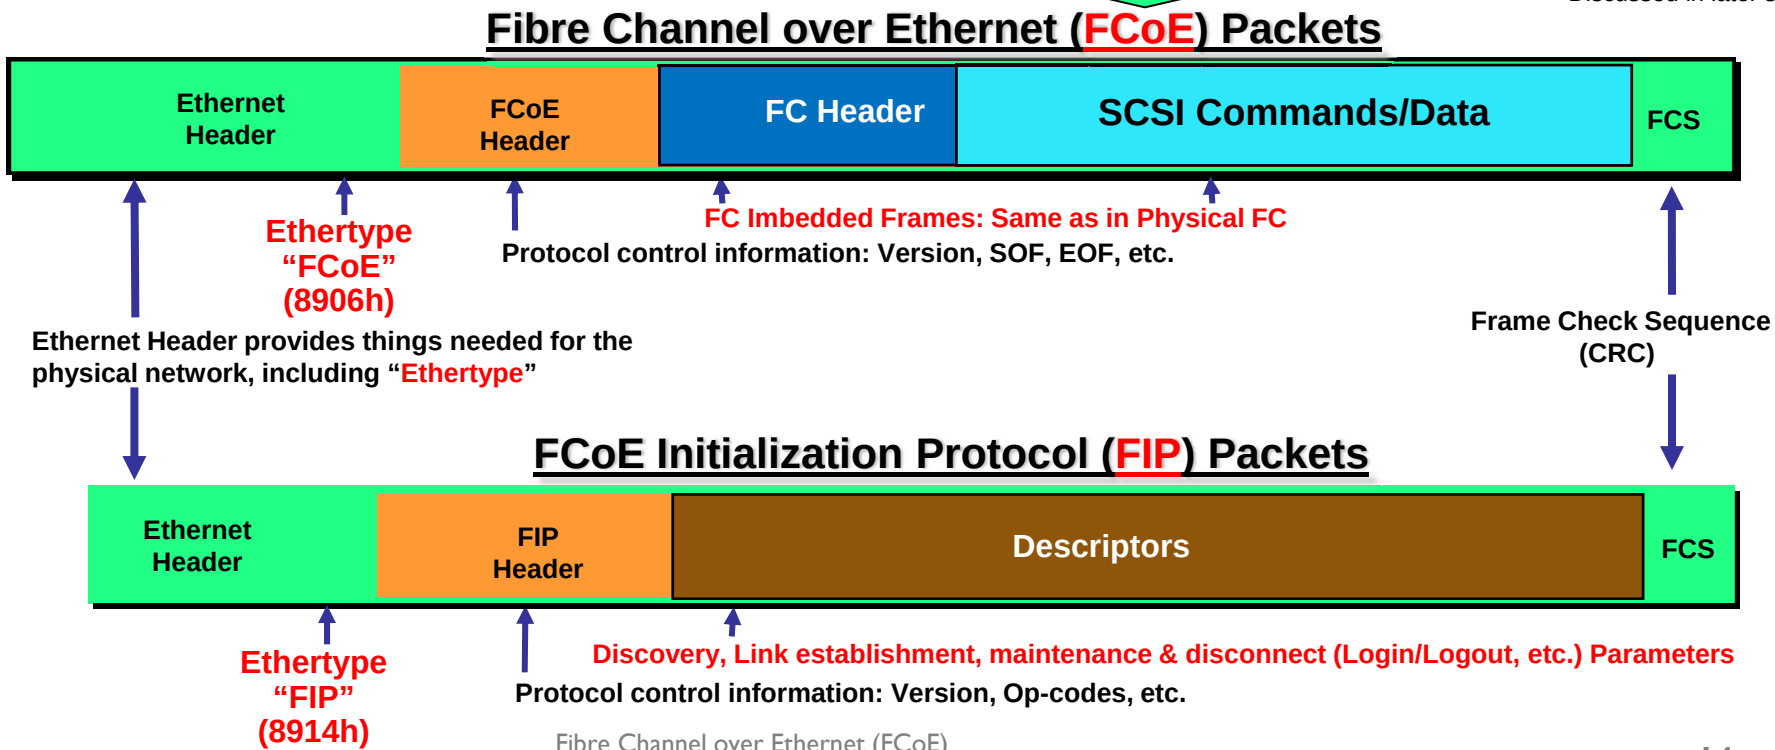

FIP Protocol and FCoE Protocol

- **Discovery Phase**
  - FCFs Discover each other, & VLANs (if any) then form a Fabric
  - ENodes Discover VLANs (if any) & then ENodes & FCFs Discover
    - Potential VN_Port ← → VF_Port pairing
    - Capabilities of Potential pairing
- **Login Phase**
  - ENodes chose among discovered FCFs' Ports
  - Creates association between ENode Ports and FCF Ports
    - VN_Port ← → VF_Port Logical FC Link
  - Two allowed methods for the ENode MAC Addressing
    - Fabric Provided MAC Addresses (FPMA) – Assigned by the Fabric
    - Server Provided MAC Addresses (SPMA) – Assigned by the Server
    - FPMA or SPMA method is chosen by FCF (based on each side's capabilities)
  - Uses: FLOGI, FLOGI ACC, LOGO, ELS, …
- **End-to-End path control & Data Transfer Phase**
  - PLOGI/PRLI
  - All other FC protocol frames (FC4 ULPs. etc.)

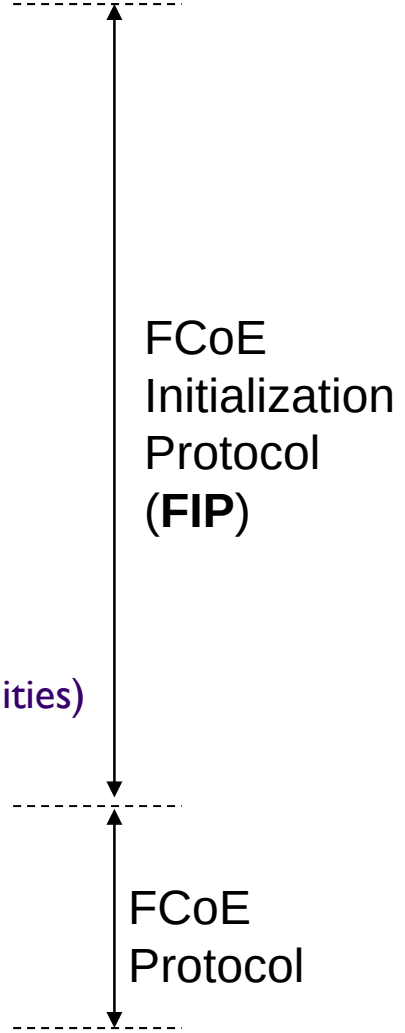

FC's Encapsulation in Ethernet (FCoE)

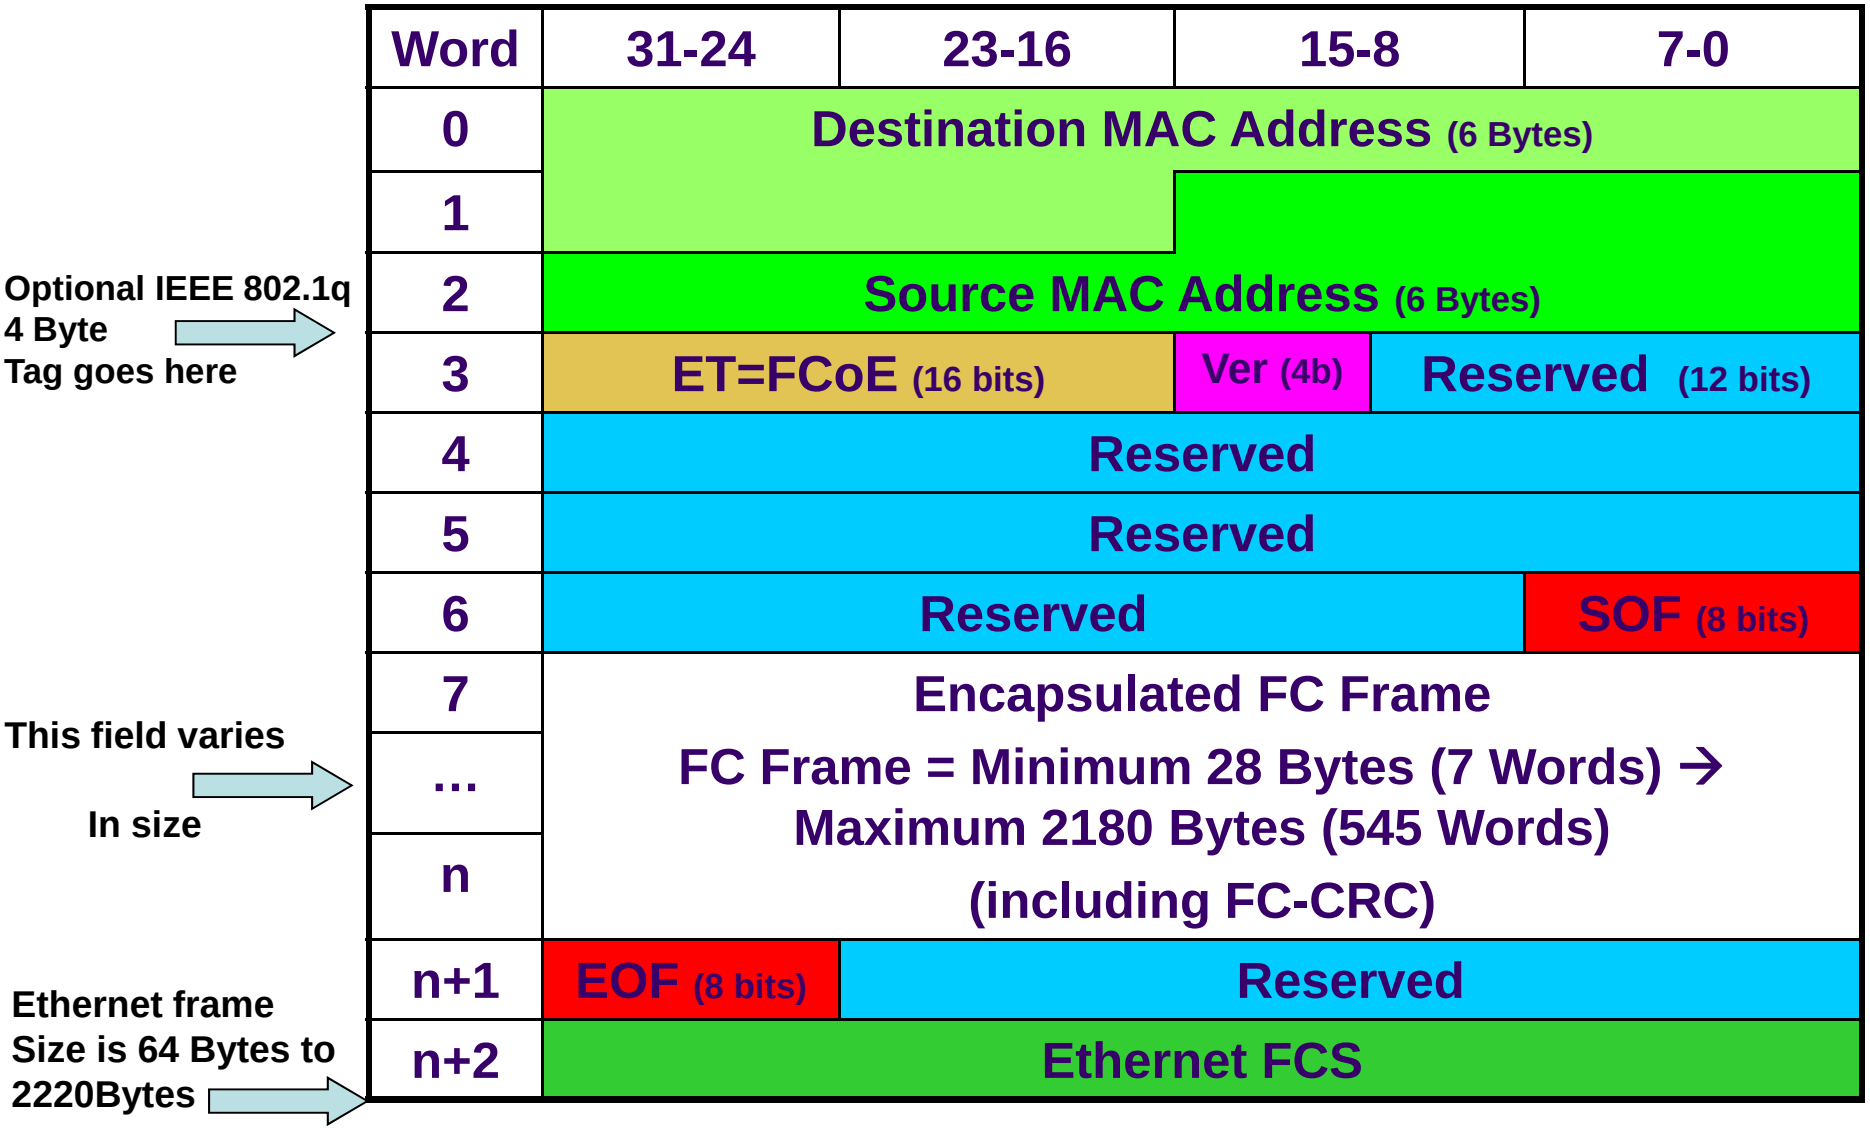

Word	31-24	23-16	15-8	7-0
0	Destination MAC Address (6 Bytes)			
1				
2	Source MAC Address (6 Bytes)			
3	ET=FCoE (16 bits)		Ver (4b)	Reserved (12 bits)
4	Reserved			
5	Reserved			
6	Reserved			SOF (8 bits)
7	Encapsulated FC Frame FC Frame = Minimum 28 Bytes (7 Words) → Maximum 2180 Bytes (545 Words) (including FC-CRC)			
…				
n				
n+1	EOF (8 bits)	Reserved		
n+2	Ethernet FCS			

