



Address Assignment Mechanism

How to select FCoE ENode MAC Addresses

Claudio DeSanti
Joe Pelissier
Landon Noll

Silvano Ce
JR Rivers
Ed Brannon

Robert Snively
John Hufferd
Suresh Vobbilisetty
Glenn Wenig
Anoop Ghanwani

T11/08-039v0

Work in Progress

Glossary (1)

- Server Provided MAC Address (SPMA)
 - Aligns with traditional Ethernet addressing mechanisms
 - MAC address of VN_Port is “world-wide” unique
 - May be installed at manufacturing time
 - May be created using local registry (note that this is a “small world”)
- Fabric Provided MAC Address (FPMA)
(Formerly named Mapped MAC address)
 - Aligns with traditional Fibre Channel addressing mechanisms
 - MAC address of the VN_Port is “fabric-wide” unique
 - Address is assigned by FCF using fabric-wide registry
 - MAC address of the VN_Port is required to have:
 - low order 24 bits = N_Port_ID assigned by FLOGI
 - high order 24 bits = FC-MAP (constant on a Fabric)

Glossary (2)

- ENode MAC Address
 - The MAC address used by an ENode during FCoE Initialization Protocol (FIP)
- VN_Port MAC Address
 - The MAC address used by an ENode for a particular N_Port_ID during Fibre Channel operations using FCoE frames

How do we facilitate address selection for FCoE?

The mechanism should:

- Be a simple change to the proposed discovery method
- Allows ENodes to suggest the VN_Port MAC address (SPMA)
- Allows FCFs to assign the VN_Port MAC address (FPMA)
- Optionally allow detection of duplicate MAC addresses
 - See T11/07-692v0.
- Use a new Ethertype for FIP (FCoE Initialization Protocol)

Additional information

- General Principles:
 - FIP frames do not use the FCoE frame formats
 - FIP is used for protocol steps where snooping by intermediate switches may be useful, including:
 - FLOGI Exchange with Request S_ID = 0
 - FDISC Exchange with Request S_ID = 0
 - LOGO Exchange with D_ID or S_ID = F_Port Controller
- NPIV support capability
 - NPIV supported using an FDISC Exchange in FIP using normal FC conventions. VN_Port MAC Addresses are assigned as during FLOGI
 - An analogous function may be performed using an FLOGI Exchange in FIP with a different VN_Port_Name
- The ENode MAC address used during discovery may be used for subsequent FIP operations if desired
- The behavior of FCF to FCF discovery and initialization uses FIP and uses the FIP Discovery frame formats, but is otherwise unchanged
- FIP frame formats are described later in this presentation
- Recommended to use separate MAC Addresses for IP devices, VF_Port capable FCF-MACs, and VE_Port capable FCF-MACs

Overview of mechanism

Discovery is enhanced with addressing capabilities information

- FPMA and/or SPMA supported

Discovery is performed in FIP, similar to T11/07-572v1

- An ENode sends a multicast solicitation, announcing its addressing capabilities
- VF_Port capable FCF-MACs supporting the ENode's capabilities replies with a unicast advertisement, announcing its addressing capabilities as well
- VF_Port capable FCF-MACs not supporting the ENode's capabilities do not reply

FLOGI is performed in FIP, according to the exchanged addressing capabilities

– VN_Port MAC Address for SPMA

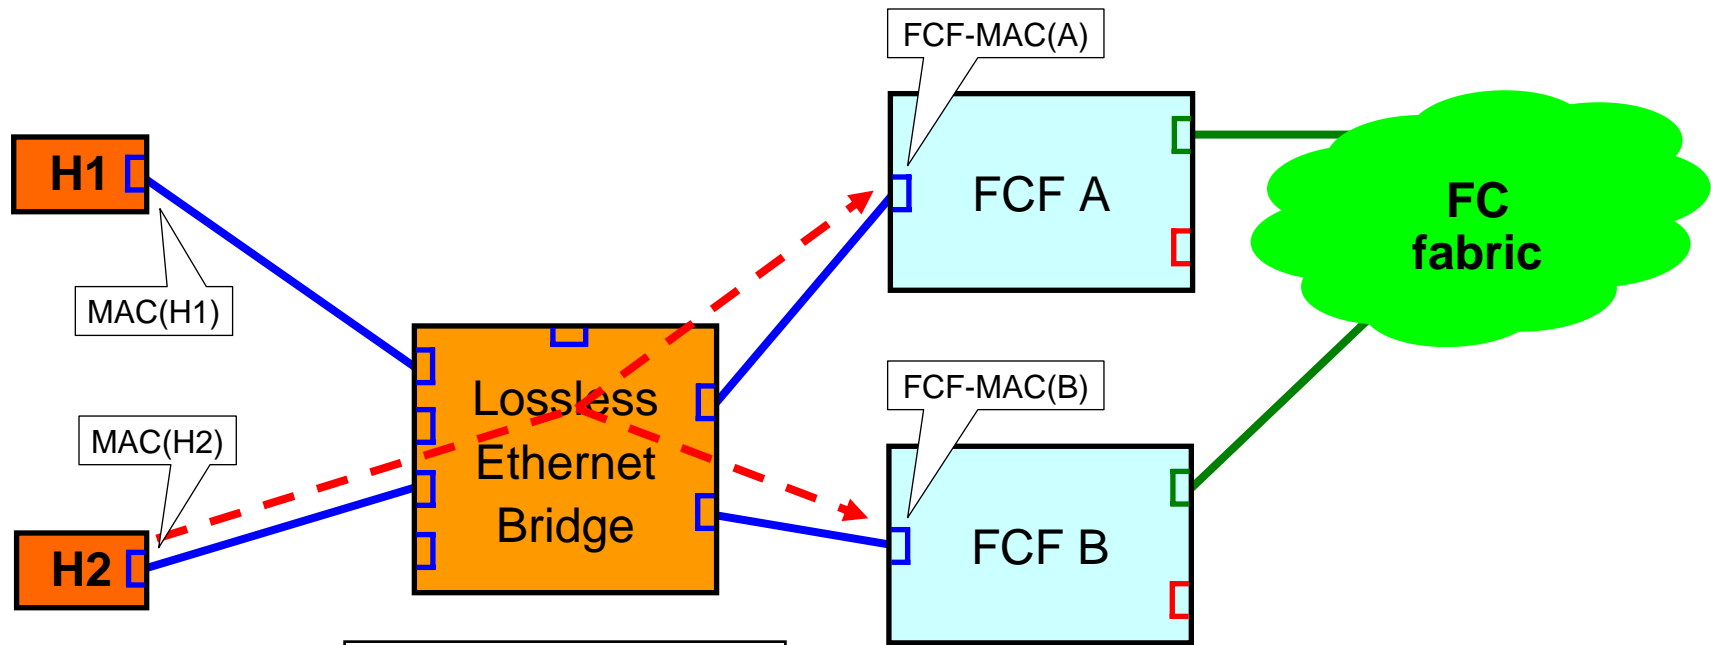
- FLOGI Response in FIP frame provides N_Port_ID information for the ENode
- ENode MAC address used during discovery is returned in the FIP FLOGI Response for use as the VN_Port MAC address for all FCoE frames for that N_Port_ID

