



Additional FIP Functions

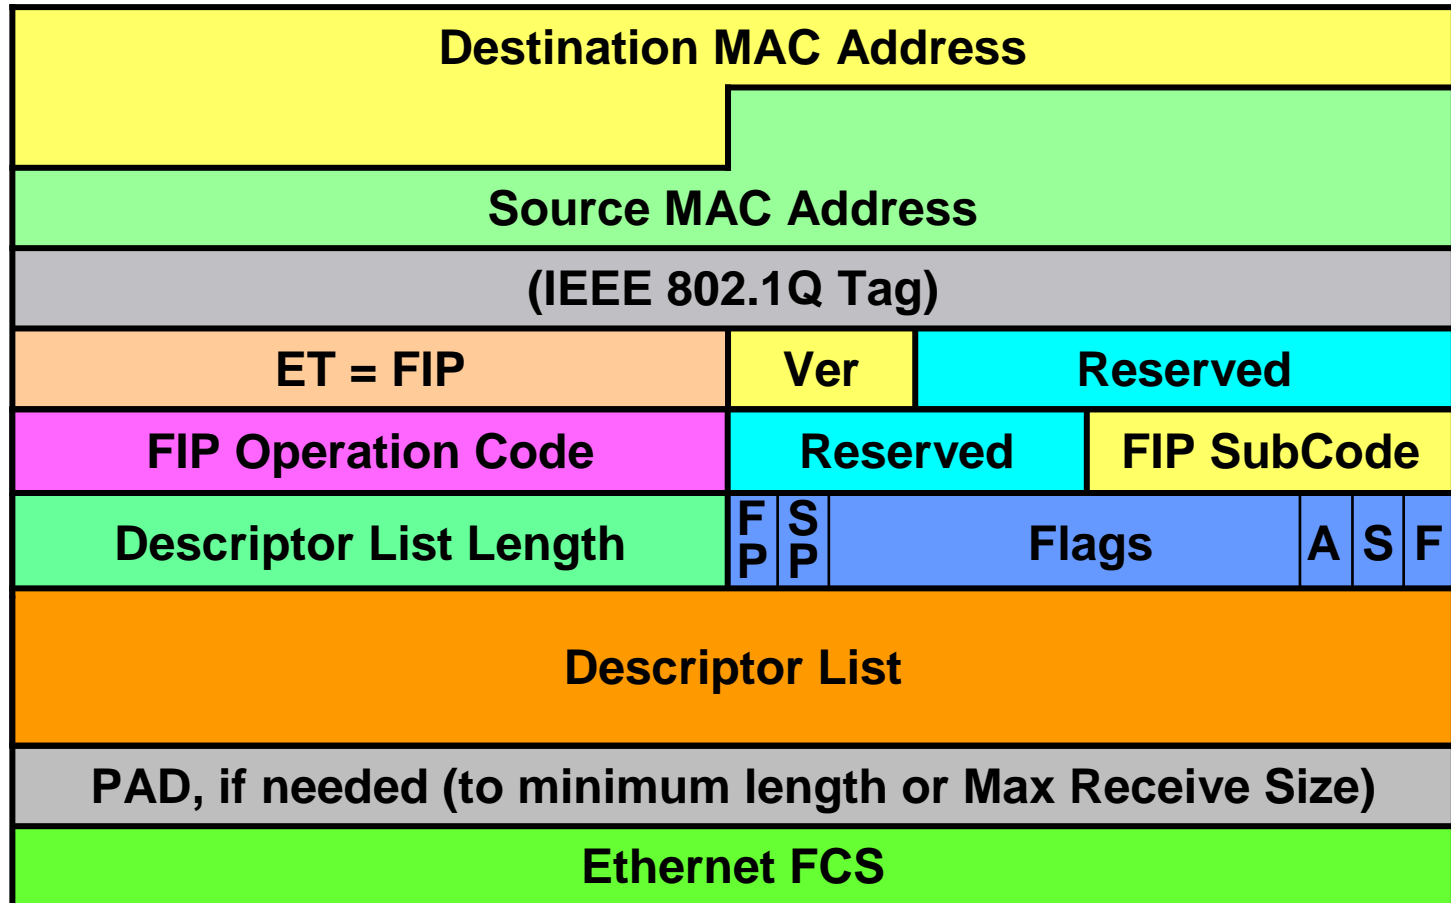
T11/08-263v0, May 2008

Claudio DeSanti

Agenda

- **FIP Discovery Summary**
- **FIP Link Reset function**

FIP Frame Format



FIP Discovery Protocol (1)

- **VF_Port capable FCF-MACs periodically send multicast Advertisements**
E.g., every 30 seconds
Not jumbo!
- **On receiving Advertisements, ENodes' FCoE Controllers create an entry per FCF-MAC in a FCF List**
By default ordered on the basis of the value carried in the Priority descriptor
Each entry has some flags:
 - ‘Max Receive Size verified’ (set to zero for entries created from multicast Advertisements, set to one when a unicast jumbo Advertisement is received)
 - ‘Available for Logins’ (set to one when the FCF is able to accept additional FLOGI/FDISC requests, set to zero otherwise)
- **ENodes' FCoE Controllers select a subset of the Available FCF-MACs for Login (the “FCF Login Set”)**
On the basis of a “local policy”
By default the one(s) with higher priority (i.e., lower priority value)

