

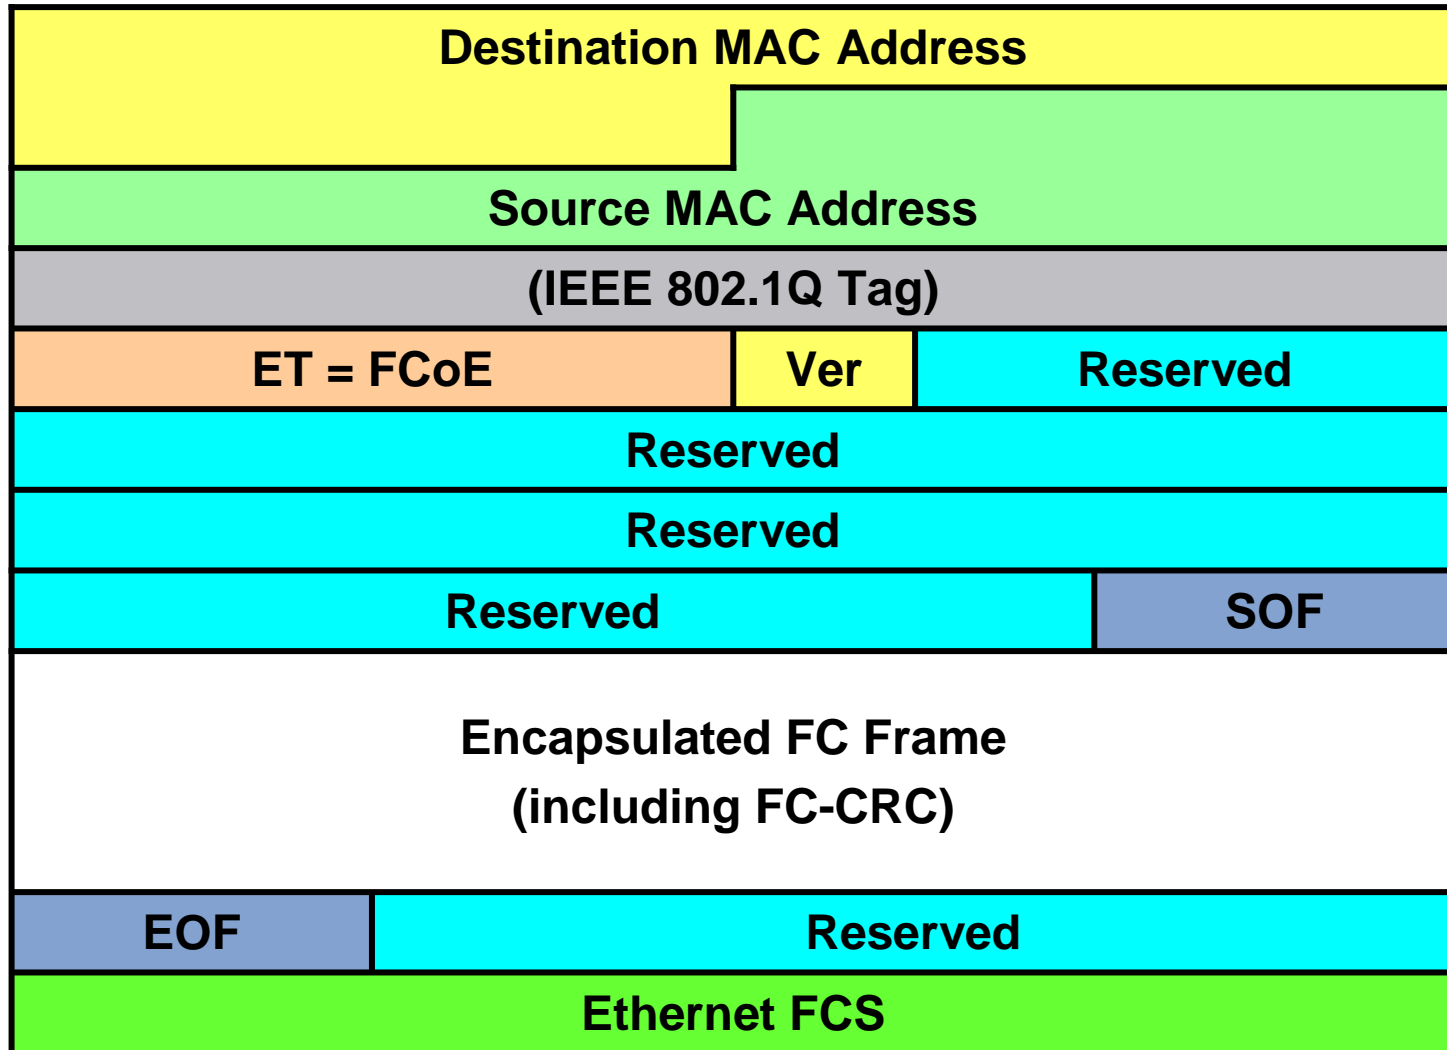
FCoE Initialization Protocol (FIP)

T11/08-208v1, April 2008

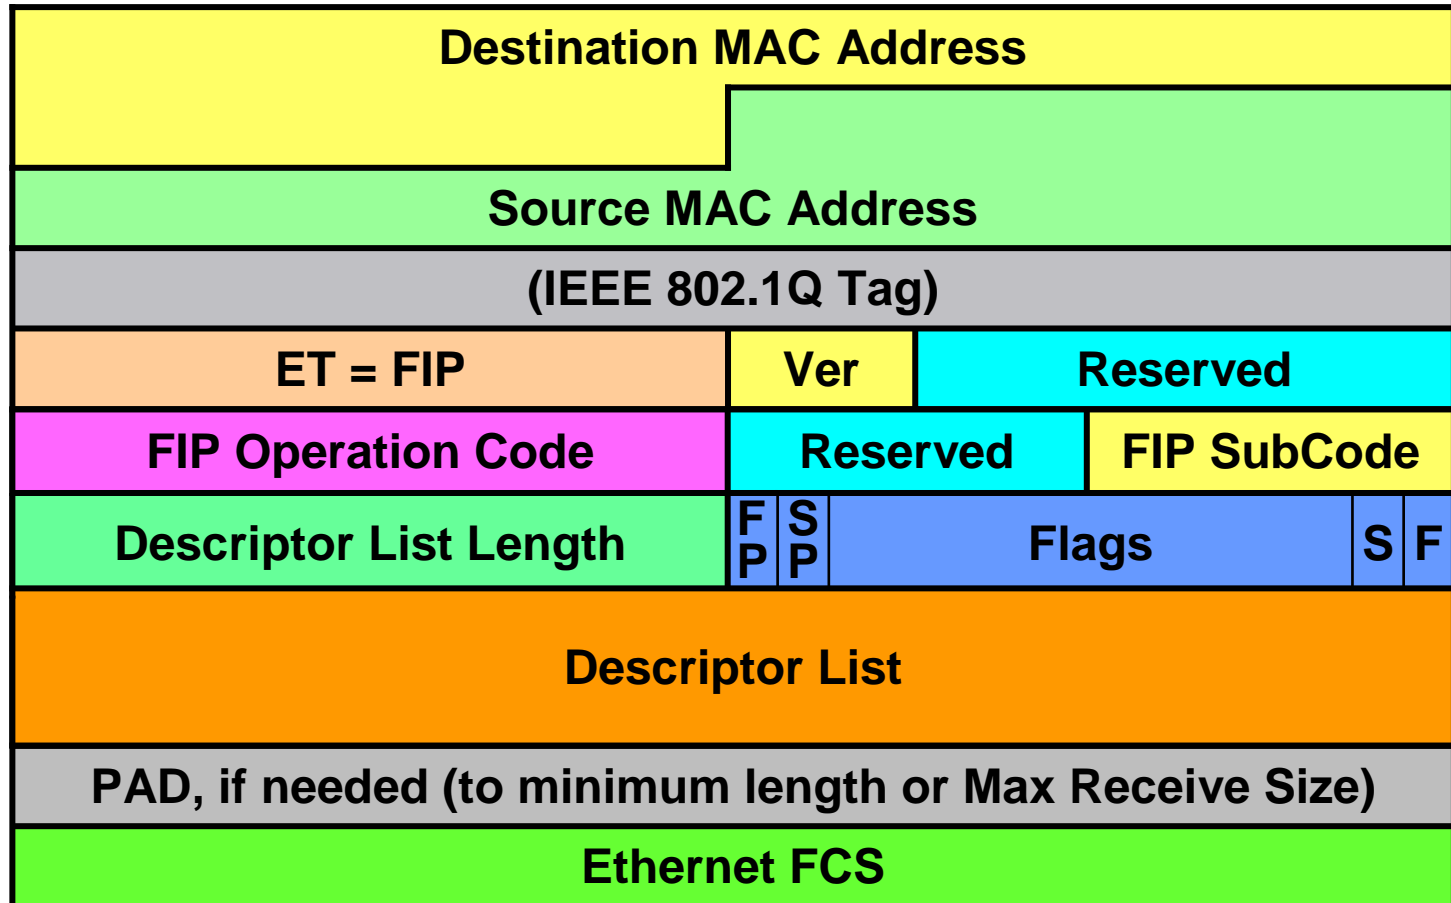
Claudio DeSanti, Cisco
Bob Nixon, Emulex
Silvano Gai, Nuova
Craig Carlson, QLogic
Ed McGlaughlin, QLogic