– VN_Port MAC Address for FPMA

- FLOGI Response in FIP frame provides N_Port_ID information for the ENode
- The MAC address provided in FIP FLOGI Response frame is used as the VN_Port MAC address for all FCoE frames for that N_Port_ID

Subsequent Fibre Channel behavior uses Ethertype = FCoE frames on the VN_Port MAC address established during FLOGI

Multicast Solicitation from H2

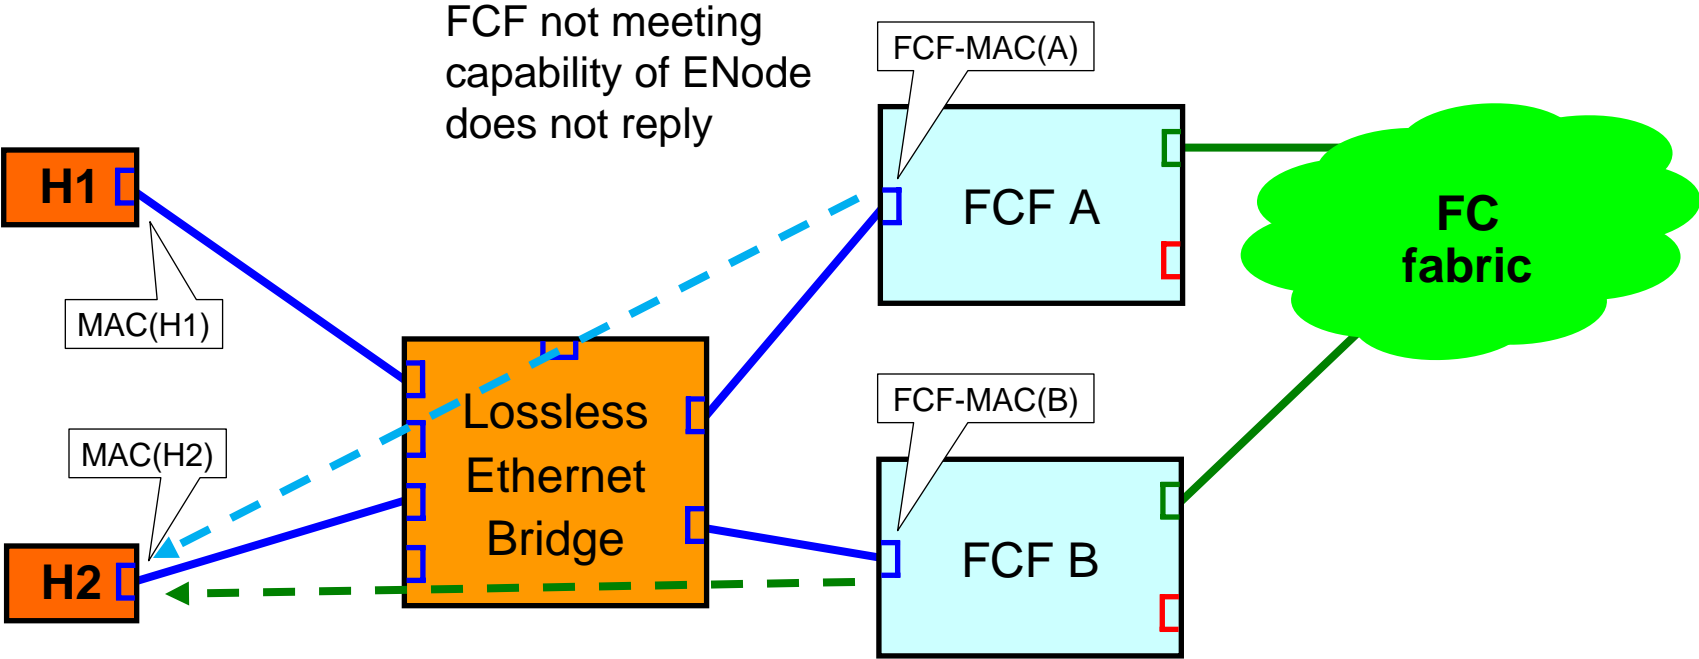


All-FCF-MACs
MAC(H2)
Solicitation (FIP)
[F=0, S=0, MAC(H2), Capability, Other]

Solicitation identifies VF_Port capable FCF-MACs with compatible addressing capabilities

Other parameters may include ENode's Port_Name for optional duplicate MAC address detection

Unicast Advertisements from A and B

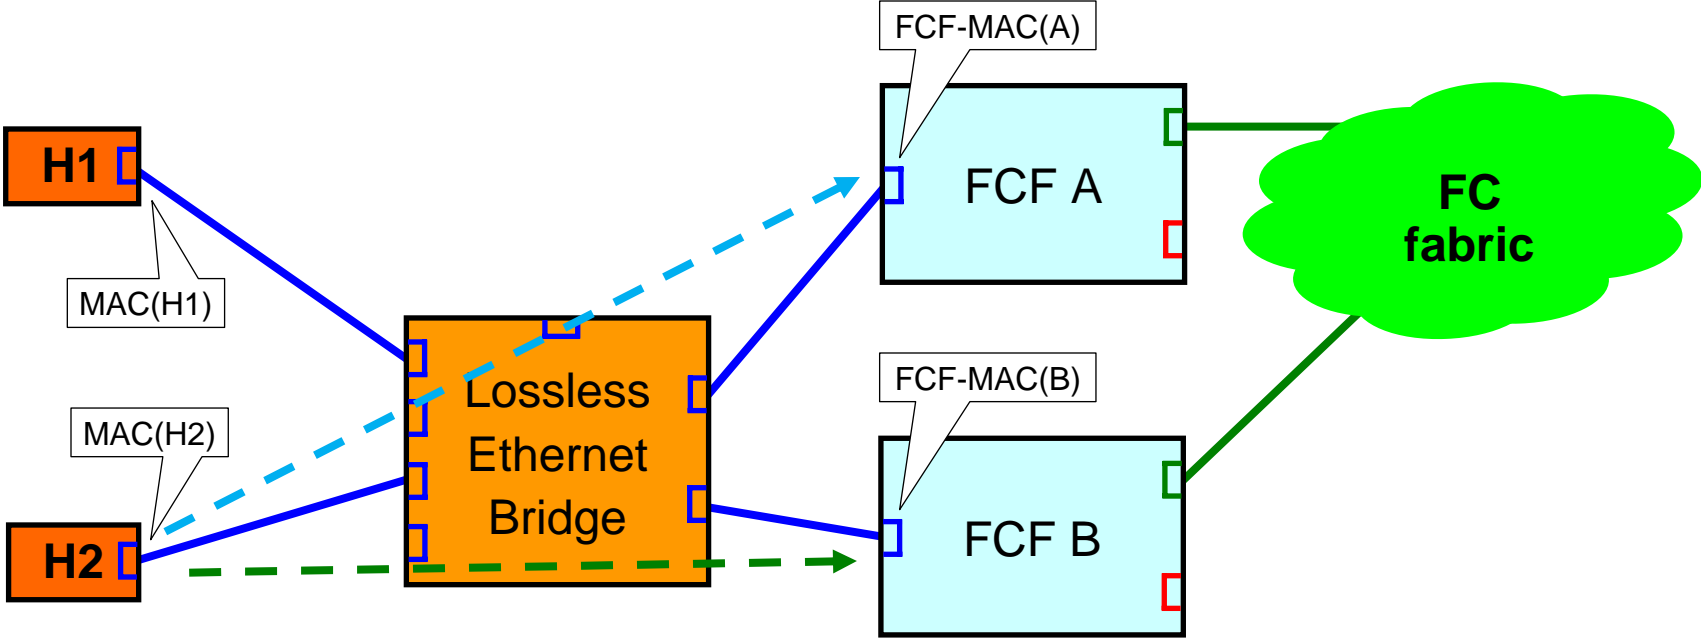


H2's FCF list:
 FCF-MAC(A) [J]
 FCF-MAC(B) [J]