FIP Discovery Protocol (2)

- **Each FCF-MAC of the FCF Login Set needs to be verified for Max Receive Size support before performing the Login**
 - By sending a unicast Solicitation and receiving a unicast Advertisement**
 - Unicast Advertisements may carry updated information specific for the soliciting ENode**
 - e.g., a more appropriate value of Priority
- **The periodic reception of multicast Advertisements allows ENode's FCoE Controllers to continuously verify the FCF-MACs reachability**
- **The 'A' flag in FIP Advertisements carries the information "Available for Login"**
 - This flag is updated on the FCF list / FCF Login Set on reception of Advertisements**

FIP Discovery Protocol (3)

- **When an ENode's FCoE Controller becomes operational it can just wait to receive the Advertisements and operate as described in (1) and (2)**

This may take some time

- **Shortcut: the ENode's FCoE Controller sends a multicast Solicitation when it becomes operational**
- **FCF-MACs respond with Unicast Advertisements**
Jumbo padded
- **The ENode's FCoE Controller may then directly proceed to the Login phase**

FIP Discovery Protocol (4)

- **VE_Port capable FCF-MACs periodically send multicast Advertisements**
 - E.g., every 30 seconds
 - Not jumbo!
- **On receiving Advertisements, VE_Port capable FCF-MAC's FCoE Controllers create an entry per FCF-MAC in a FCF List**

By default ordered on the basis of the value carried in the Priority descriptor

Each entry has some flags:

'Max Receive Size verified' (set to zero for entries created from multicast Advertisements, set to one when a unicast jumbo Advertisement is received)

'Available for ELP' (set to one when the FCF is able to accept additional ELP requests, set to zero otherwise)

FIP Discovery Protocol (5)

- **Each FCF-MAC of the FCF List needs to be verified for Max Receive Size support before performing the ELP**

By sending a unicast Solicitation and receiving a unicast Advertisement

Unicast Advertisements may carry updated information specific for the soliciting FCF-MAC

- **The periodic reception of multicast Advertisements allows VE_Port capable FCF-MAC's FCoE Controllers to verify the reachability of other VE_Port capable FCF-MACs**
- **The 'A' flag in FIP Advertisements carries the information "Available for ELP"**

This flag is updated on the FCF list on reception of Advertisements

FIP Discovery Protocol (6)

- **When a VE_Port capable FCF-MAC's FCoE Controller becomes operational it can just wait to receive the Advertisements and operate as described in (4) and (5)**

This may take some time

- **Shortcut: the VE_Port capable FCF-MAC's FCoE Controller sends a multicast Solicitation when it becomes operational**
- **VE_Port capable FCF-MACs respond with Unicast Advertisements**

Jumbo padded

- **The VE_Port capable FCF-MAC's FCoE Controller may then directly proceed to the ELP phase**

Dealing with Multicasts and Periodicity

- **Periodic transmission of multicast messages as well as responses to multicast messages may create undesirable synchronizations**
- **The interval between subsequent transmissions of multicast Advertisements should be randomized**
 - When the clock expires, wait a random time
- **Advertisements sent in response to a multicast Solicitation should be also delayed by a random time**
- **The random time should be uniformly distributed between 0 and MAX_RANDOM_DELAY seconds**
 - 2 seems a reasonable value for MAX_RANDOM_DELAY

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 with the 'Max Receive Size verified' flag set to one**
 - Unsolicited Advertisements create entries with the 'Max Receive Size verified' flag set to zero**
- **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 FCF-MACs are preferred to establish Logins**
 - 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**

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 = 11		FP	SP
		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

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

Solicitations Processing Rules

- **The Switch_Name descriptor is used by a VE_Port capable 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 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 network

