FCoE: Ethernet Direct-Attached Fabrics (EDAF)

An EDAF is an FCoE fabric that consists of:

- An Ethernet fabric consisting of any number of enhanced ethernet switches, hubs, bridges, etc.
- One of more ENodes with VN_Ports attached directly to the ethernet fabric.
- One or more FCFs with VF_Ports attached to the ethernet fabric.

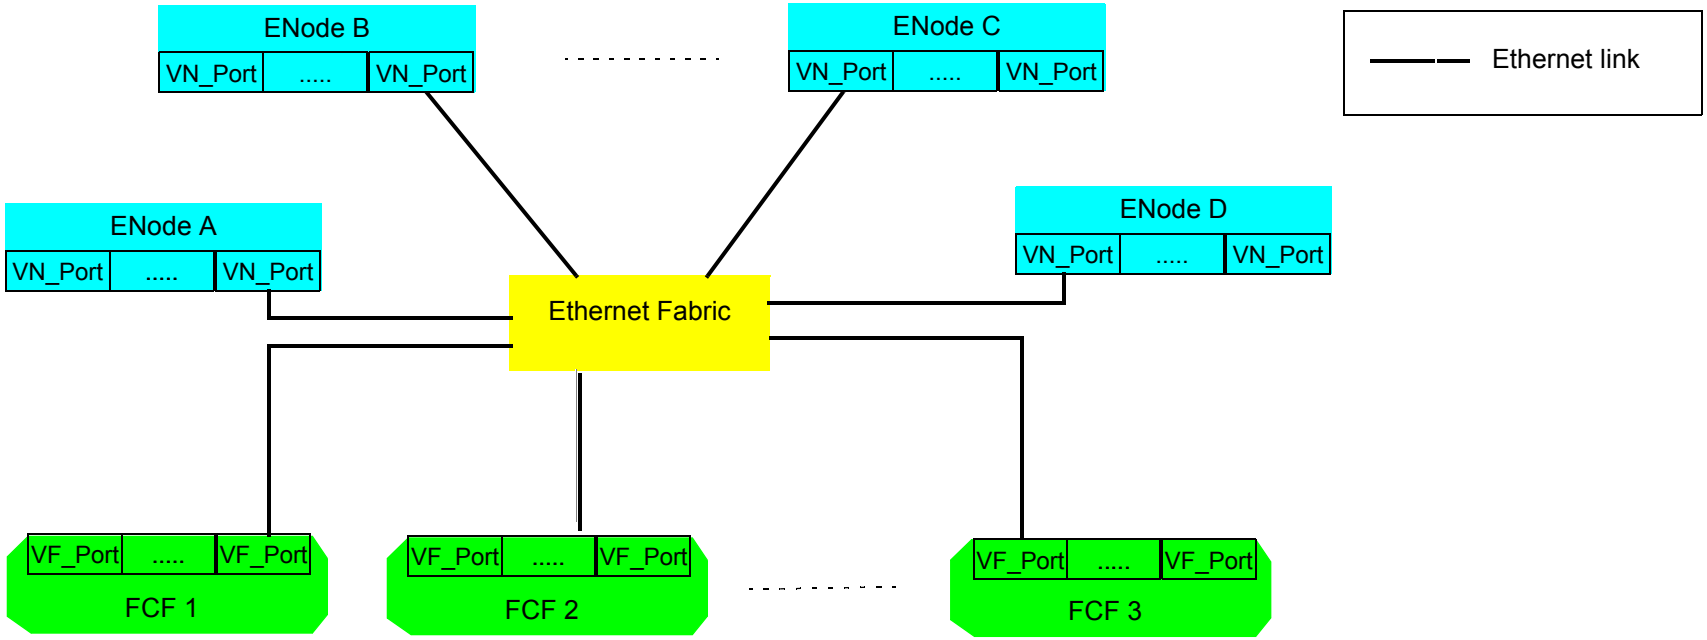

Figure 1. Ethernet Direct-Attached Fabric (EDAF)

Ethernet Direct Communication (EDC) Between VN_Ports

- Allows FCoE frames to be exchanged between two VN_Ports directly over an ethernet fabric.
- Does not require FCoE traffic to be routed through an FCF - reduces traffic through FCFs.
- Takes advantage of the shortest available path.
- Requires new FIP operations to allow creation of dynamic ACLs via snooping:
  - Ethernet Direct Communication, Permit
  - Ethernet Direct Communication, Prohibit
  - Ethernet Direct Communication, Prohibit-all
- Two proposed methods for enabling direct communication:
  - Proposal 1 - FLOGI approach
  - Proposal 2 - PLOGI approach

Ethernet Direct Communication

An example of ethernet direct communication is show below.