MAC(H2)
FCF-MAC(A)
Mini-jumbo Advertisement (FIP)
[S=1, F=1, Priority, FC-MAP, FCF-MAC(A), Switch_Name, Fabric_Name, Capability, Other]

MAC(H2)
FCF-MAC(B)
Mini-jumbo Advertisement (FIP)
[S=1, F=1, Priority, FC-MAP, FCF-MAC(B), Switch_Name, Fabric_Name, Capability, Other]

FLOGI Request

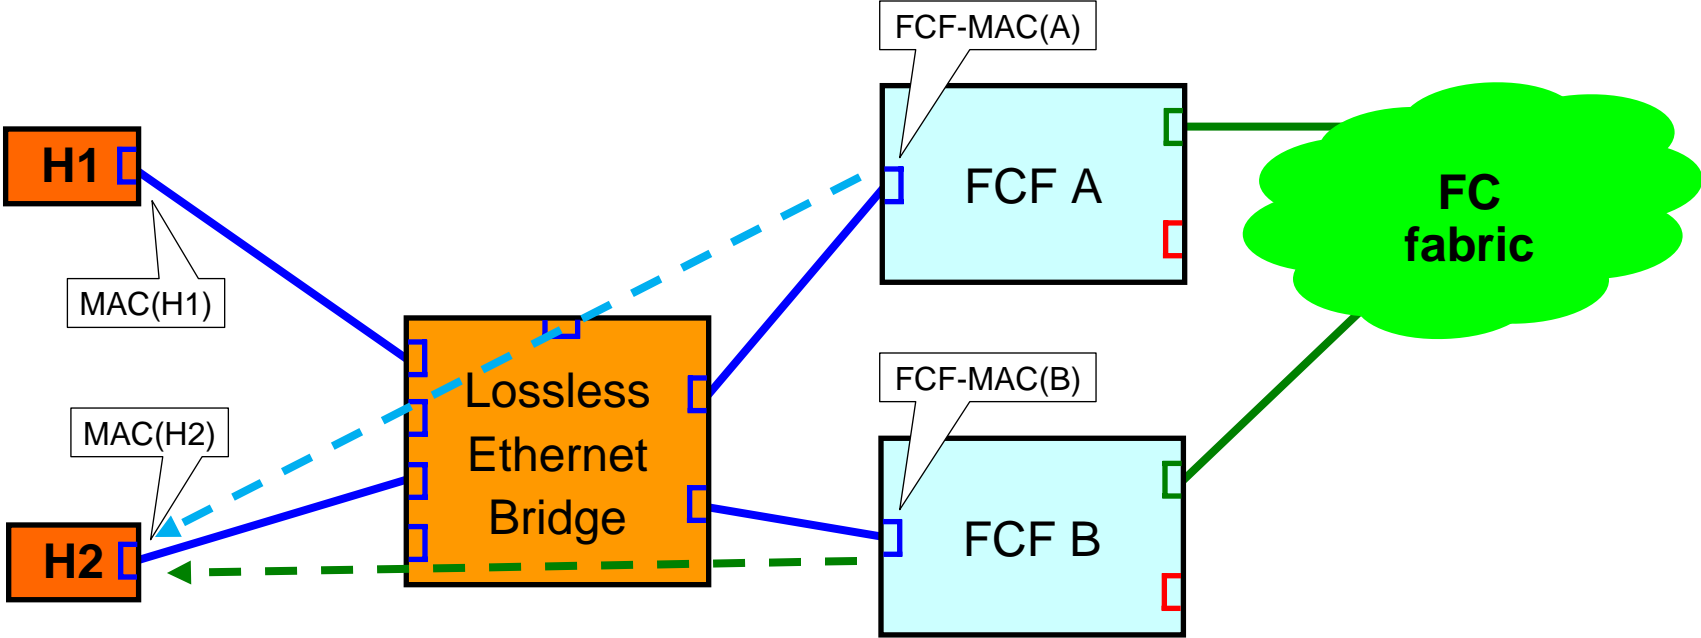


FCF-MAC(A)
MAC(H2)
FLOGI Request (FIP)
[FC Header, FLOGI data, Proposed MAC'(H2)]

FCF-MAC(B)
MAC(H2)
FLOGI Request (FIP)
[FC Header, FLOGI data, Proposed MAC''(H2)]