FIP Operation Format

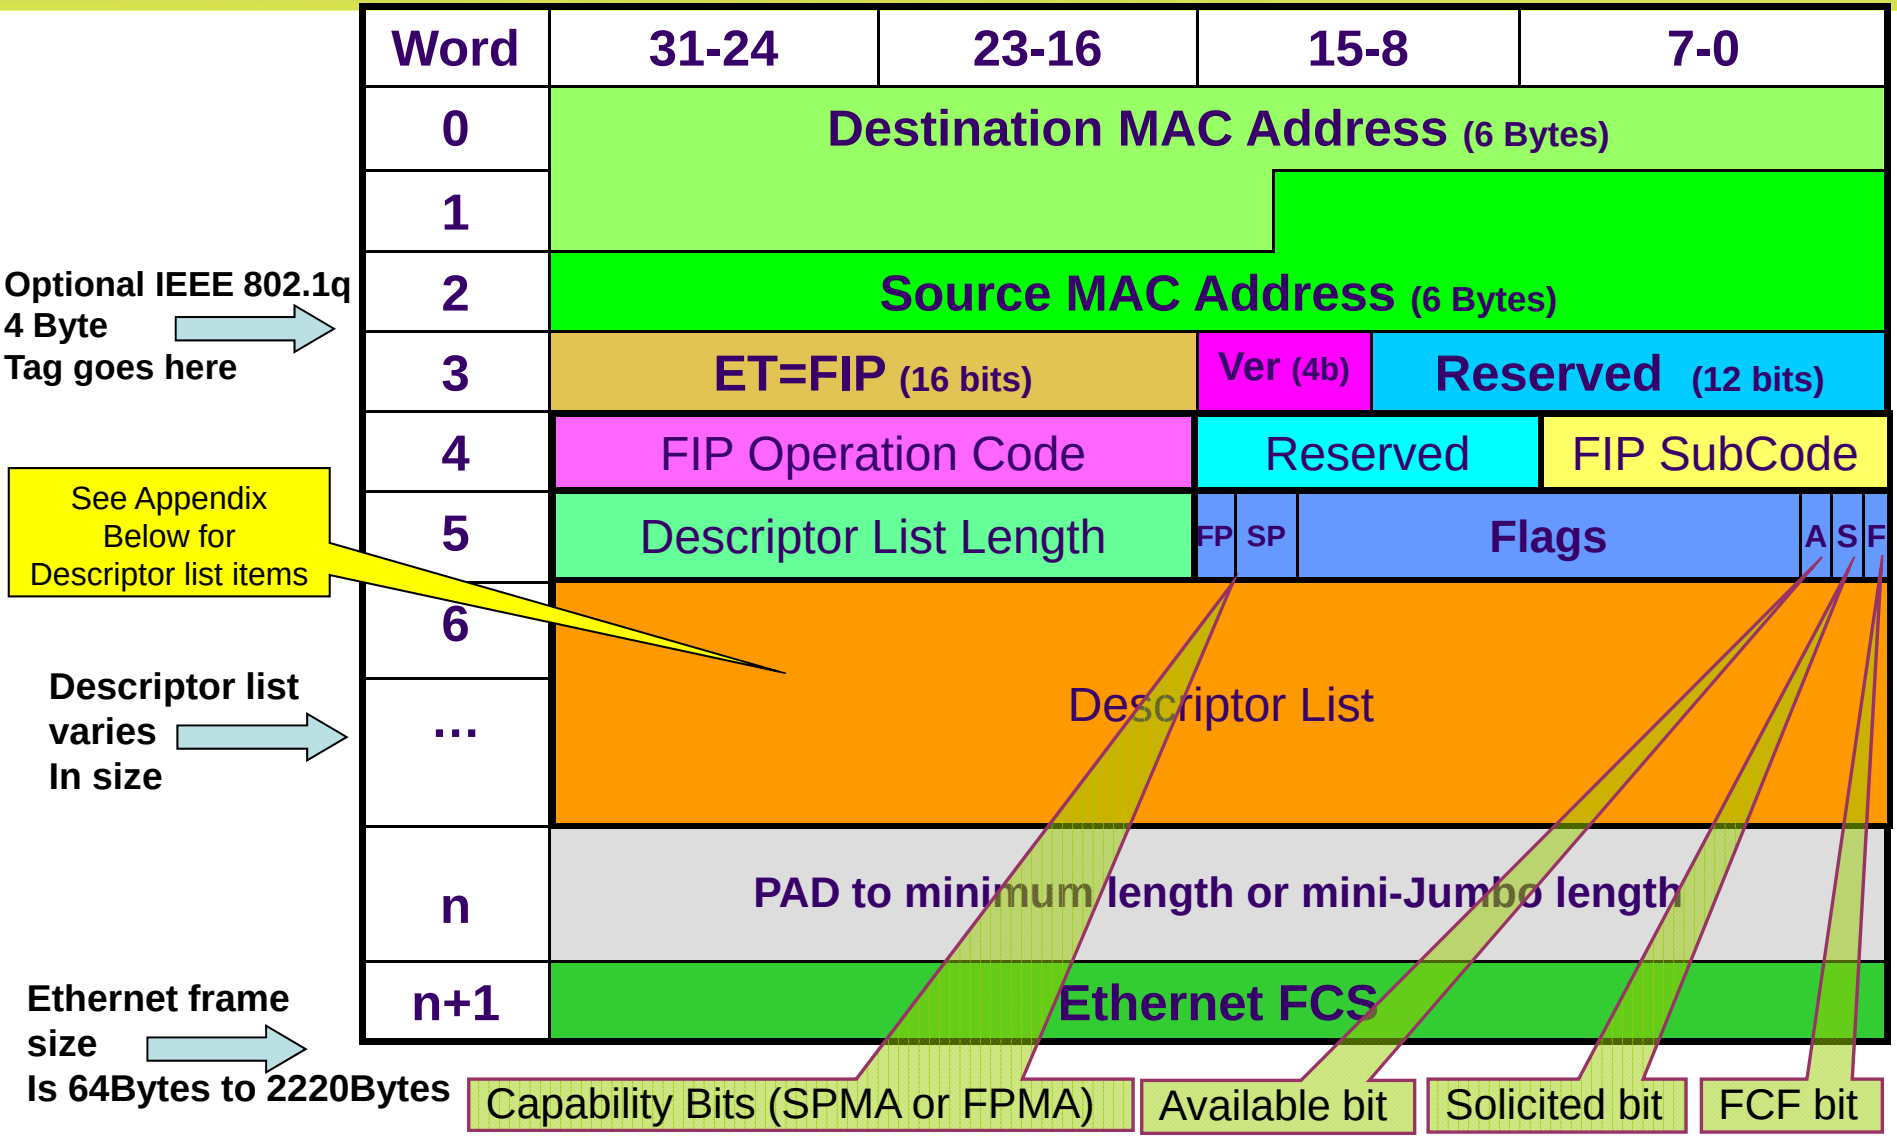

FCF Model

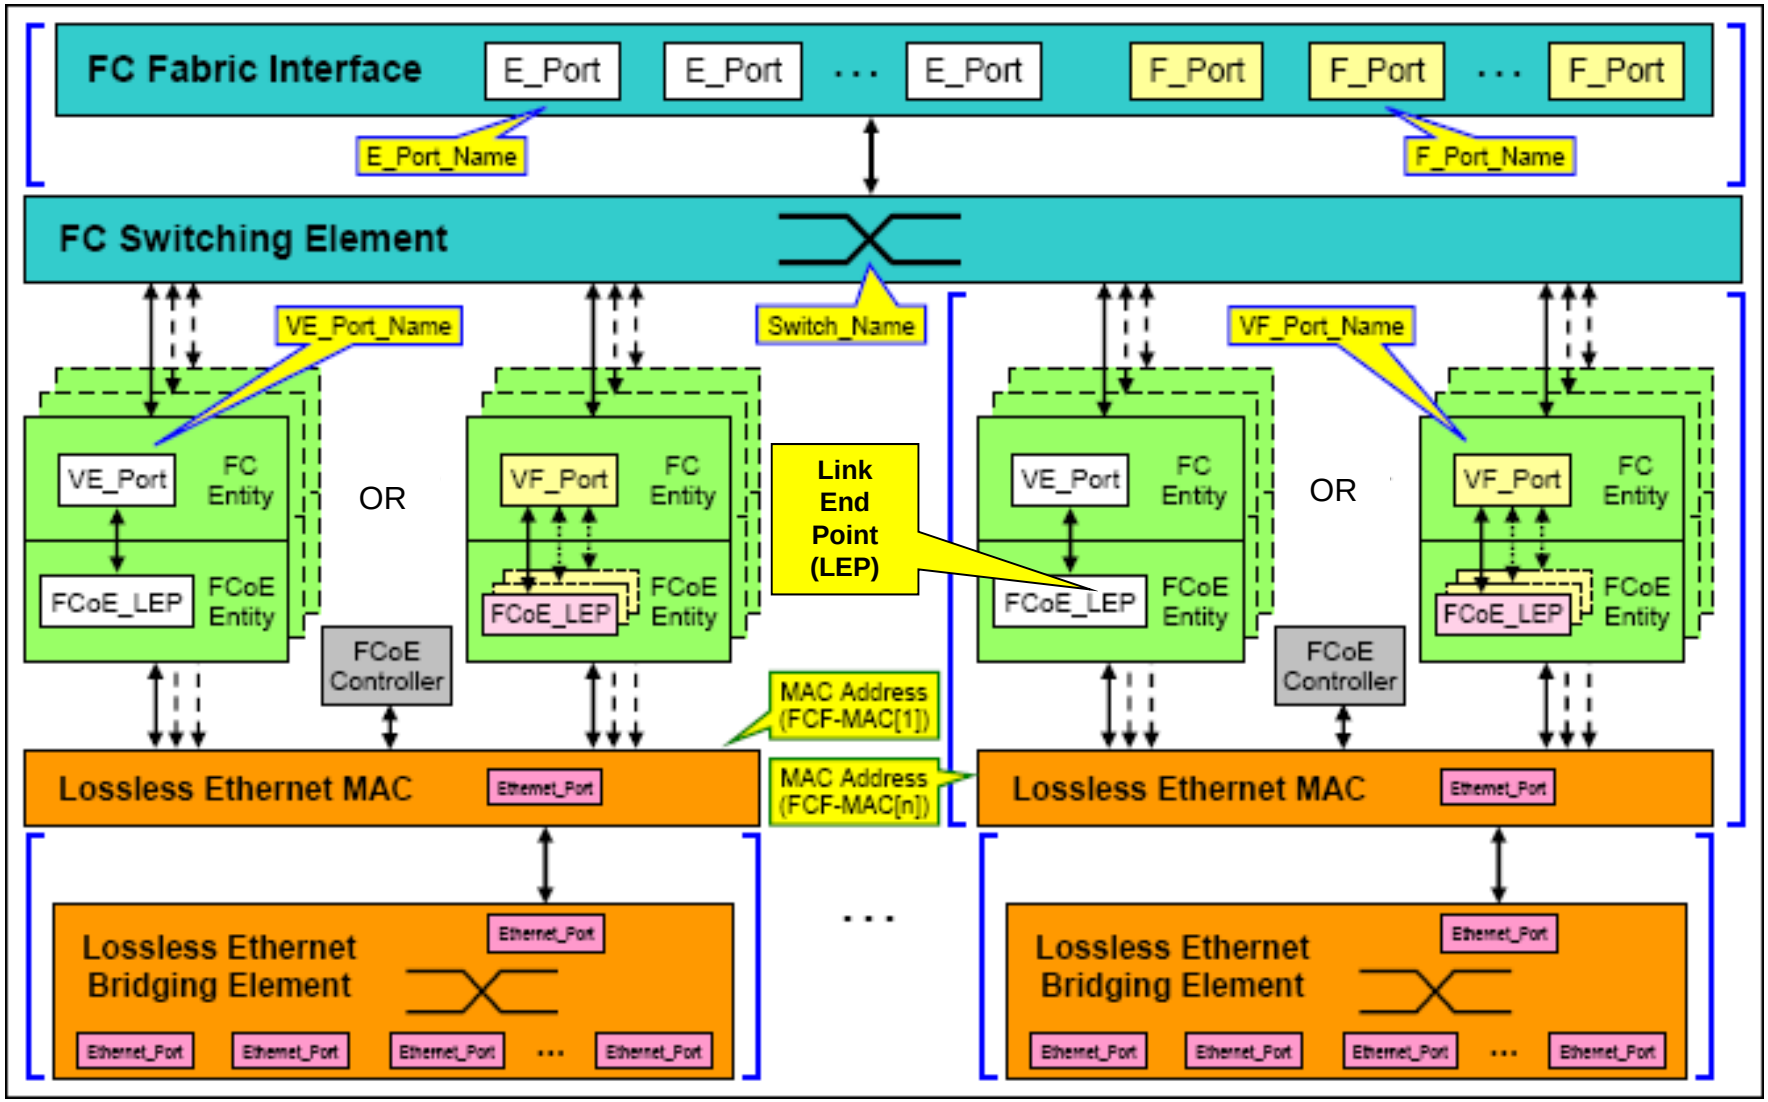

ENode (HBA) Model

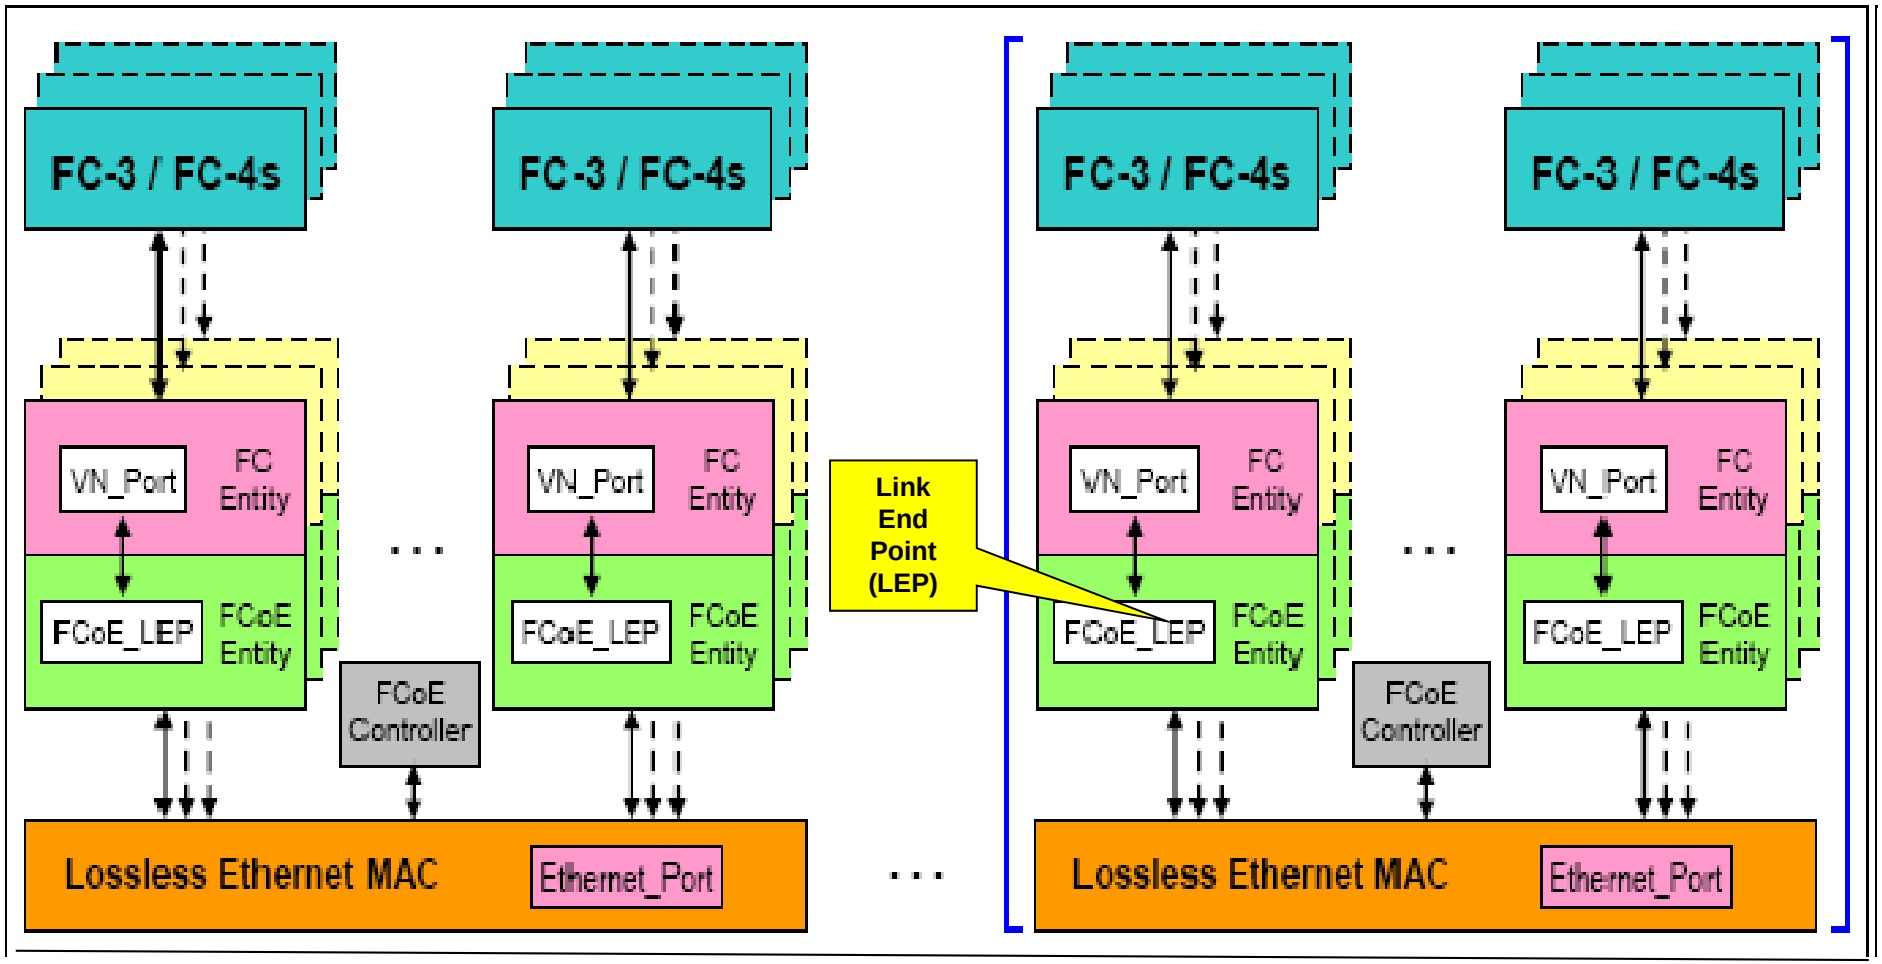

Each ENode (HBA/CNA) may have multiple Physical Ethernet Ports
Each Physical Port may have multiple Logical VN_Ports

Model of the ENODE with Multiple Logical FC interfaces

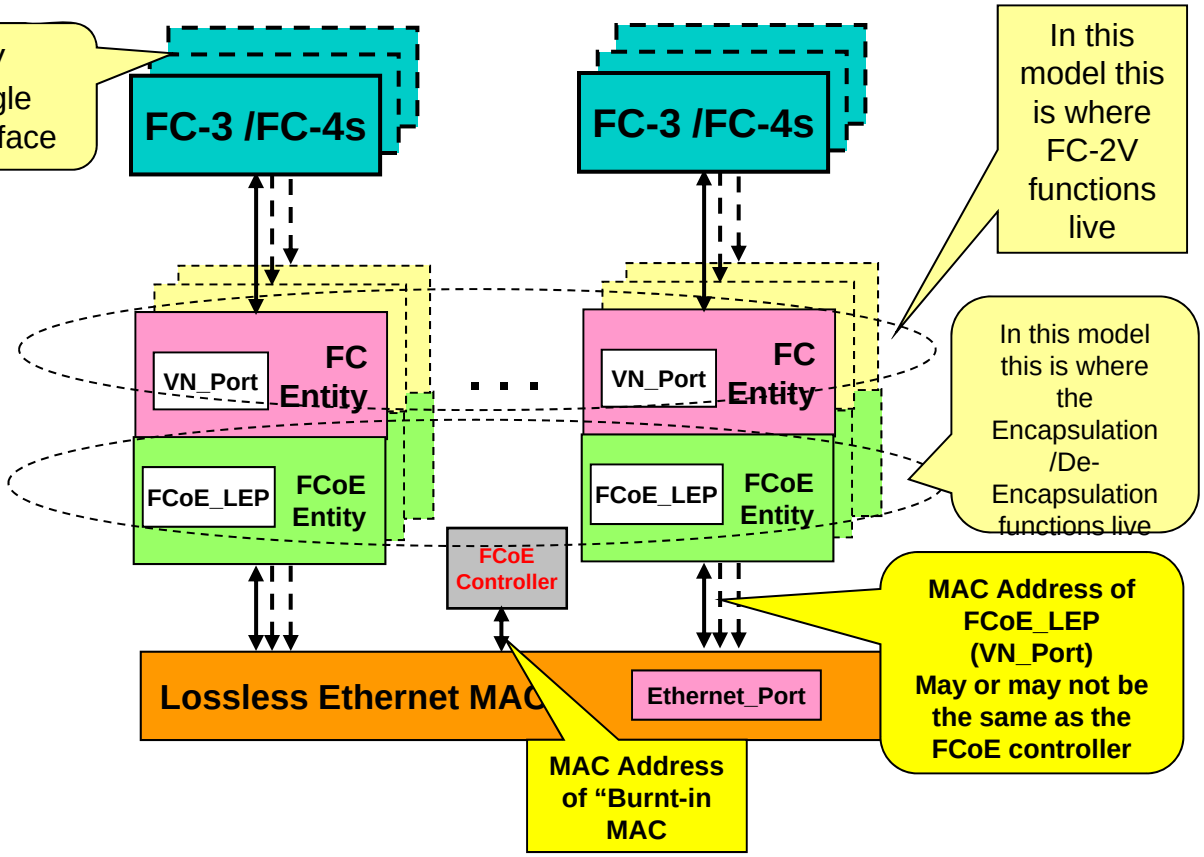

- For each logical N_Port (VN_Port) there is one FLOGI and perhaps 100's of FDISC
- Each VN_Port is seen by the Host as a separate (logical) FC Login/FDISK connection
- The number of (logical) FC connections is implementation dependent

- Only one MAC Address is required for the FCoE Controller and the VN_Ports on a single physical MAC (aka Server Provided MAC Address – SPMA)
- FCF may chose to specify new MAC addresses for each VN_Port (aka Fabric Provided MAC Address – FPMA)

Multiple Logical FC connections via a single Ethernet MAC

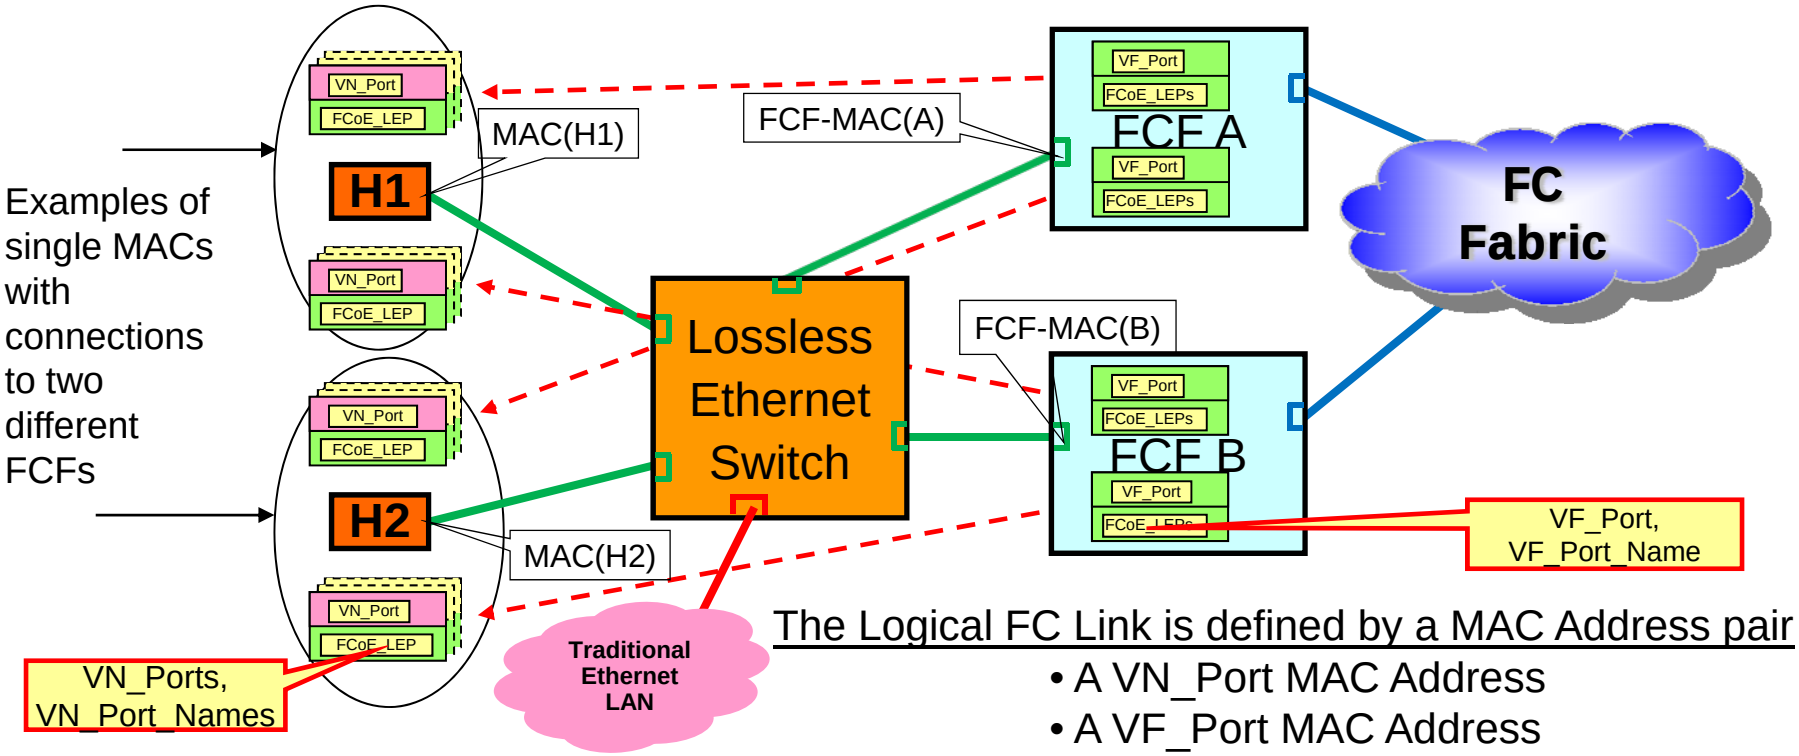

The Logical FC Link is defined by a MAC Address pair

- A VN_Port MAC Address
- A VF_Port MAC Address

For a logical FC link the FCoE Frames are always sent to and received from a specific FCF's MAC Address

- Therefore, pathing to and from the FC driver is always defined by the MAC Address of the partner FCF's VF_Port

Note: VF ports get created off of the FLOGI from the VN port and multiple VF ports can sit behind a single FCF physical interface with a single MAC address

Functions of an FCoE Converged Network Adapter (CNA) in the Initiator

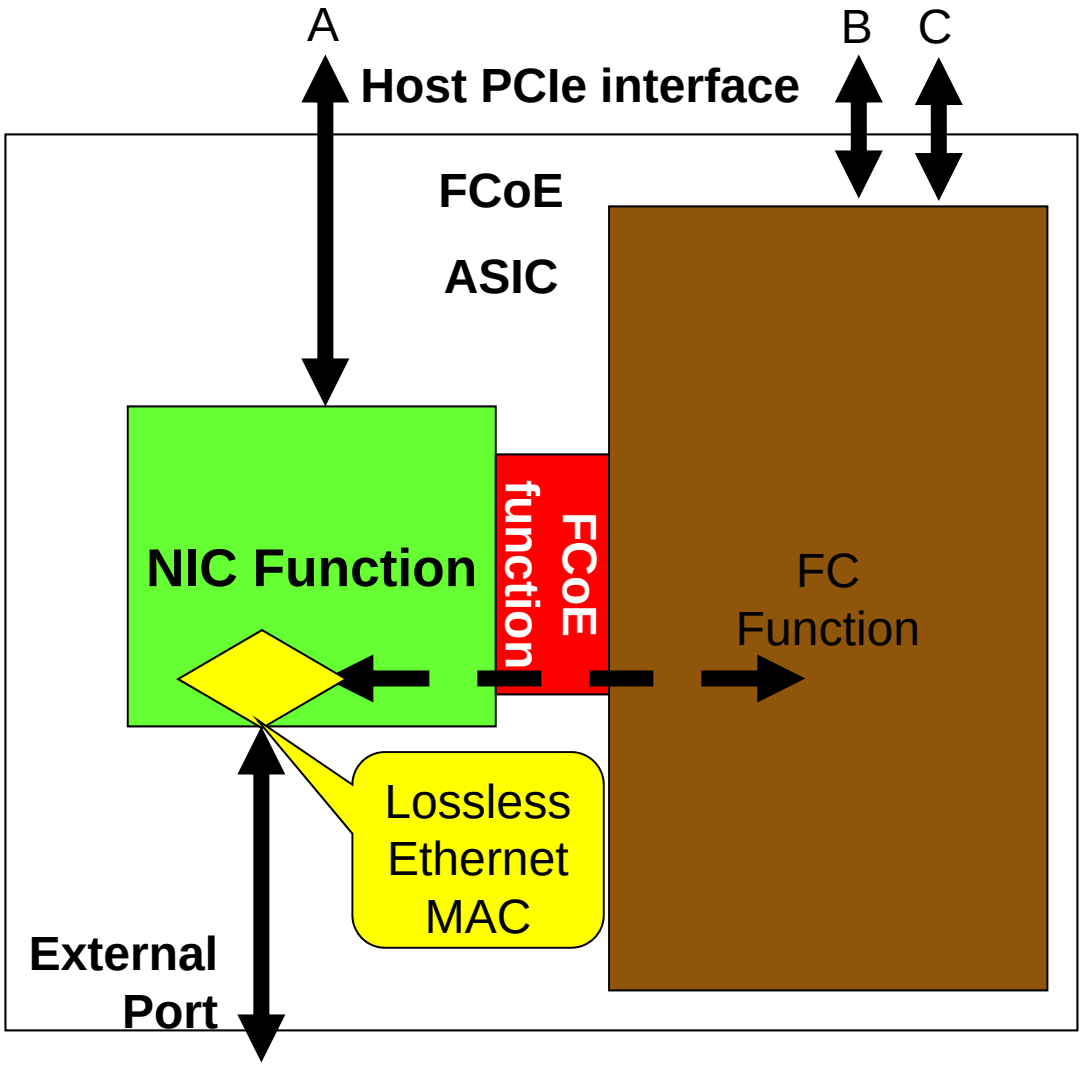

- Has a Normal NIC interface (A) to the Host
- Has one or more Normal FC interfaces (B,C) to the Host
- FCoE functions not seen by the Host
- FCoE functions perform the Encapsulation and De-encapsulation
- The FCoE function Instantiates Logical FC N_Ports, called VN_Ports

CNA with Multiple Logical Interfaces

- The FCoE controller will perform the FIP functions and will instantiate new VN_Ports as FCoE Link End Points (LEP)
  - With the same MAC address as the FCoE Controller (SPMA)
  - Or with a new MAC address specified by the FCF (FPMA)

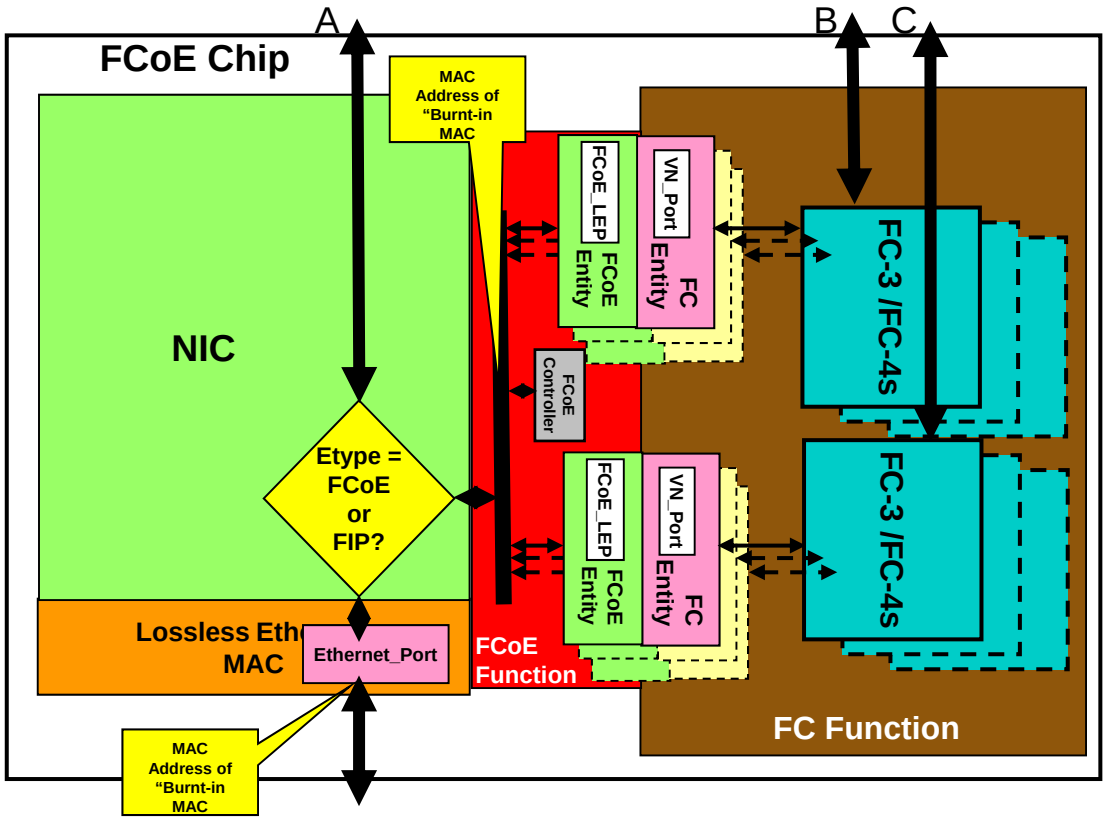

- MAY have one "burnt-in MAC address for both IP and FCoE/FIP packets

Or

- MAY have different "Burnt-in" MACs for IP and FCoE/FIP packets
  - Used to separate HW based FCoE from other Ethernet Traffic
  - Most NICs come with several "Burnt-in" MAC Addresses

Discovery and Link Instantiation (FIP -- FCoE Initiation Protocol)

ENode to FCF Logical Links

An ENode must Locate FCFs with a Discovery protocol and then establish a Logical FC Link with an FCF (VN_Port ← → VF_Port) before a Normal FCP frame flow

