VA_Ports:
FDF / Controlling FCF Protocols

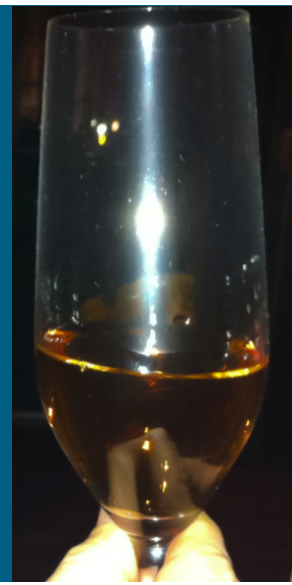

Claudio DeSanti

T11/10-271v3, October 2010

Agenda

- **Protocols**
- **Primary to Secondary Controlling FCF**
- **Controlling FCF to FDF**

The Controlling FCF

- **Controlling FCFs should operate in primary/secondary pairs for high availability**

  If more than two controlling FCFs are present on a network the two with the higher Switch_Name are selected as the Primary/Secondary Controlling FCF pair

- **Controlling FCFs instantiate VE_Port to VE_Port Virtual Links between themselves**

  The VE_Port to VE_Port Virtual Link is used to keep the N_Port_ID allocation state synchronized

- **Controlling FCFs instantiate VA_Port to VA_Port Virtual Links with the controlled FDFs**

- **The Primary Controlling FCF perform N_Port_ID allocations for all its controlled FDFs and synchronize this state with the Secondary Controlling FCF**

Finnish Defense Forces
Image Credit:
Flickr user georgerigato
Creative Commons License

The FDF (1)

- **The FDF instantiates VF_Ports and VA_Ports**
- **An FDF encapsulates the parameters of a received FIP FLOGI or FIP NPIV FDISC Request in an FCoE N_Port_ID Allocation SW_ILS sent to the Primary Controlling FCF**
- **The Primary Controlling FCF allocates an N_Port_ID, updates the FC Name Server, and generates the appropriate RSCN(s)**
- **The Primary Controlling FCF also re-computes the zoning "ACLs"**
  - **The list of N_Port_IDs allowed to communicate with the newly allocated one**
- **The Primary Controlling FCF distributes the allocated N_Port_ID and the updated zoning ACLs to all its FDFs and to the Secondary Controlling FCF**
- **On receiving the zoning ACLs for the allocated N_Port_ID, an FDF is able to enforce these zoning rules and reply to the FIP FLOGI or FIP NPIV FDISC Request**

The FDF (2)

- **The FDF performs all ENode/FCF FIP functions**

  It implements the VF_Ports

  In particular it terminates FIP for Virtual Link maintenance

- **When a VF_Port to VN_Port Virtual Link is deinstantiated, the FDF sends a message to the Primary Controlling FCF to communicate that VN_Port is not anymore reachable**
- **The Primary Controlling FCF deallocates that N_Port_ID, updates the FC Name Server, generate the appropriate RSCN(s), recompute the zoning ACL ,and sends the updated zoning ACL to the affected FDFs**

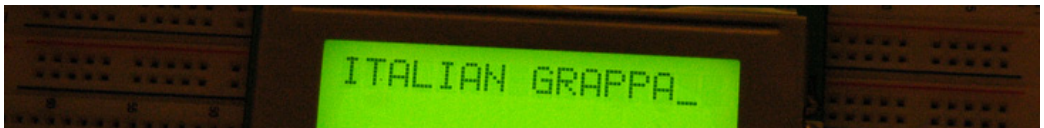

FIP Discovery

Fairly Important Person Room

Image Credit: Flickr user badjonni
Creative Commons License

- **From a FIP perspective, an FDF operates as an FCF discovering other FCFs**

  FDFs and controlling FCFs send periodic advertisements to All-FCF-MACs to discover other controlling FCFs and other FDFs connected to the same lossless Ethernet network

- **Two new flags are defined in FIP advertisements**

  FDF flag, to indicate "I am an FDF"

  Controlling FCF flag, to indicate "I am a controlling FCF"