Capability agreed during discovery

FLOGI LS_ACC

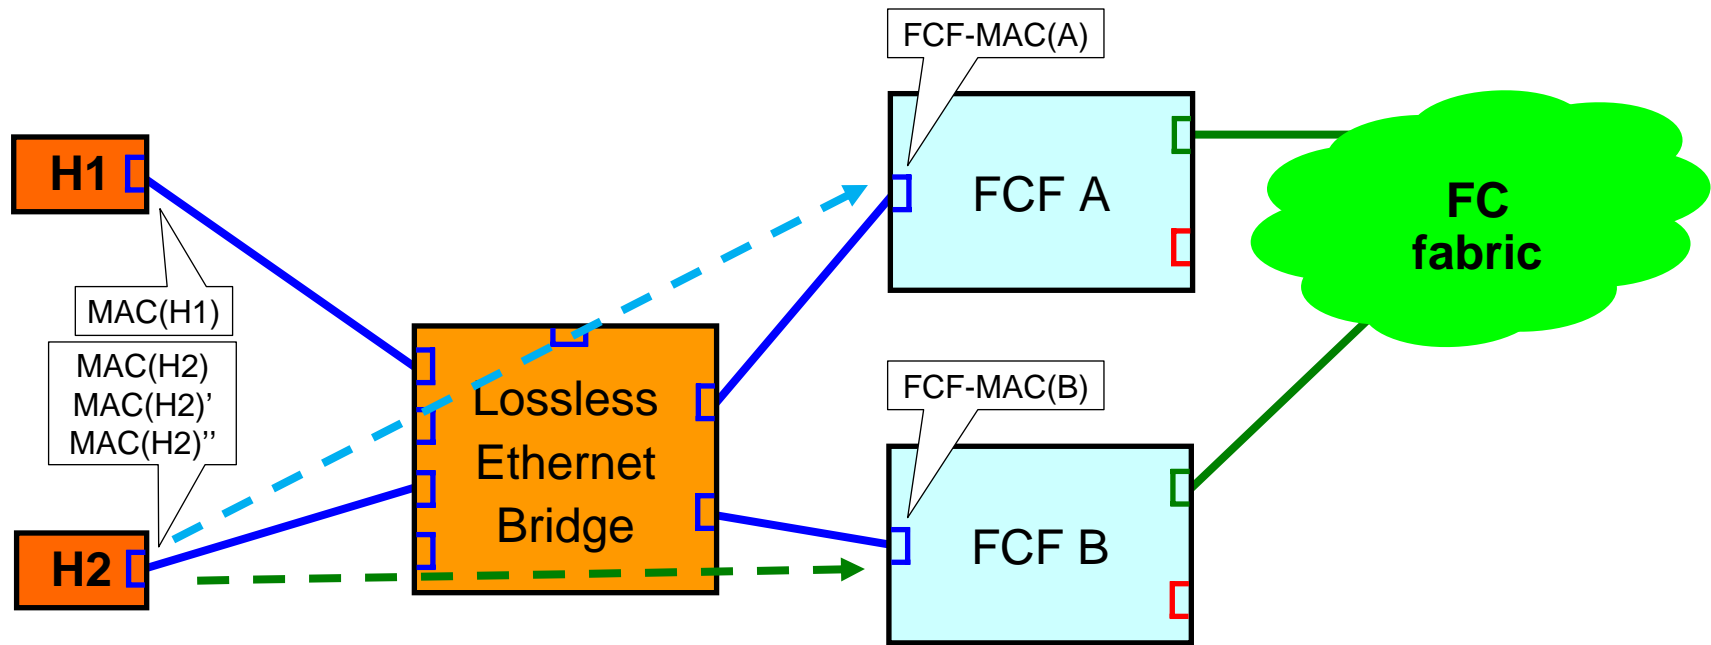


ENode uses MAC address in FIP FLOGI LS_ACC as the VN_Port MAC address for the N_Port_ID contained in the FLOGI data for subsequent FCoE frames

MAC(H2)
FCF-MAC(A)
FLOGI LS-ACC (FIP) [FC Header, LS_ACC data, Approved MAC(H2)']

MAC(H2)
FCF-MAC(B)
FLOGI LS-ACC (FIP) [FC Header, LS_ACC data, Approved MAC(H2)']

All subsequent FCoE frames use granted MAC address and assigned N_Port_ID



FCF-MAC(A)
MAC(H2)'
Fibre Channel Frame (FCoE) [FC SOF, FC Header, FC Data, FC CRC FC EOF]

FCF-MAC(B)
MAC(H2)''
Fibre Channel Frame (FCoE) [FC SOF, FC Header, FC Data, FC CRC, FC EOF]

FIP frames continue to use MAC(H2)
 For SPMA, MAC(H2)' = MAC(H2)'' = MAC(H2)
 For FPMA, MAC(H2)' and MAC(H2)'' use N_Port_IDs as low order 24 bits and FC-MAP for upper 24 bits

Subsequent FLOGI/FDISC

For all subsequent FLOGI or FDISC Exchanges intended to assign an additional N_Port_ID (i.e., with S_ID of Request = 0):

- FLOGI or FDISC is transmitted with SA = ENode MAC address
- Request and Reply Sequence use FIP format
- ENode and FCF assume SPMA or FPMA method as identified during discovery and used in the initial FLOGI
- LS_ACC Sequence provides the proper VN_Port MAC address and Fibre Channel N_Port_ID (D_ID of reply frame)
- If the Node_Name/N_Port_Name information carried in the FLOGI Request are the ones of an already logged in VN_Port, the original login session is replaced with a new login session

FDISC Exchanges intended to verify parameters for an existing login session use FCoE frames with the assigned VN_Port MAC address

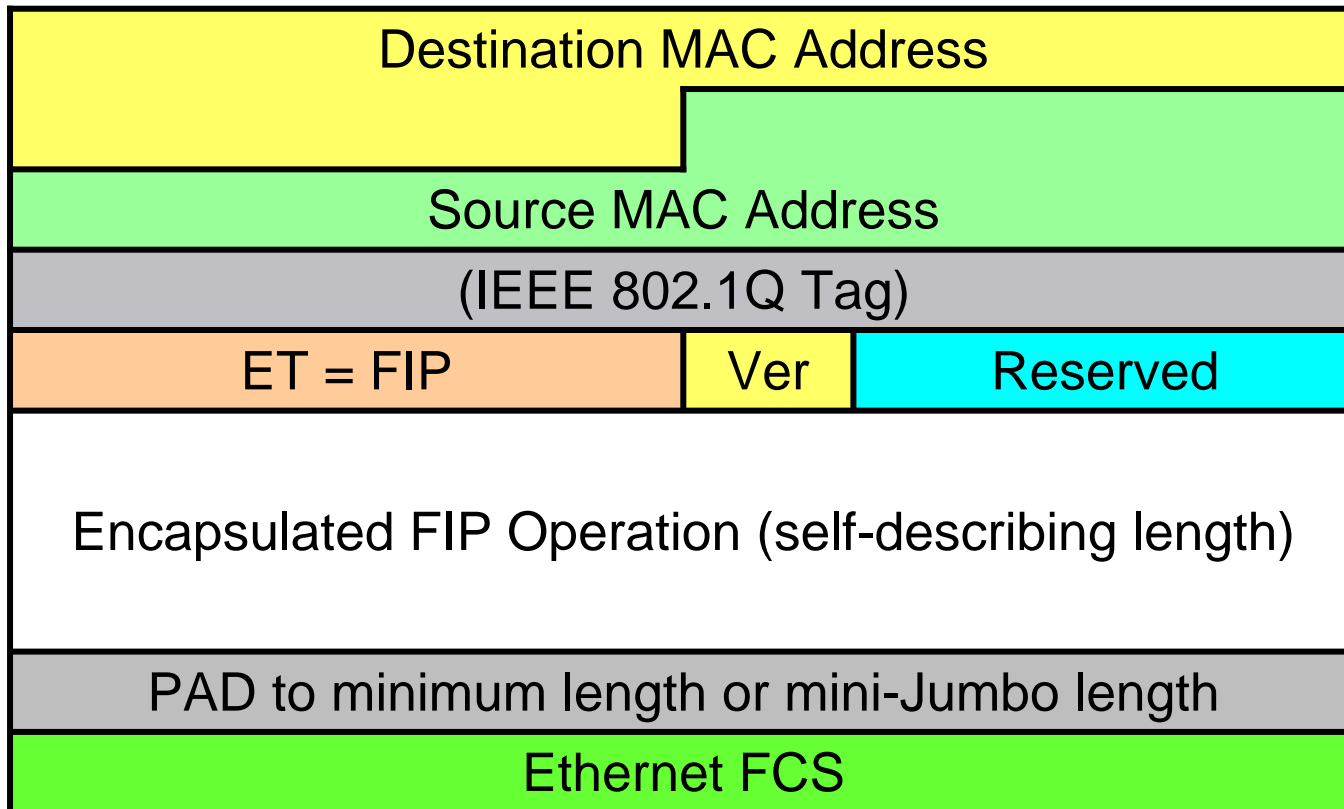
LOGO

For all LOGO Exchanges performing a logout involving a VF_Port (i.e., with either the Fibre Channel D_ID or S_ID = the F_Port Controller address of FFFFFFFEh):