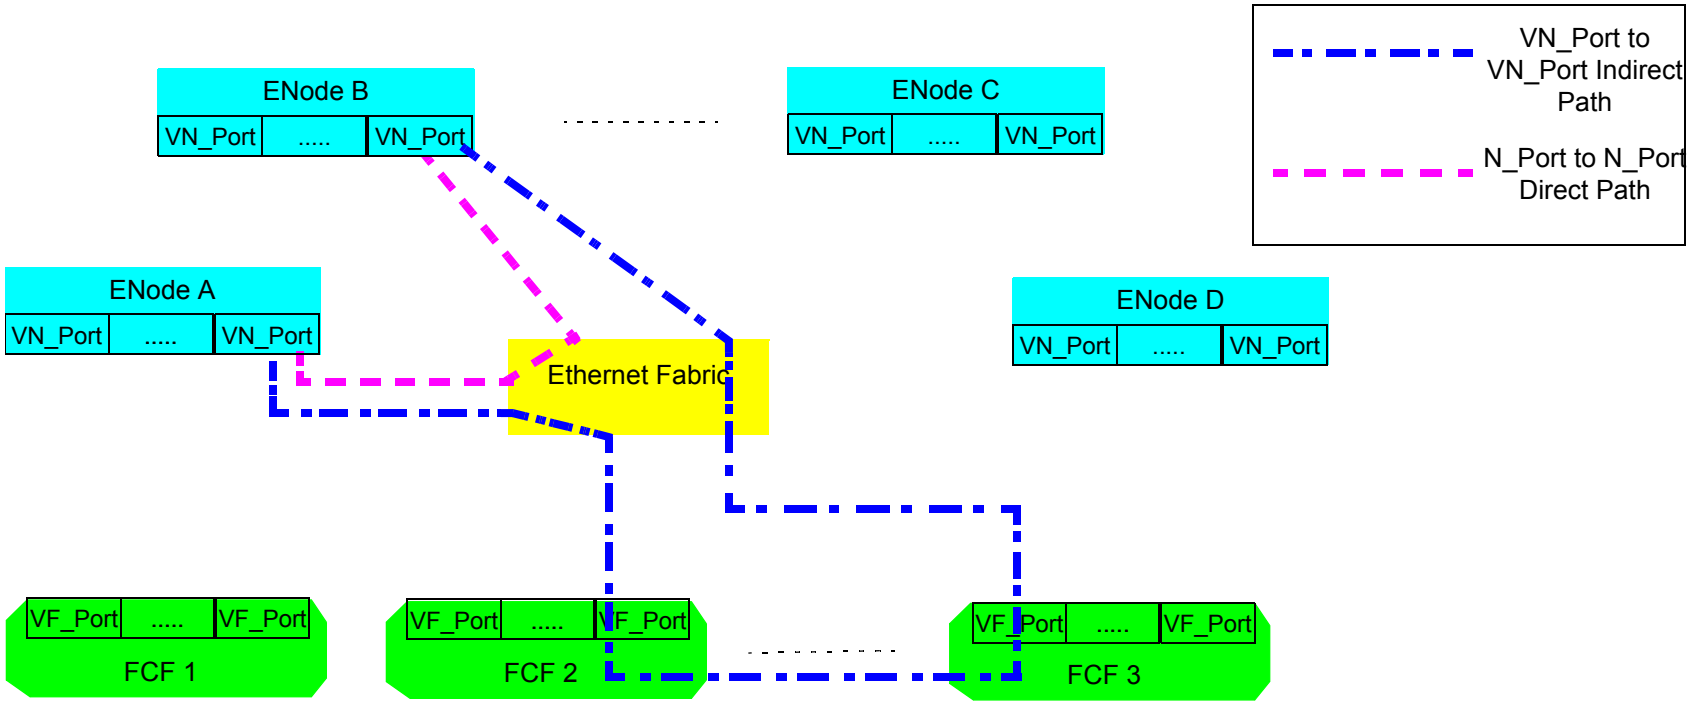

Figure 2. Example of FCoE Direct and Indirect Paths through an Ethernet Fabric

Background - FCF Discovery

- Enode port sends a mutlicast ethernet frame with the ethertype set to FIP and the operation set to discovery-solicitation to determine the set of FCFs that are available to the Enode for fabric login.

- Each FCF in the multicast group sends a unicast ethernet frame with the ethertype set to FIP and the operation code set to discovery-advertisement.