- **This enables:**

  Controlling FCFs to discover other controlling FCFs and FDFs

  FDFs to discover controlling FCFs and other FDFs

FIP Solicitation

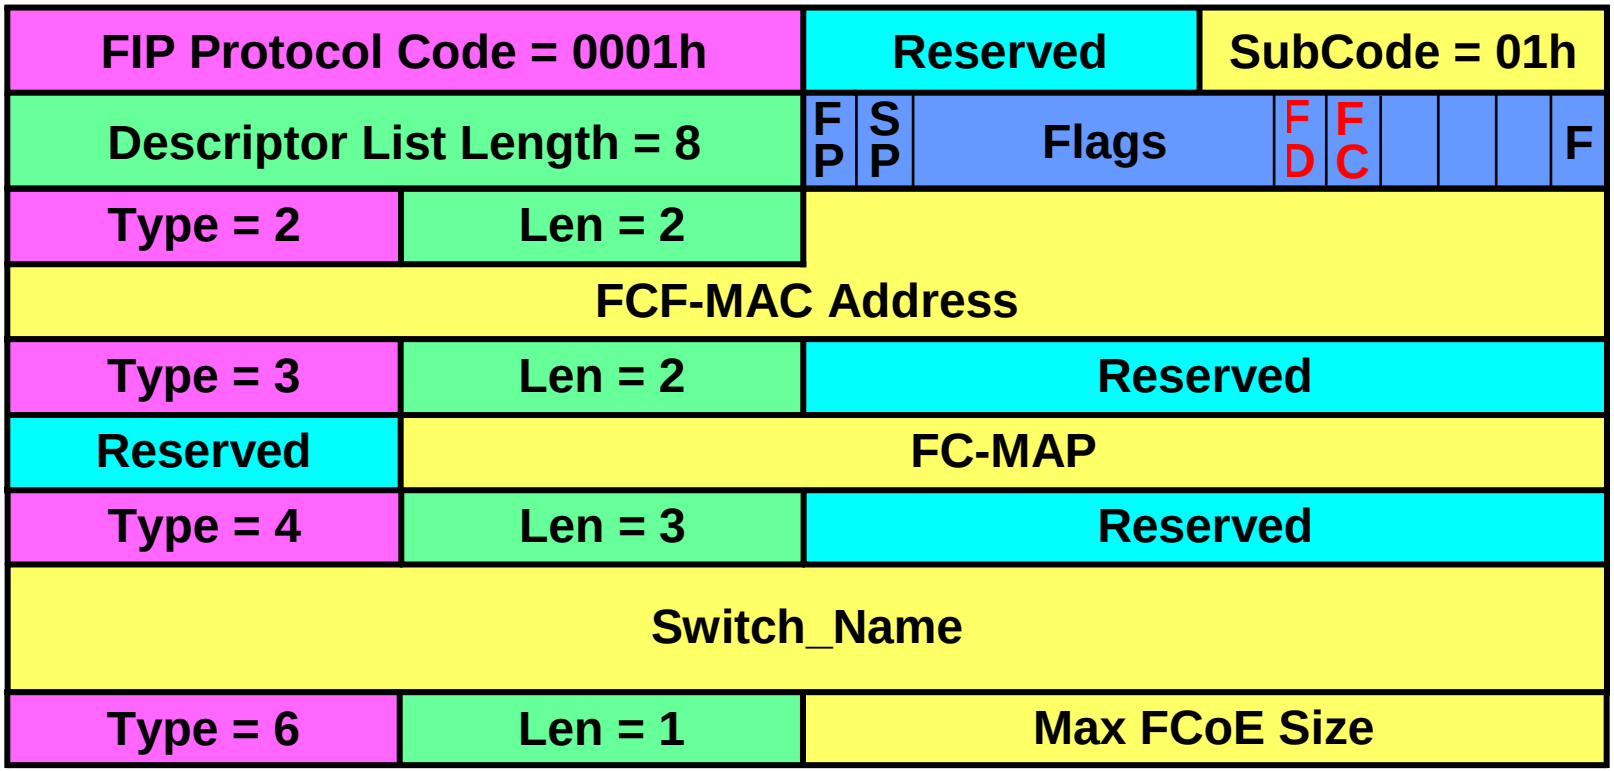

FIP Advertisement

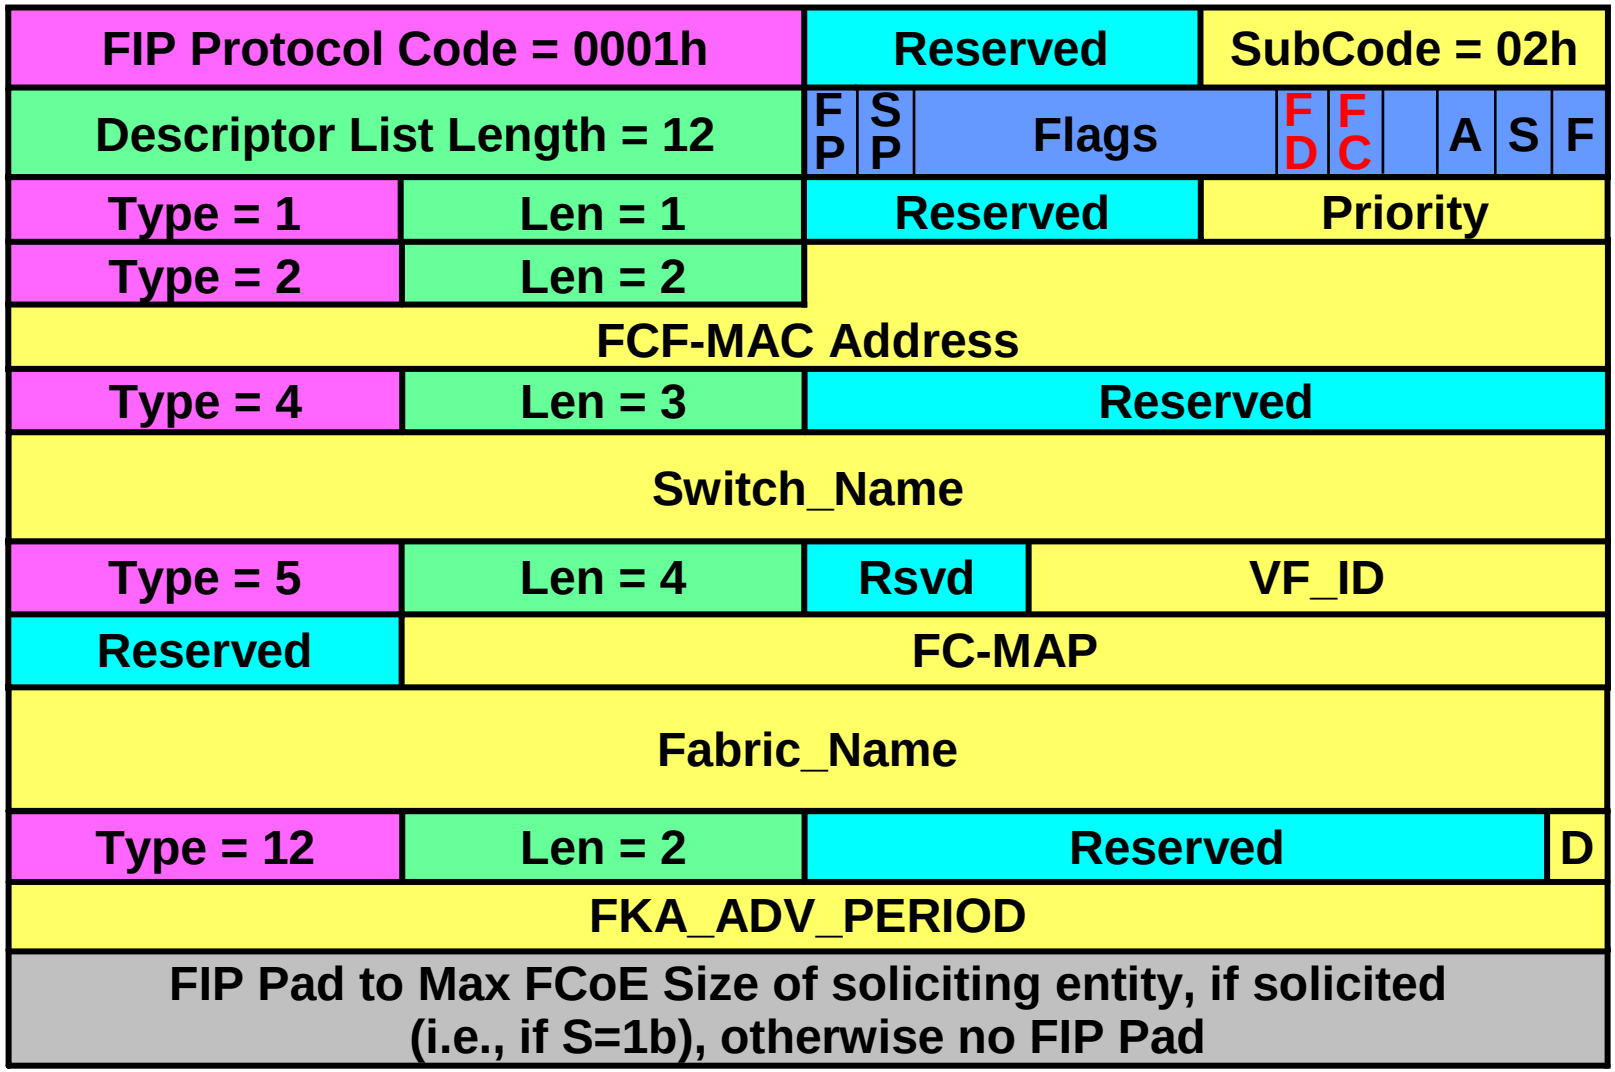

Virtual Links

- **Two new flags are defined in ELP Requests and SW_ACC**
  - **FDF flag, to indicate “I am an FDF”**
  - **Controlling FCF flag, to indicate “I am a controlling FCF”**
- **These flags enable the following Virtual Links:**
  - **Controlling FCF to Controlling FCF: VE_Port to VE_Port**
  - **Controlling FCF to FDF: VA_Port to VA_Port**
  - **FDF to FDF: VA_Port to VA_Port**
- **When a VA_Port to VA_Port Virtual Link is established, the primary controlling FCF provides to the newly connected FDF the full list of allocated N_Port_ID ranges and a summary of the Domain_IDs reachable through that FCF**
  - **Enabling the newly connected FDF to set up its forwarding tables**

Reaching other Domain_IDs

- **FDFs rely on the controlling FCF to reach other Domain_IDs**

  i.e., only a controlling FCF can be connected to other Domain_IDs (other FCFs or native FC Fabric)

Agenda

- **Protocols**
- **Primary to Secondary Controlling FCF**
- **Controlling FCF to FDF**