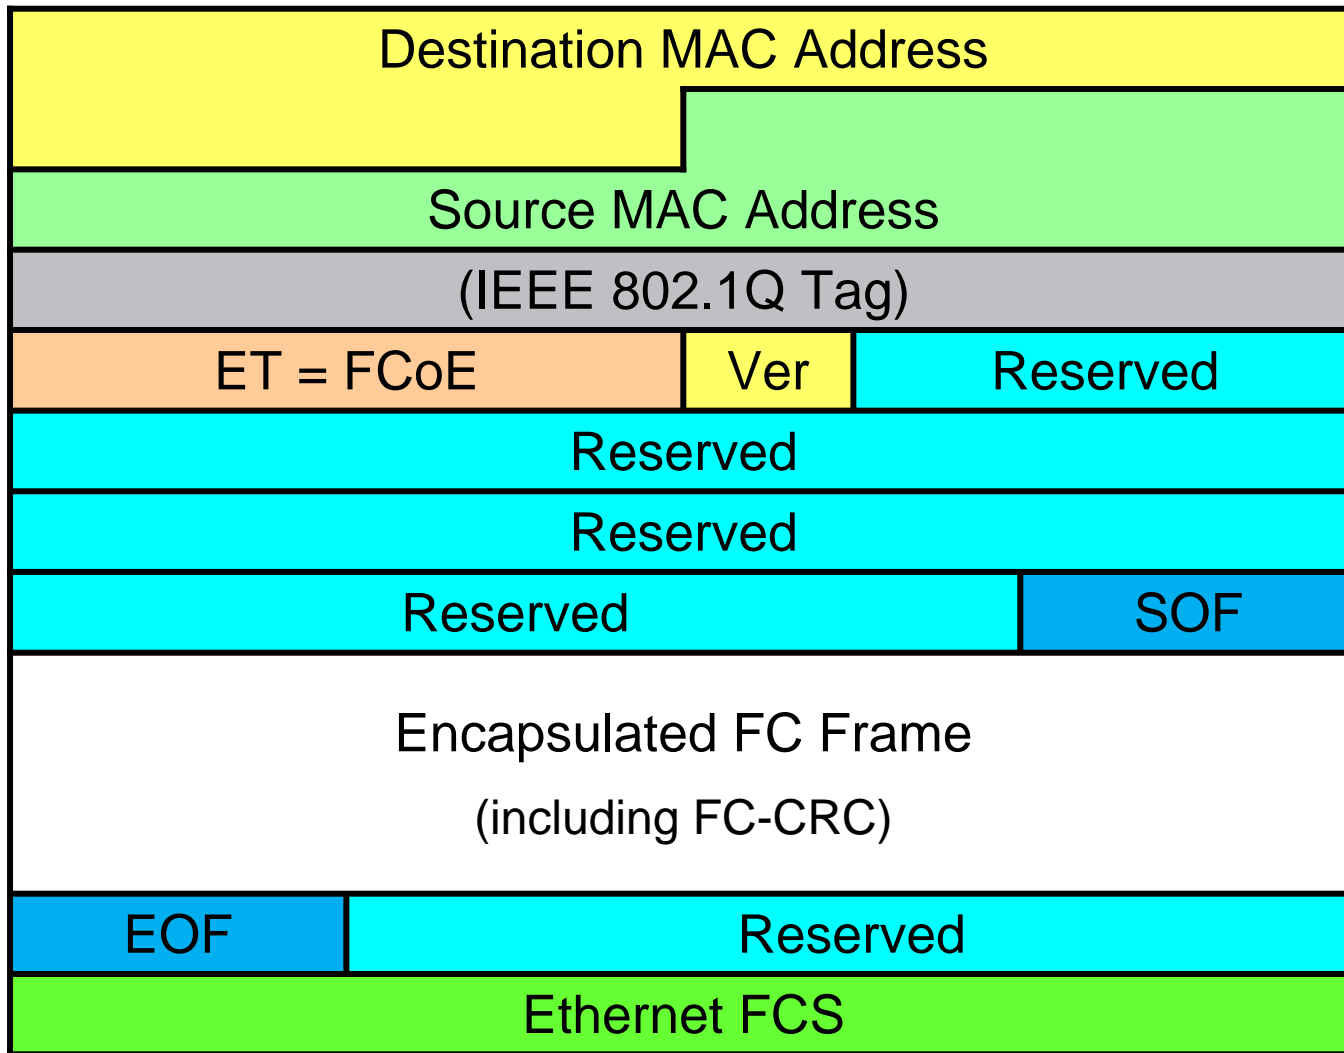
- Request and Reply Sequence use the FIP format
- The ENode use the ENode MAC Address

LOGO Exchanges involving any other pair of ports use FCoE frames with the assigned VN_Port MAC address