Joe Pelissier, Cisco
Bob Snively, Brocade
John Hufferd, Brocade
Steve Wilson, Brocade
Dave Peterson, Brocade

FCoE Frame Format



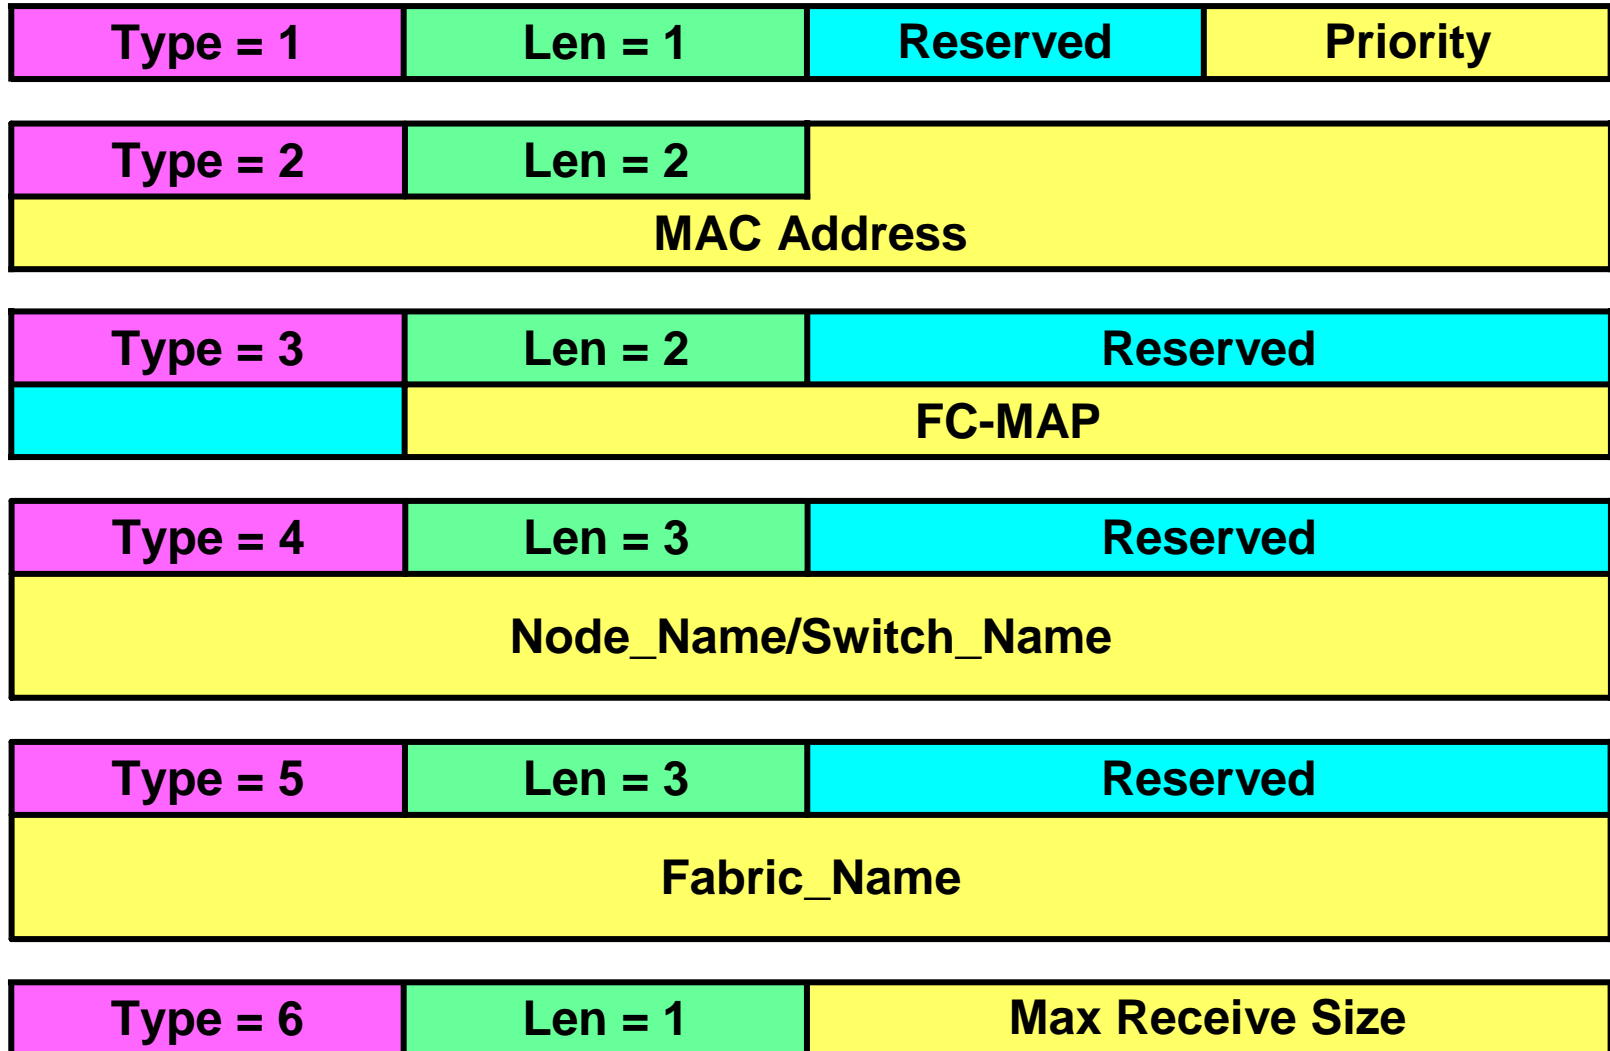
FIP Frame Format



Operation Codes and SubCodes

Operation Code	SubCode	Operation
0001h	01h	Discovery, Solicitation
	02h	Discovery, Advertisement
0002h	01h	FLOGI/FDISC/LOGO/ELP, Request
	02h	FLOGI/FDISC/LOGO/ELP, Reply
All others	All others	Reserved