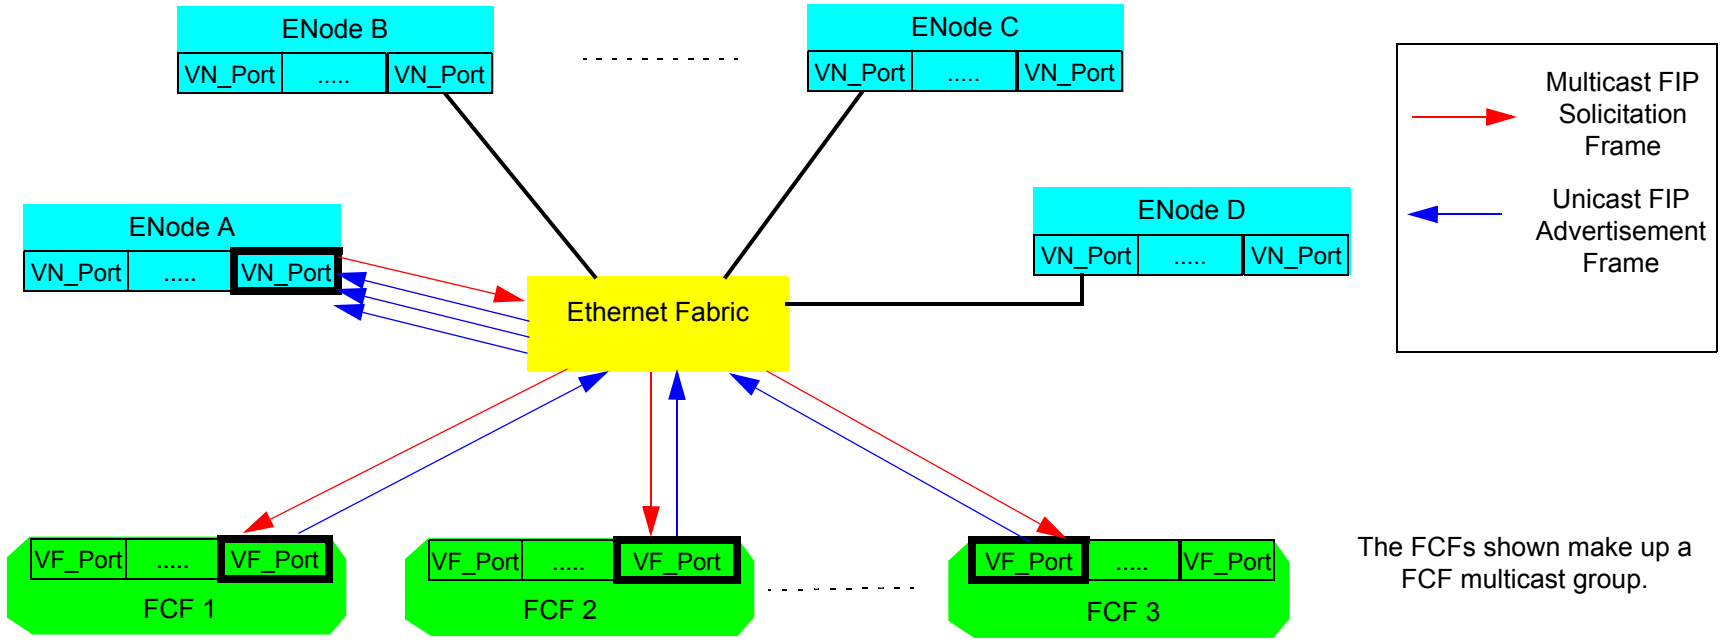

Figure 3. Solicitation and Advertisement FIP Etherframes

Background - FCF Discovery (cont.)

FIP Solicitation

	Source Address	Destination Address
FIP Frame Header	Enode MAC address	FCF Multicast Address

Figure 4. MAC Addresses for a FIP Solicitation

Descriptor List includes:

- Enode MAC address
- Node Name
- Maximum frame receive size

FIP Advertisement

	Source Address	Destination Address
FIP Frame Header	FCF MAC Address	Enode MAC address

Figure 5. MAC Addresses for a FIP Advertisement

Descriptor List includes:

- FCF MAC address
- 3-byte **FC-MAP** that is to be used as the first 3 bytes of the MAC address for FPMA addressing.
- 8-byte switch name
- 8-byte fabric name.

Background - VN_Port Fabric Login

- A VN_Port performs fabric login by sending a FIP frame containing an encapsulated FLOGI request to one VF_Port out of the set of VF_Ports specified in the FIP advertisement.
- NPIV VN_Port IDs are established by sending encapsulated FDISC requests in a FIP frame to the selected VF_Port.