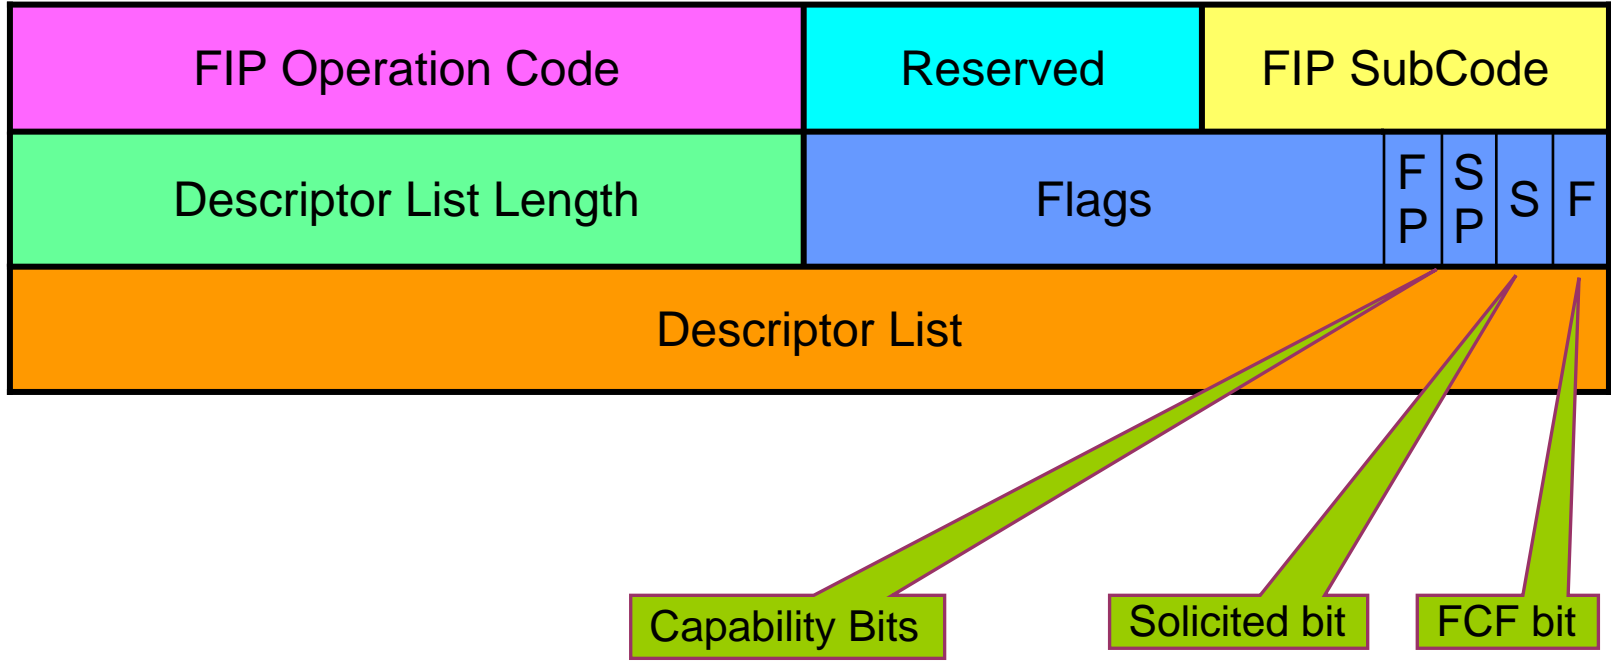
FIP Frame: Contains FIP Operation



FCoE Frame: Contains standard FC Frame







FIP Operation Format



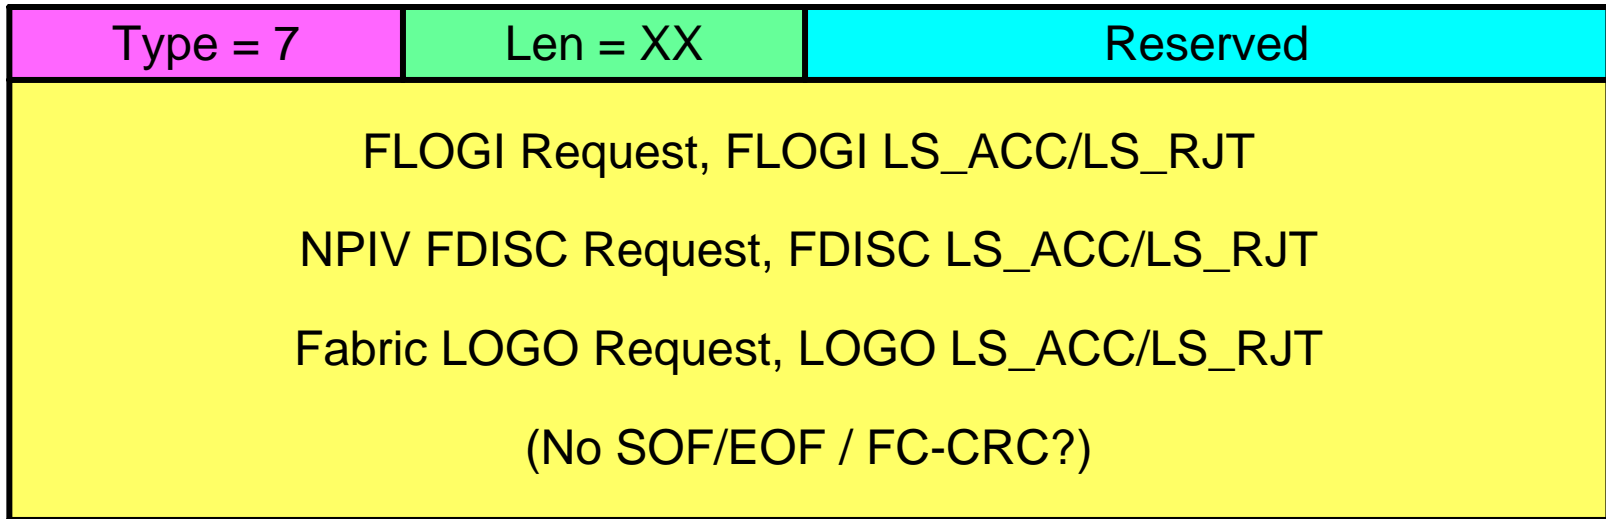

FIP Operation Codes and FIP SubCodes

FIP Operation Code	FIP SubCode	Operation
00000001h	01h	Discovery, Solicitation
	02h	Discovery, Advertisement
00000002h	01h	FLOGI/FDISC/LOGO, Request
	02h	FLOGI/FDISC/LOGO, Reply
All others	All others	Reserved

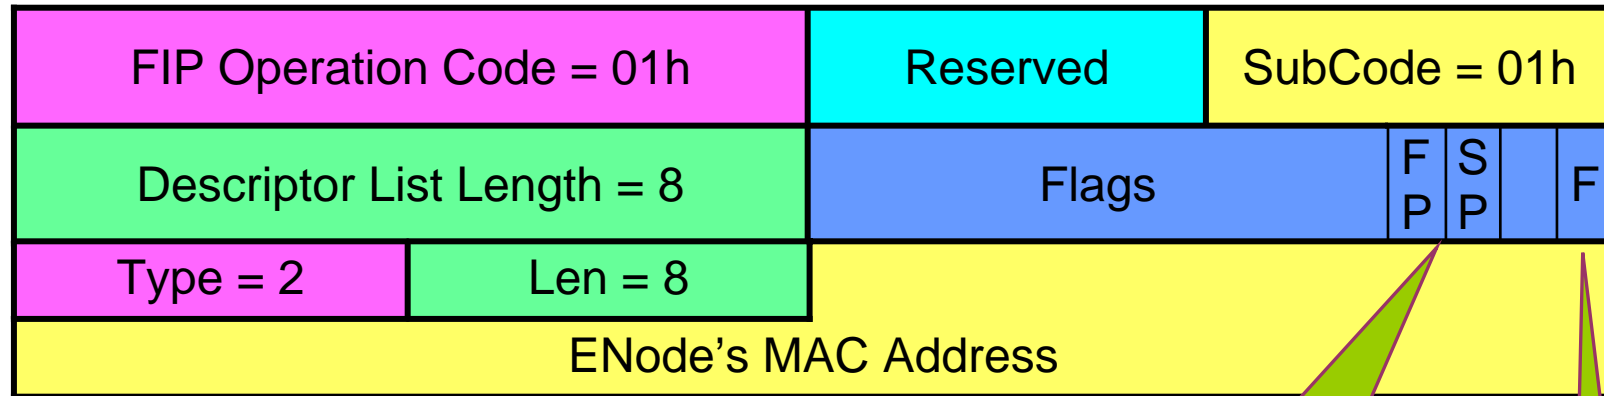
FIP Descriptors (1)



FIP Descriptors (2)