FIP Solicitation

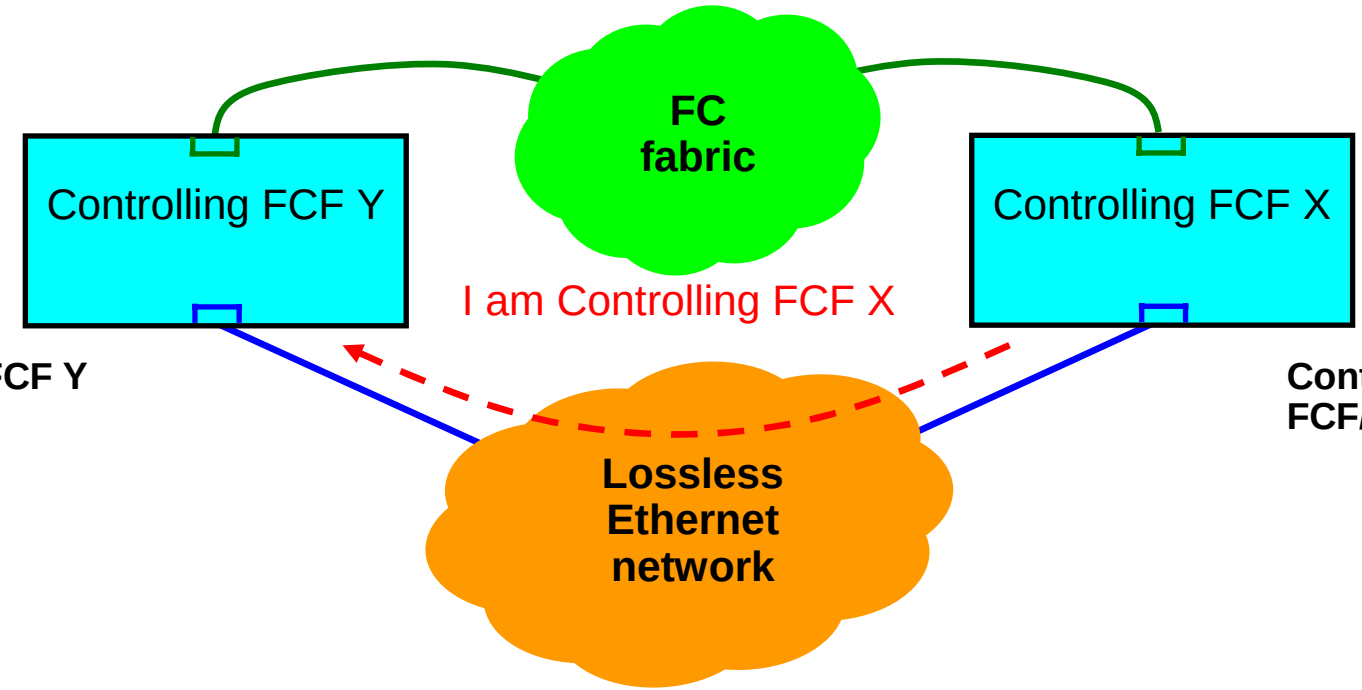

FIP Solicited Advertisement

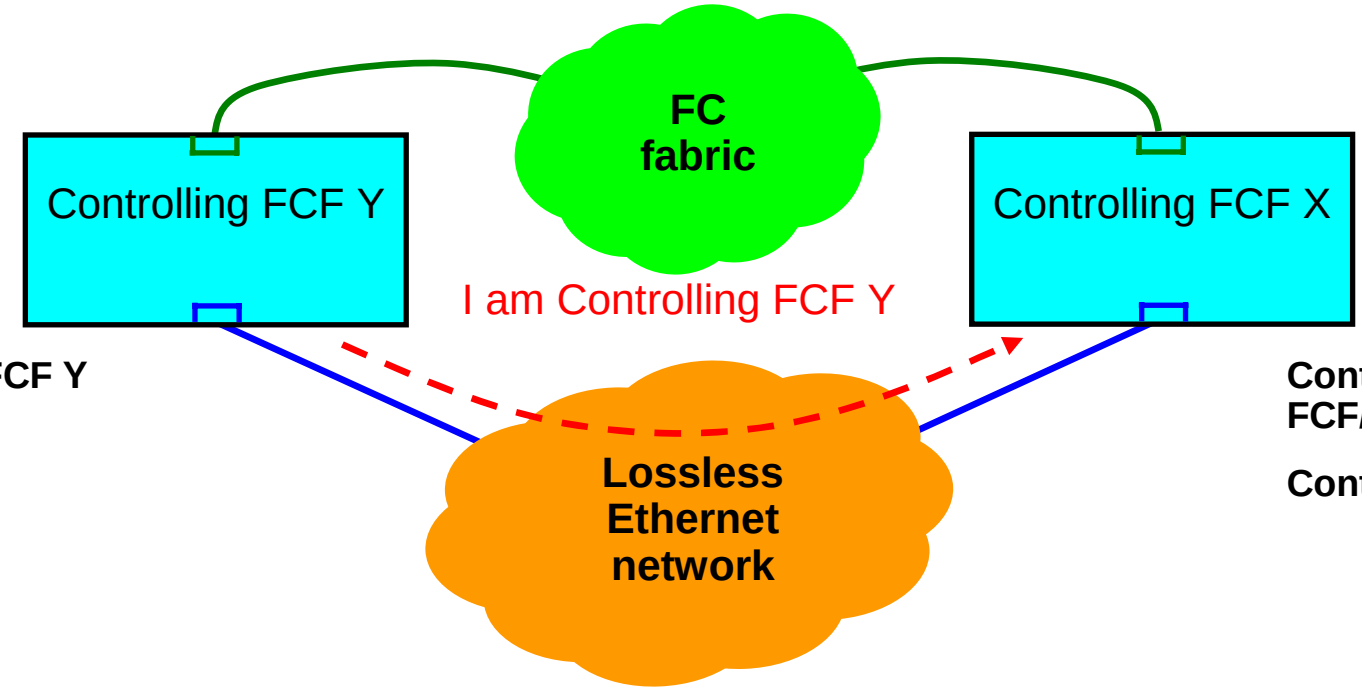

FIP ELP Request