FIP Descriptors (1)



Lengths are measured in 32-bit words

Max Receive Size Descriptor

- **Answers the question: “how big should be a jumbo FCoE Advertisement?”**
- **Used in Solicitations**
- **Allows an ENode MAC or a VE_Port capable FCF-MAC to communicate the Max Receive Size they support for FCoE**
 - The responding FCF-MACs use this information to pad solicited Advertisements**
- **Possible FC limitations use the FLOGI parameters and DO NOT USE the Max Receive Size for limiting length**
 - A limited Max Receive Size should be properly reflected in the FLOGI Request Payload**

FIP Descriptors (2)

Type = 7	Len = 36/9	Reserved
FLOGI Request, FLOGI LS_ACC/LS_RJT (no CRC, SOF, nor EOF)		

Type = 8	Len = 36/9	Reserved
NPIV FDISC Request, FDISC LS_ACC/LS_RJT (no CRC, SOF, nor EOF)		

Type = 9	Len = 11/8/9	Reserved
Fabric LOGO Request, LOGO LS_ACC/LS_RJT (no CRC, SOF, nor EOF)		

Type = 10	Len = 33/9	Reserved
ELP Request, ELP SW_ACC/SW_RJT (no CRC, SOF, nor EOF)		

No CRC, SOF, nor EOF?

- **Simplifications for FIP (to facilitate software processing):**

- Removal of FC-CRC**

- Removal of EOF**

- Assuming implicitly that it is always EOFt

- No support for multi frame Sequences in FIP

- Removal of SOF**

- FIP FLOGI is performed in Class 3

- Class 2 capability for FCoE is allowed by the FLOGI

- Class of Service bits

- FLOGI “profiling”**

- Do not support the FLOGI extension data

- **At least for the initial FIP protocol**

- Keep it simple!**

- **If more is needed, we will add later**

A Descriptor Value per ELS

- **Allows an easier processing of Accepts/Rejects**

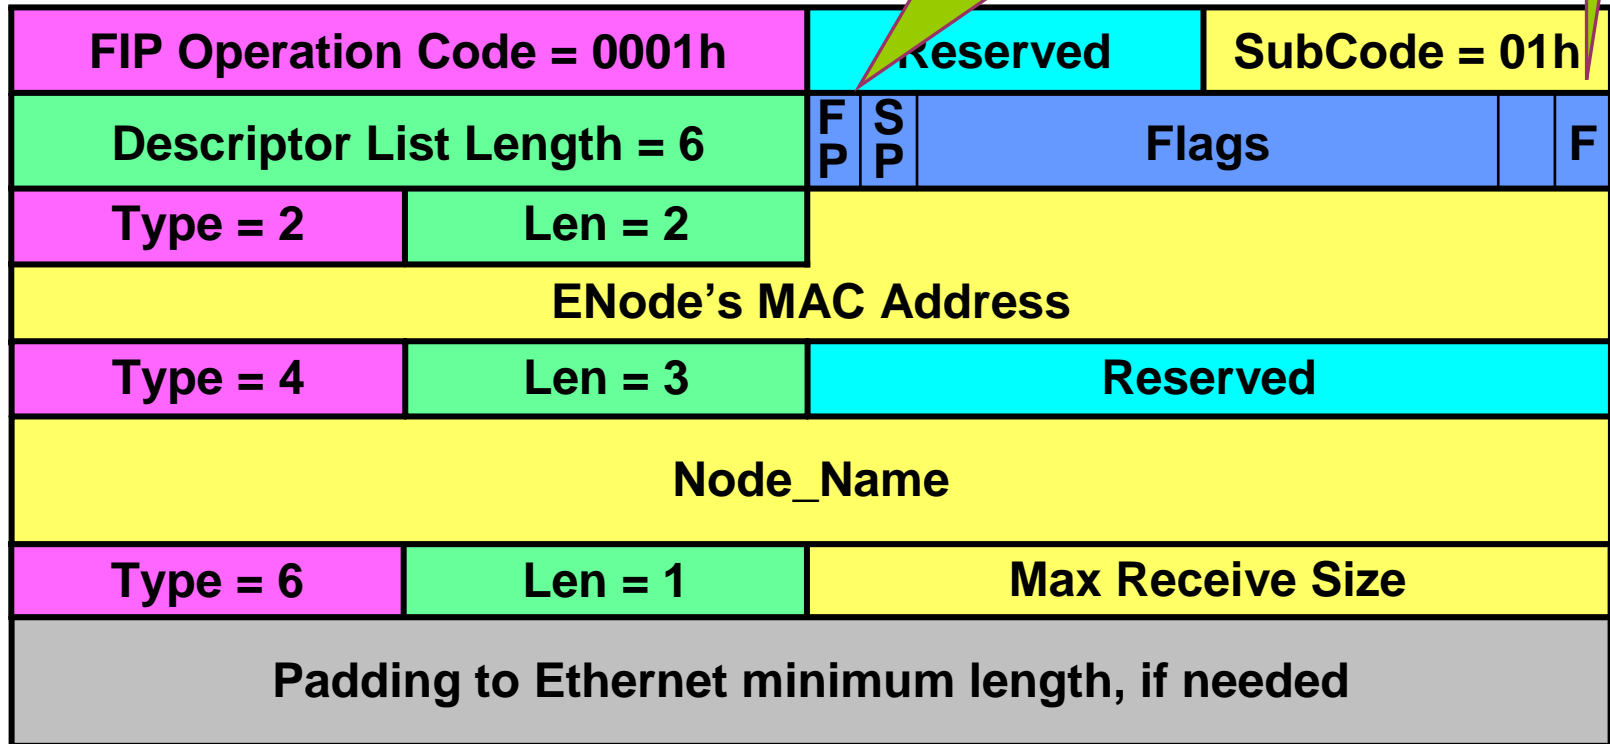
They have all the same opcode

- **Without a descriptor value per ELS the only way to process Accepts and Rejects by an intermediate Ethernet switch is by capturing the Request and keeping track of the OX_ID**
- **With a descriptor value per ELS the type value tells if the Accept/Reject is related to FLOGI, or FDISC, or LOGO, or ELP**