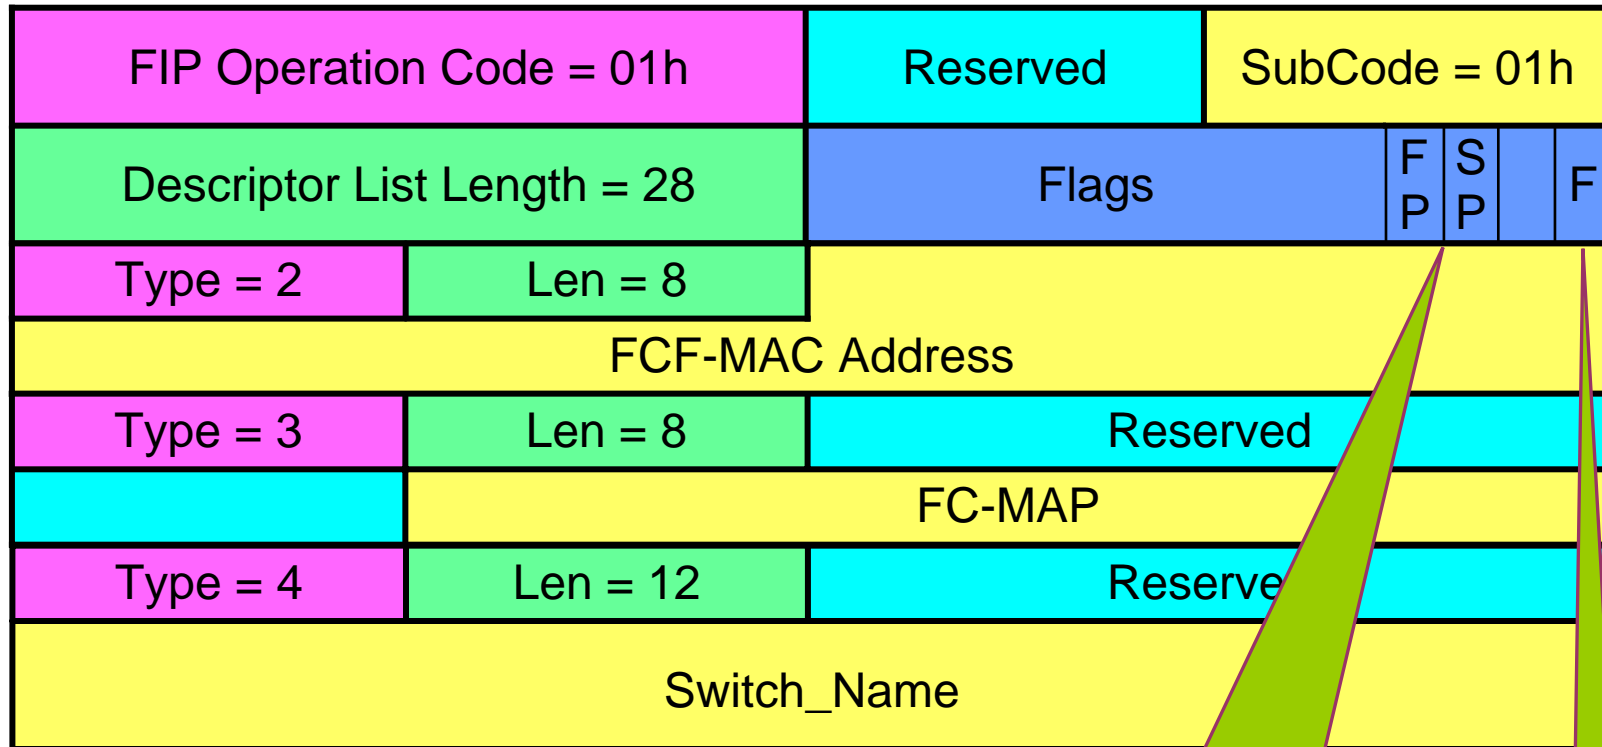
Solicitation from ENode



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

F = 0b

Solicitation from FCF



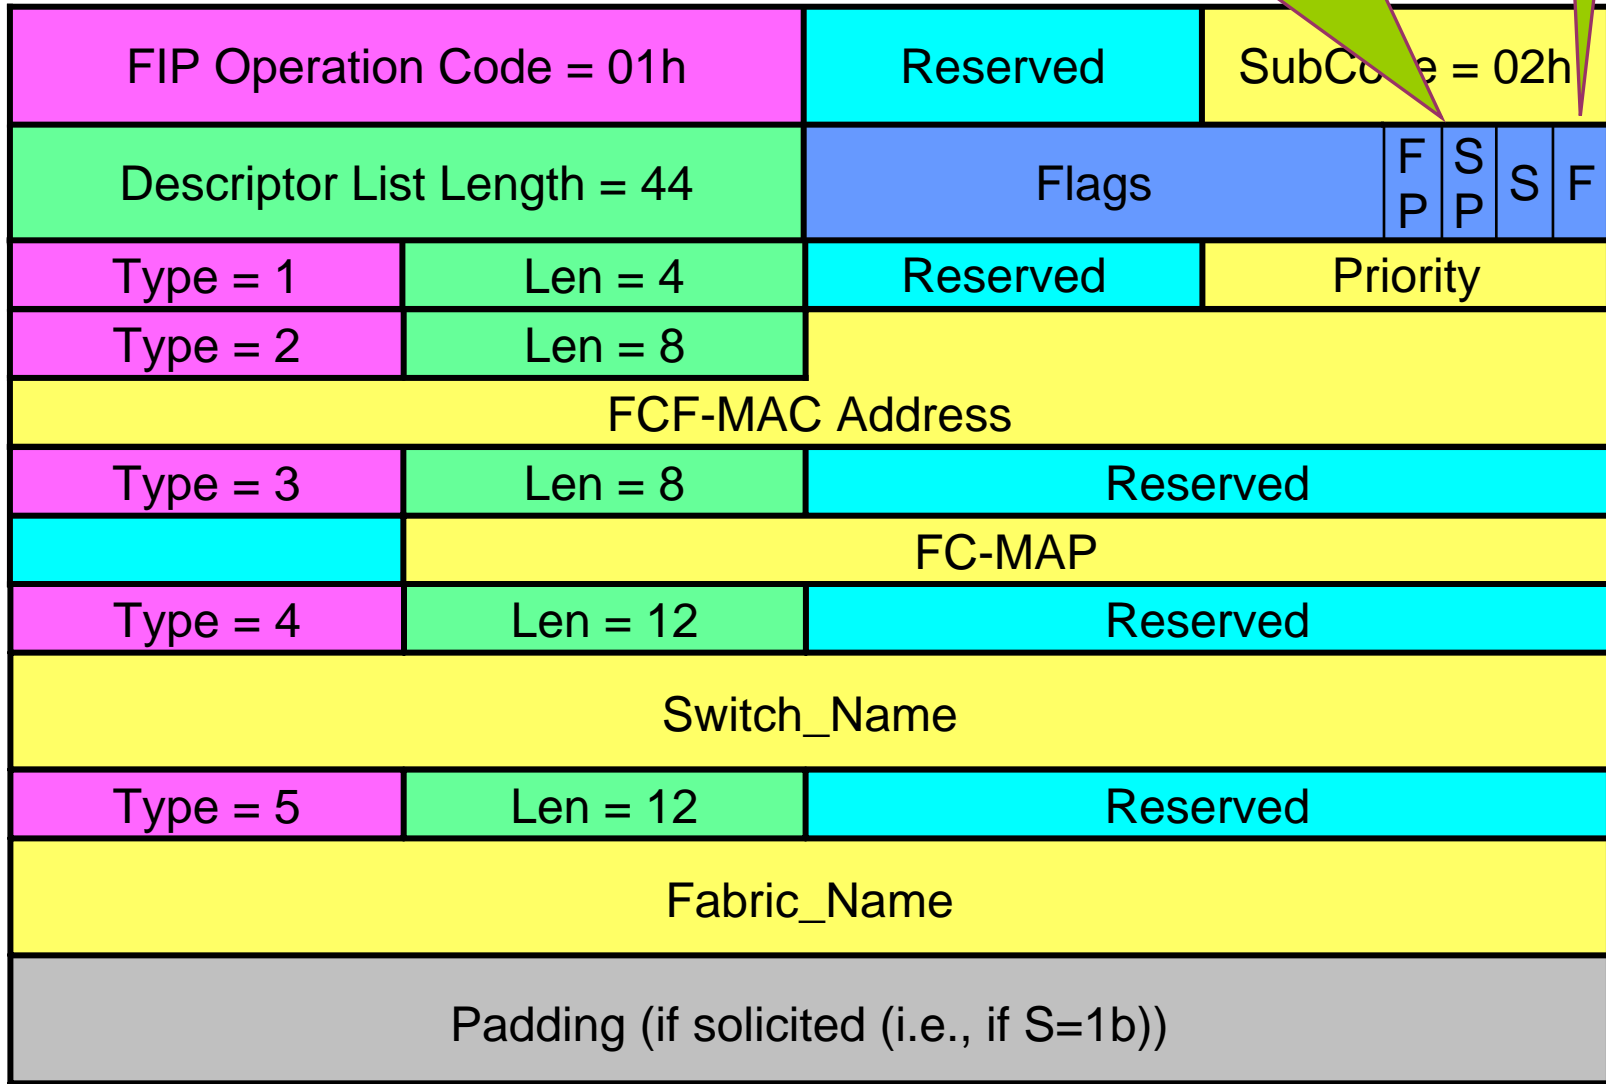
FCF Capabilities:
 FP = 1 if FPMA supported
 SP = 1 if SPMA supported