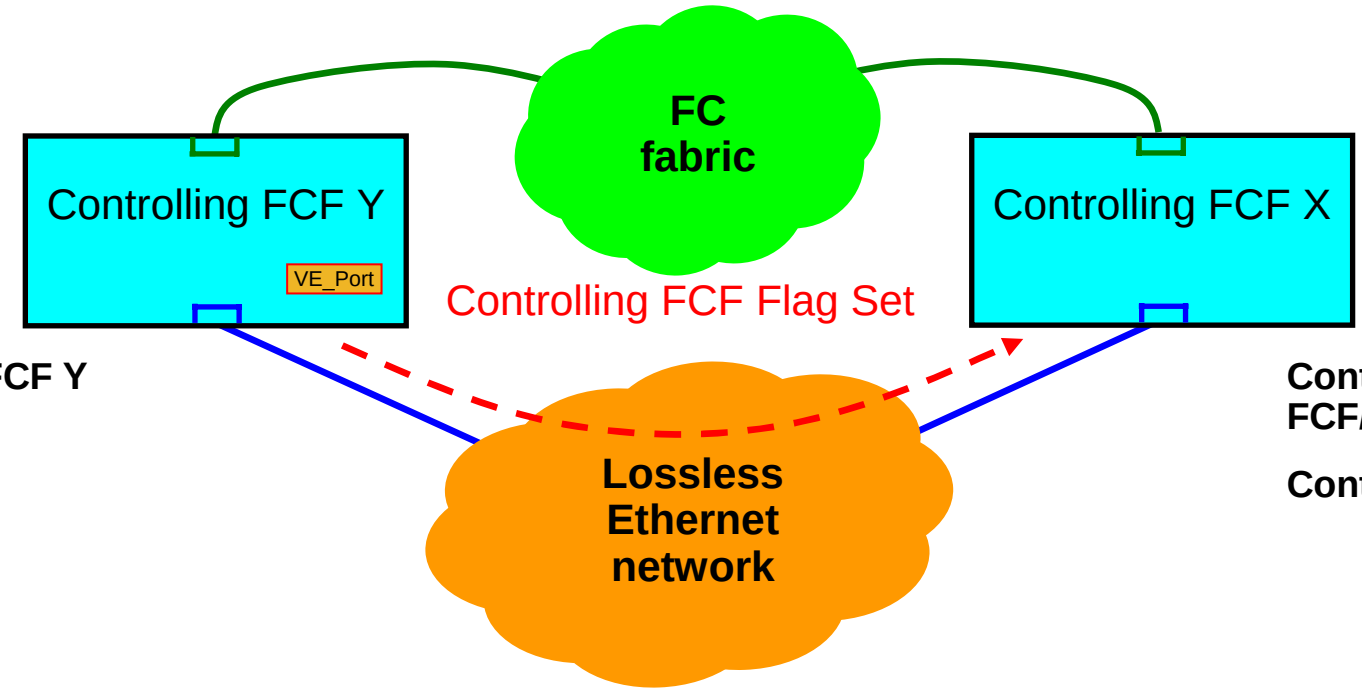

Controlling FCF Y
FCF/FDF list:

Contr. FCF X

Controlling FCF X
FCF/FDF list:

Contr. FCF Y