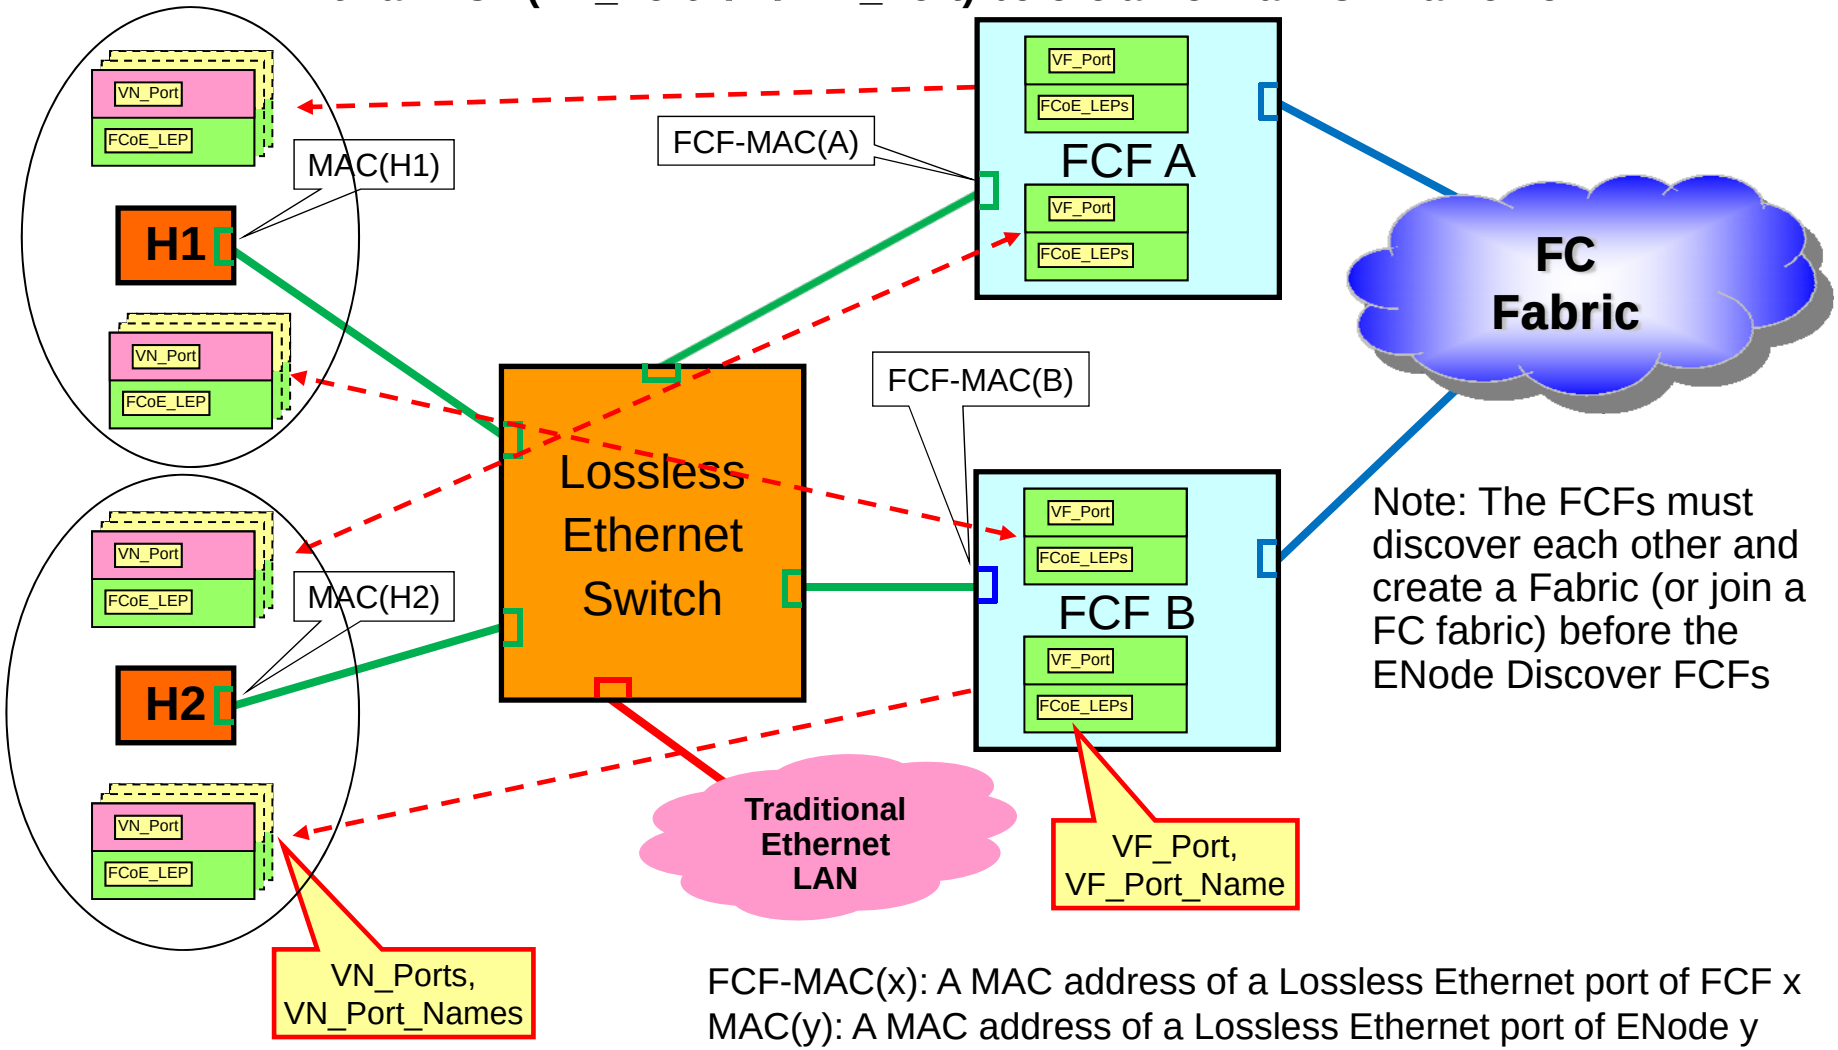

Note: The FCFs must discover each other and create a Fabric (or join a FC fabric) before the ENode Discover FCFs

FCF-MAC(x): A MAC address of a Lossless Ethernet port of FCF x
MAC(y): A MAC address of a Lossless Ethernet port of ENode y

Initial Login Flow Ladder (2 FIP Phases)

(within any specific VLAN)

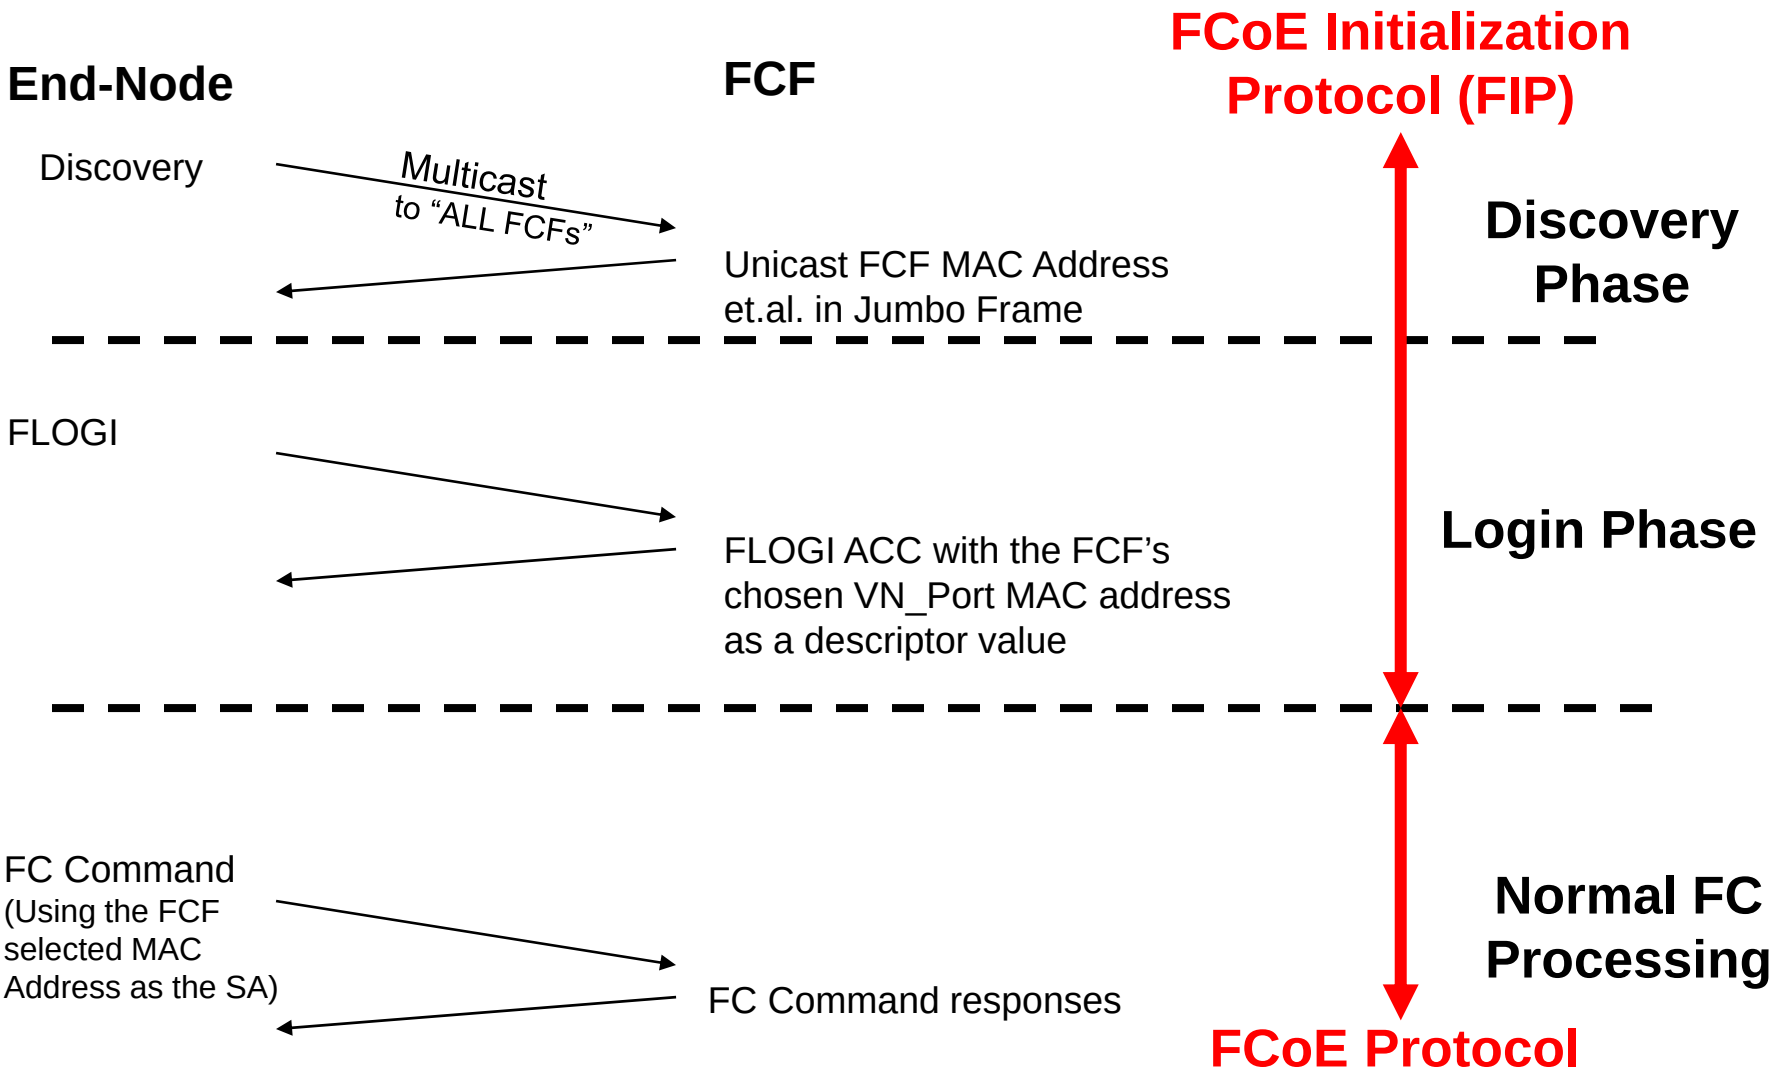

See Appendix for more details on Discovery Consideration and Actions, including discovering VLANs

Subsequent Login Flow Ladder (1 FIP Phase)

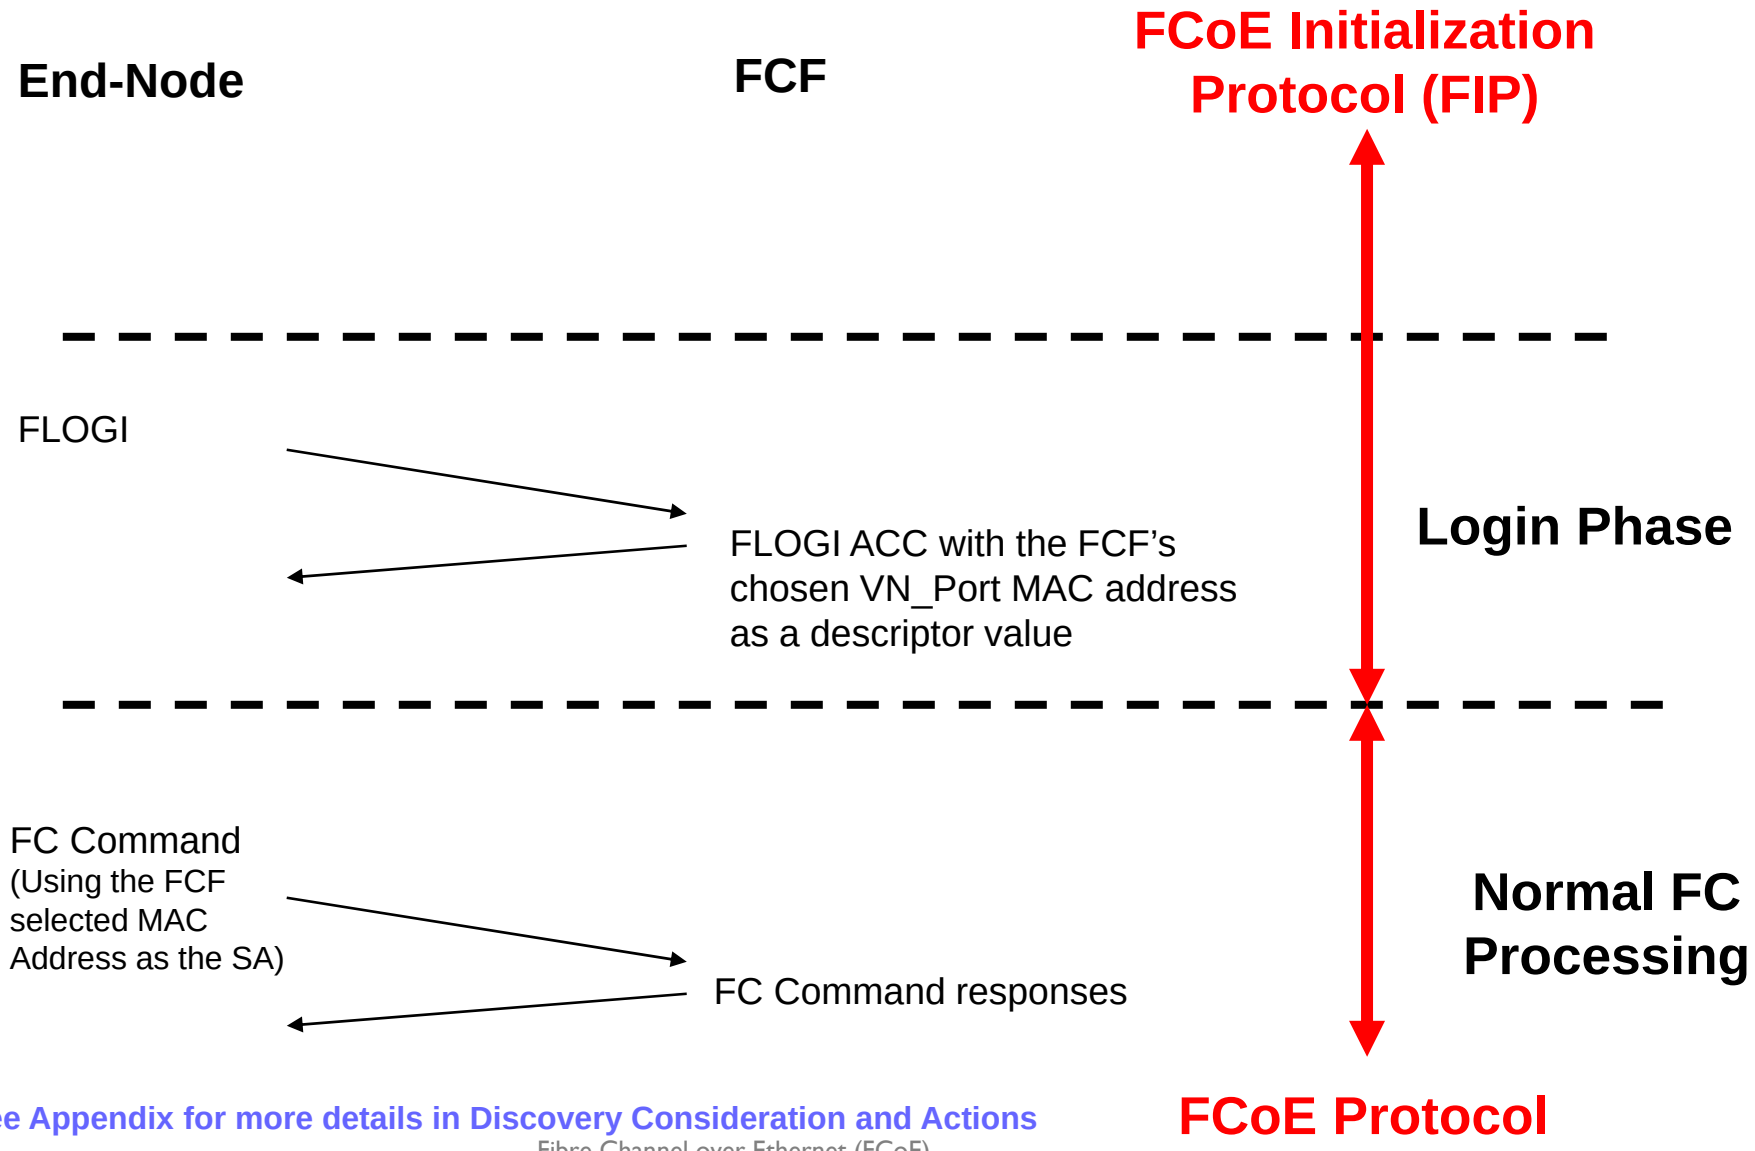

See Appendix for more details in Discovery Consideration and Actions

Topologies

A Simple CEE/FCoE Fabric

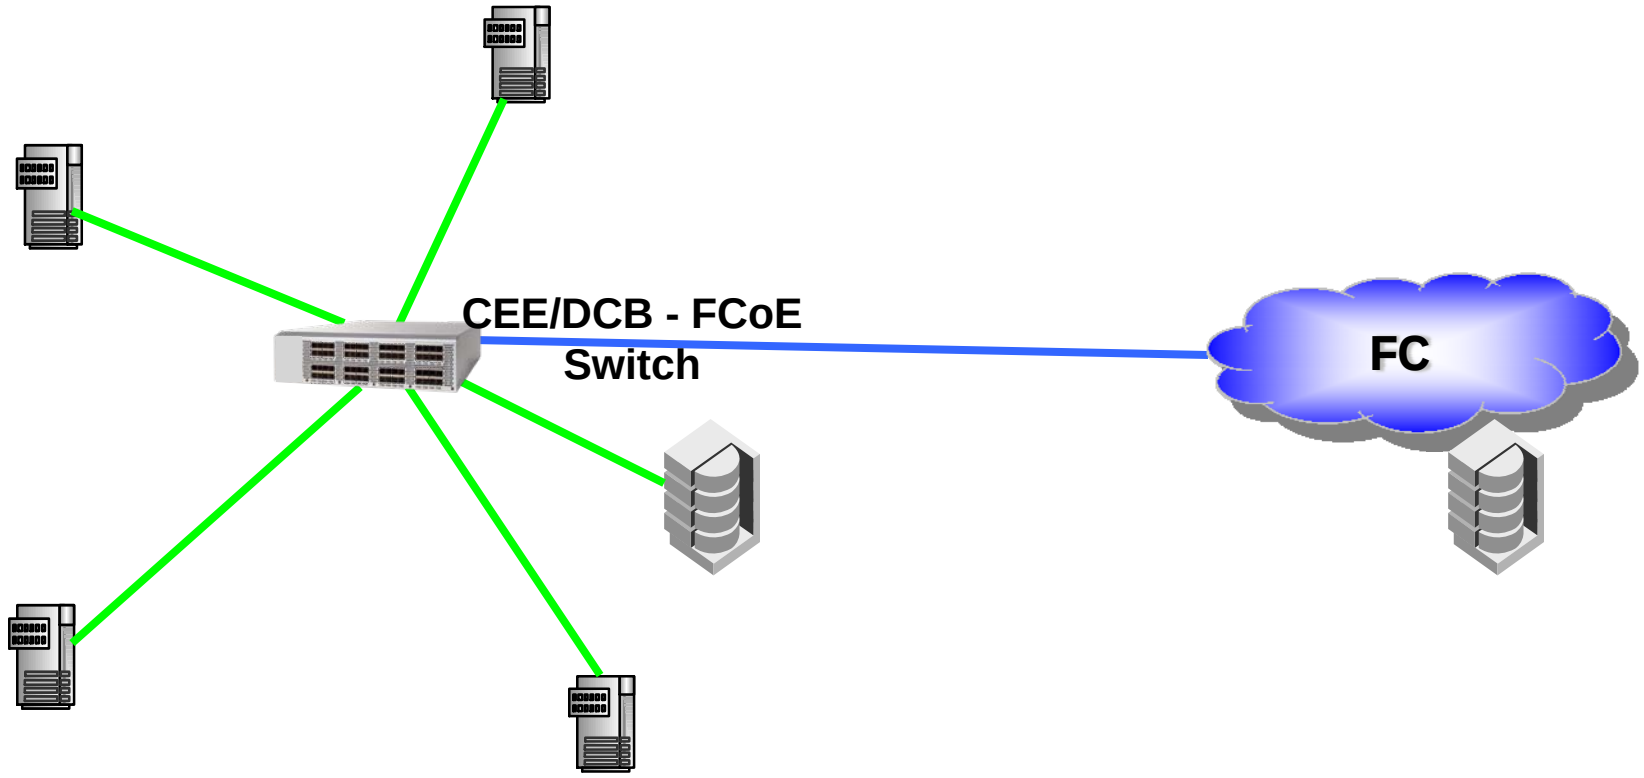

An CEE/DCB - FCoE (Integrated) Switch may connect to a Traditional FC switch/fabric

Via the FC E-Port

Ethernet
FC

Multiple Topologies using FCoE Switches (FCFs)

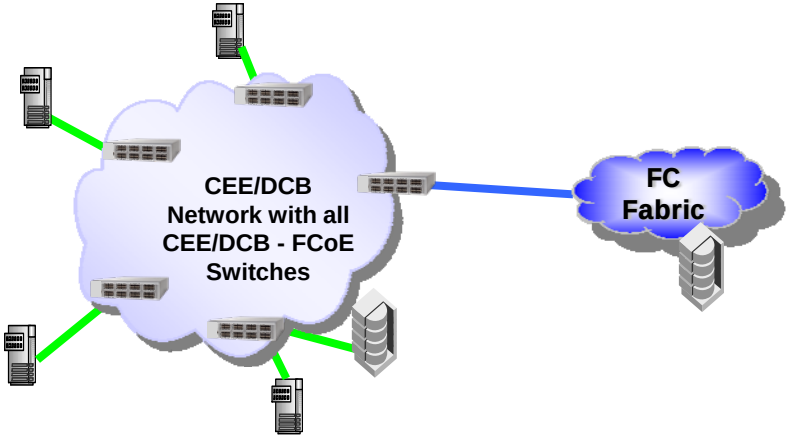

A Lossless Ethernet Network can be made up of all CEE/DCB - FCoE (Integrated) Switches

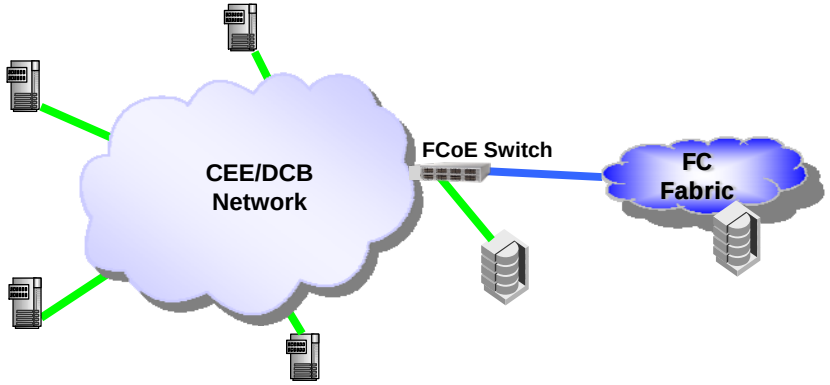

Lossless Ethernet (CEE/DCB) switches configured into a Lossless Ethernet (CEE/DCB) Network can Front the FCoE Switch

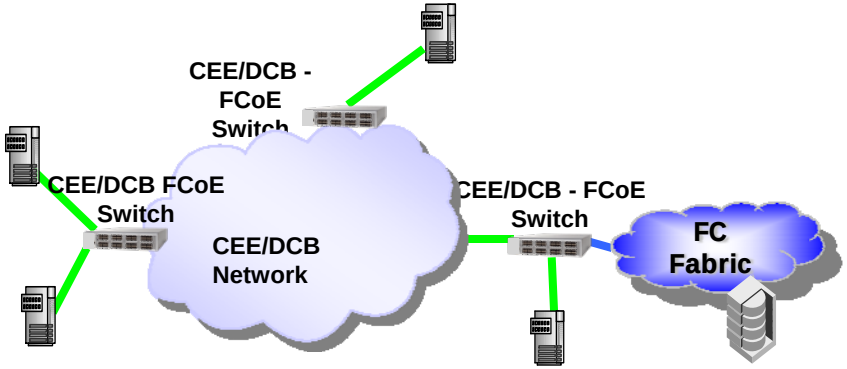

CEE/DCB - FCoE (Integrated) Switches deployed at the edges of the Lossless Ethernet CEE/DCB Network
CEE/DCB - FCoE Switches connected via VE_Ports and a Lossless Ethernet CEE/DCB Network