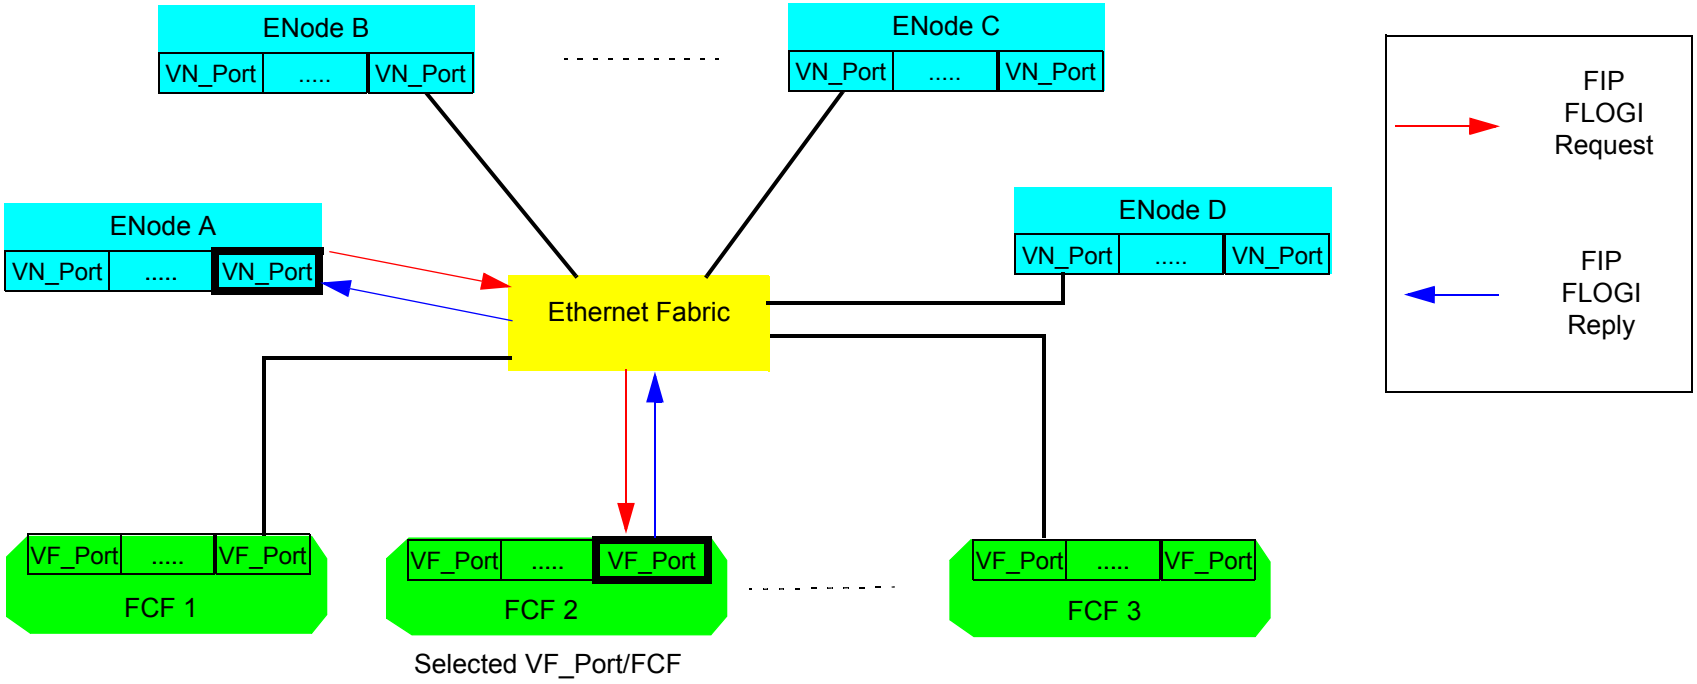

Figure 6. FLOGI FIP Requests and LS_ACC Replies

Background - VN_Port MAC Addressing

Two types of addressing are supported by VN_Ports and FCFs:

- **Server Provided MAC addressing (SPMA) -** is a “universally administered” MAC address according to ethernet standards. It is generally set by the manufacturer at installation and may be created using a local registry. The SPMA address to be used by a VN_Port is provided in the FLOGI LS_ACC reply and is equal to the address used during discovery or as specified by the Enode in the FLOGI request.

- **Fabric Provided MAC address (FPMA)** - is a fabric-unique address assigned by the fabric. The low order 24 bits are equal to the N_Port ID assigned during fabric login and the high order 24 bits are equal to the FC-MAP associated with the fabric. The FPMA address to be used by the VN_Port is provided in the FLOGI response.