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

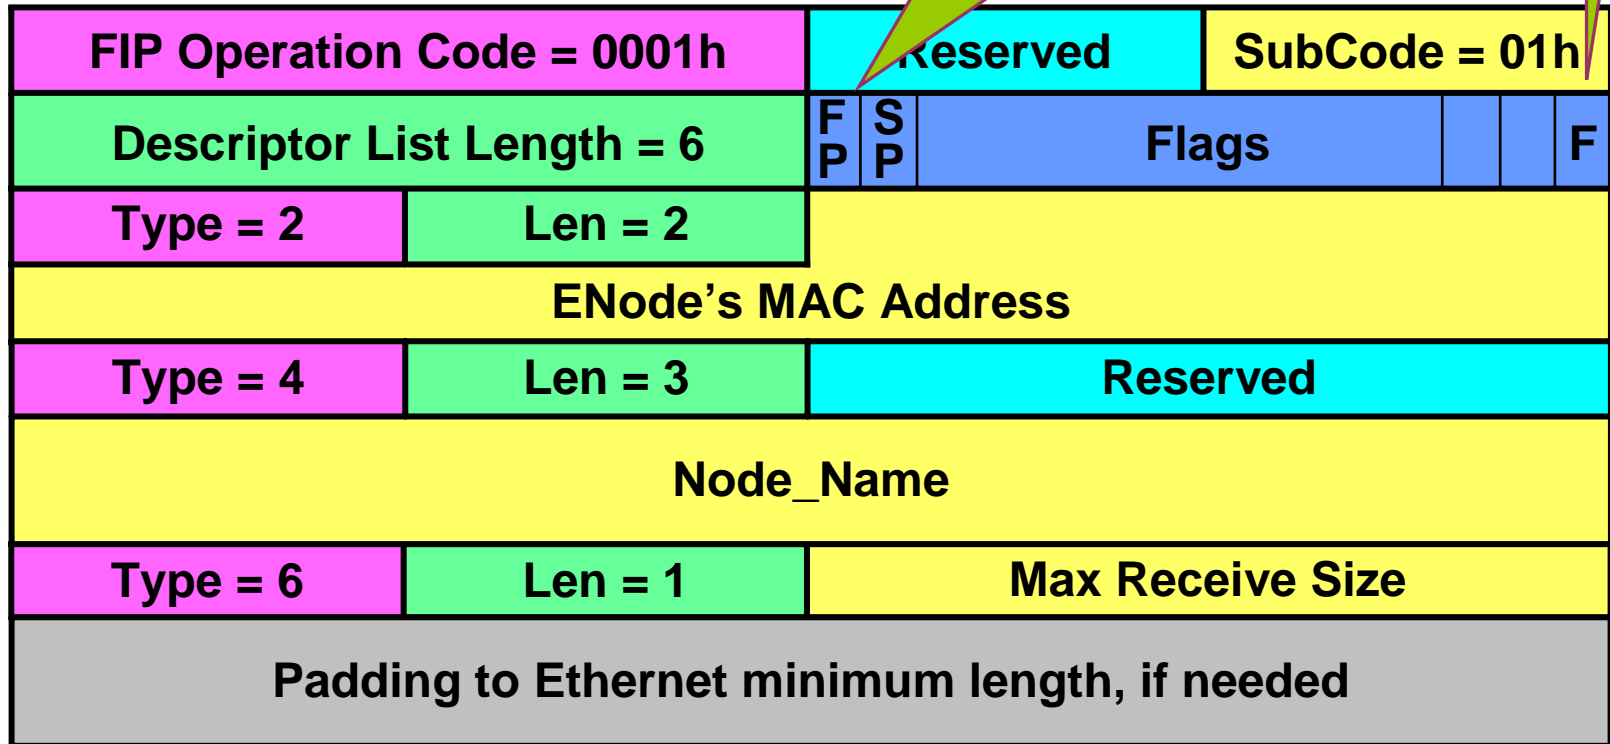
An exception may be signaled

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

Solicitation from ENode

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

F = 0b

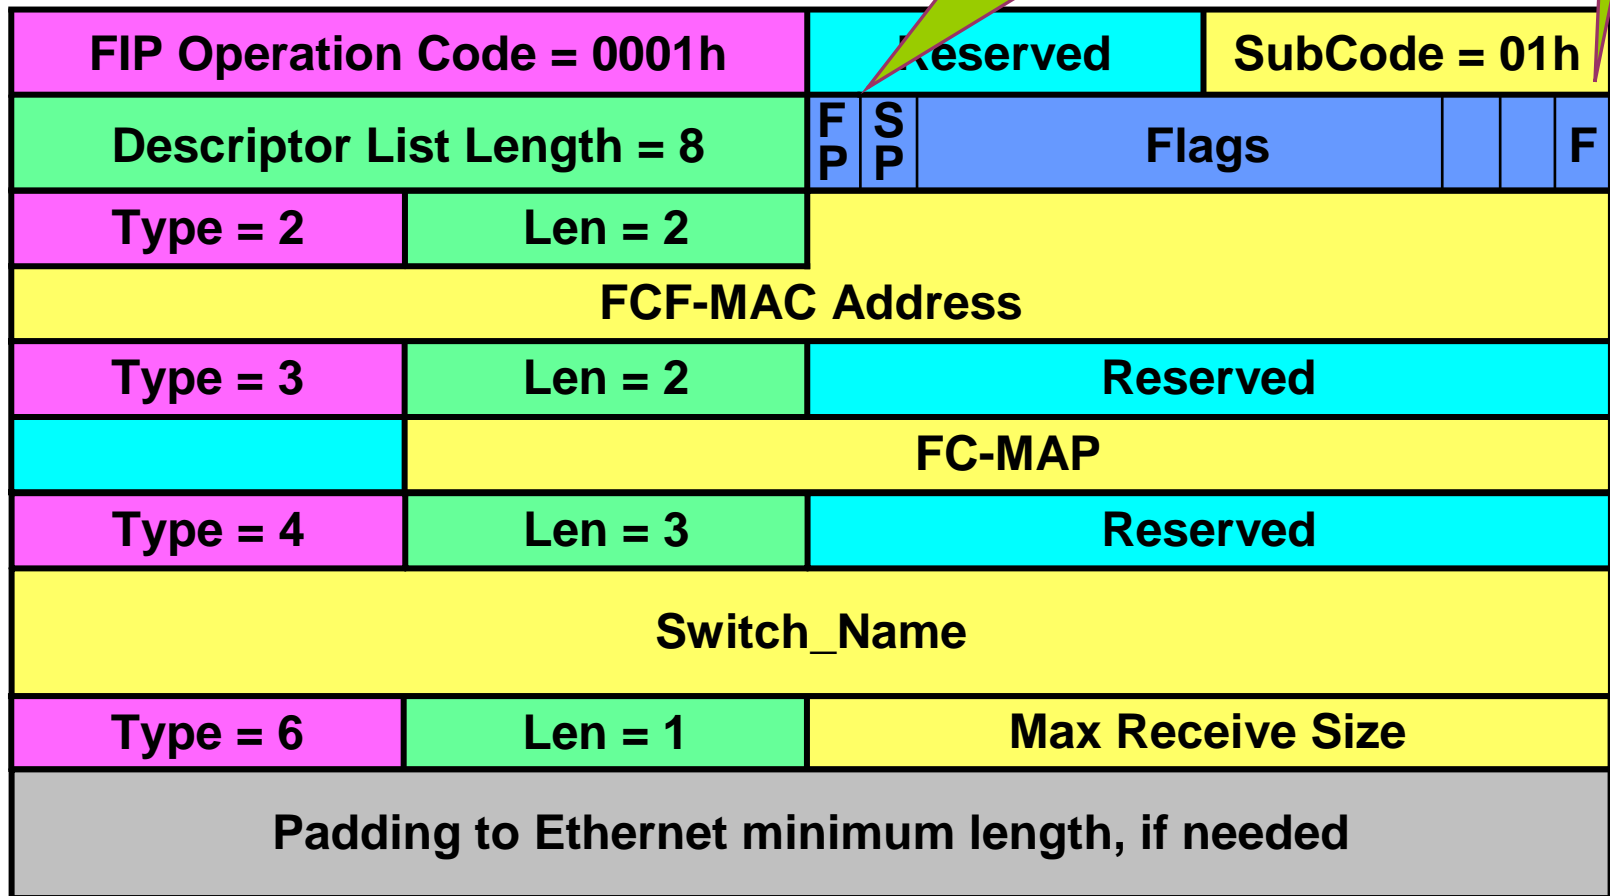


Solicits VF_Port capable FCF-MACs

Solicitation from FCF

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

F = 1b



Solicits VE_Port capable FCF-MACs

Agenda

- **FIP Discovery Summary**
- **FIP Link Reset function**

A FIP Clean-up Function

- **Native Fibre Channel ports may perform a link initialization**
On a N_Port to F_Port link this removes all N_Port_IDs (VN_Ports)
- **It is desirable to have something similar also on the virtualized FCoE environment**

A **FIP Link Reset** function

When invoked, it removes all VN_Ports, VF_Port and Virtual Links between an ENode's MAC and a VF_Port capable FCF-MAC

When invoked, it removes all VE_Ports and Virtual Links between two VE_Port capable FCF-MACs

Allows a better cleaning of state (also on intermediate Ethernet bridges) when a loss of connectivity is detected

- **However use of this function need to be limited**
In order to reduce possibilities for denial of service attacks
- **Limitation: only FCFs may invoke the FIP Link Reset function**

FIP 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
0003h	01h	FIP Link Reset
All others	All others	Reserved

FIP Link Reset

- Can be generated only by an FCF
- Does not have a reply message
- Carries the FCF-MAC Address and the Switch_Name of the generating FCF
- When sent to an ENode's MAC, carries the list of VN_Port MAC addresses to be removed