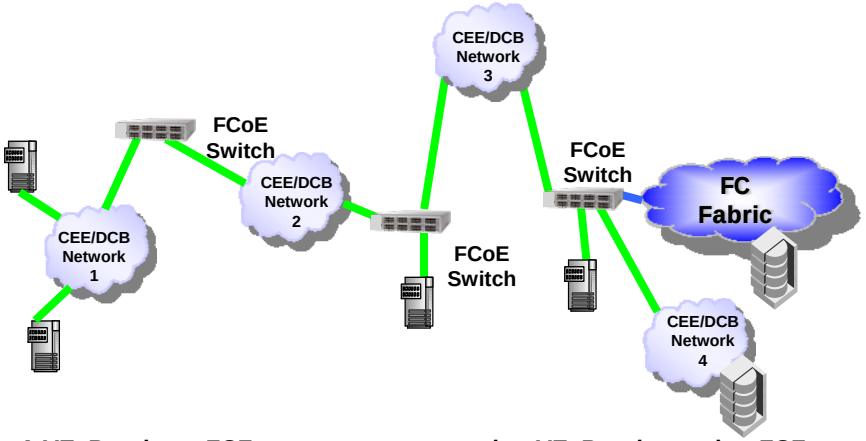

A VE_Port in an FCF may connect to another VE_Port in another FCF
And an FCF FC E_Port may connect to an FC switch E_Port in the FC Fabric

Scenarios

Scenario 1: FCoE & IP Flows

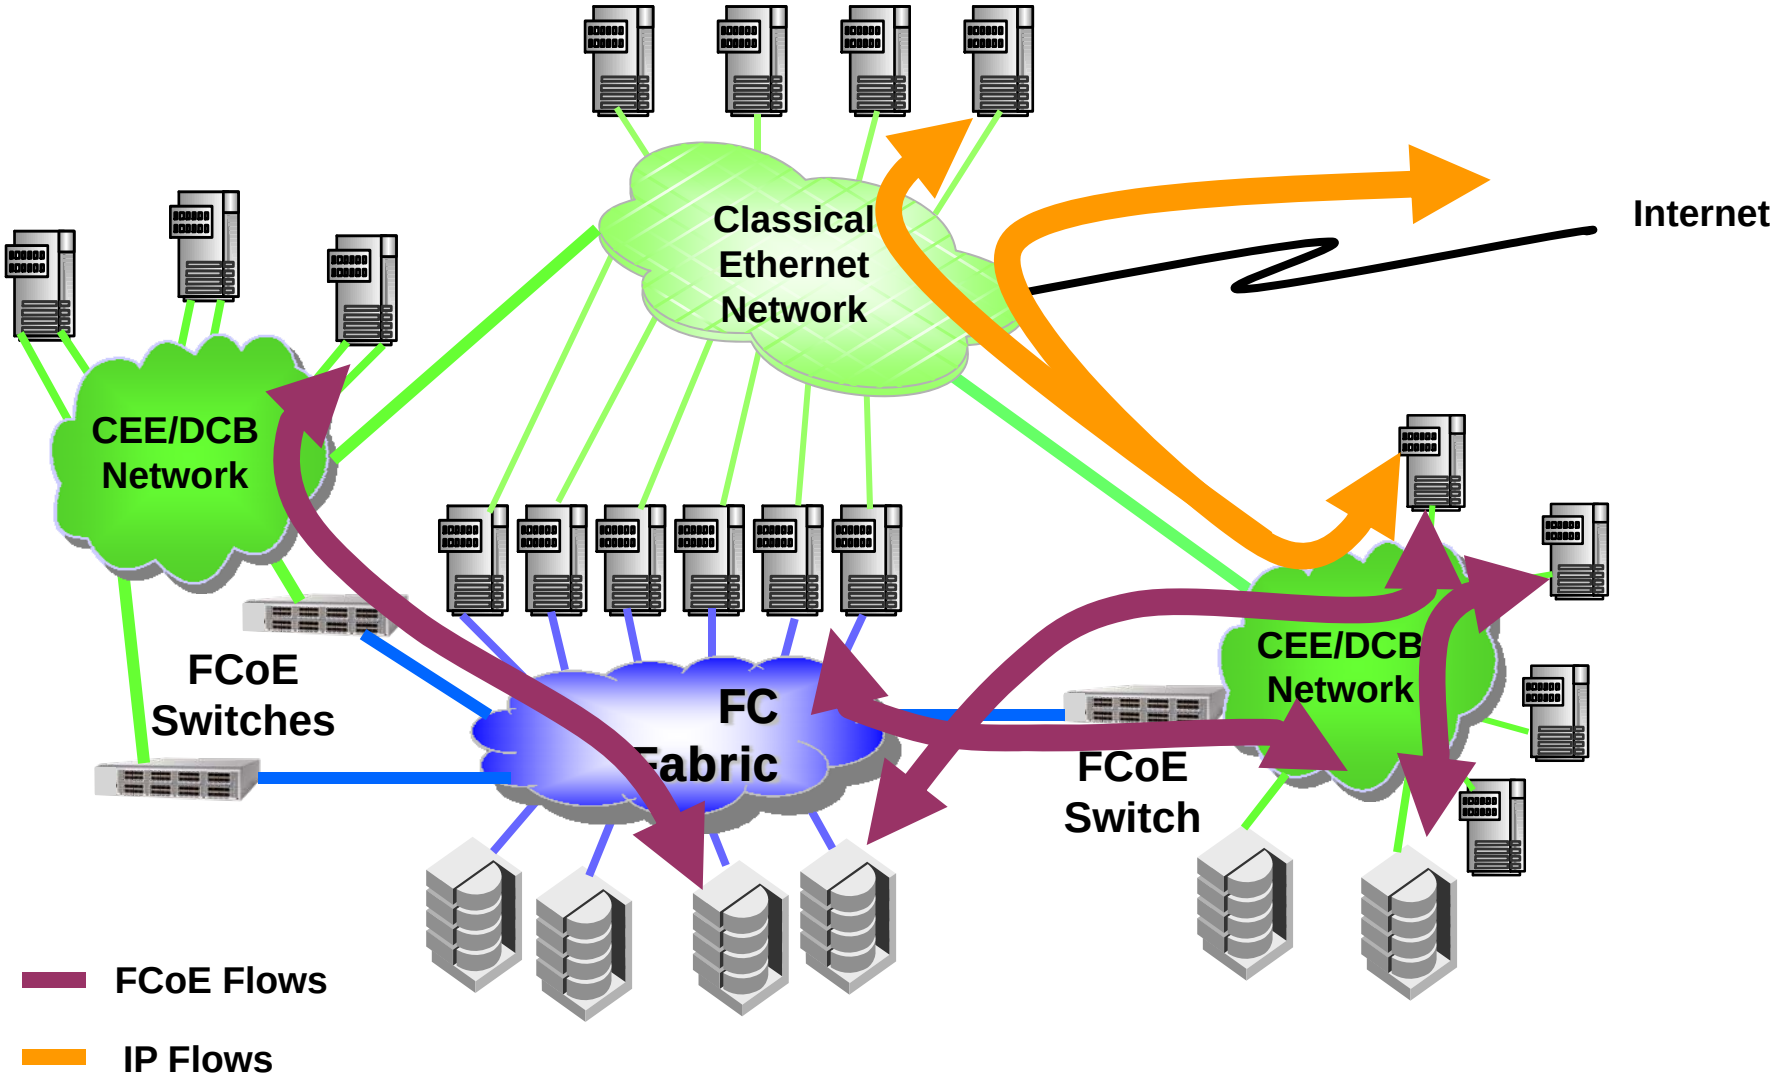

Scenario 2: FCoE Right & Wrong

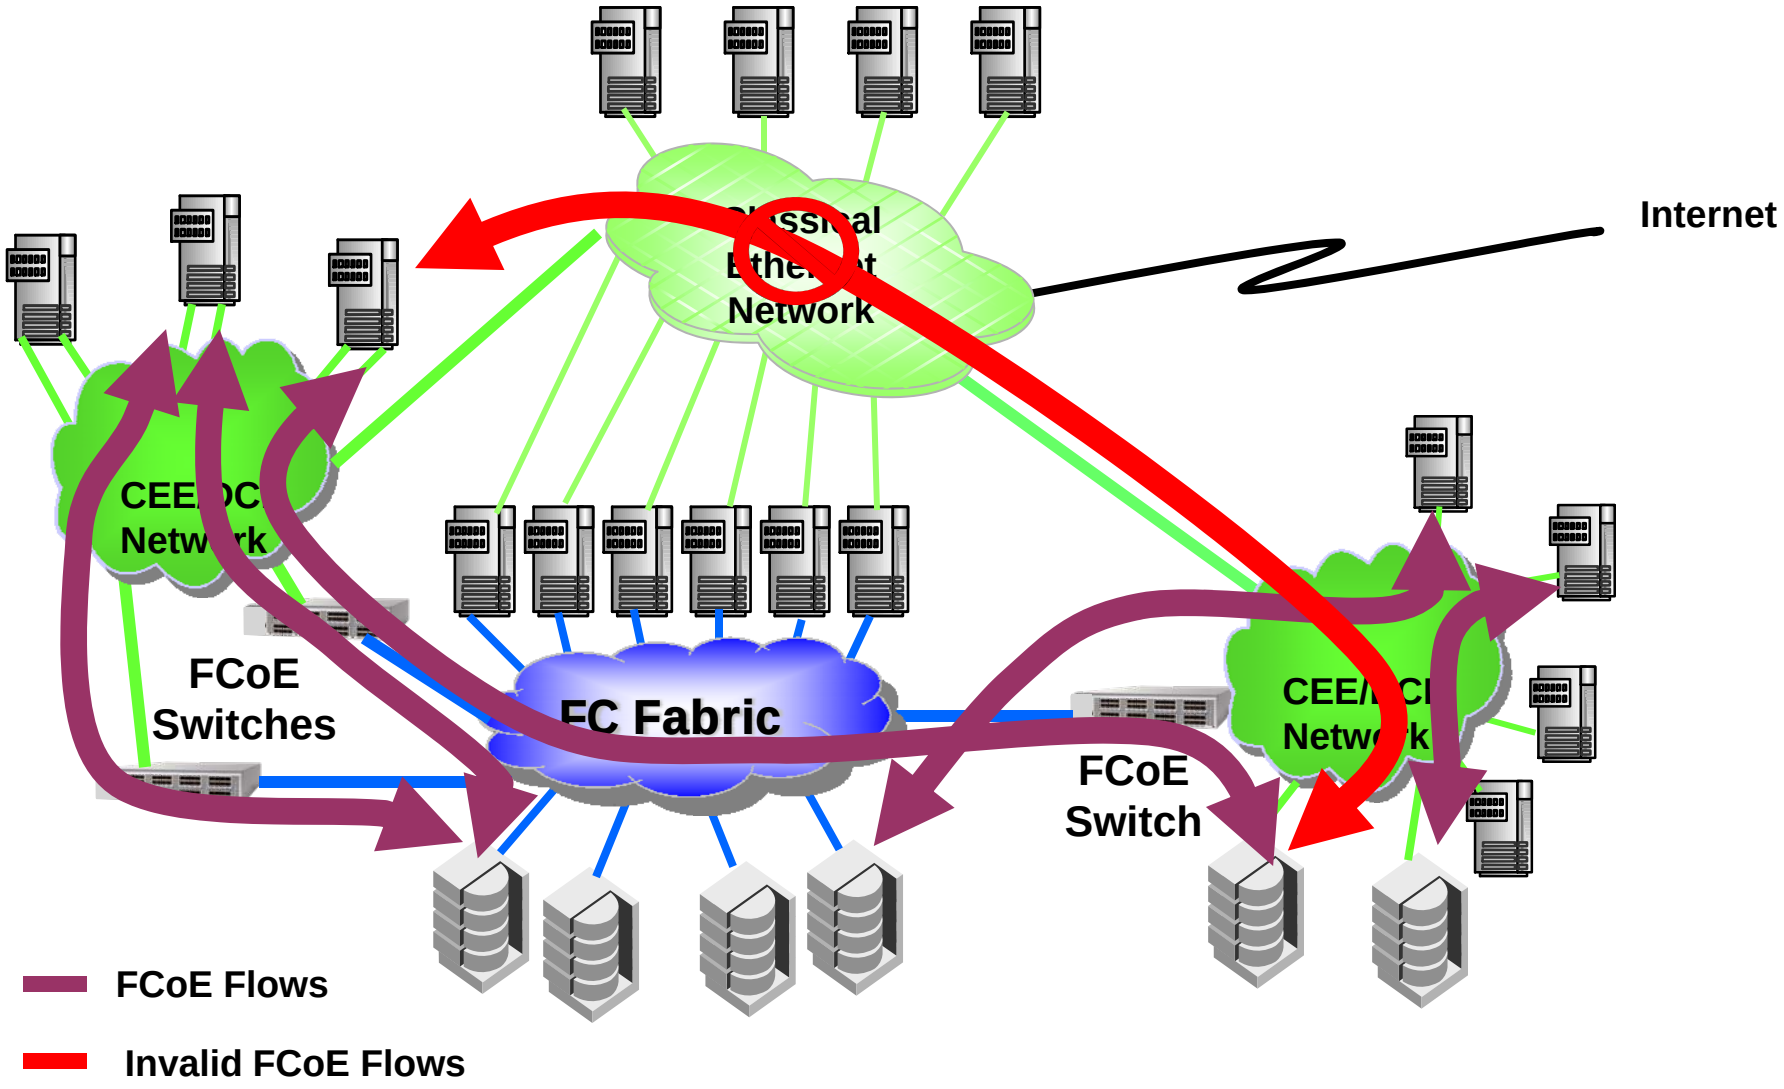

Summary

FCoE Summary

- T11.3's FC-BB-5 Ad-Hoc Working Group completed the specification in June 2009. The document is now at INCITS where it is being published.
  - Volume product ramp expected in 2010
- FCoE is a simple, efficient mechanism for encapsulating Fibre Channel in Ethernet frames on a New Ethernet type Network
  - Not a traditional Ethernet Interface or fabric
  - A New Network – the Converged Enhanced Ethernet (CEE) Network
  - CEE (also called DCB) being defined in the IEEE 802.1 standards working group
  - FC protocols frames will just be inserted into these Ethernet frames
- An evolutionary deployment model was designed into FCoE
  - Specification permits the installation to evolve from FC to FCoE
    - Any Fabric mix of FC and FCoE is possible
    - May only need FCoE at the Server Edge with a Converged Interface
    - But a total FCoE SAN is also possible (using FCFs)
    - Value in reduced Server Edge Cables, Adapters, Power, and Cooling
  - All FCoE devices should interoperate with Real FC devices
- FCoE is made for a Data Center Fabric – Not applicable for the Outfacing Network

Q&A / Feedback

- **Please send any questions or comments on this presentation to SNIA: tracknetworking@snia.org**

Many thanks to the following individuals for their contributions to this tutorial.

SNIA Education Committee

Claudio DeSanti	Howard Goldstein	Walter Dey
Robert Snively	Suresh Vobbilisetty	Silvano Gai
Joe Pelissier	John Hufferd	

- **For additional information refer to http://www.t11.org/fcoe**

Thank You!

Education

Appendix

- Additional Info on FCoE Fabrics
- FCoE Relation to ISO Layers
- Flows
- Additional Topologies
- FSPF and STP
- FIP Considerations and Actions
- FIP Descriptors
- Pause vs. BB_Credit

FCoE Fabrics (part 3)

- FCoE is a direct mapping of Fibre Channel over an Ethernet network
- FCoE is layered on top of a Lossless Ethernet
  - FSPF used to route FCoE packets
  - Ethernet Spanning Tree type protocol, RSTP, MSTP, etc, is at a layer below
- FCoE allows an evolutionary approach towards consolidation of fabrics
  - The Fibre Channel N_Port, F_Port and E-Port constructs must be retained
    - With FCoE, ports may be connected with Logical Ethernet Links
      - May pass through Lossless Ethernet NICs & Switches
      - Identified by pairs of end point MAC addresses
  - Physical Ethernet Links can replace physical FC Links
  - Physical Ethernet Links can carry all Ethernet traffic, including FCoE, but combined traffic needs the CEE/DCB capabilities

FCoE Fabrics (part 4)

- "Integrated FCoE Switches" are being built that support traditional Ethernet traffic, FCoE traffic , & FC traffic
- The FCoE solutions appear as a Fibre Channel to a Fibre Channel experienced customer
- FCoE keeps the Fibre Channel operations independent from Ethernet forwarding
  - Keeps Management /Troubleshooting simple
  - Common physical structures, different logical structures
    - Based on Ethertype (Ethertype = FCoE, or FIP)
- Storage Management should be unchanged

FCoE Relation to ISO Layers

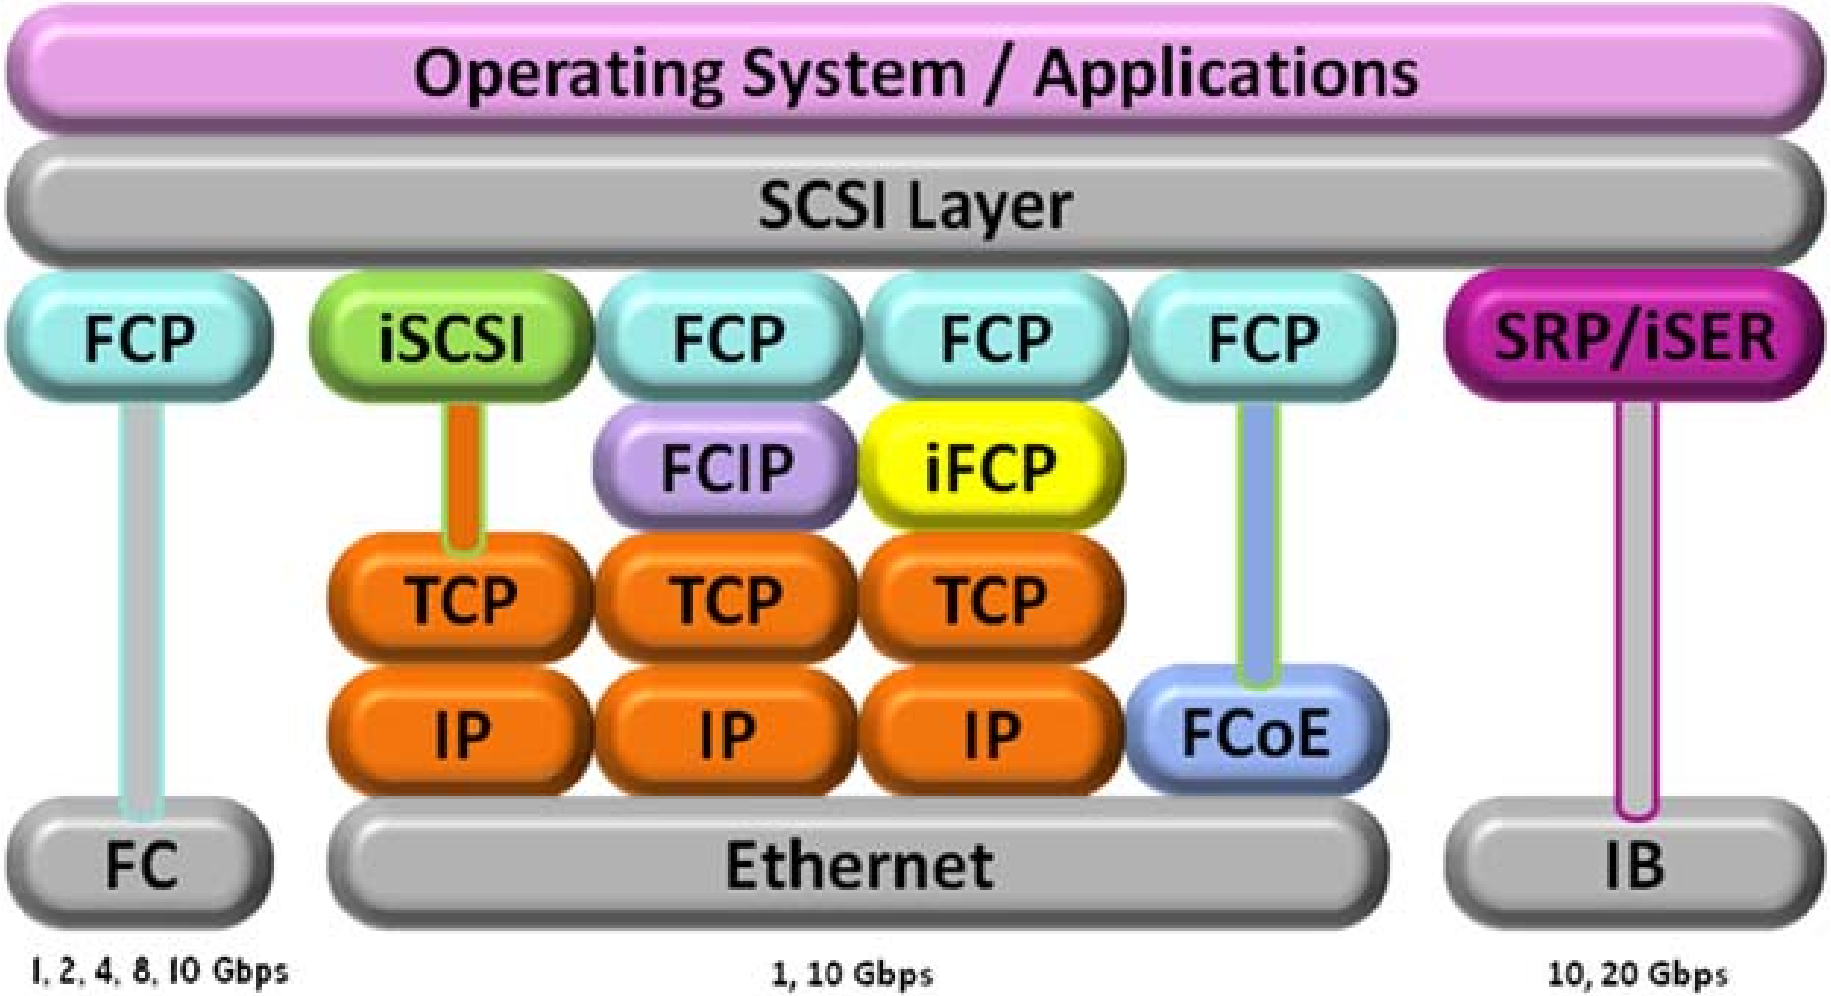

Flows

Logical Fabric Topology

An FCoE Switch receives FCoE frames addressed to its FC-MAC address and forwards them based on the D_ID of the encapsulated FC frame

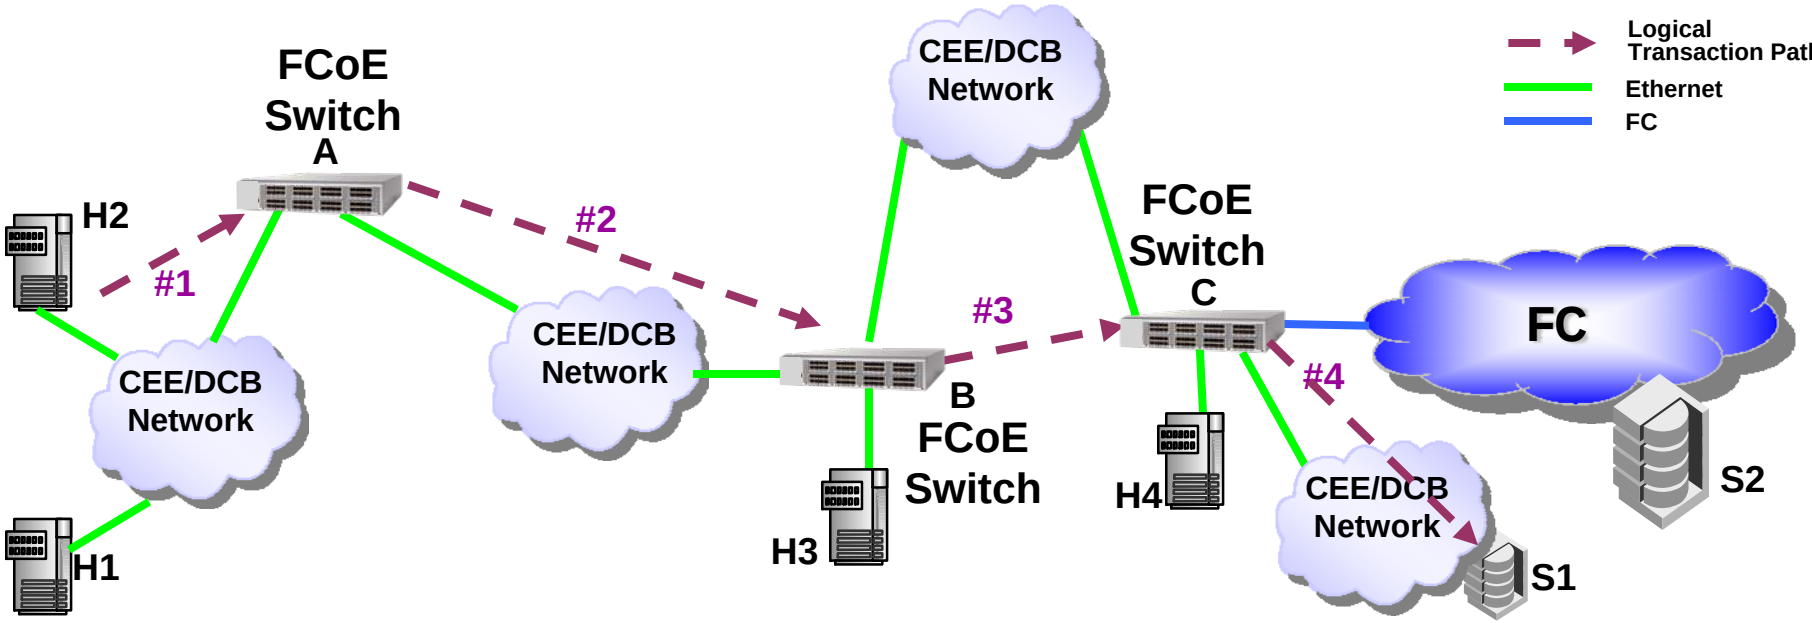

An FCoE Switch rewrites the SA and DA of an FCoE frame

	Path #1	Path #2	Path #3	Path #4
Ethernet Destination & Source	FCoE-A MAC	FCoE-B MAC	FCoE-C MAC	FCoE-S1 MAC
	FCoE-H2 MAC	FCoE-A MAC	FCoE-B MAC	FCoE-C MAC
Encapsulated FC Frame D_ID	FC_ID for S1	FC_ID for S1	FC_ID for S1	FC_ID for S1
S_ID	FC_ID for H2	FC_ID for H2	FC_ID for H2	FC_ID for H2