The type of MAC address (SPMA or FPMA) that is to be used by a VN_Port is based on the addressing capabilities of the server and the FCF as specified in the FIP frames. This address is used in all FCoE frames.

The MAC addresses used in FIP frames between Enodes and FCFs are always equal to the Enode MAC address and either a FCF multicast or FCF unicast MAC address.

Ethernet Direct Communication FIP Operations

A new FIP operation, the “ethernet direct communication” operation, is defined as shown below. The operation is sent by a FCF to a VN_Port and contains a list of the VN_Port IDs and VN_Port MAC addresses with which the VN_Port is permitted/prohibited to have ethernet direct communication. As part of the operation, the ethernet edge switch at the destination VN_Port snoops the FIP EDC frame and updates its ACL accordingly to permit/prohibit ingress access to the fabric for all of the MAC address listed in the FIP frame.

FIP Operation Code	FIP Subcode	Operation
01h	01h	Discovery, Solicitation
	02h	Discovery, Advertisement
02h	01h	FLOGI/FDISC/LOGO/ELP Request
	02h	FLOGI/FDISC/LOGO/ELP Reply
03h	01h	Ethernet Direct Communication (EDC), Permit
	02h	Ethernet Direct Communication (EDC), Prohibit
	02h	Ethernet Direct Communication (EDC), Prohibit-all
all others	all others	reserved

New FIP Operation - Ethernet Direct Communication (cont.)

The descriptor list for EDC frames may contain one or more type X descriptors as described below.

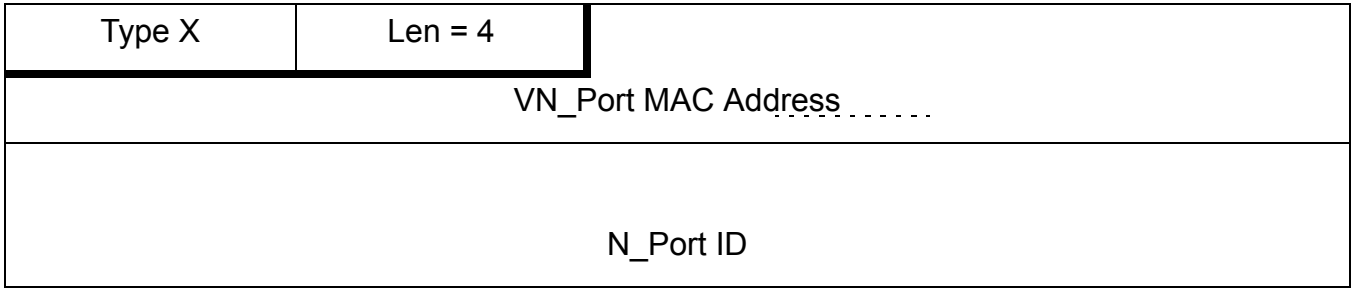

The ACL at the edge ethernet switch for the VN_Port receiving FIP EDC frame as follows:

- When a FIP EDC-permit frame is received, an entry is added for each MAC address in the descriptor list: **SA = MAC{VN_Port}, DA=MAC{Descriptor Entry}, Eth=FCoE, permit**
- When a FIP EDC-prohibit frame is received, the following entry is deleted from the ACL, if present: **SA = MAC{VN_Port}, DA=MAC{Descriptor}, Eth=FCoE, permit**
- When a FIP EDC-prohibit--all frame is received, the following entries are deleted from the ACL: **SA = MAC{VN_Port}, DA=MAC{non-FCF}, Eth=FCoE, permit**

Ethernet Direct Communication Mapping Table

A VN_Port that supports EDC maintains a table that identifies the VN_Ports that it is permitted to directly communicate with over the ethernet fabric. The mapping identifies a N_Port ID and the associated MAC address for that N_Port ID. The VN_Port MAC address may be SPMA or FPMA.

N_Port ID	Associated MAC Address
First Directly Accessible N_Port ID	MAC Address
Second Directly Accessible N_Port ID	MAC Address
. . .	. . .
Last Directly Accessible N_Port Id	MAC Address

Table 1: EDC Mapping Table (EMT) for each VN_Port

The table is updated when a VN_Port receives a FIP EDC frame as follows:

- When a FIP EDC-permit frame is received, an entry is added for each entry in the descriptor list.
- When a FIP EDC-prohibit frame is received, entries corresponding to each entry in the descriptor list are deleted.
- All entries in the table are deleted when a FIP EDC-prohibit-all frame is received.

Ethernet Direct Communication PLOGI Proposal

The PLOGI proposal is based on information provided by the name server to identify the VN_Ports that are capable of EDC. The FIP EDC-Permit frame is issued only when a VN_Port that supports EDC intends to perform a PLOGI with another VN_Port that also supports EDC. The requirements and process are described below.