Solicitation from ENode

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

F = 0b



Solicits VF_Port capable FCF-MACs

Solicitation from FCF

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

F = 1b

FIP Operation Code = 0001h		Reserved	SubCode = 01h	
Descriptor List Length = 8		F P	SP	Flags
Type = 2	Len = 2	FCF-MAC Address		
Type = 3	Len = 2	Reserved		
		FC-MAP		
Type = 4	Len = 3	Reserved		
		Switch_Name		
Type = 6	Len = 1	Max Receive Size		
Padding to Ethernet minimum length, if needed				

Solicits VE_Port capable FCF-MACs

Advertisement

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

F = 1b

FIP Operation Code = 0001h		Reserved	SubCode = 02h
Descriptor List Length = 11		F P	S P
		Flags	
Type = 1	Len = 1	Reserved	Priority
Type = 2	Len = 2		
FCF-MAC Address			
Type = 3	Len = 2	Reserved	
FC-MAP			
Type = 4	Len = 3	Reserved	
Switch_Name			
Type = 5	Len = 3	Reserved	
Fabric_Name			
Padding to Max Receive Size of soliciting entity, if solicited (i.e., if S=1b), otherwise no padding			

Solicitations Generation Rules

- **Solicitations are sent:**

- By ENodes' MACs to solicit VF_Port capable FCF-MACs

- F=0b

- By VE_Port capable FCF-MACs to solicit other VE_Port capable FCF-MACs

- F=1b

- **Solicitations are sent to the multicast address 'All-FCF-MACs'**

Solicitations Processing Rules

- **The Switch_Name descriptor is used by an FCF to detect Solicitations coming from its own FCF-MACs**

If the value carried in the Switch_Name descriptor is the FCF's own Switch_Name, the FCF does not respond

- **The FC-MAP descriptor is used by an FCF to detect anomalous network configurations**

Because the FC-MAP shall be the same for FCFs belonging to the same network

If the value carried in the FC-MAP descriptor is not the FCF's own FC-MAP, the FCF does not respond to the Solicitation

An exception may be signaled

- **The FC-MAP descriptor may be ignored if set to zero**

Advertisements Generation Rules

- Advertisements are sent by FCF-MACs
- Advertisements may be Solicited (i.e., sent in response to a Solicitation) or unsolicited (i.e., sent periodically)
- Solicited Advertisements are unicast messages, have the Solicited bit set to one and are padded to the Max Receive Size specified in the Solicitation they are responding to
- Solicited Advertisements are sent to the MAC address carried in the MAC address descriptor contained in the Solicitation they are responding to
- Unsolicited Advertisements are multicast messages, have the Solicited bit set to zero and are not padded
- Unsolicited Advertisements are periodically sent:
 - to the multicast address 'All-ENode-MACs' by VF_Port capable FCF-MACs
 - to the multicast address 'All-FCF-MACs' by VE_Port capable FCF-MACs

Advertisements Processing Rules (1)

- **The FCoE Controllers of ENode's MACs and VE_Port capable FCF-MACs maintain an FCF list**

- **Each Advertisement creates an entry in the FCF list**

Solicited Advertisements (Jumbo) create entries immediately usable for connections

Unsolicited Advertisements create entries that need to be verified with a unicast Solicitation/Advertisement pair before being used for connections

- **An FCF having multiple FCF-MACs may respond to a multicast Solicitation with multiple unicast Advertisements**

Having the same value in the Switch_Name descriptor

The priority descriptor indicates which Advertisements are preferred to establish connections with the advertising FCF

They have higher priority in the FCF List

The MAC address carried in the MAC address descriptor is used to establish virtual links

Advertisements Processing Rules (2)

- **The FC-MAP descriptor is used by an ENode's MAC supporting Fabric Provided MAC Addresses to configure such addresses**
- **The FC-MAP descriptor is used by a VE_Port capable FCF-MAC to detect anomalous network configurations**

Because the FC-MAP shall be the same for FCFs belonging to the same Fabric

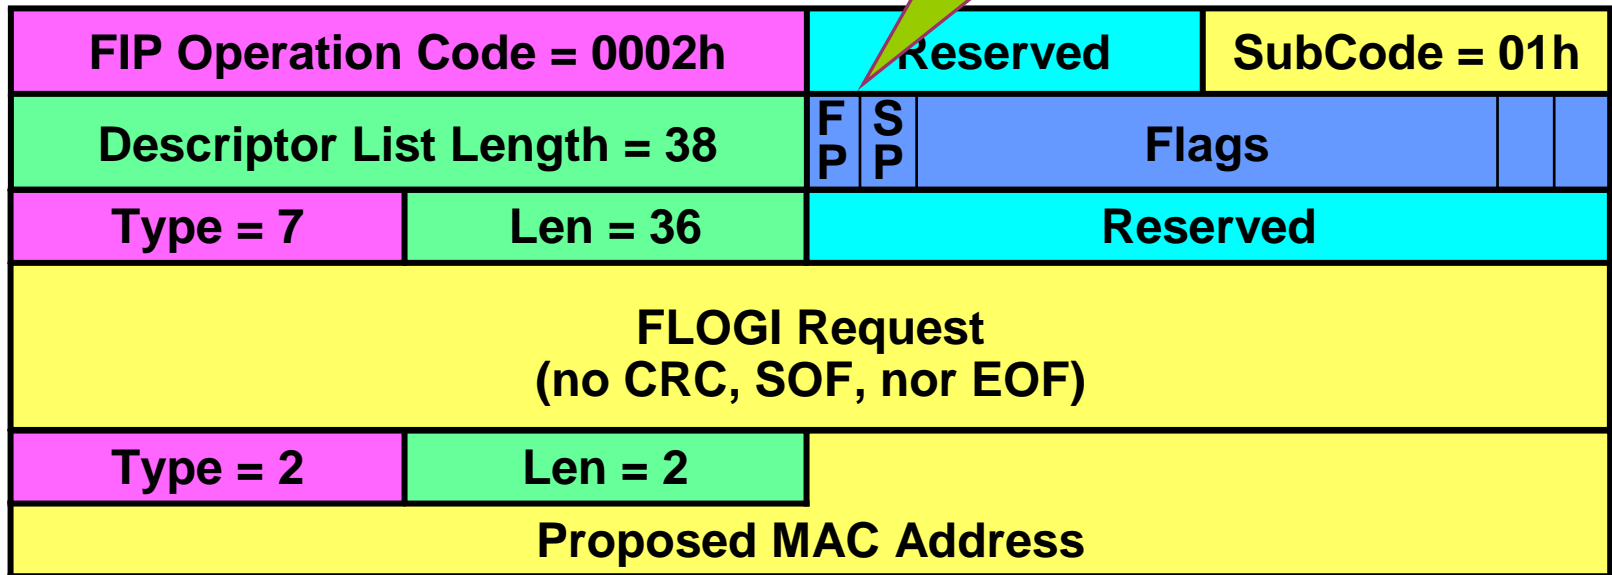
An FCF ignores an Advertisement if the value carried in the FC-MAP descriptor is not equal to its configured FC-MAP

An exception may be signaled

- **The FC-MAP descriptor may be ignored if set to zero**
- **An ENode's MAC should detect an anomalous configuration if:**
 - different Advertisements carry different FC-MAP values**
 - different Advertisements carry different Fabric_Names**
- **The Fabric_Name descriptor may be ignored by VE_Port capable FCF-MACs**


FIP FLOGI Request (ENode → FCF)