Single Ethernet Fabric with FCoE Switches

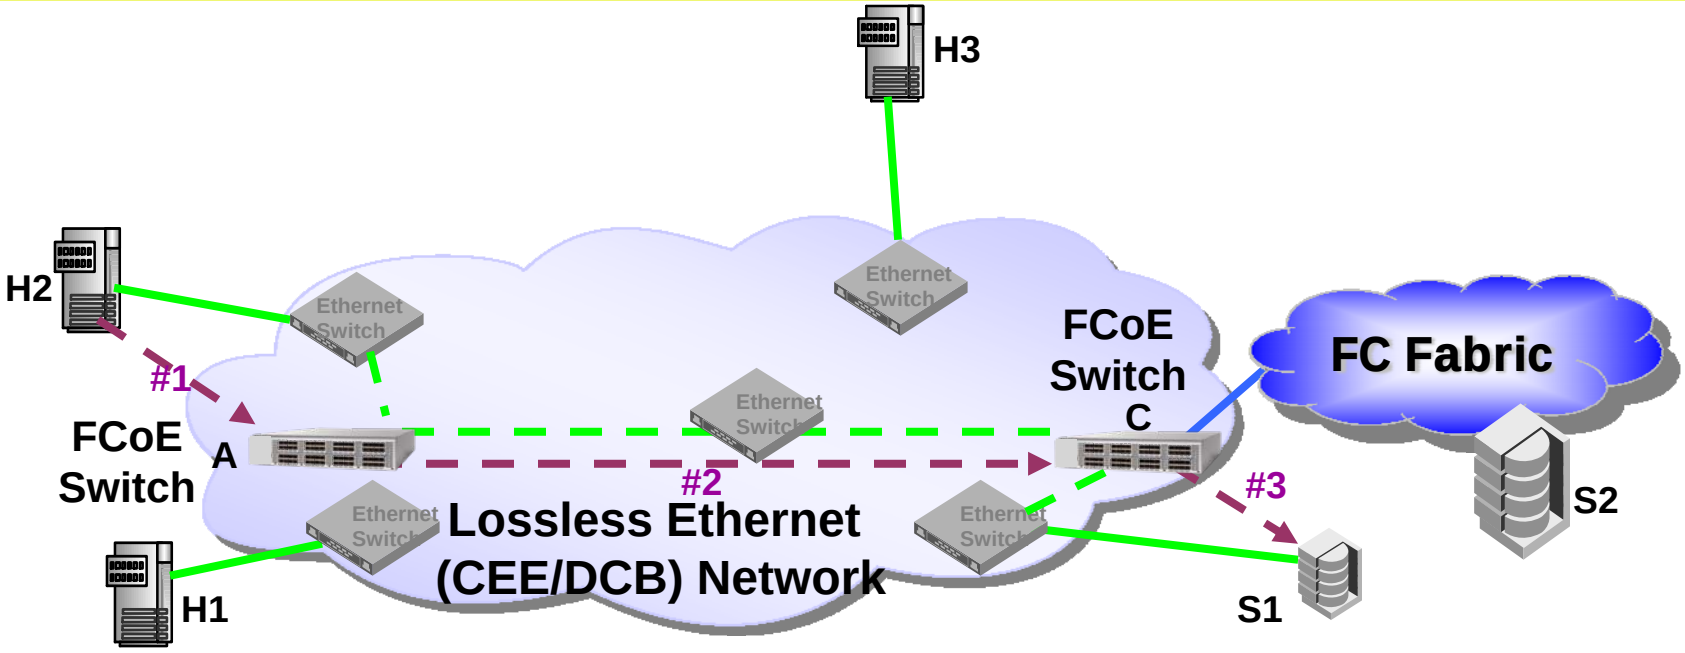

	Path #1	Path #2	Path #3
Ethernet Destination & Source	FCoE-A MAC FCoE-H2 MAC	FCoE-C MAC FCoE-A MAC	FCoE-S1 MAC FCoE-C MAC
Encapsulated FC Frame			
D_ID	FC_ID for S1	FC_ID for S1	FC_ID for S1
S_ID	FC_ID for H2	FC_ID for H2	FC_ID for H2

Logical Transaction Path
Ethernet
FC

FC Host to FCoE Storage

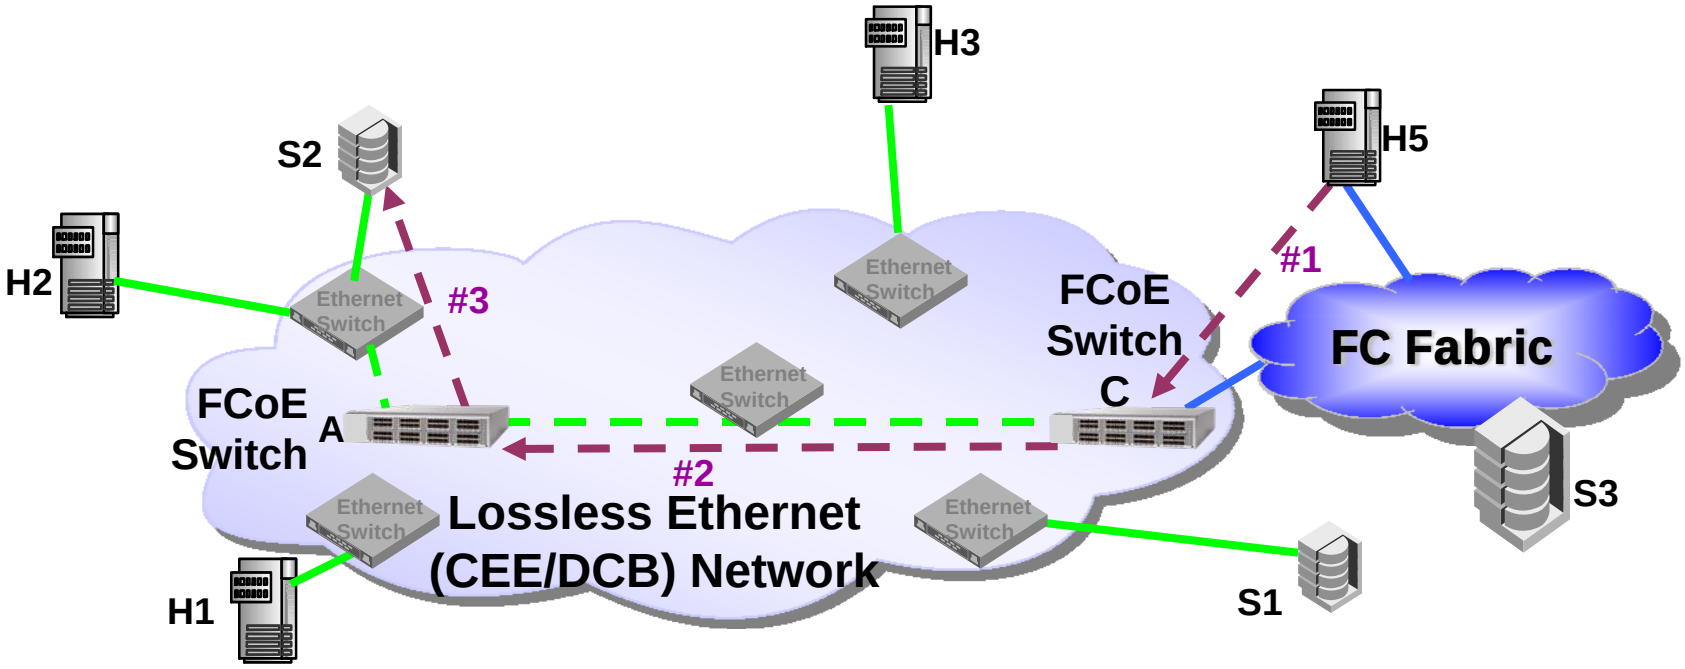

	Path #3	Path #2	Path #1
Ethernet Destination & Source	FCoE-S2 MAC FCoE-A MAC	FCoE-A MAC FCoE-C MAC	
Encapsulated FC Frame			
D_ID	FC_ID for S2	FC_ID for S2	FC_ID for S2
S_ID	FC_ID for H5	FC_ID for H5	FC_ID for H5

Logical Transaction Path
Ethernet
FC

Additional Topologies

Additional Topologies (1)

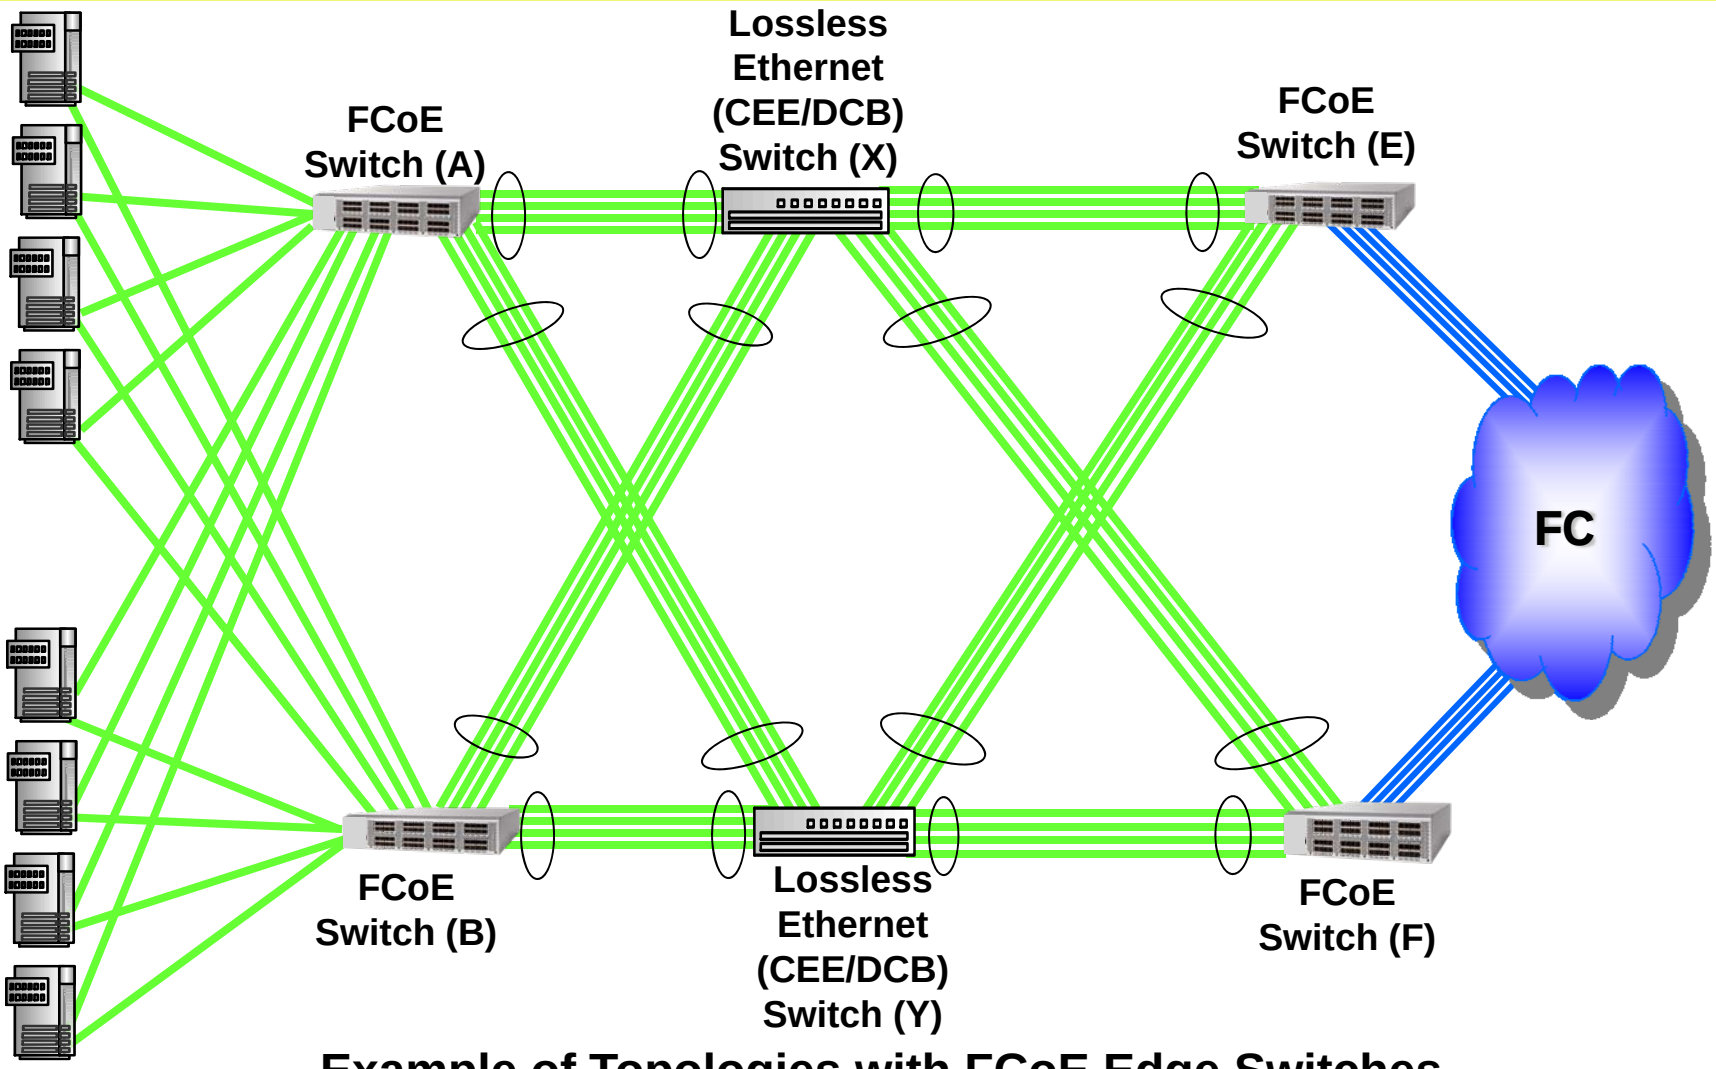

Example of Topologies with FCoE Edge Switches

Additional Topologies (2)

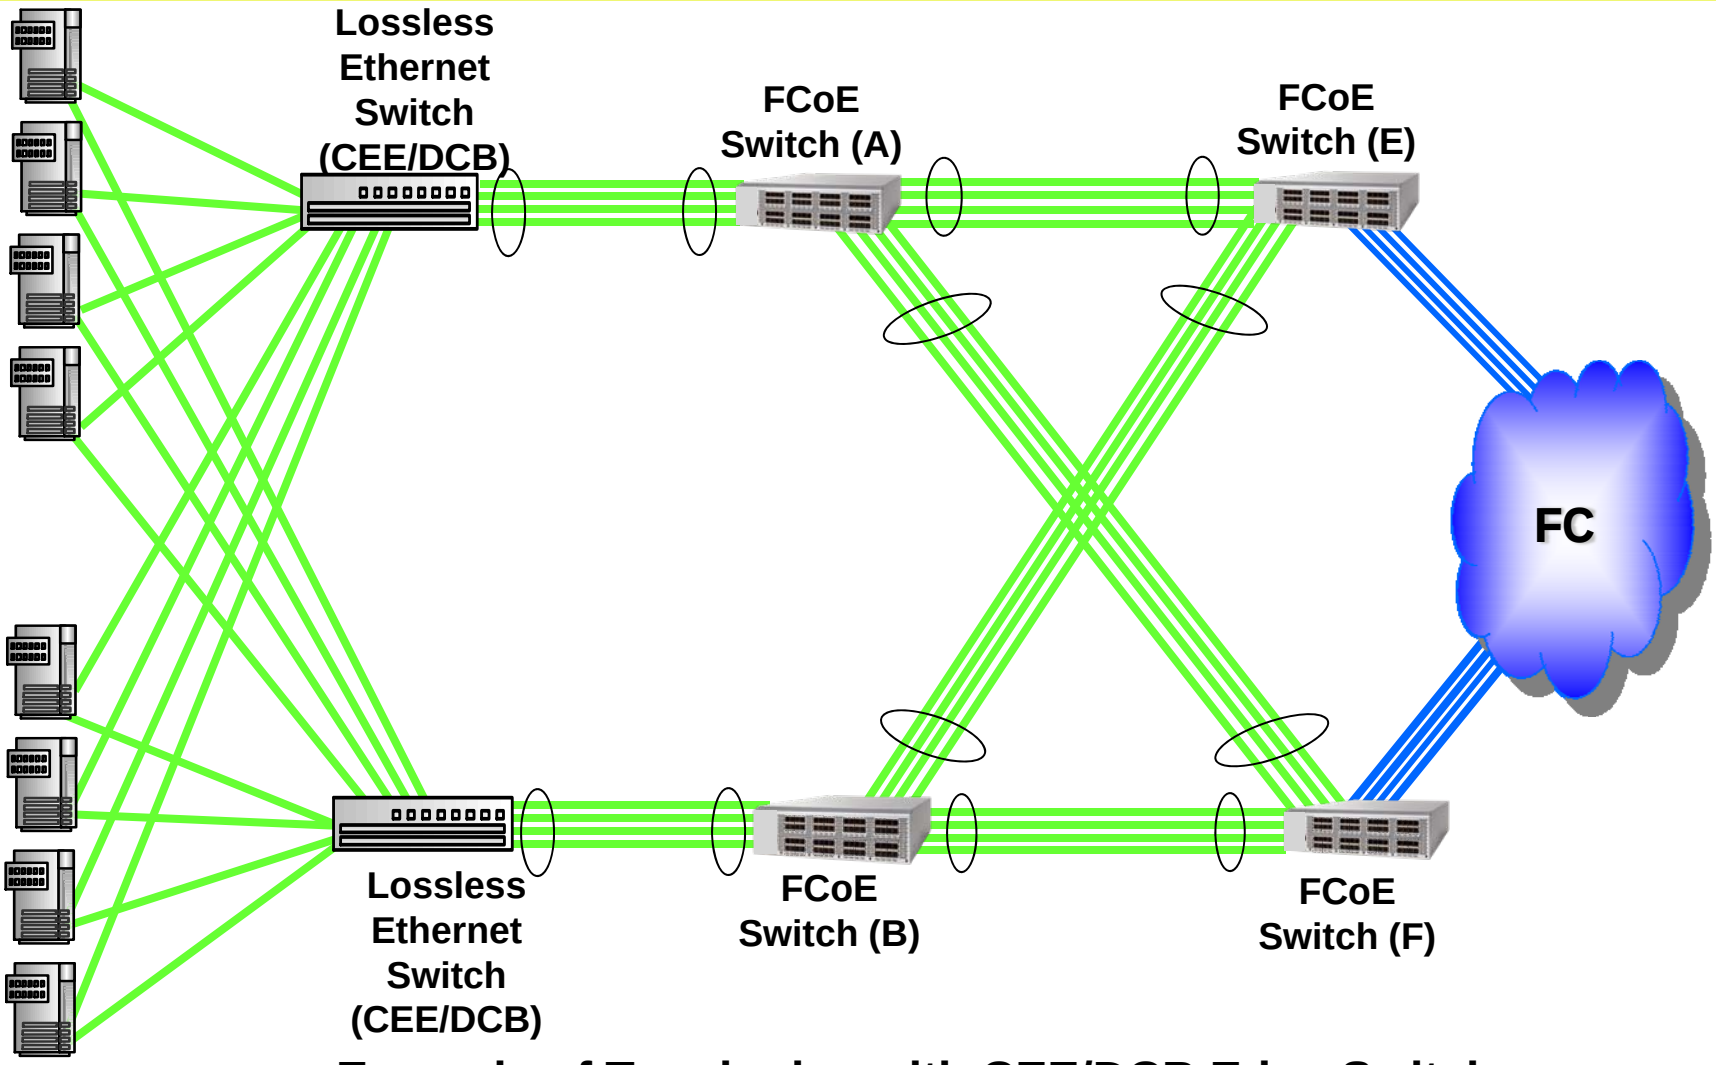

Example of Topologies with CEE/DCB Edge Switches

FSPF and STP

FSPF & STP Concepts with FCoE (basic)

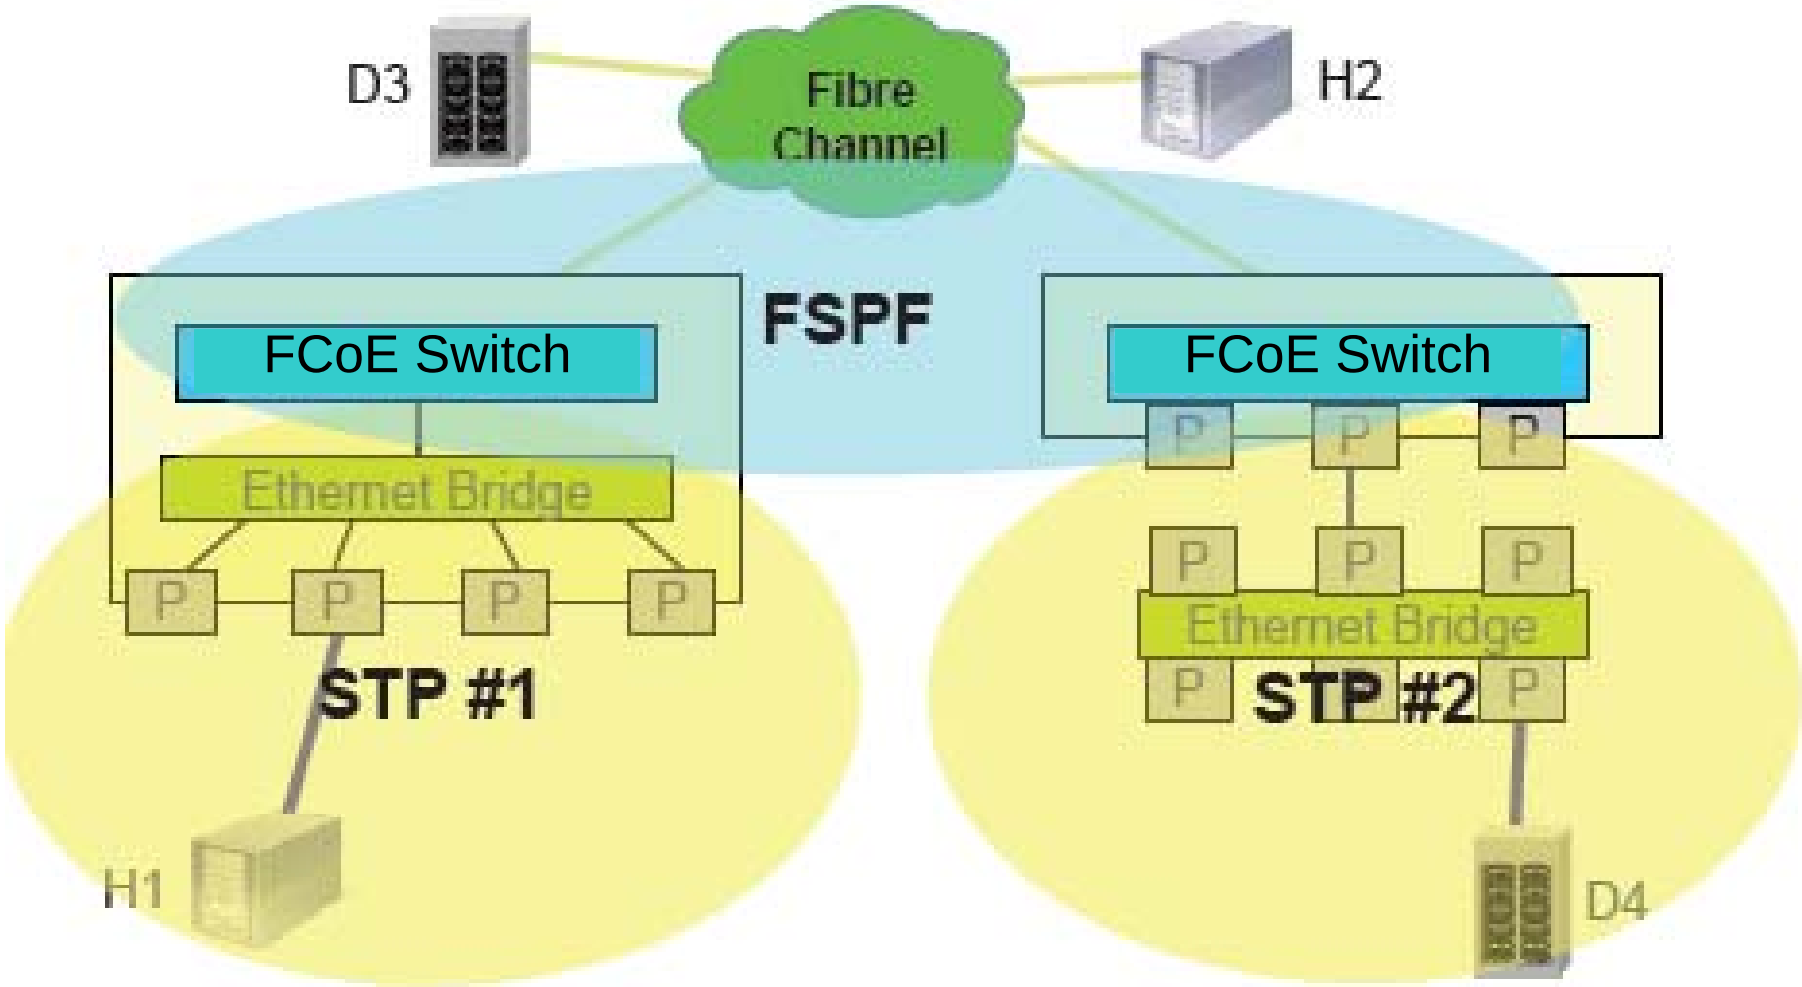

FSPF & STP Concepts with FCoE (Interconnected)