- A new name server FC-4 feature bit is defined to indicate that a VN_Port supports EDC, as specified by the VN_Port at name server registration.
- The MAC address associated with a VN_Port is added to the information maintained by the name server for each VN_Port, as provided by the VN_Port at name server registration.
- A new FIP operation is added to request that an EDC-permit or EDC-prohibit frame be issued to two VN_Ports.

FIP Operation Code	FIP Subcode	Operation
04h	01h	Request EDC, Permit
	02h	Request EDC, Prohibit

The descriptor list contains the two type-X descriptors that specify the VN_Ports that a FIP EDC request/prohibit frame is to be issued to. The resulting FIP EDC-request/prohibit specifies the opposite VN_Port in the descriptor list.

EDC PLOGI Proposal (cont.)

- Name server “Get” commands are defined/modified to provide the MAC address associated with a WWPN or N_Port ID.

- Following registration with the name server, a VN_Port that supports EDC performs a name server query to get the WWPN, VN_Port ID and MAC address of each VN_Port that supports EDC.

- For each WWPN retrieved from the name server that is in the configuration of the VN_Port (i.e. with which it intends to perform PLOGI), the VN_Port issues a FIP request-EDC frame. The descriptor list contains the following:

  - A type-X descriptor specifying the VN_Port that is intending to perform the PLOGI.
  - A type-X descriptor specifying the opposite VN_Port of the PLOGI pair.

- The FCF issues a FIP EDC-permit frame to each of the VN_Ports specified in the FIP-request EDC frame specifying the opposite port of the PLOGI pair.

- The VN_Port that issued the FIP request-EDC performs the PLOGI using an FCoE frame over the EDC path. If a PLOGI timeout is detected (indicating no direct path), a request-EDC-frame is issued to remove the access permissions between the two ports and to remove the corresponding entries from the EDC mapping table

- When a VN_Port performs a VN_port logout, the VN_Port issues a FIP request-EDC frame with descriptor list that specifies itself and the opposite VN_Port of the LOGO pair.

- When a VN_Port logs out from the fabric, explicitly or implicitly, the FCF issues a FIP EDC-prohibit to each port that supports EDC with a descriptor specifying the logged out port.

EDC PLOGI Proposal Example

VN_Port A has performed a name server query that indicates that VN_Port C supports EDC and is in VN_Port A's configuration. FIP frames are issued as shown below.