FIP ELP SW_ACC

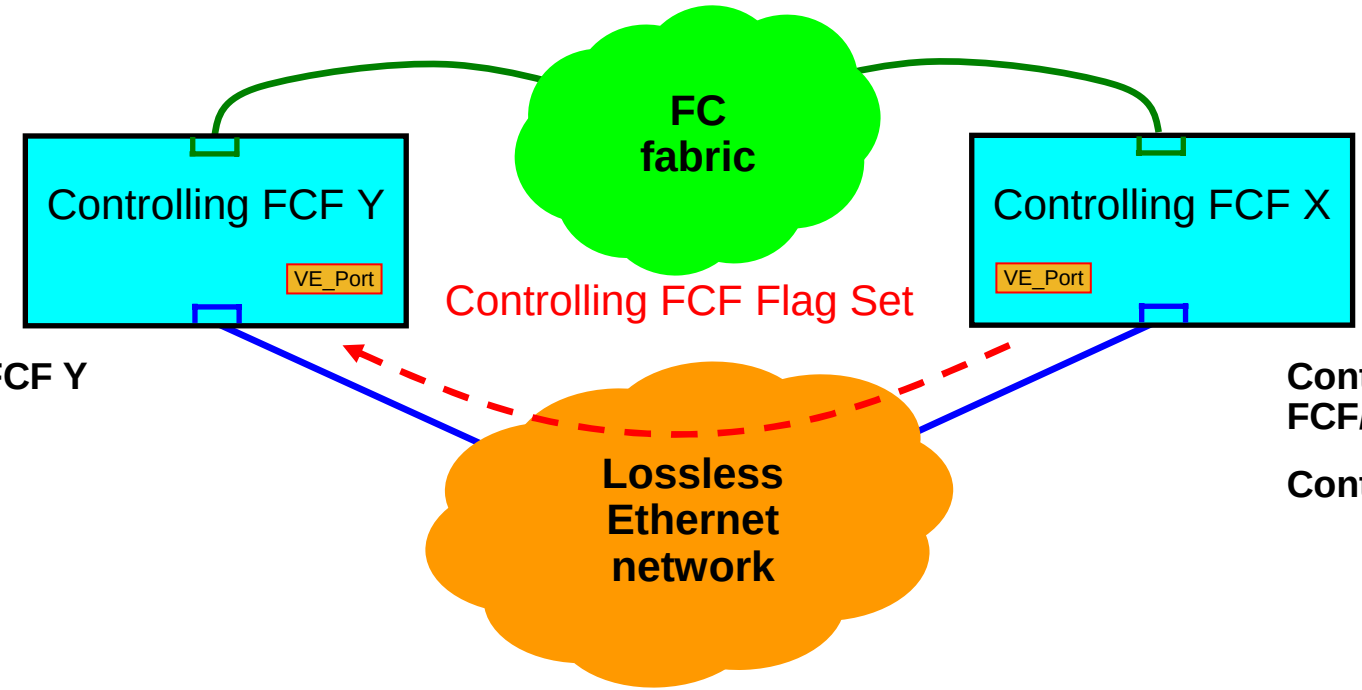

Controlling FCF Y
FCF/FDF list:

Contr. FCF X

Controlling FCF X
FCF/FDF list:

Contr. FCF Y

VE_Port to VE_Port Virtual Link

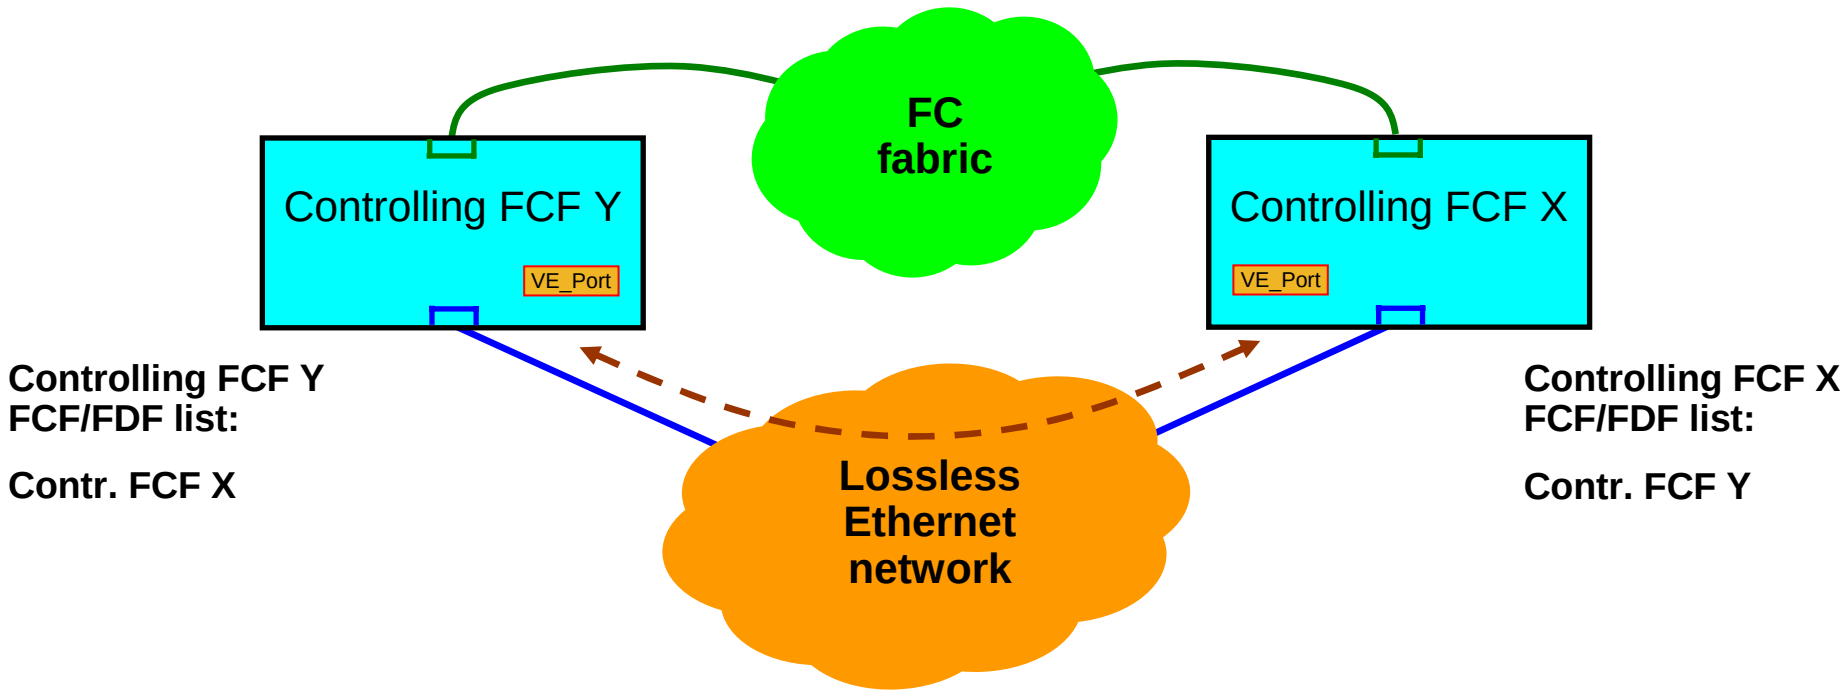

- **The VE_Port to VE_Port Virtual Link between two Controlling FCFs is used to establish a high available configuration**
- **The Controlling FCF with the higher Switch_Name is the Primary, the other one is the Secondary**

High Availability

- **The Primary Controlling FCF obtains an additional Domain_ID from the Principal Switch**
  - **Used as Virtual Domain_ID to allocate N_Port_IDs to FDFs**
- **The Primary Controlling FCF communicates the Virtual Domain_ID to the Secondary Controlling FCF**
  - **Through a FCoE encapsulated SW_ILS**
- **Then the Primary Controlling FCF generates the LSR describing the Virtual Domain Z**
- **Resulting FSPF Topology:**