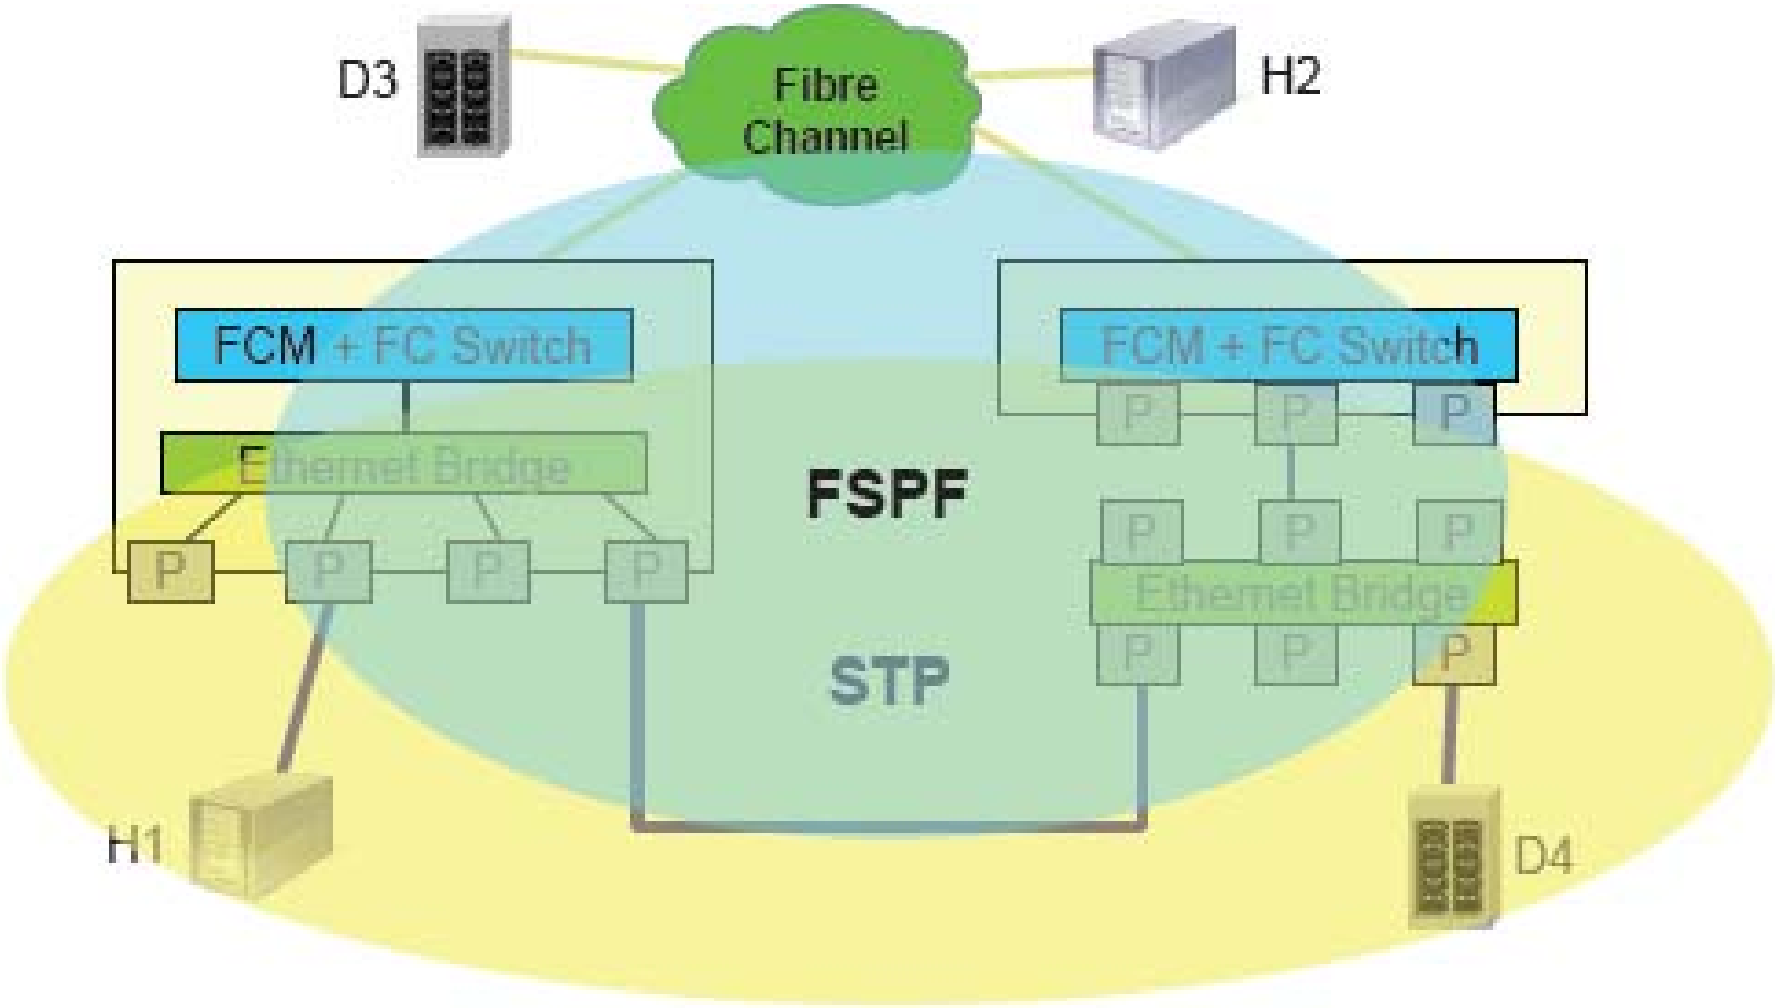

Equivalent FC Topology

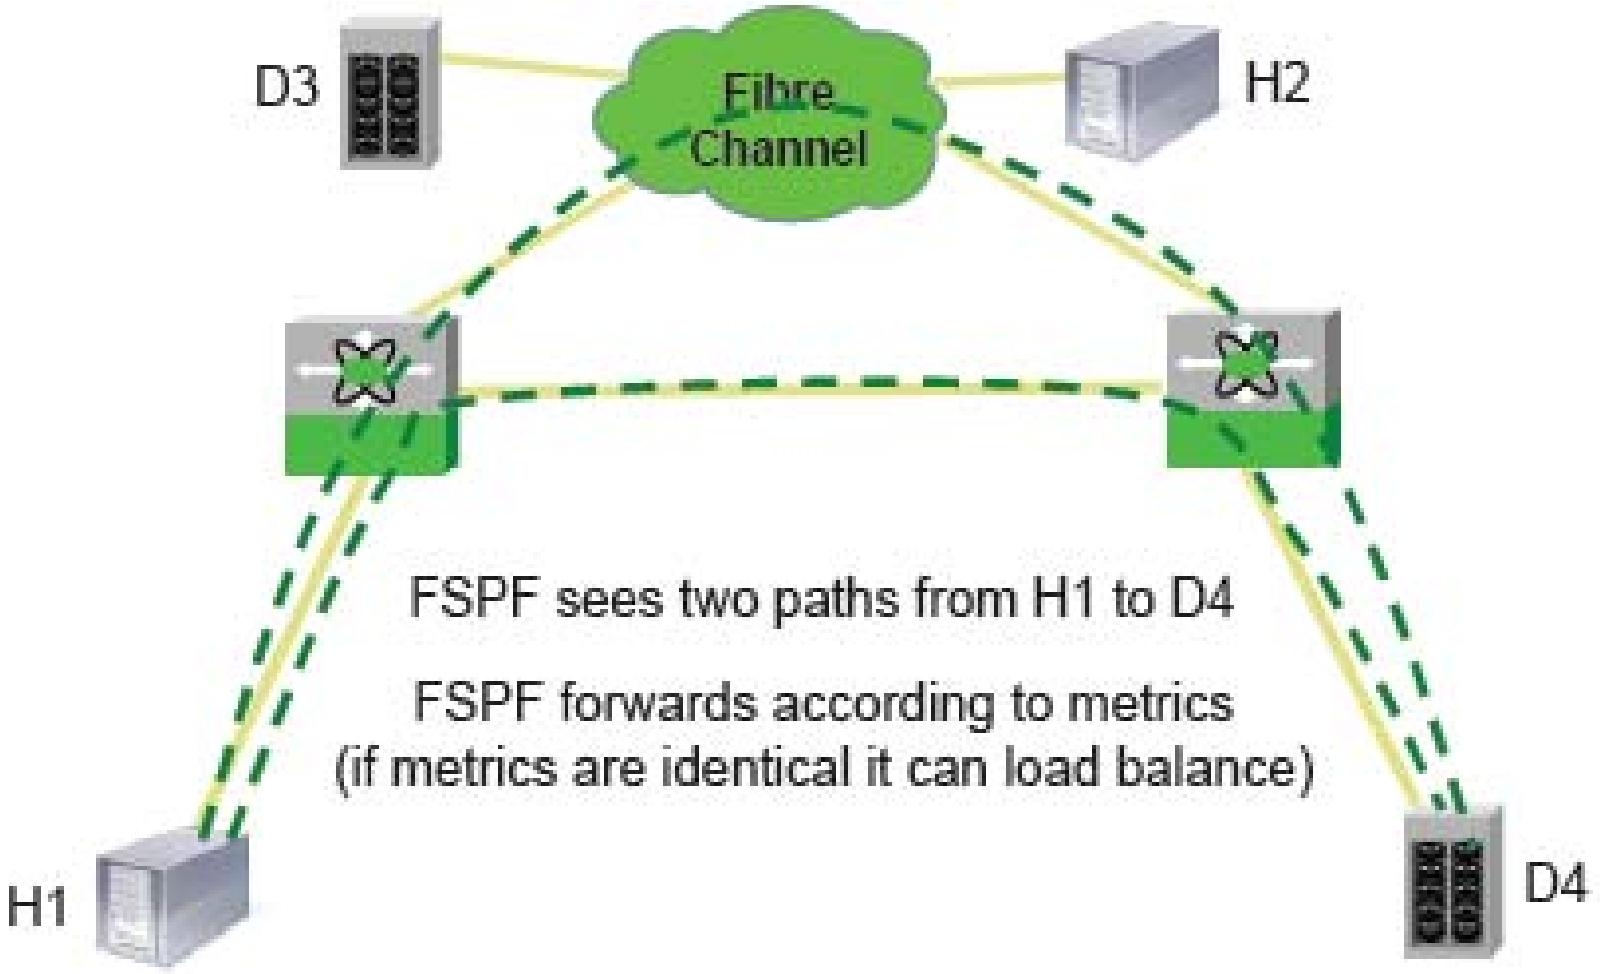

FIP Considerations and Actions

VLAN Determination

- There **may** need to be a dynamic determination of what VLANs can be used for FCoE operations
  - But this is optional -- the FIP VLAN discovery protocol is not needed if these VLANs are already known or if VLANs are not used
  - Expect this to be administratively determined
- But if dynamic VLAN information is needed:
  - FCFs will issue a VLAN information "Request" to other FCFs
    - Sent via a Multicast to "All_FCF_MACs"
    - Receiving FCF VE_Port capable MACs will respond with a Unicast "Notification" which contains a list of VLANs that can be used
      - The receiving FCF VE_Port may use one or more of these VLANs
  - ENodes will issue a VLAN information "Request"
    - Sent via a Multicast to "All_FCF_MACs"
    - Receiving FCF VF_Port capable MACs will respond with a Unicast "Notification" which contains a list of VLANs that can be used
      - The ENode may use one or more of these VLANs

FCoE Discovery Considerations

- After VLANs, if any, are determined, the discovery process uses two types of messages, **Solicitations** and **Advertisements**
- This part of the FIP Discovery Phase helps define the FCF Ports that are available for the Link instantiation Phase
- The ENodes discover the FCF ports that can become VF_Ports and FCFs discover other FCF ports that can become VE_Ports
  - ENodes Solicit (via Multicast of "All-FCF-MACs") Advertisements from FCFs while specifying their capabilities
  - In response FCFs Advertise (via Jumbo Unicast) their VF_Port capable MACs availability and capabilities back to the solicitating ENode
  - FCF Multicast (to "All-FCF-MACs") their VE_Port capable MACs existence to other FCFs
- Sometimes a New FCF will come on line and Multicast its availability to "ALL_ ENode_MACs" and "All_FCF_MACs"
- The FIP Discovery phase exchanges solicitation and/or Advertisements between (HBA's and/or FCF's) "FCoE Controllers"

FCF Discovery Actions

An FCF supporting VE_Ports:

- Discovers other VE_Port capable FCF-MACs, connected to the same Lossless Ethernet segment, by:
  - Transmitting a multicast Solicitation to 'All-FCF-MACs'
    (with the FCF bit set to one)
  - Receiving back Jumbo Unicast Advertisements from VE capable MACs
    - Which also verifies the support of Ethernet Jumbo frames in the path
  - Sending Jumbo Unicast Advertisements from its own VE capable MACs
    - In response to receipt of a Multicast to "All-FCF-MACs" from another FCF VE_Port
- Instantiates VE_Port to VE_Port connections and Exchanges FC ELP (Extended Link Protocol) and Fabric configuration (using Ethertype=FIP) with the other FCFs VE_Port capable MACs

ENode Actions

- When an ENode becomes operational:
  - The ENode discovers the VF_Port capable FCF-MACs connected to the same Lossless Ethernet segment by:
    - Transmitting a multicast Solicitation to 'All-FCF-MACs' (with the FCF bit set to zero)
    - Receiving back Jumbo Unicast Advertisements from compatible VF_Port capable FCF-MACs
      - (May store the discovered FCF-MACs in an FCF port list)
- When an ENode receives an Advertisement that a new FCF port is available, it may send a Unicast Solicitation to it and receive a Jumbo Unicast Advertisement in reply
  - To verify the support of Ethernet Jumbo frames in the path
- May then perform FLOGIs (with Ethertype=FIP) to a vendor specific subset of the FCF-MACs in the FCF port list

FIP Keep Alive

- FCoE connections may not be directly attached to the FCF
  - May flow through intermediate Lossless Ethernet switches,
  - Needs to be a method to detect that something was wrong in the path
    - Between VN_Ports and VF_Ports as well as between VE_Ports
- Timers connected to the ENodes & FCFs can determined that a port is sending messages
- Need to determine if a port is just dormant, but still alive
  - Therefore, a periodic message needs to be sent.
- The periodic unsolicited Advertisements from the FCF can be used for Keep Alive in the following directions
  - FCF → ENode
  - FCF ←→ FCF
- But something else is needed for the ENode → FCF direction
  - Because there is no periodic unsolicited Advertisement from a ENode
- So -- a special FIP Keep Alive Message was created to just inform the FCF that the ENode and its VN_Ports are still alive

FIP KAs (Keep Alives) & the FCoE Controllers

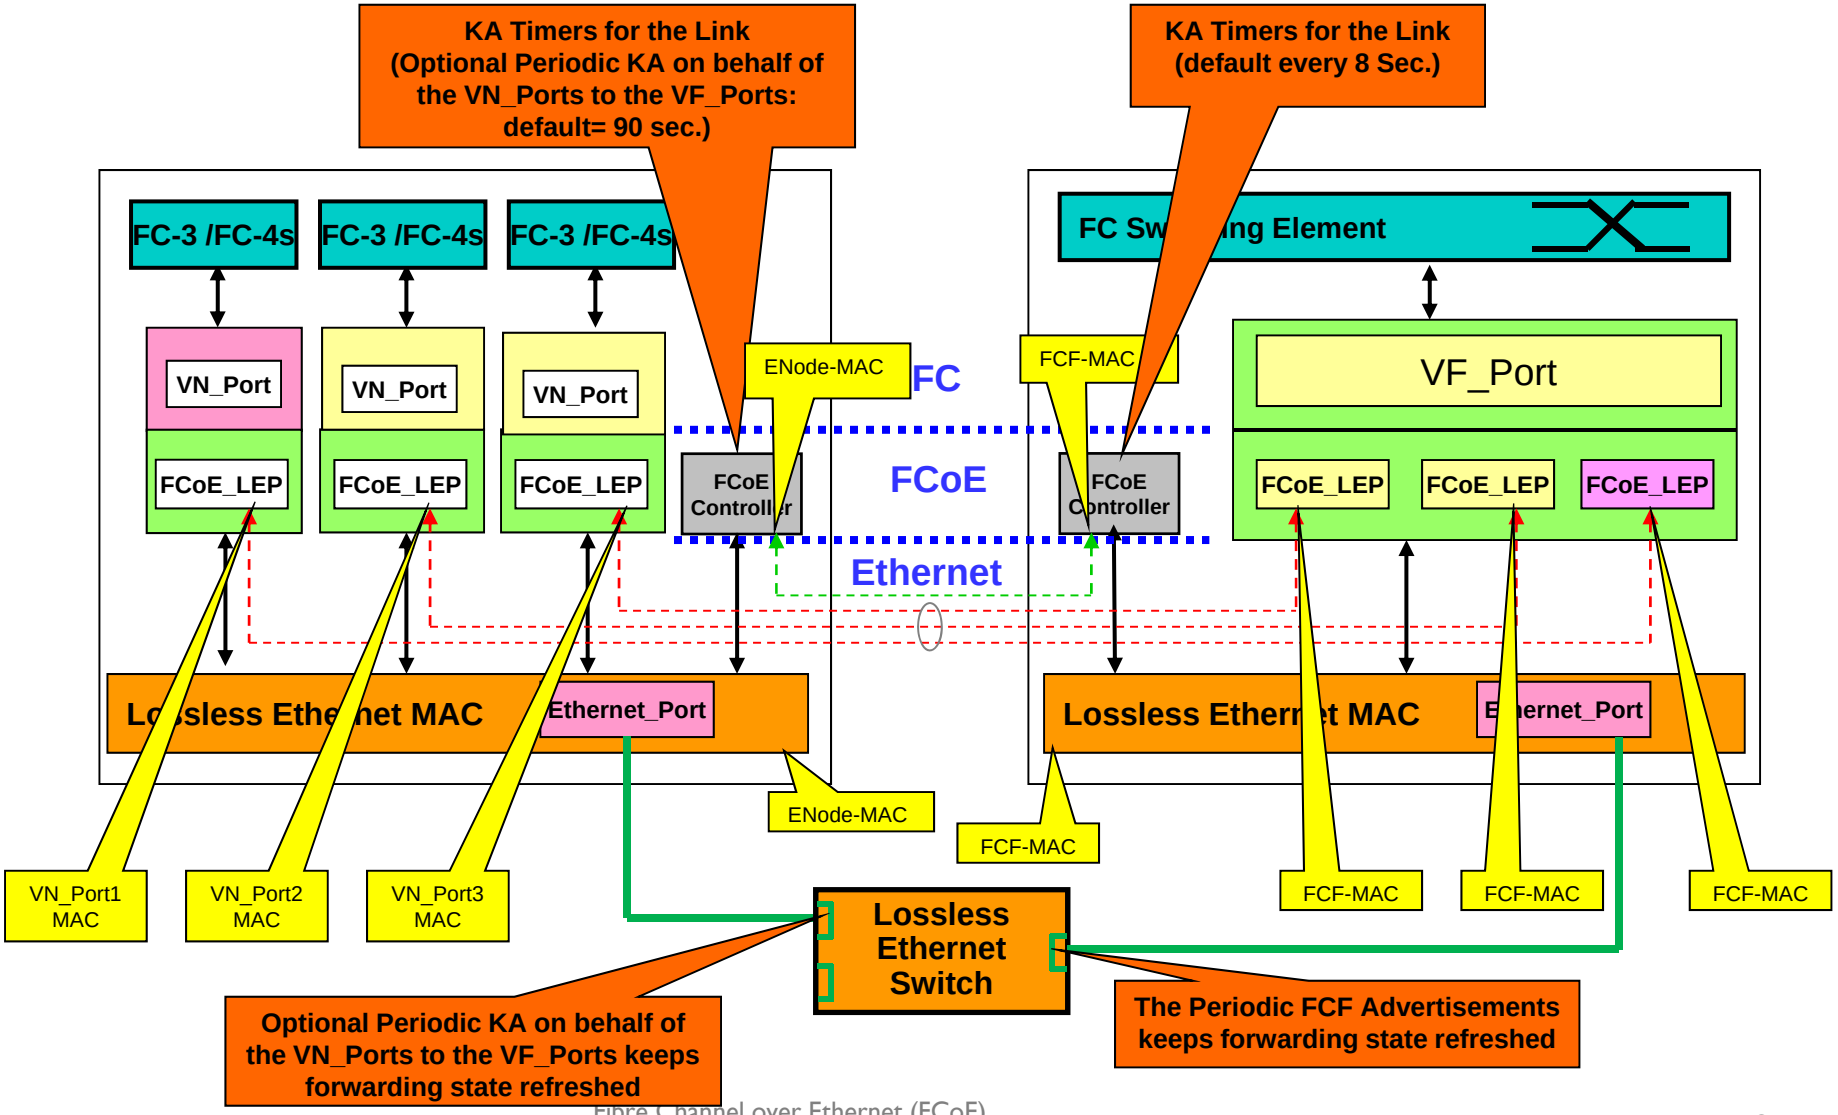

FIP Clear Virtual Link

- When a port connection is determined by a side to be not functional
  - There is a requirement for the side to attempt to clear all state,
    - (Especially state which may have been established in any intermediate switches)
- The VN_Port can issue a FIP Logout (LOGO)
- But if the ENode is dead or its link busted, the FCF needs to clear things on its own
- But there is no FC capability for a FC switch to Logout a N_Port
  - So -- a special FIP function was created called **Clear Virtual Link**
    - Can be issued by the FCF
  - When intermediate switches see this frame they should cleanup their ACLs etc.