F = 1b

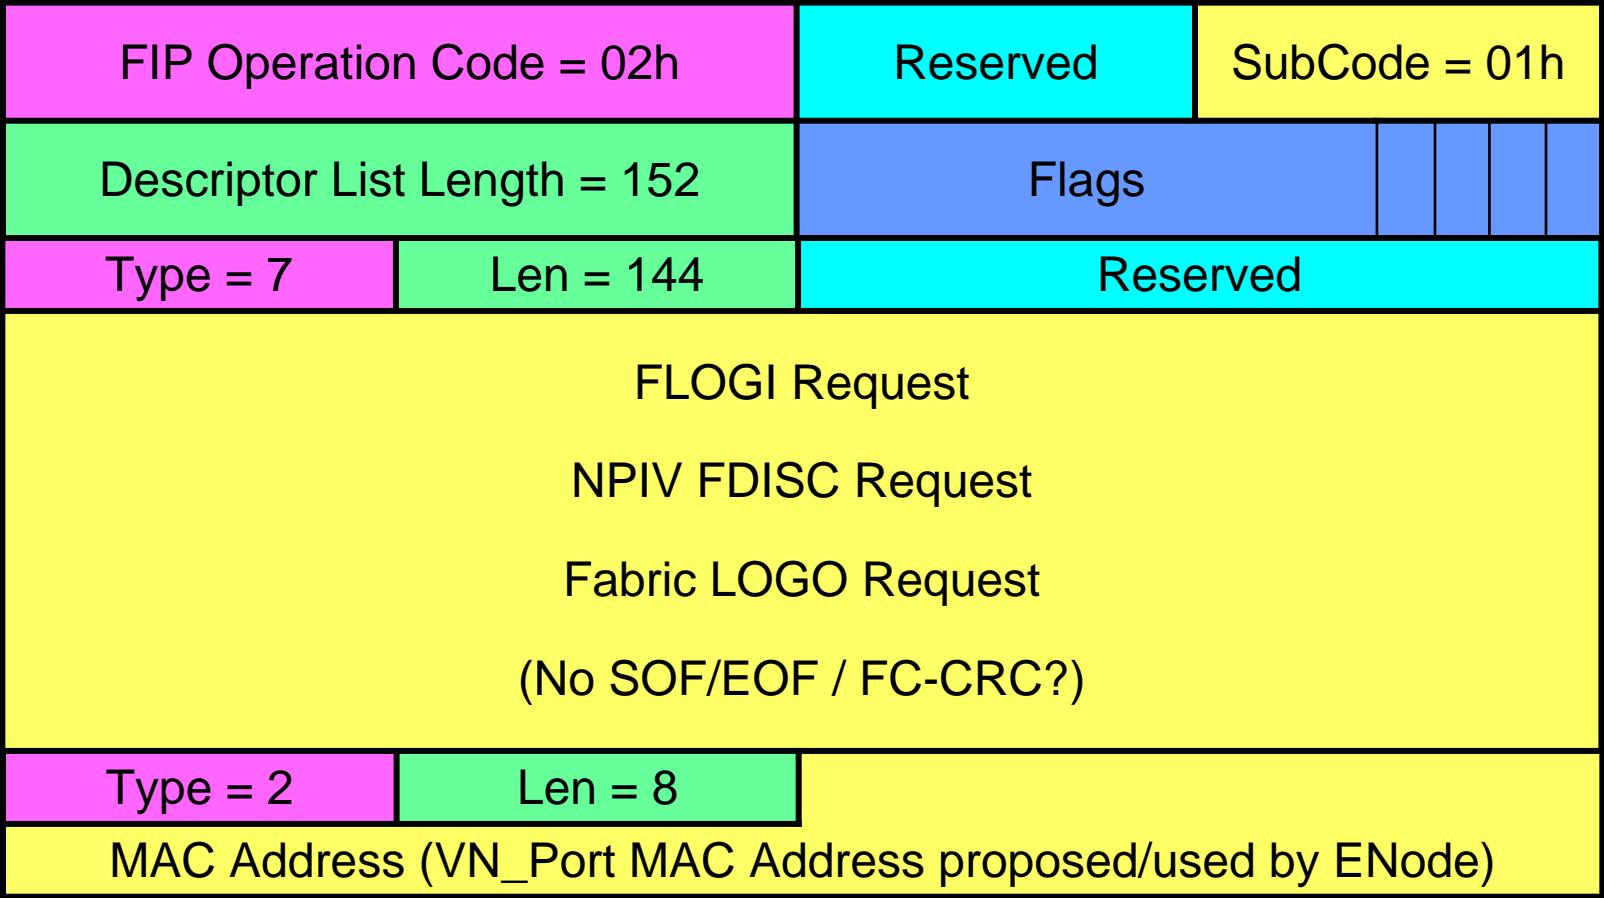
Advertisement

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

F = 1b



FIP FLOGI/FDISC/LOGO Request



FIP FLOGI/FDISC/LOGO LS_ACC