FP = 1 if FPMA requested, or
SP = 1 if SPMA requested

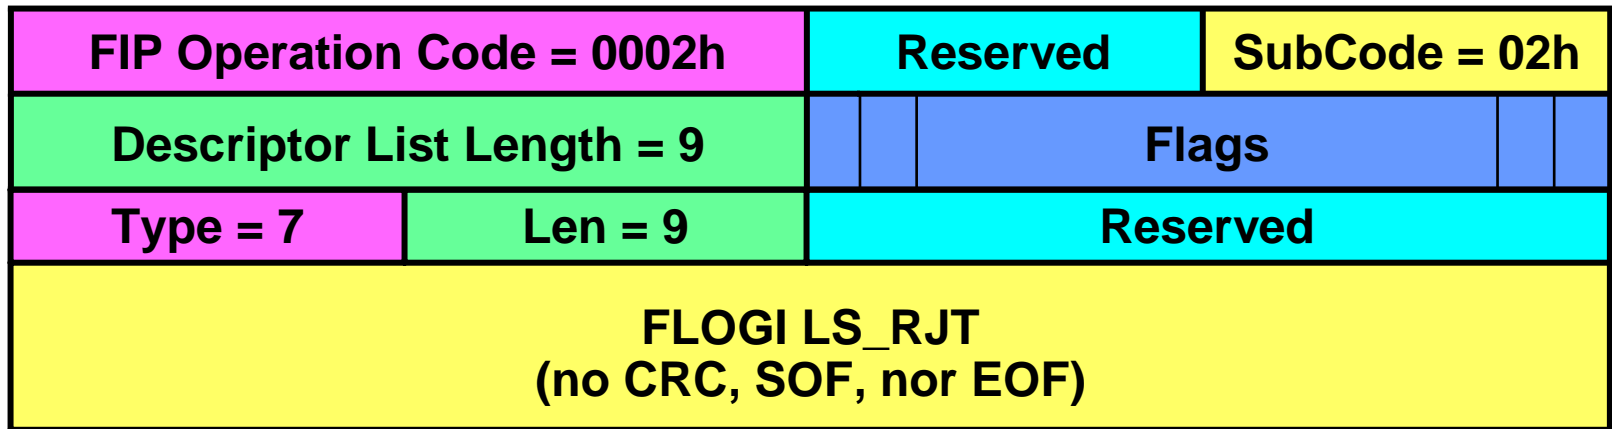


FIP FLOGI Accept (FCF → ENode)

FP = 1 if FPMA granted, or
SP = 1 if SPMA granted




FIP FLOGI Reject (FCF → ENode)



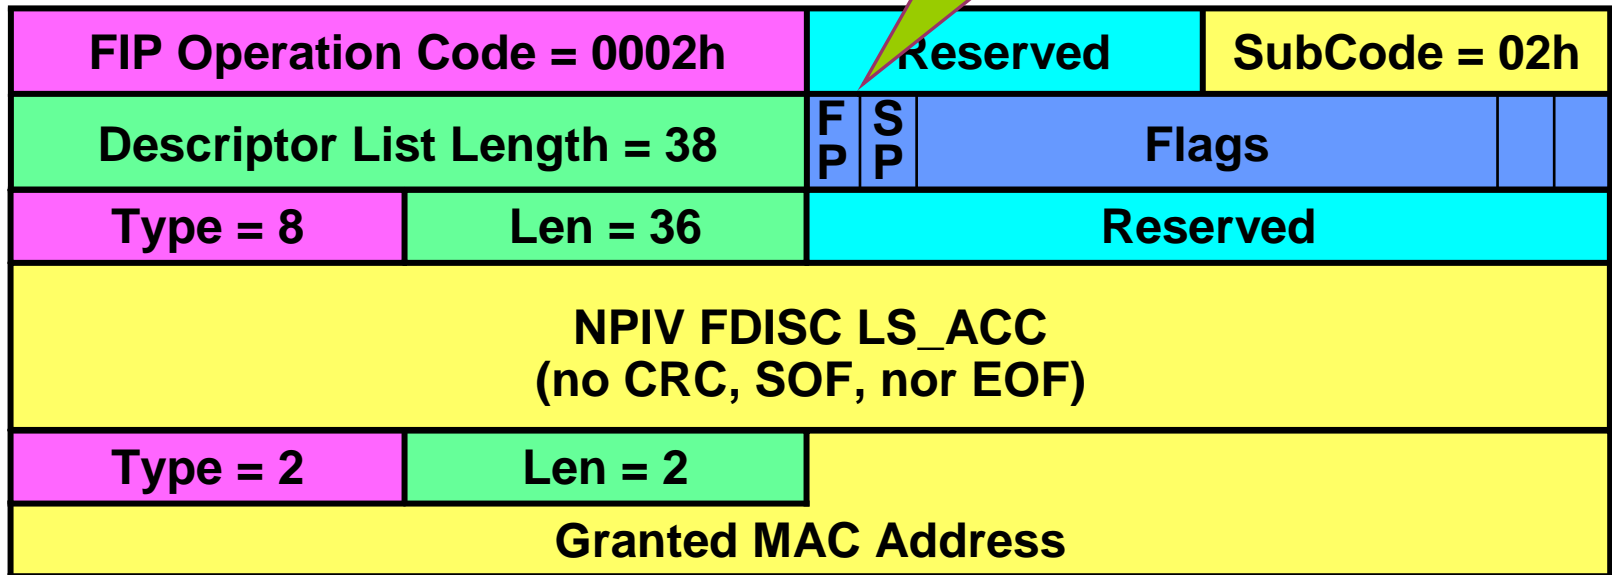
FIP NPIV FDISC Request (ENode → FCF)

FP = 1 if FPMA requested, or
SP = 1 if SPMA requested




FIP NPIV FDISC Accept (FCF → ENode)

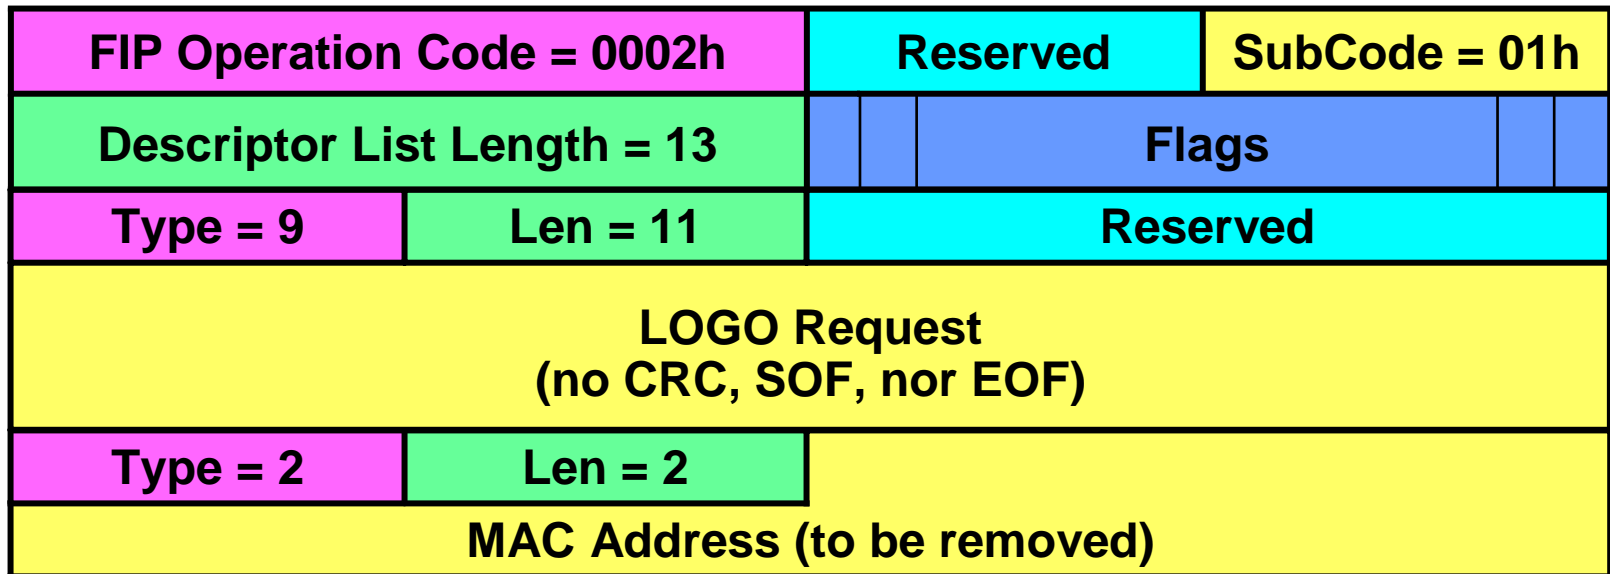
FP = 1 if FPMA granted, or
SP = 1 if SPMA granted