For a “full” FIP Link Reset, all VN_Port MAC addresses

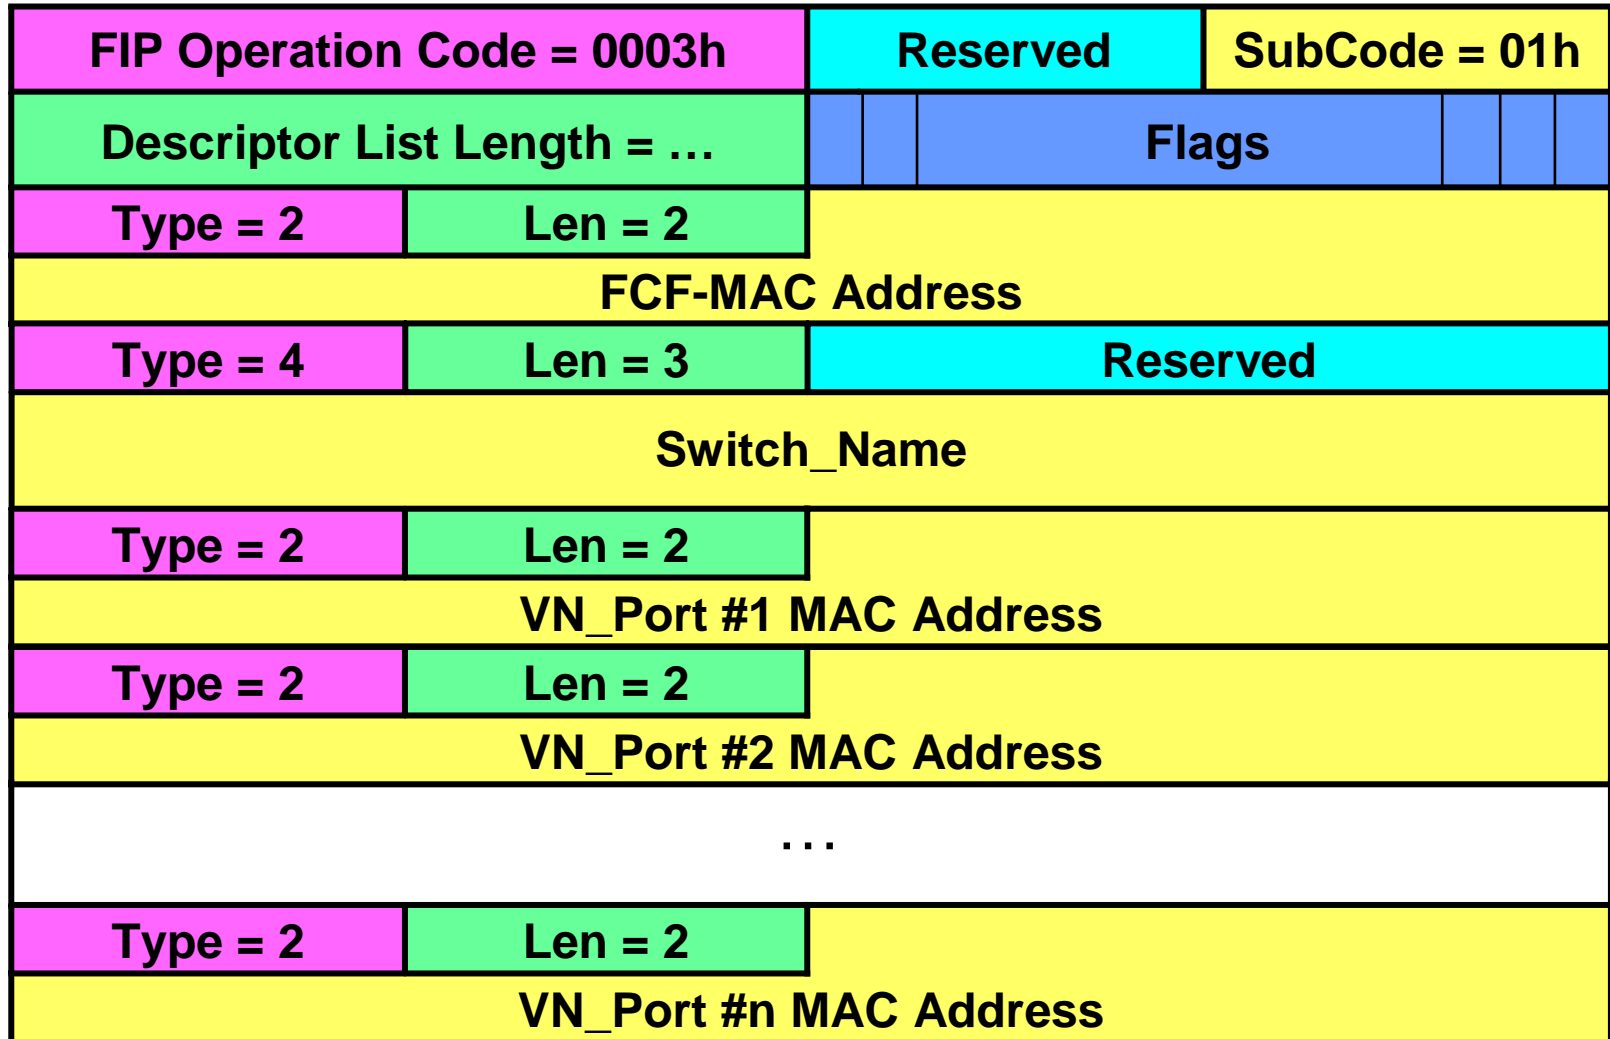
To clean-up the state of some non-responding VN_Ports, the MAC addresses of the non-responding VN_Ports