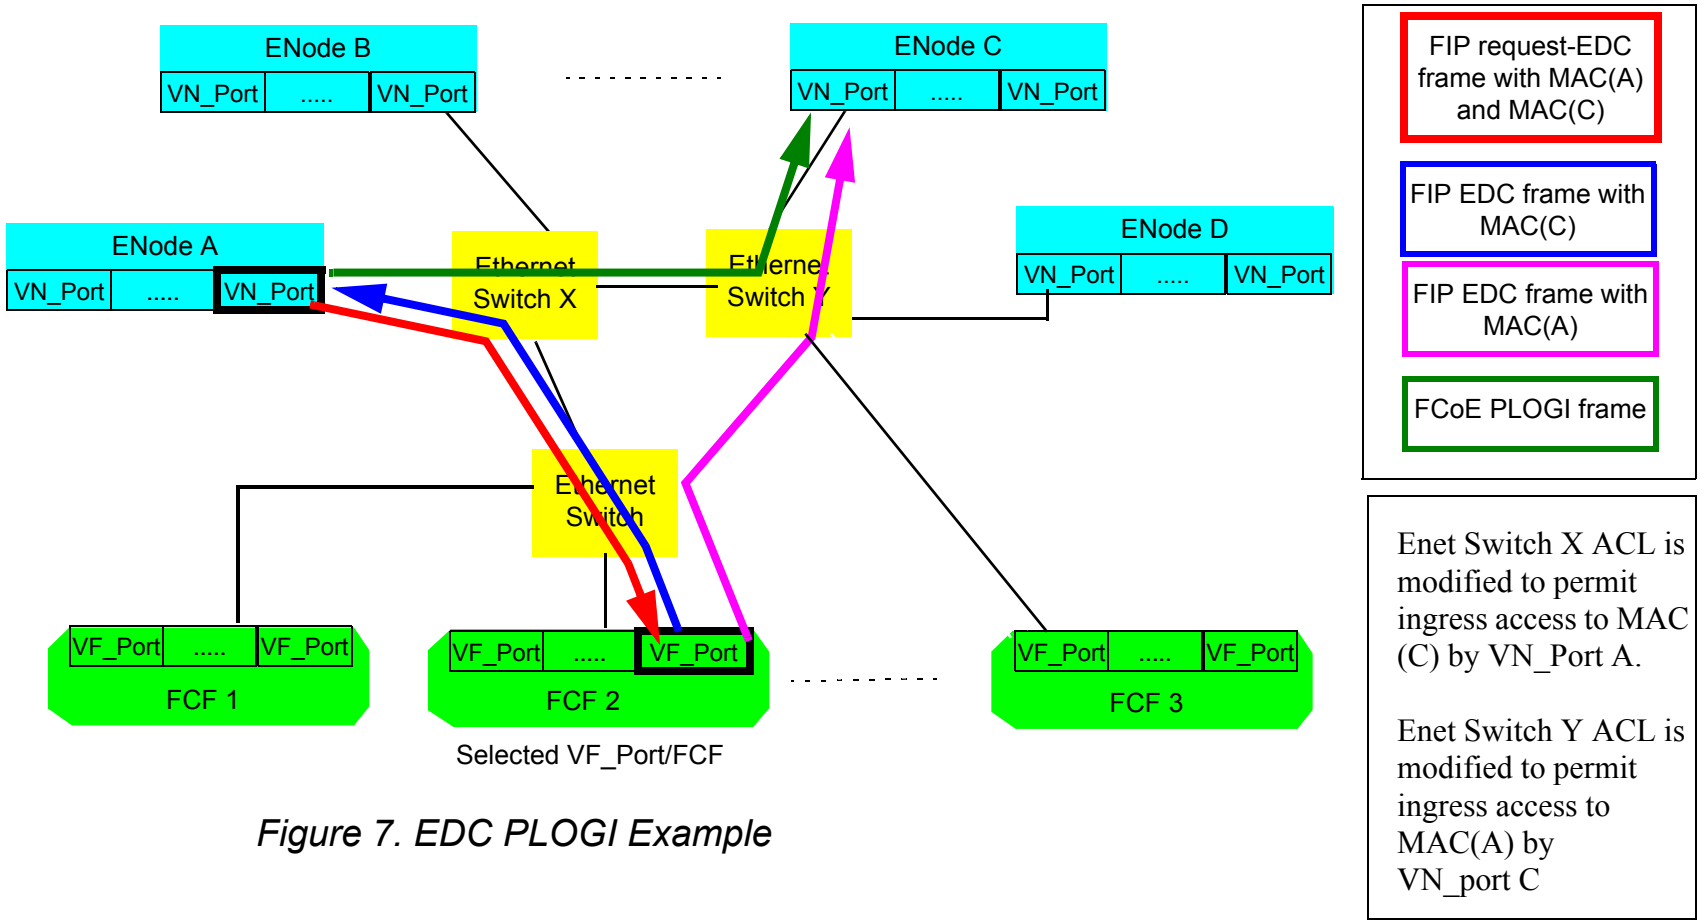

Figure 7. EDC PLOGI Example

Ethernet Direct Communication - FLOGI proposal

The FLOGI proposal is based on issuing FIP EDC-permit frames when a VN_Port that supports EDC logs in to the fabric. The process for enabling EDC is described below.

- The FLOGI service parameters are modified to include a bit indicating that a VN_Port supports EDC.
- When a VN_Port logs in to the fabric, the FCF determines if there are other VN_Ports that also support EDC. If so, the FCF issues FIP EDC-permit operations as follows:
  - The FCF issues a FIP EDC-permit frame to each other VN_port that supports EDC. The descriptor list contains a single type-8 descriptor specifying the VN_Port performing the FLOGI.
  - The FCF issues a FIP EDC-permit frame to the port that performed the FLOGI with a descriptor list specifying each VN_Port in the fabric that supports EDC.
- Each VN_Port updates its EDC mapping table with the VN_Port IDs and MAC addresses provided in the descriptor list of the FIP EDC frame.
- When a VN_Port performs a VN_port logout, the VN_Port issues a FIP request-EDC frame with a descriptor list that specifies itself and the opposite VN_Port of the LOGO pair.
- When a VN_Port that support EDC logs out from the fabric, explicitly or implicitly, the FCF issues a FIP EDC-prohibit to each other VN_Port in the fabric that supports EDC. The descriptor contains a single type-X descriptor specifying the port that has logged out.

Ethernet Direct Communication FLOGI Proposal Example

In this example,.VN_Ports A and B are logged in to the fabric with EDC enabled. VN_Port C logs in to the fabric with service parameters that indicate EDC is supported. The FCF issues FIP EDC frames as shown below.