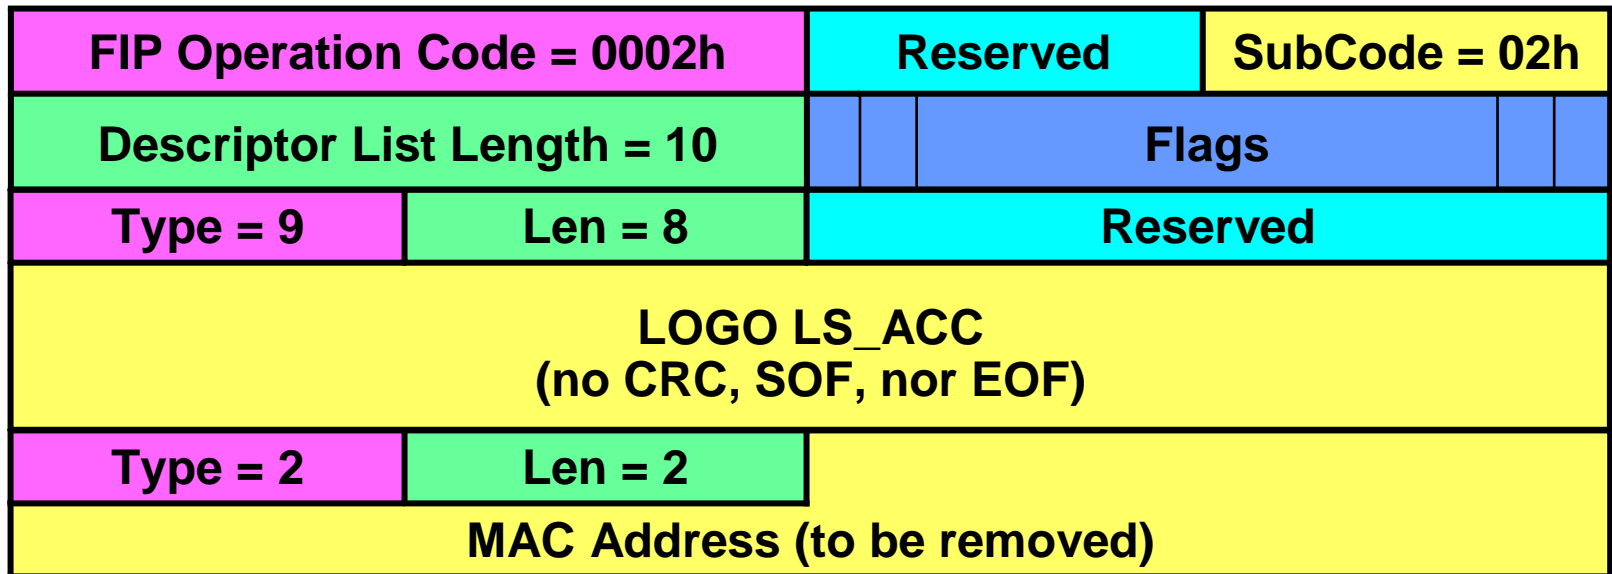
FIP NPIV FDISC Reject (FCF → ENode)



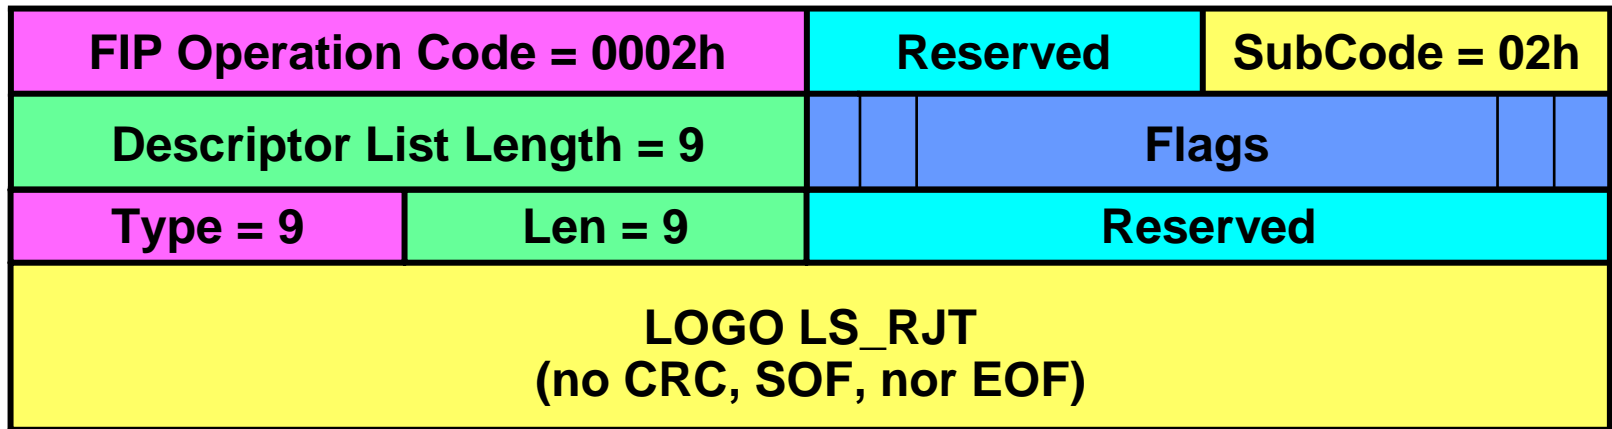
FIP Fabric LOGO Request (ENode \leftrightarrow FCF)




FIP Fabric LOGO Accept (ENode \leftrightarrow FCF)




FIP Fabric LOGO Reject (ENode \leftrightarrow FCF)