After a “full” FIP Link Reset, operations restarts from the Discovery protocol

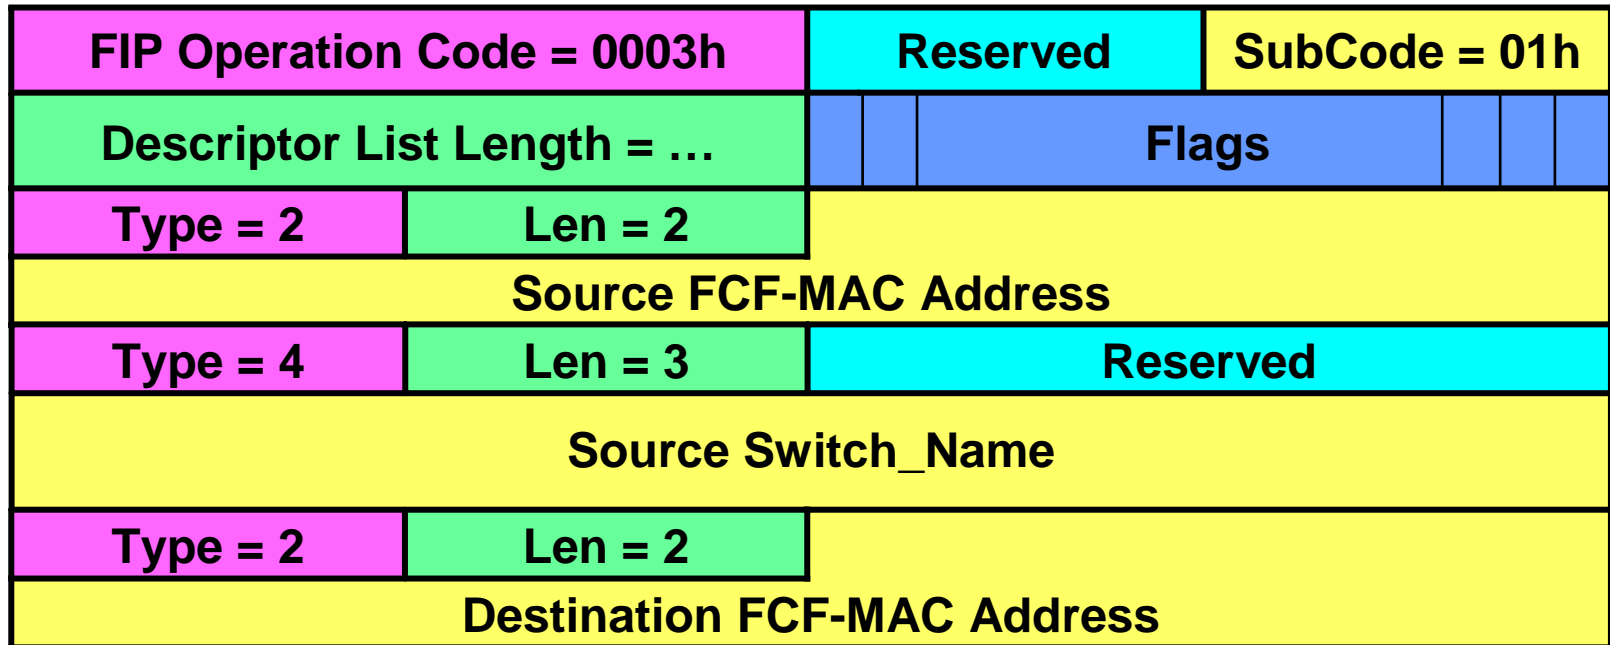
- When sent to an FCF-MAC, carries the destination FCF-MAC address