FIP Descriptors

FIP Operation Codes

Operation Code	Subcode	Operation
0001h	01h	Discovery, Solicitation
	02h	Discovery, Advertisement
0002h	01h	FLOGI/FDISC/LOGO/ELP, Request
	02h	FLOGI/FDISC/LOGO/ELP, Reply
0003h	01h	FIP Keep Alive
	02h	FIP Clear Virtual Link
0004h	01h	FIP VLAN Request
	02h	FIP VLAN Notification
FFF8h .. FFFEh	00h .. FFh	Vendor Specific
All others	All others	Reserved

Type and Values of FIP Descriptors

Type	Value
00	Reserved
01	Priority
02	MAC Address
03	FC-MAP
04	Node_Name / Switch_Name
05	Fabric
06	Max FCoE Size
07	FLOGI
08	NPIV FDISK

Type	Value
09	Fabric Log Out
10 (0Ah)	ELP
11 (0Bh)	VN_Port ID
12 (0Ch)	FKA_ADV_Period
13 (0Dh)	Vendor ID
14 (0Eh)	VLAN
15 (0Fh) – 240 (F0h)	Reserved (128-240 are Non Critical)
241 (F1h) – 254 (FEh)	Vendor Specific
255 (FFh)	Reserved

FIP Descriptors (1)

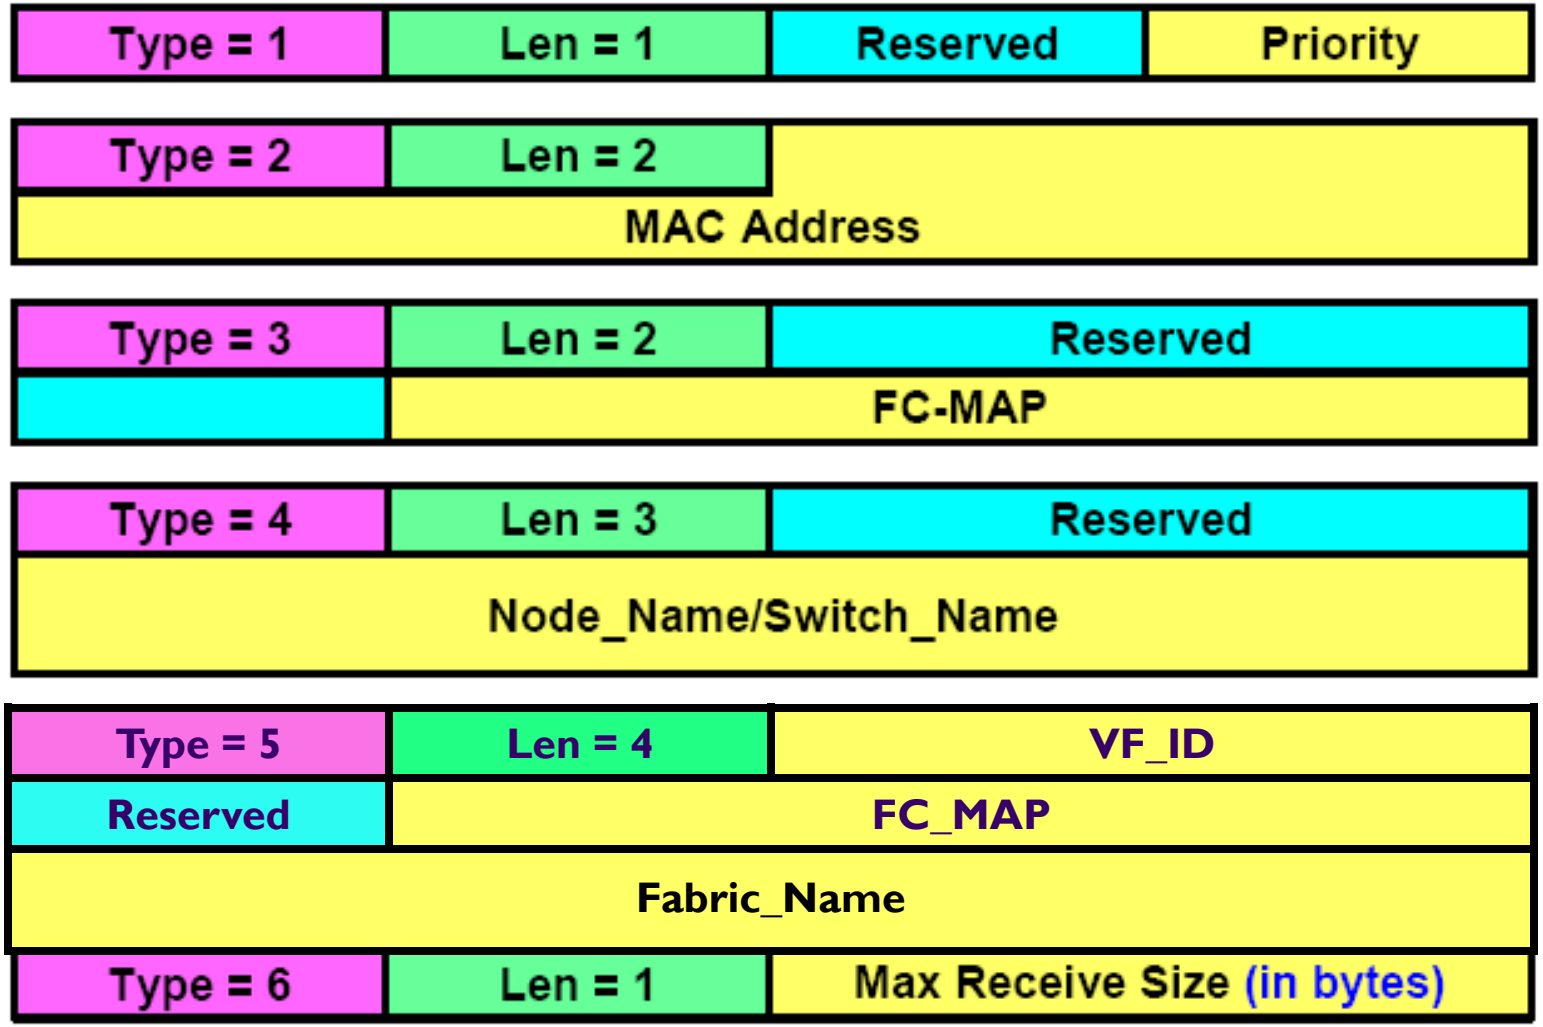

Lengths are measured in 32-bit words

FIP Descriptors (2)

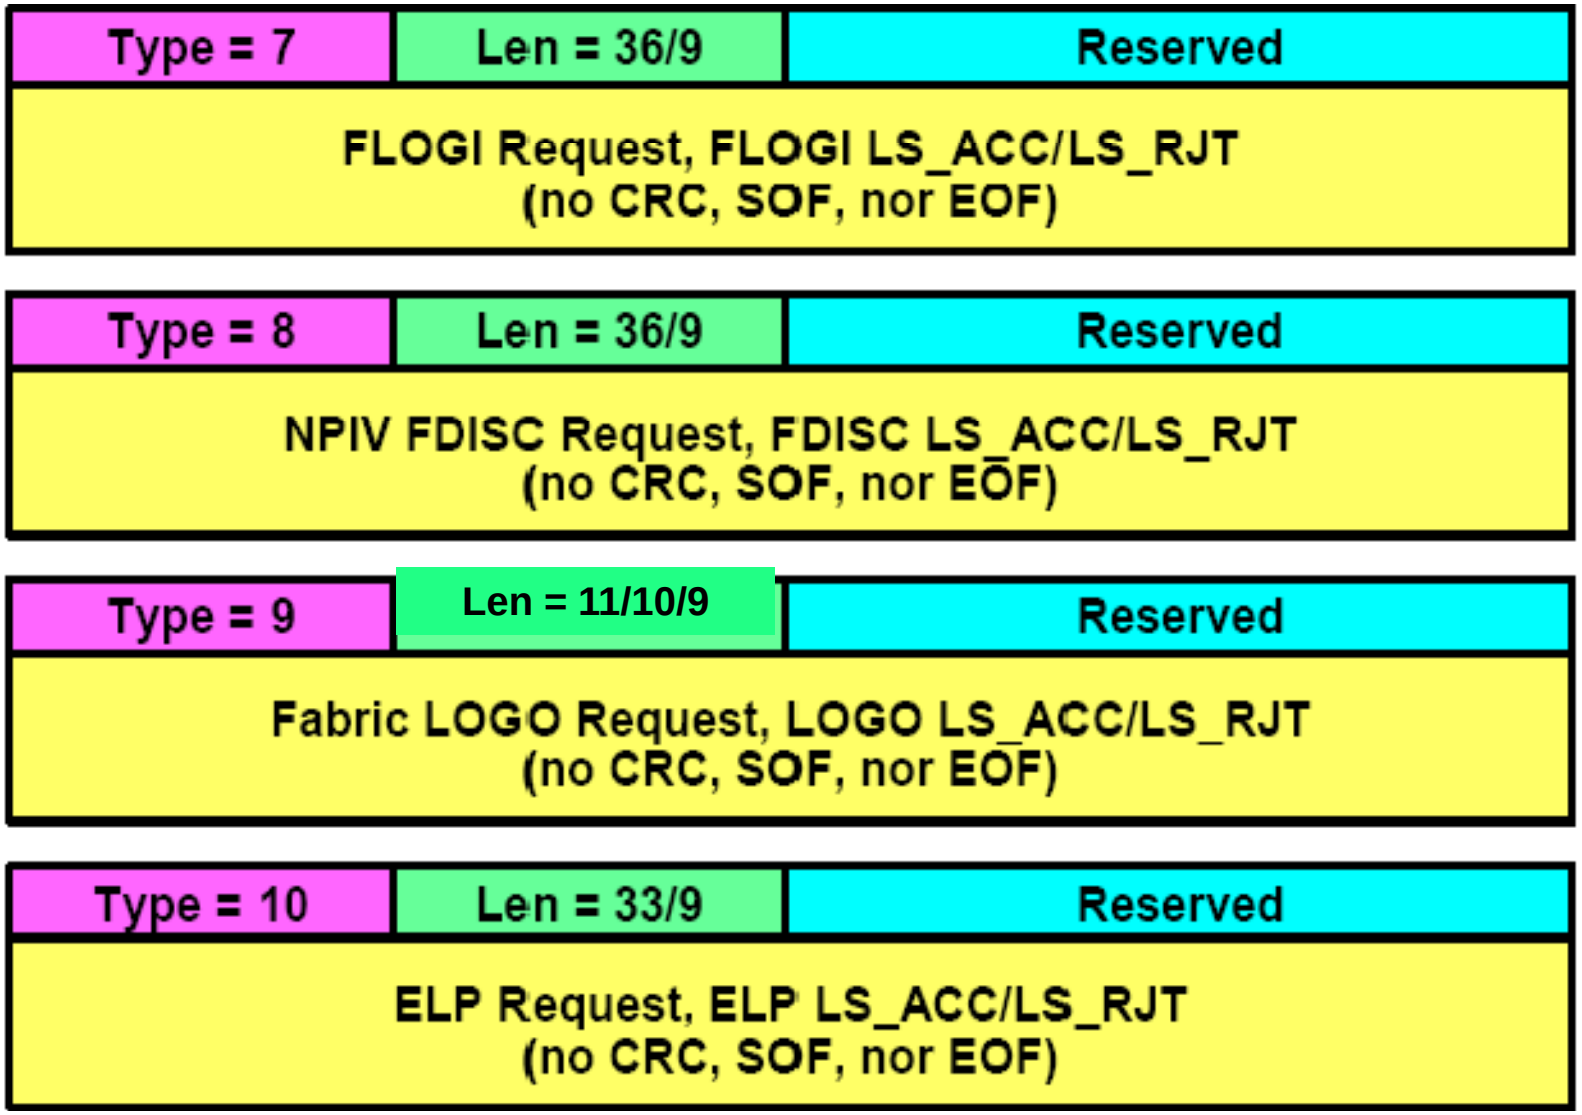

FIP Descriptors (3)

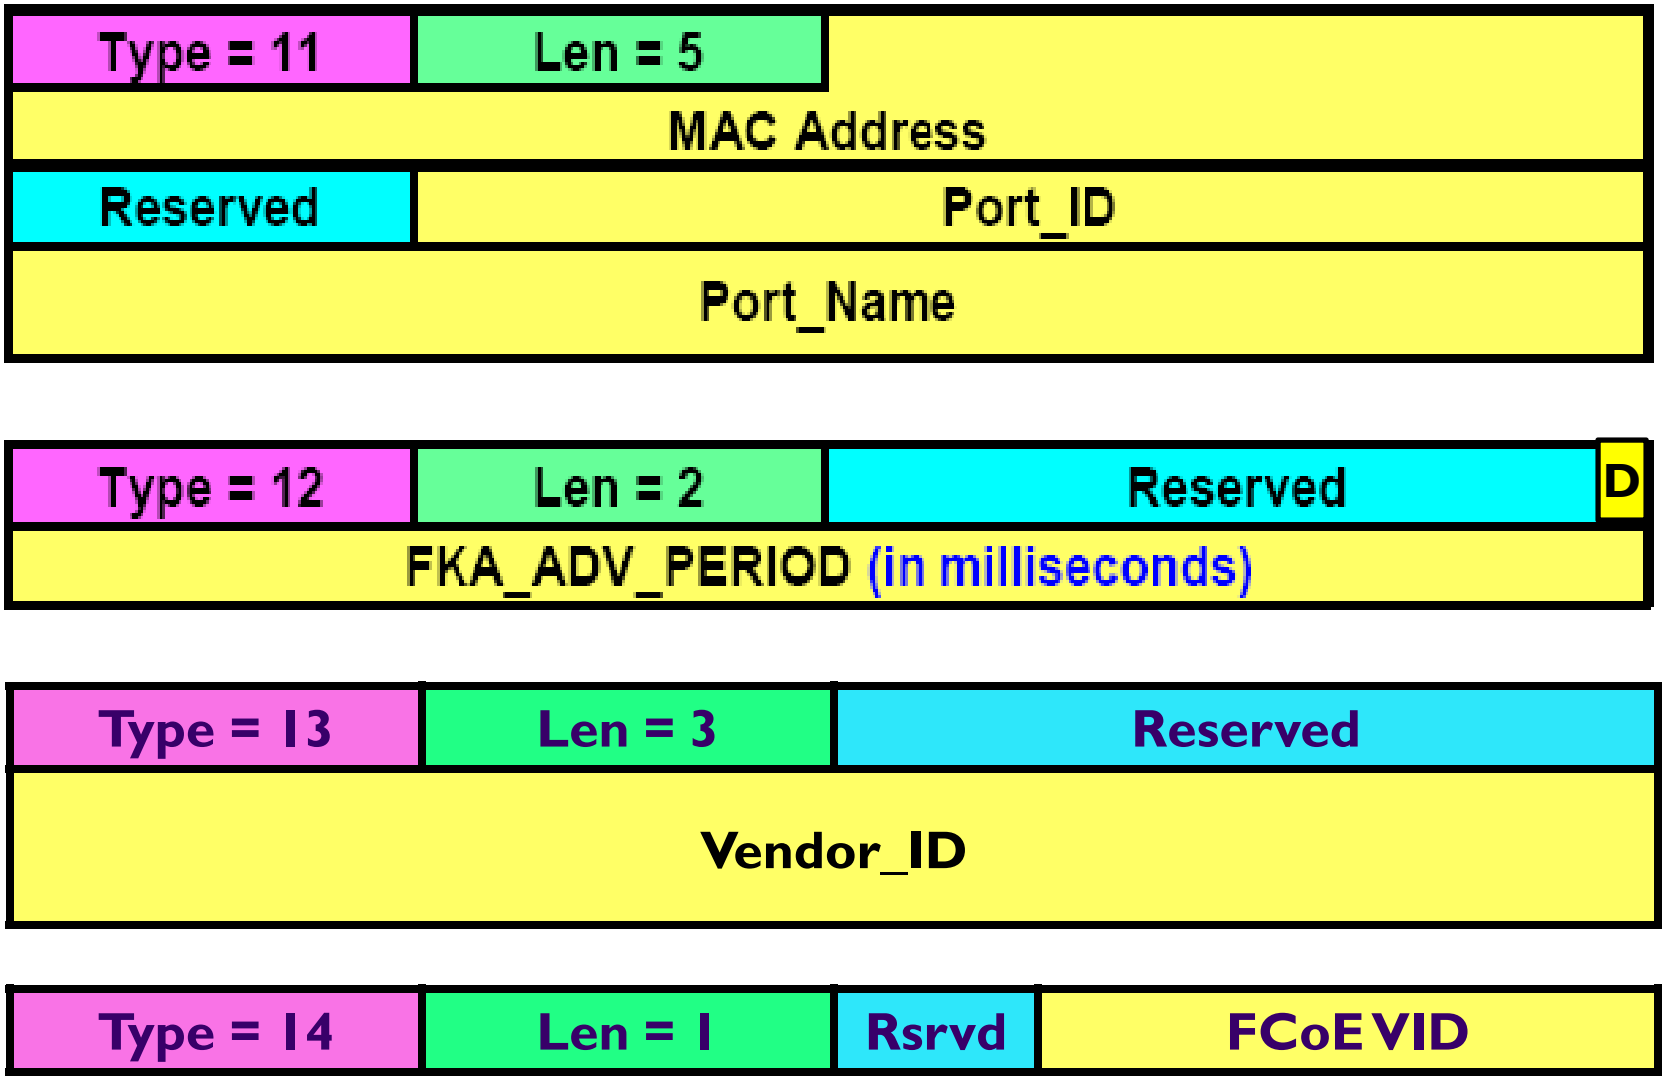

VLAN Request & Notification

VLAN Request
ENode/FCF → FCF

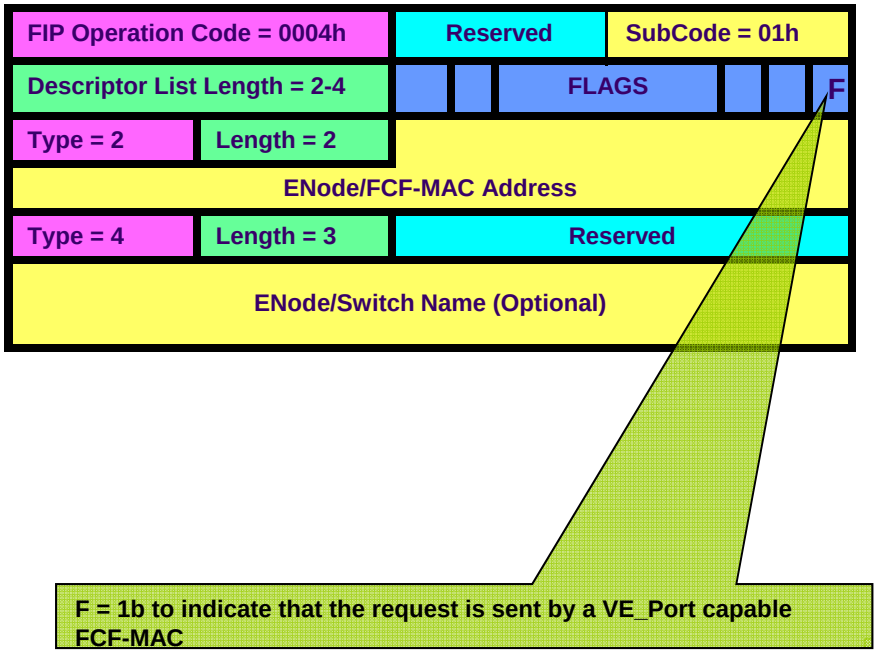

VLAN Notification
FCF → FCF/ENode

FIP Operation Code = 0004h		Reserved	SubCode = 02h
Descriptor List Length = 2-n		FLAGS	
Type = 2	Length = 2		
FCF-MAC Address			
Type = 14	Len = 1	Rsrvd	FCoE VID -1
.	.	.	.
.	.	.	.
Type = 14	Len = 1	Rsrvd	FCoE VID-n

Discovery Solicitation & Advertisement

Education SNIA

Solicitation ENode → FCF

ENode Capabilities:
FP = 1 if FPMA supported
SP = 1 is SPMA supported

F = 0b

FIP Operation Code = 0001h				Reserved			SubCode = 01h	
Descriptor List Length = 6				FP	SP	Flags		F
Type = 2	Len = 2							
ENode's MAC Address								
Type = 4	Len = 3	Reserved						
Node_Name								
Type = 6	Len = 1	Max Receive Size						

Solicitation FCF → FCF

FIP Operation Code = 0001h				Reserved			SubCode = 01h	
Descriptor List Length = 8						FLAGS		F
Type = 2	Length = 2							
FCF-MAC Address								
Type = 3	Length = 2	Reserved						
Reserved	FC-MAP							
Type = 4	Length = 3	Reserved						
Switch Name								
Type = 6	Length = 1	Max Receive Size						

Advertisement

FCF Capabilities:
FP = 1 if FPMA supported
SP = 1 if SPMA supported

A = 1b if FCF available for Login/ELP

F = 1b

FIP Operation Code = 0001h				Reserved			SubCode = 02h	
Descriptor List Length = 12				FP	SP	Flags	A S	F
Type = 1	Len = 1			Reserved			Priority	
Type = 2	Len = 2							
FCF-MAC Address								
Type = 4	Len = 3	Reserved						
Switch_Name								
Type = 5	Len = 4	VF_ID						
Reserved	FC_MAP							
Fabric_Name								
Type = 12	Len = 2	Reserved						
FKA_ADV_Period								
Padding to Max Receive Size of Soliciting Entity, if Solicited (i.e. if S=1b, otherwise no padding								

F = 1b to indicate that the Solicitation is sent by a VE_Port capable FCF-MAC (i.e., it is soliciting VE_Port capable FCF-MACs)

FIP FLOGI

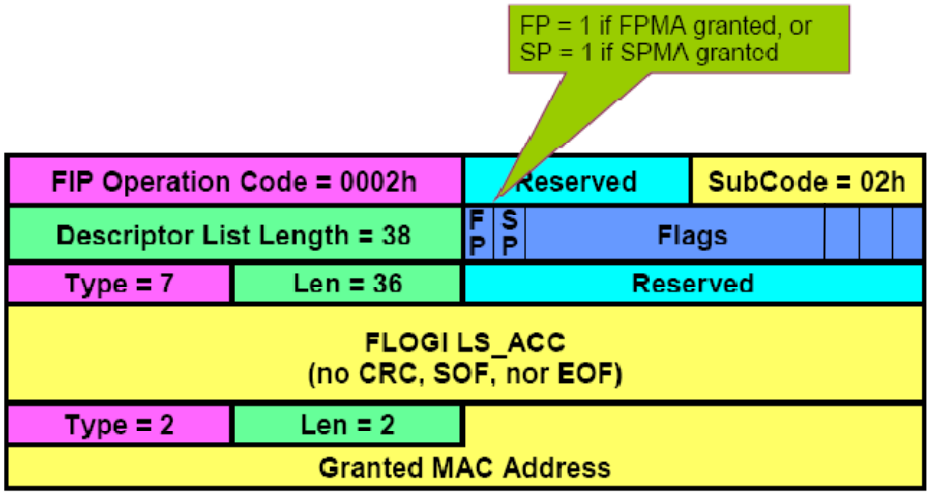

Accept (FCF → ENODE)

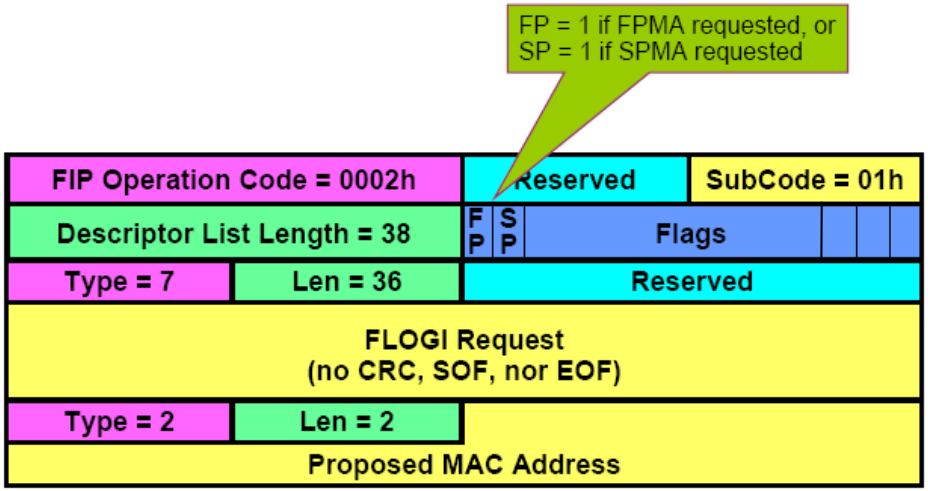

Request (ENODE → FCF)

Reject (FCF → ENODE)

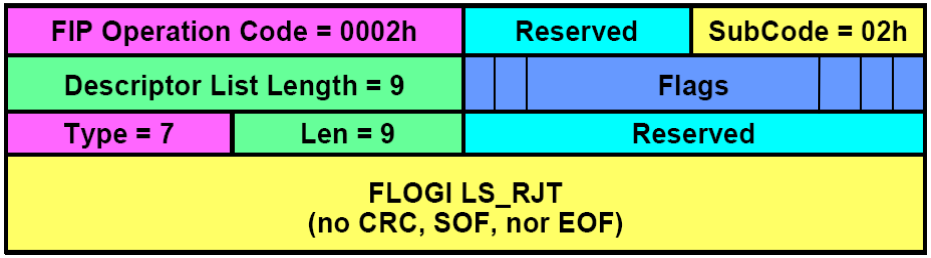

FIP NPIV FDISC

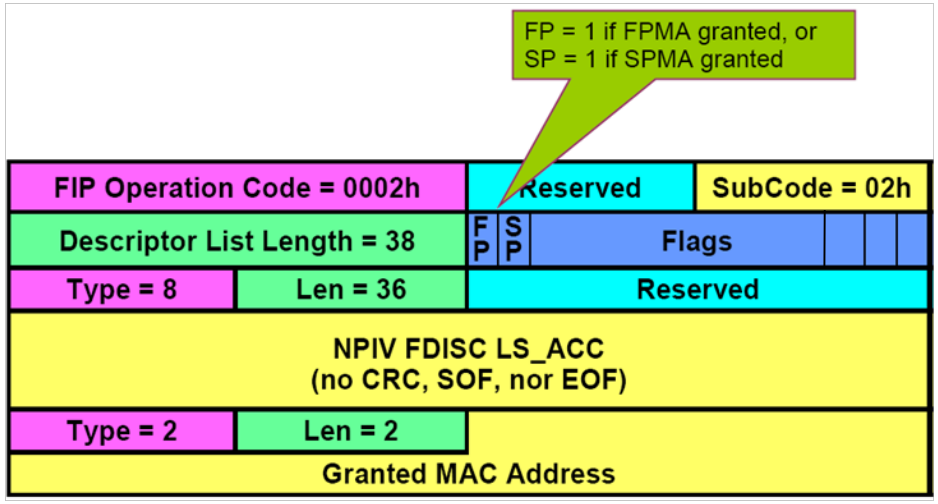

Accept (FCF → ENODE)