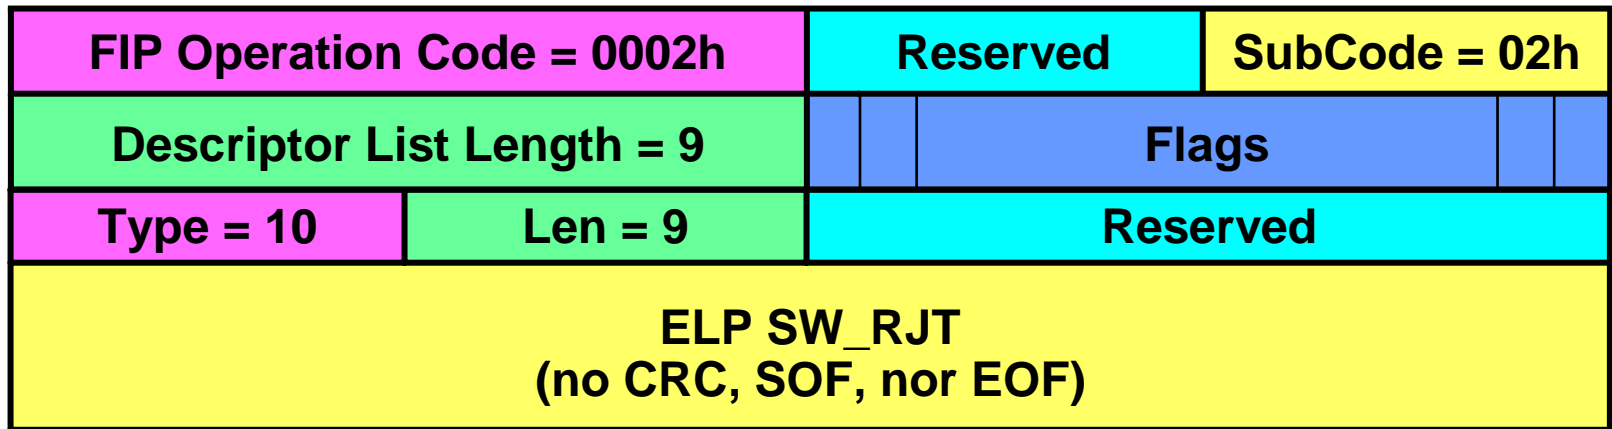
FIP ELP Request (FCF \leftrightarrow FCF)



FIP ELP Accept (FCF \leftrightarrow FCF)



FIP ELP Reject (FCF \leftrightarrow FCF)



Motion

- **To accept the FIP protocol specified in this document as the basis for the FIP protocol in FC-BB-5**

Detection of Loss of Connectivity

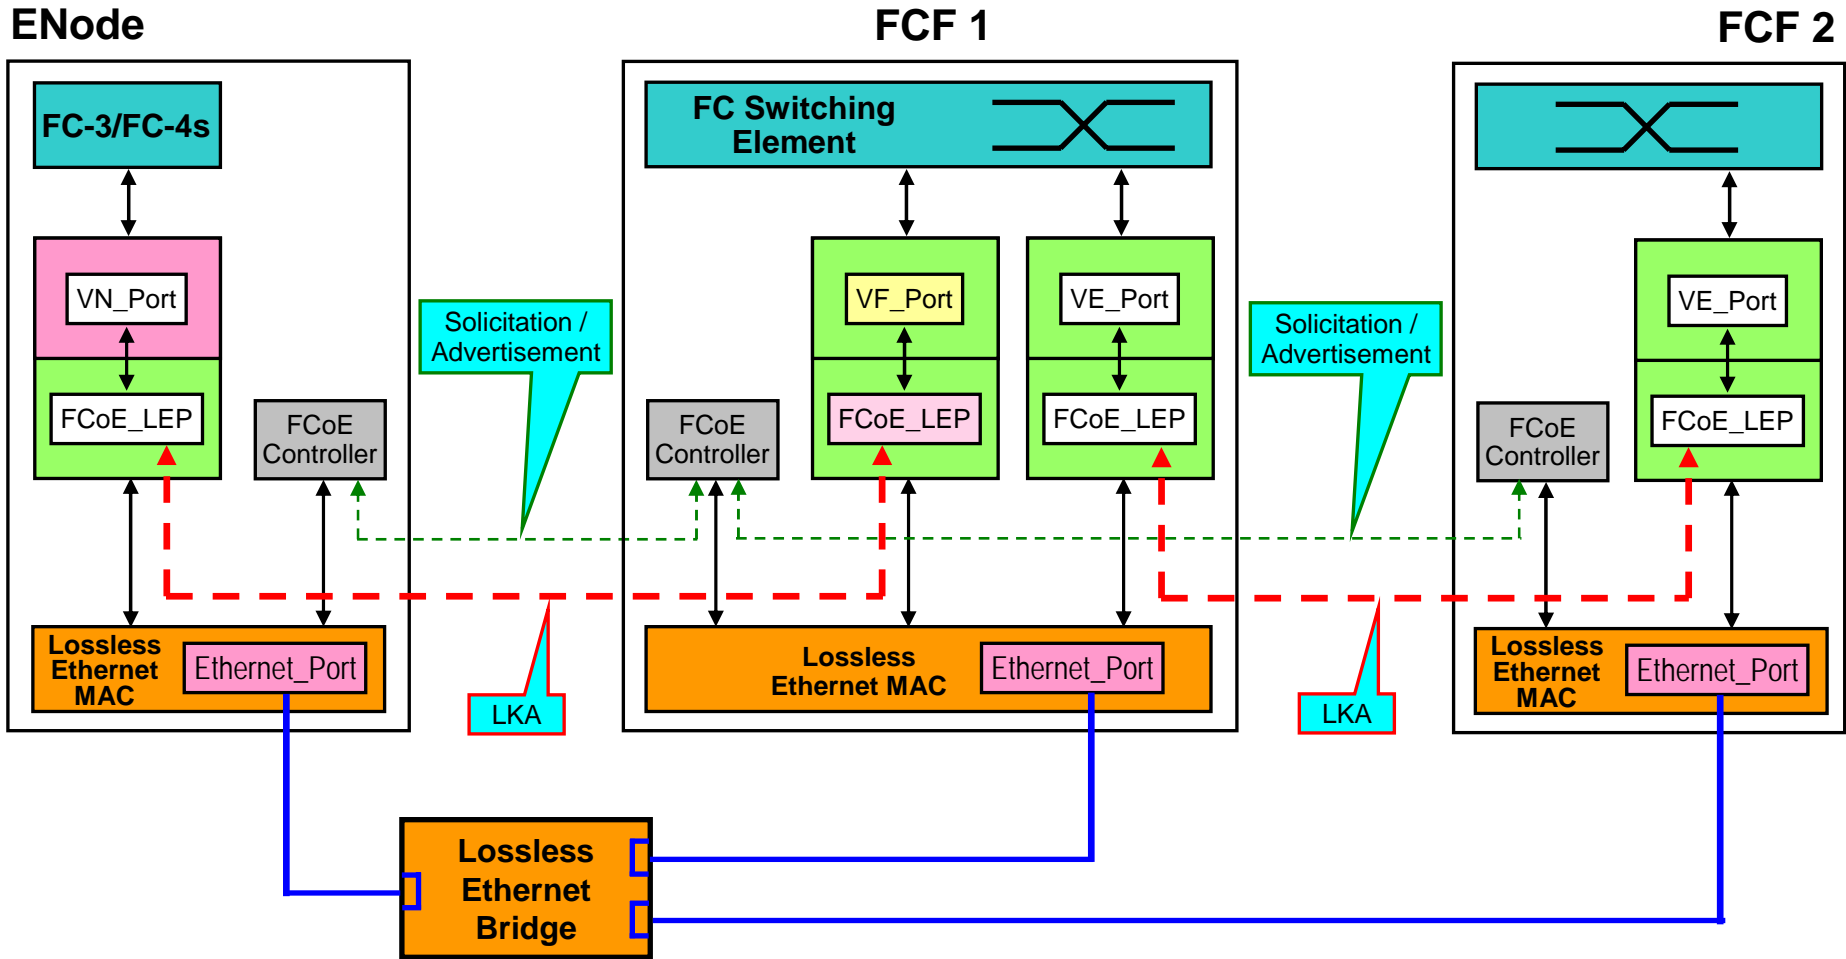
- **VF_Ports, VN_Ports, and VE_Ports may generate the LKA ELS to verify Virtual Link connectivity in absence of other FC traffic**