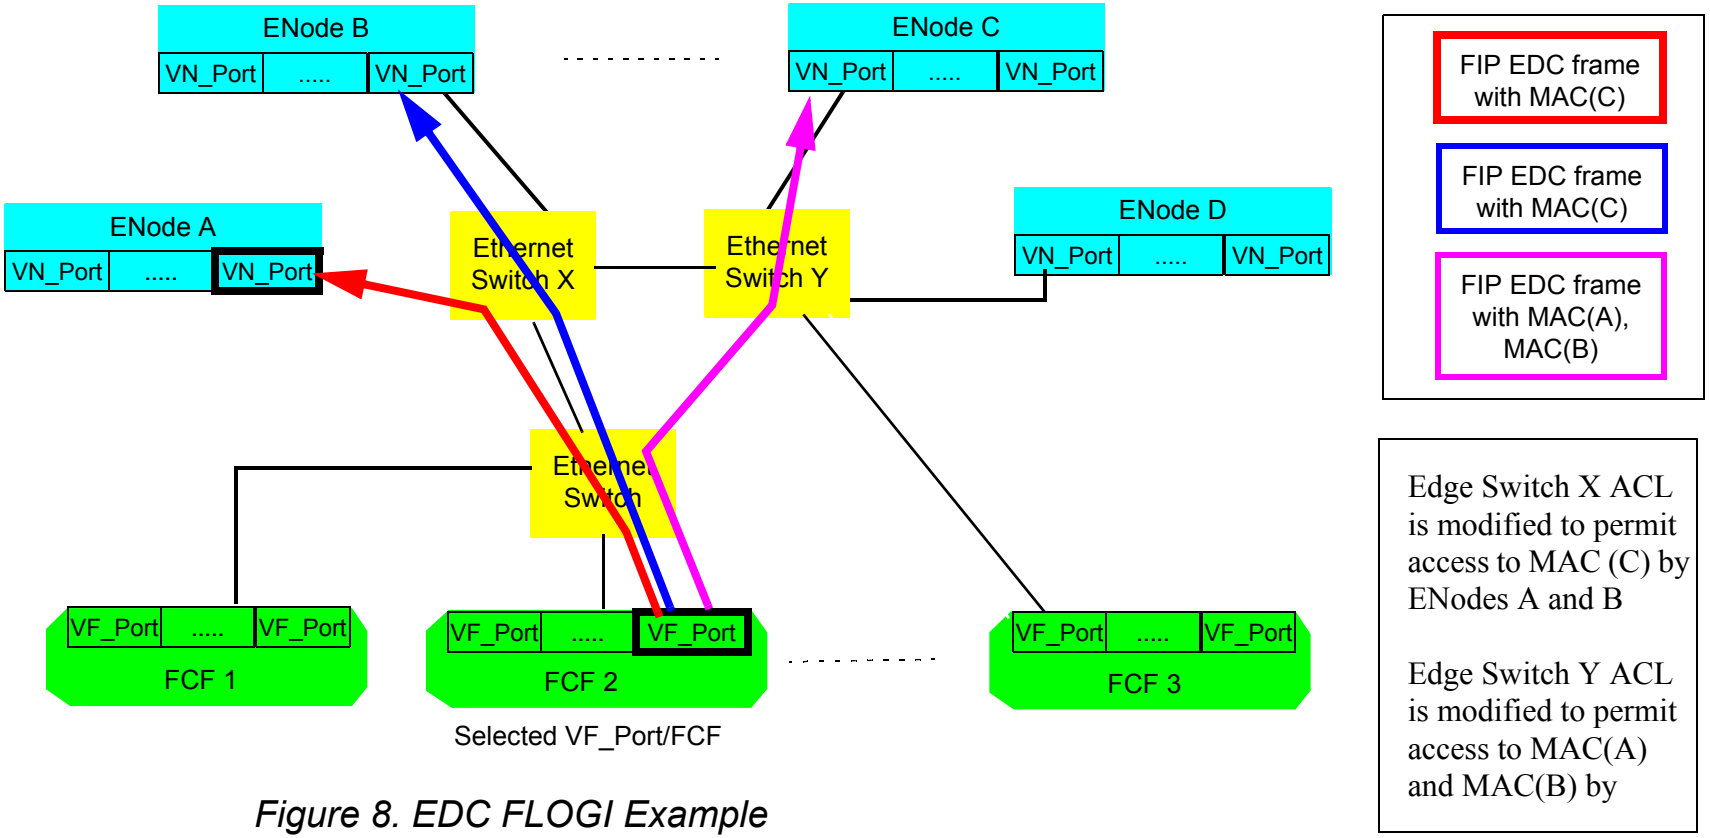

Figure 8. EDC FLOGI Example