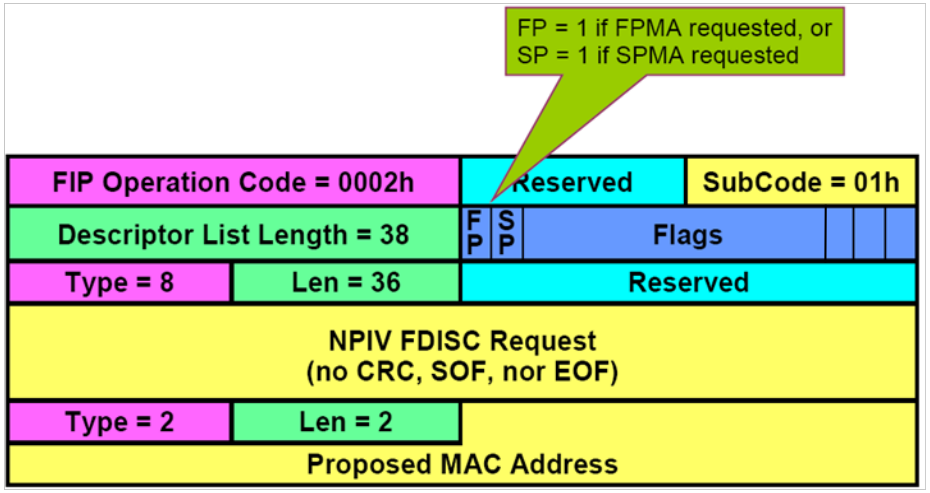

Request (ENODE → FCF)

Reject (FCF → ENODE)

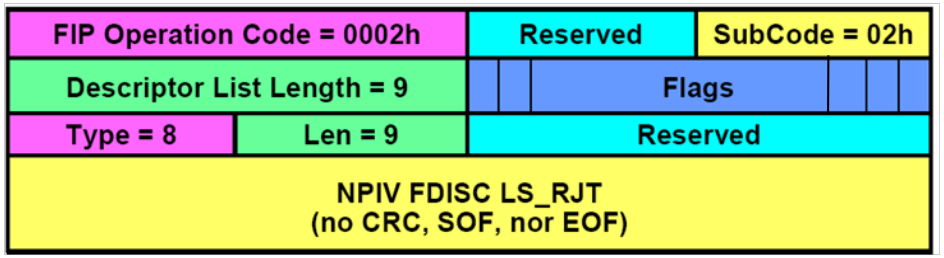

FIP Fabric LOGO

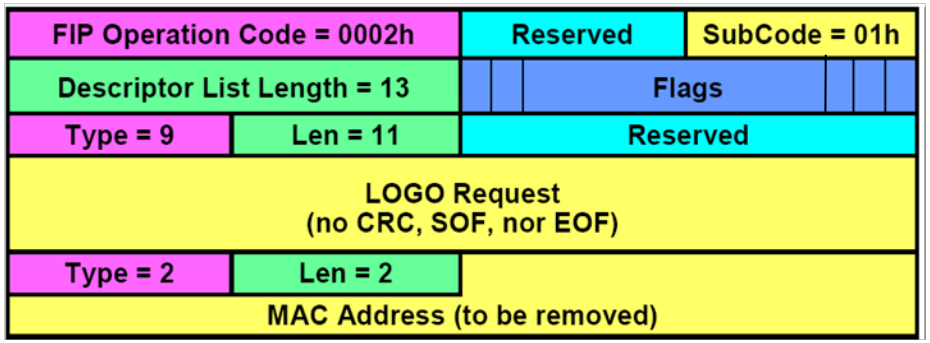

Request (ENODE → FCF)

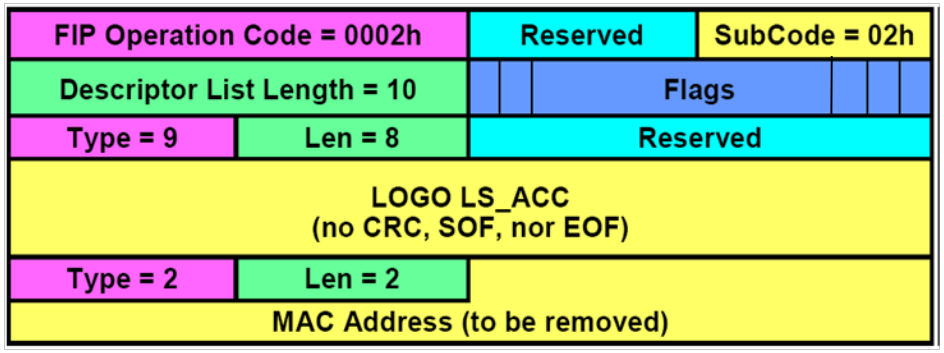

Accept (FCF → ENODE)

Reject (FCF → ENODE)

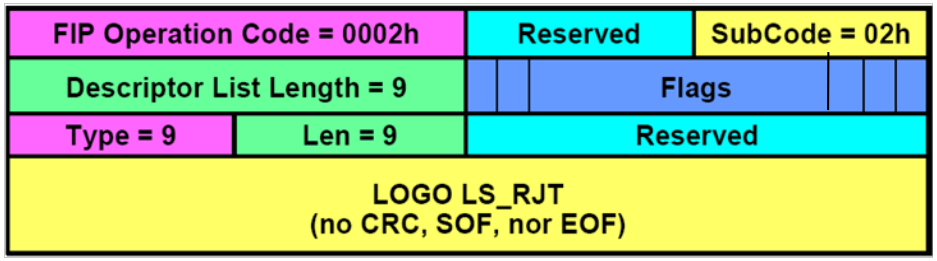

FIP ELP

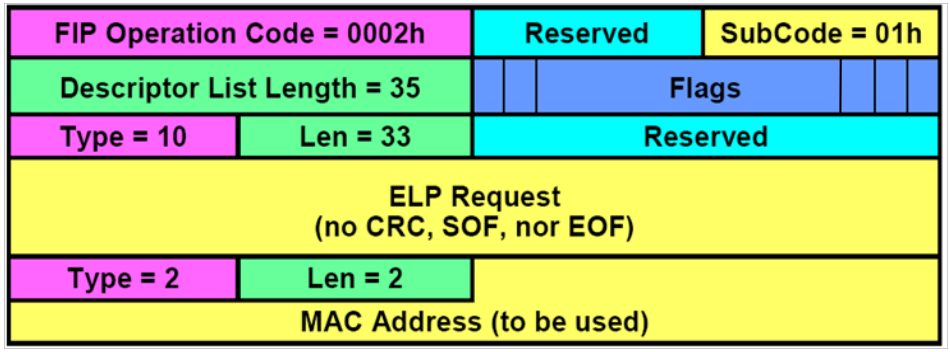

Request (FCF ← → FCF)

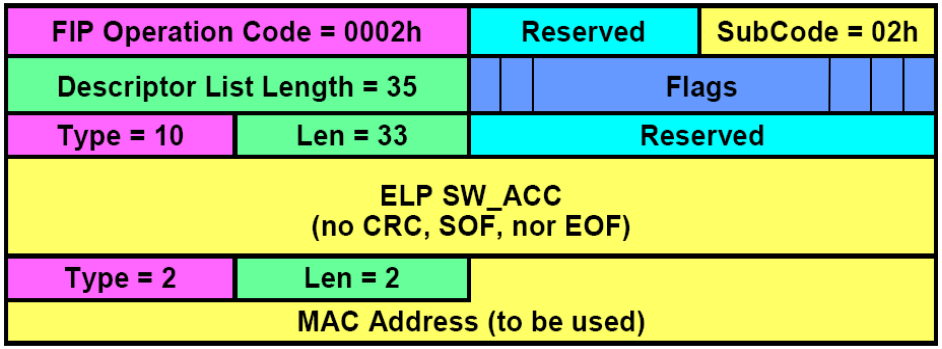

Accept (FCF ← → FCF)

Reject (FCF ← → FCF)

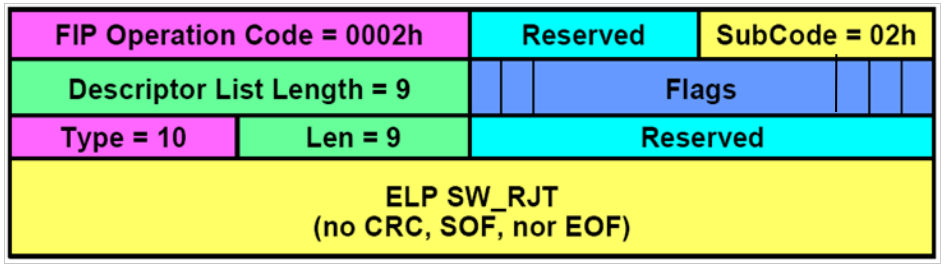

FIP Keep Alive

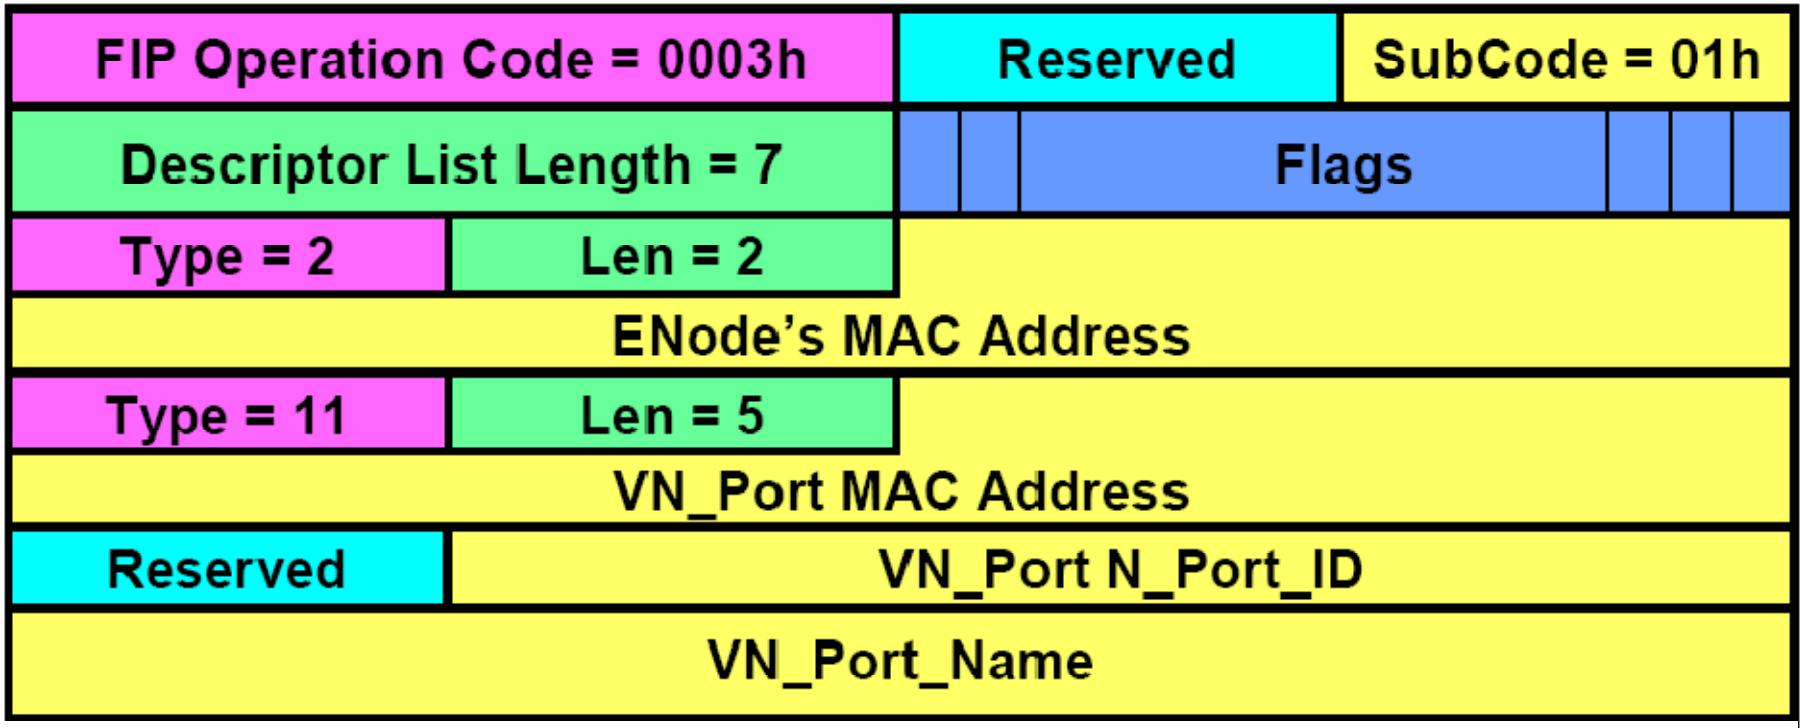

ENode → FCF

(The Keep Alive from the FCF → ENode/FCF is a normal unsolicited Advertisement)

FIP Clear Virtual Link

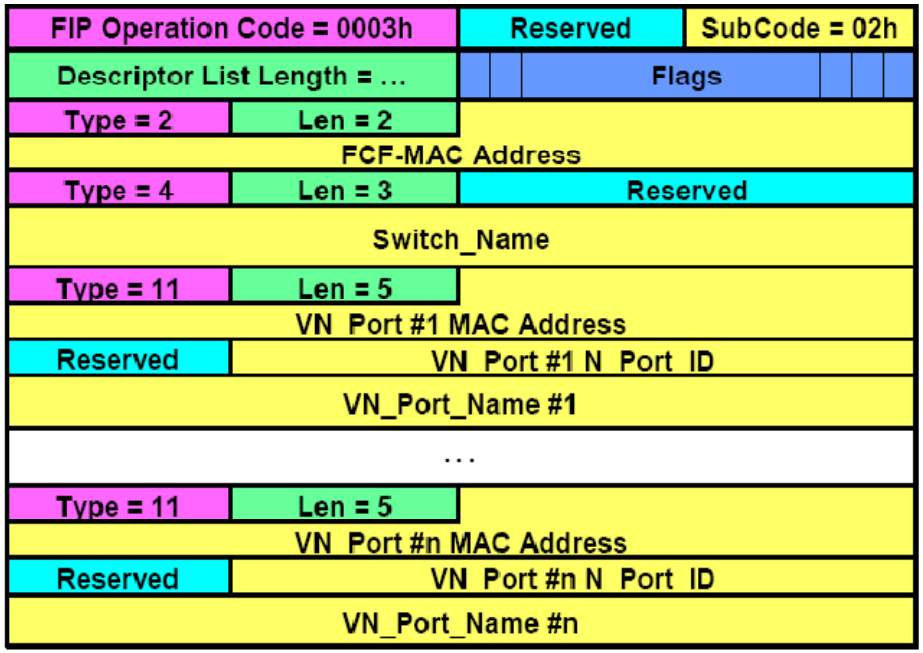

FCF → ENODE

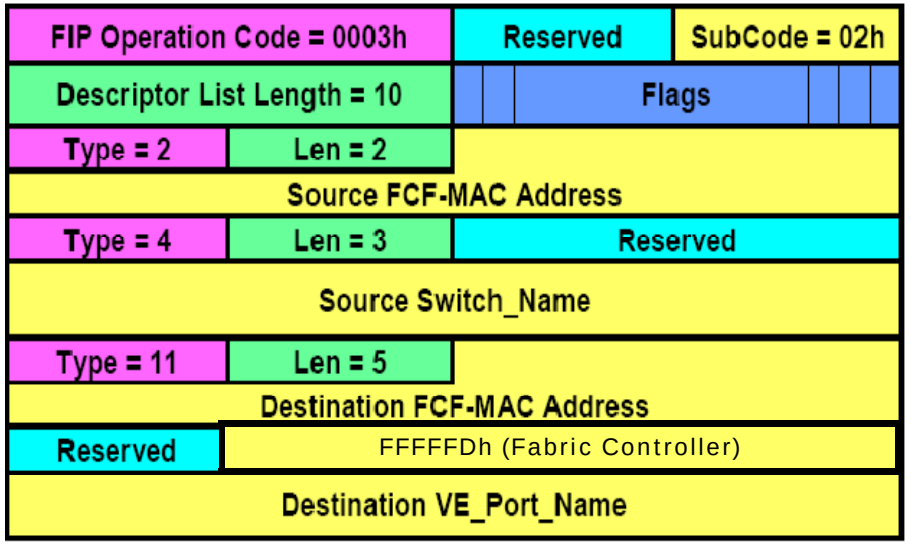

FCF ← → FCF

(Note: the ENode uses FIP Fabric LOGO to Clear the Virtual Link from its side)

Vendor Specific FIP Message

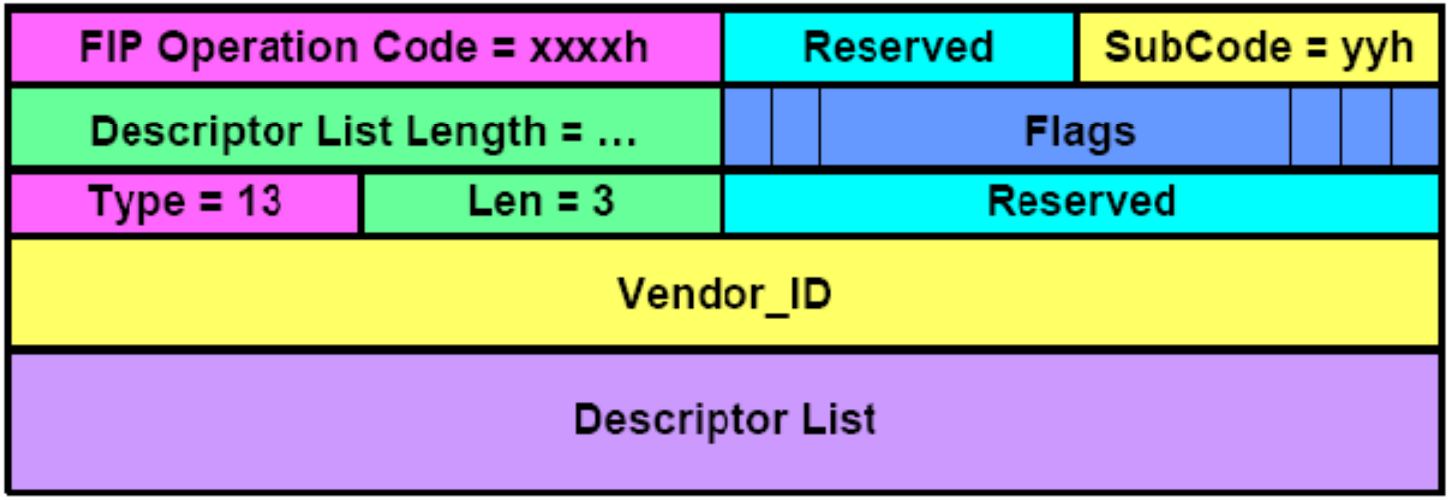

- An unknown Vendor Specific message shall be ignored

  Normal operation will be the result

- A device is never required to generate a Vendor Specific message for normal operation

Vendor Specific Operation Code =FFF8 -- FFFE with any SubCode
Vendor Specific Type (in Descriptor List) = F1h – FEh
Reserved Types = 0Fh – F0h, & FFh

Pause vs. BB_Credit

- Both mechanisms are used to avoid dropping frames
  - With different trade-offs
- The Pause mechanism requires at least the (2 x RTT x bandwidth) product on a link as buffer space
  - But allows Buffer handling in an arbitrary way
  - Well suited for networks with limited (bandwidth x delay) product (e.g. within the data center)
- The Pause frame is handled by the MAC layer
  - Similar to the R_RDY handling by the FC-1 level
- The BB_Credit mechanism prevents losing frames over any link
  - But links go under-utilized if link credits (& buffers) are < that needed for (RTT x BW)
  - Requires buffer handling in maximum frame size units

The Origins of a Data Center Fabric

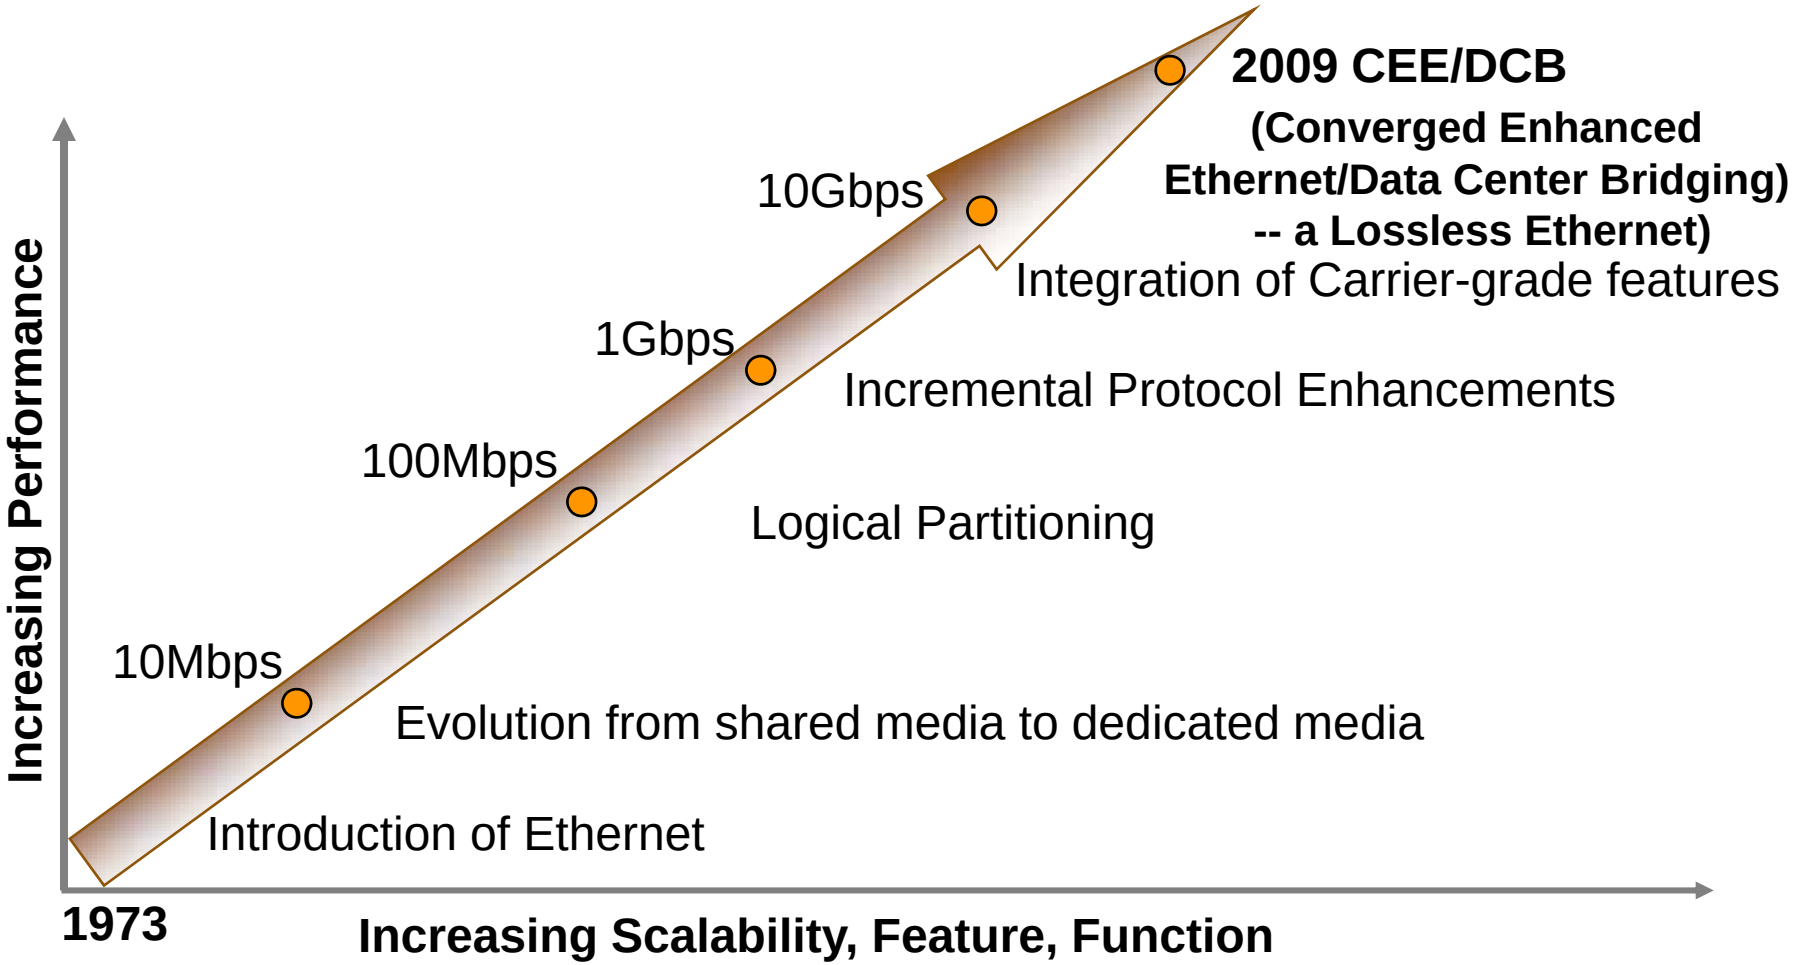

- *The technology has evolved continuously, showing a great ability to adapt to new technologies and increasing business requirements*