LKA allows the exercise of all functional components of a VN_Port to VF_Port Virtual Link or of a VE_Port to VE_Port Virtual Link

- **After a certain number of missed LKA responses a Virtual Link may be declared down**
- **A FIP LOGO Request should be generated to clean up state in possible intermediate Lossless Ethernet switches**

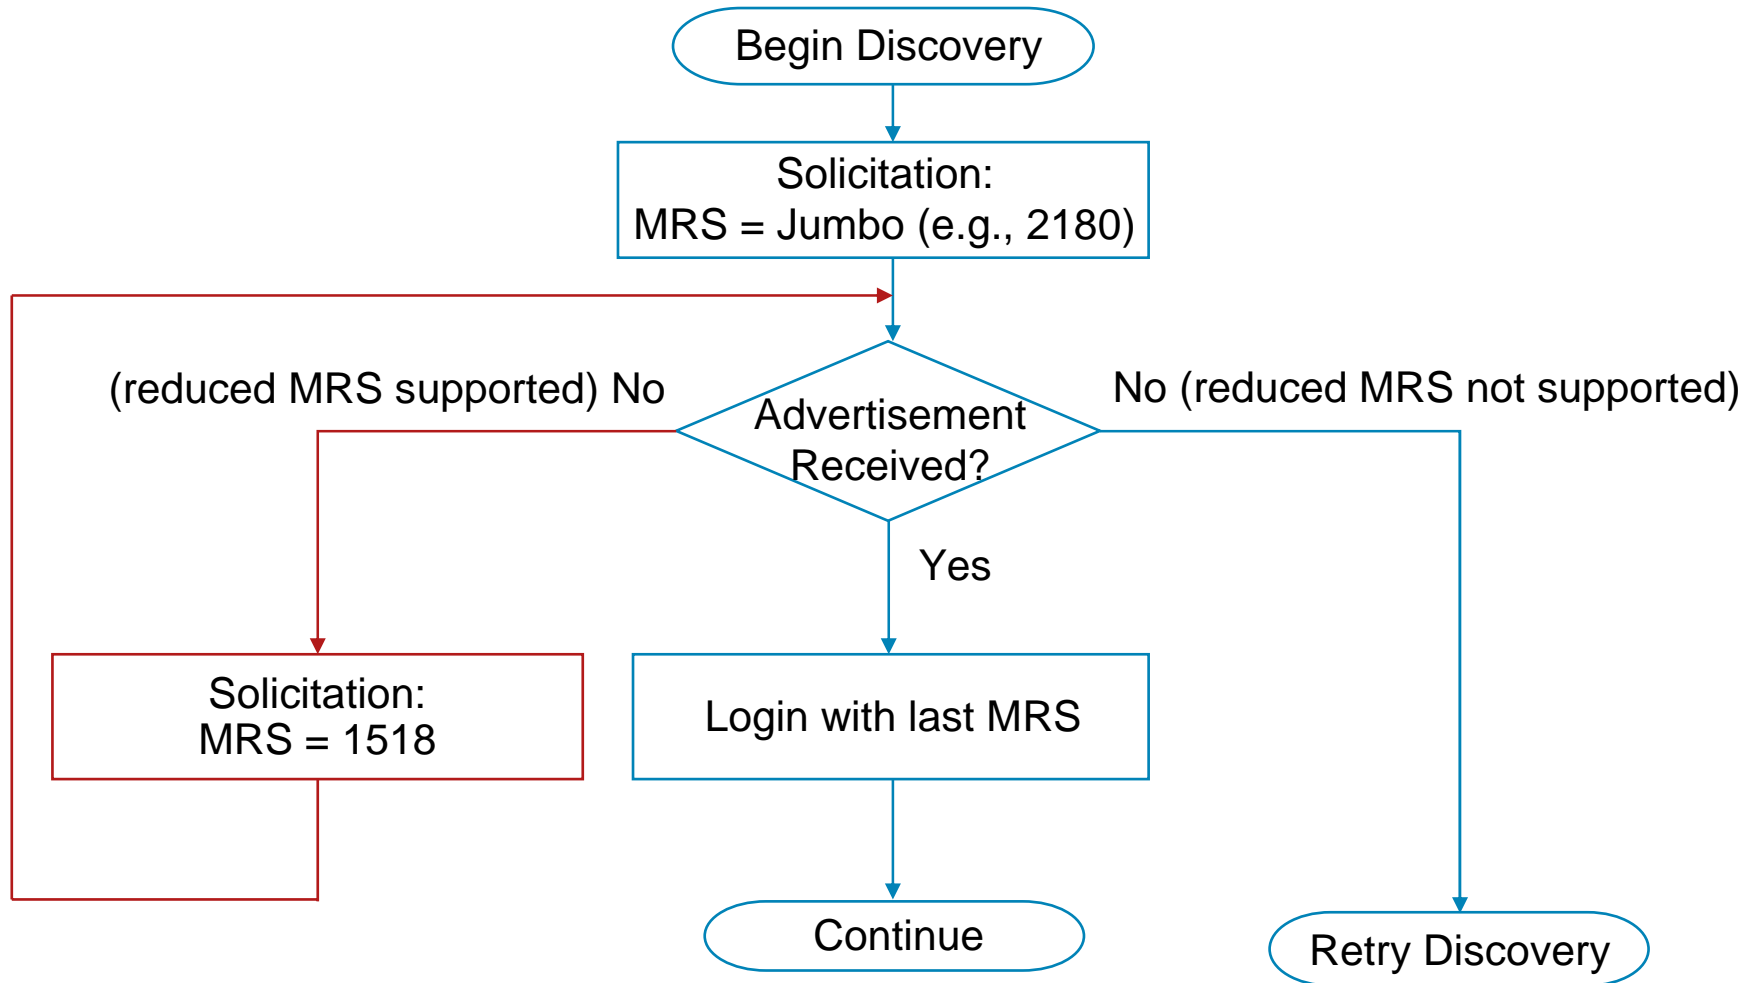
The FIP LOGO Request could be used also for VE_Port to VE_Port Virtual Links

Virtual Links Verification



Thank You

Possible ENode behavior for Max Receive Size



Support for reduced Max Receive Size is optional