The VE_Port address to be removed

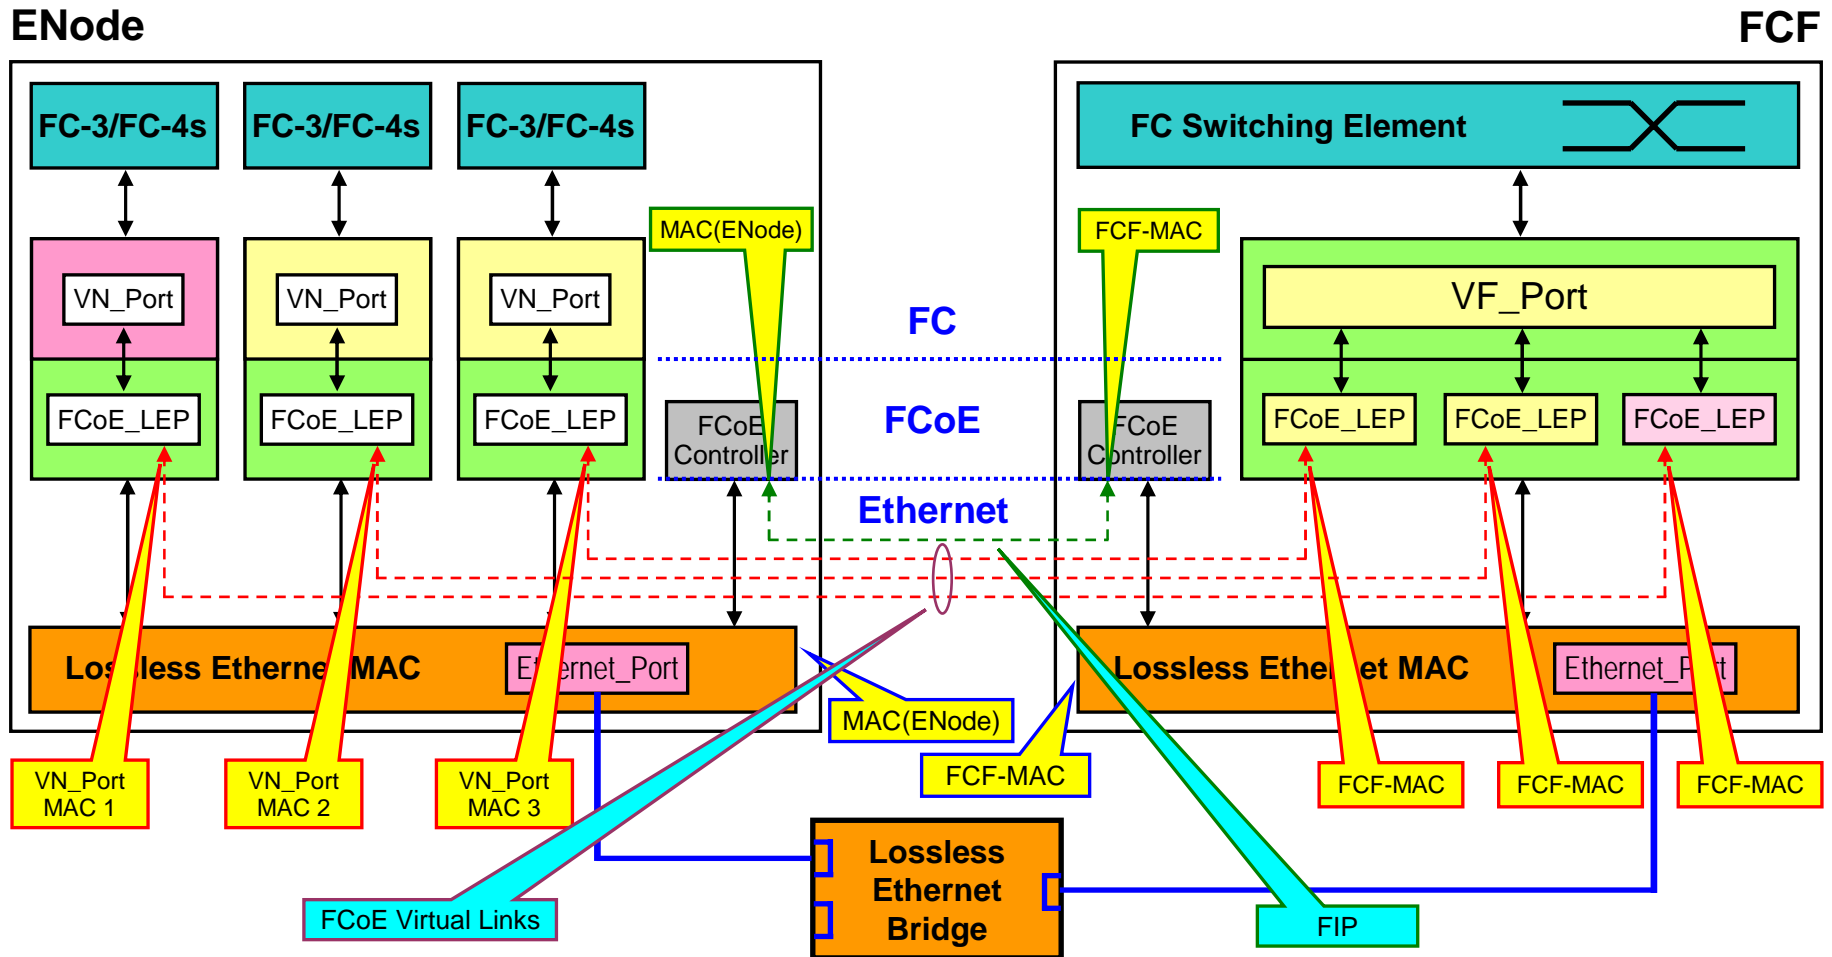
FCF to ENode FIP Link Reset Payload



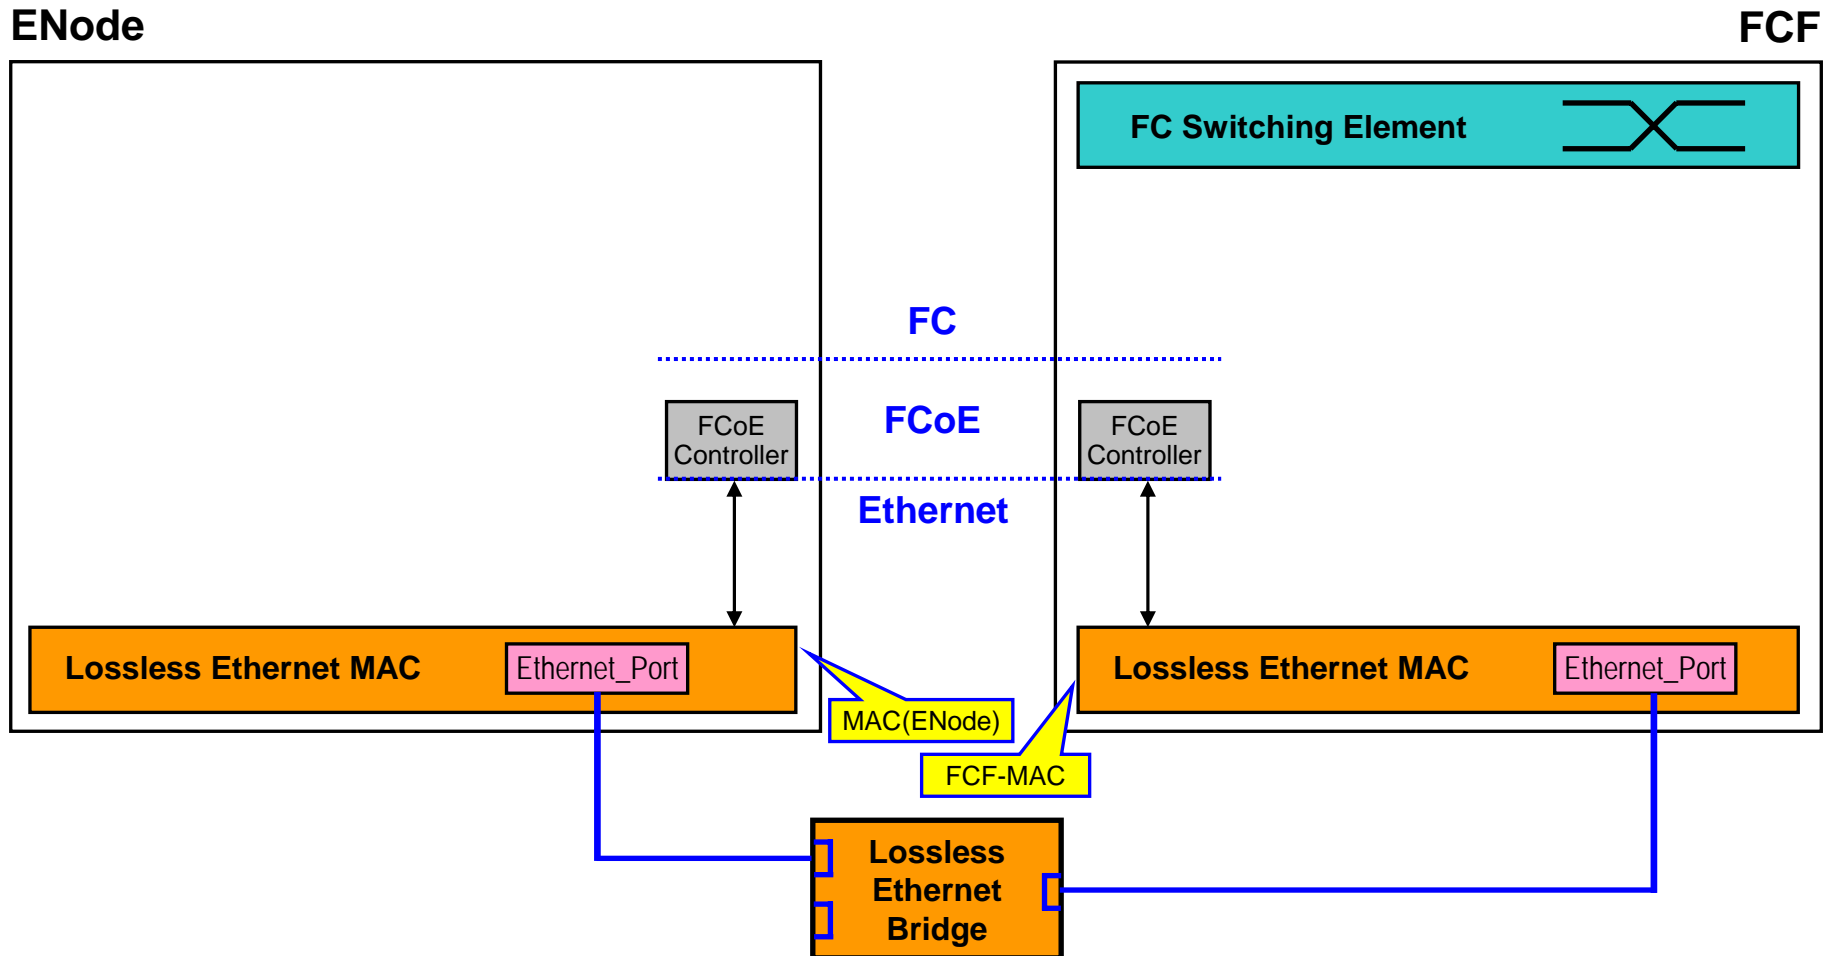
FCF to FCF FIP Link Reset Payload



FIP and FCoE Virtual Links



After a Full FIP Link Reset...



Relationships

- **NPIV associates one VF_Port to a set of VN_Ports**
- **There is a VF_Port per each logged in ENode's MAC**
- **These relationships allow to simplify the minimum set of requirements for detecting faults**

Multicast Advertisements (that are periodic) can be used by an ENode's MAC to infer VF_Port reachability

LKA ELSs generated by VN_Ports can be used by an FCF to infer ENode's MAC reachability

- **For VE_Port to VE_Port Virtual Links, the periodic multicast Advertisements may be used by VE_Port capable FCF-MACs to verify the reachability of remote VE_Port capable FCF-MACs**

Detection of Loss of Connectivity (1)

- **VN_Ports should 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**
- **An ENode's MAC should consider down a Virtual Link after a certain number (e.g., 3) of missed LKA responses**
- **An FCF-MAC should consider down a Virtual Link if no traffic has been received from it for more than a certain time**
- **An ENode's MAC may use the reception of multicast Advertisements generated by a VF_Port capable FCF-MAC to verify the reachability of the FCoE Controller of that FCF-MAC**

Detection of Loss of Connectivity (2)

- A VE_Port capable FCF-MAC should use the reception of periodic multicast Advertisements generated by a remote VE_Port capable FCF-MAC to verify the reachability of the remote FCF-MAC
- After some missing Advertisements (e.g., 3) a unicast Solicitation should be sent
- If no Advertisement is received in reply, the VE_Port to VE_Port Virtual Link should be considered down

Reaction to Loss of Connectivity (1)

- **Loss of Connectivity is detected by both an ENode's MAC and by a VF_Port capable FCF-MAC**
- **A VF_Port capable FCF-MAC reacts by sending a FIP Link Reset to the involved ENode's MAC**

The FIP Link Reset payload contains the MAC addresses of the no more operational VN_Ports if a partial fault is detected

The FIP Link Reset payload contains all VN_Port MAC addresses for a full FIP Link Reset

- **An ENode's MAC reacts by performing an implicitly logout of the involved VN_Ports**
- **The FIP Link Reset is easy to be processed by intermediate Lossless Ethernet bridges**

Reaction to Loss of Connectivity (2)

- **Loss of Connectivity is detected by both VE_Port capable FCF-MACs connected by a Virtual Link**
- **Both VE_Port capable FCF-MACs react by sending a FIP Link Reset to the other FCF-MAC**
- **The FIP Link Reset is easy to be processed by intermediate Lossless Ethernet bridges**

Thank You