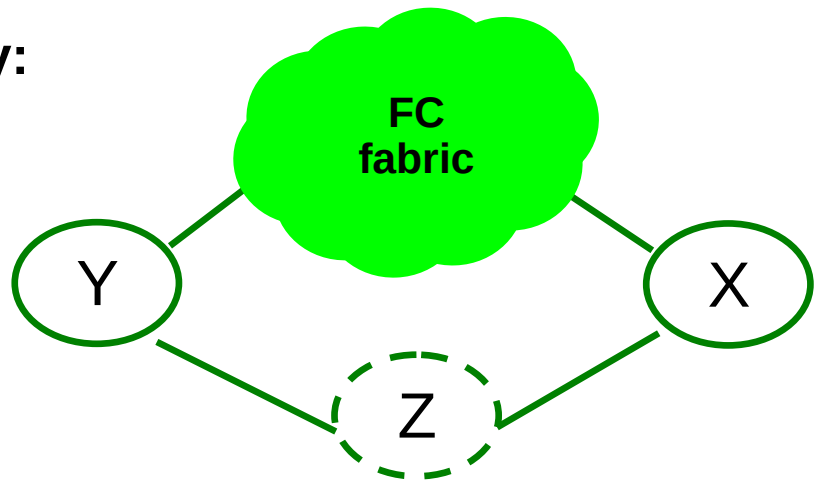

Virtual Domain_ID Z Announcement

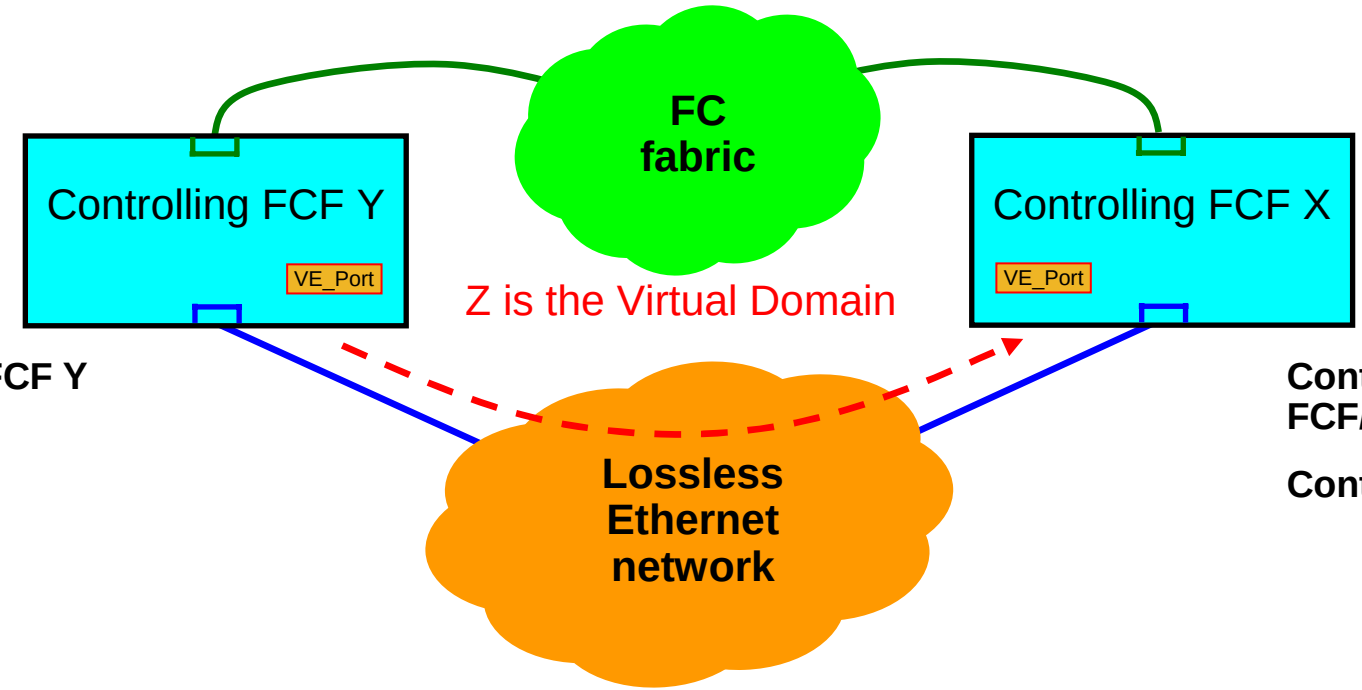

Controlling FCF Y
FCF/FDF list:

Contr. FCF X

Controlling FCF X
FCF/FDF list:

Contr. FCF Y

Virtual Domain Z Announcement SW_ACC

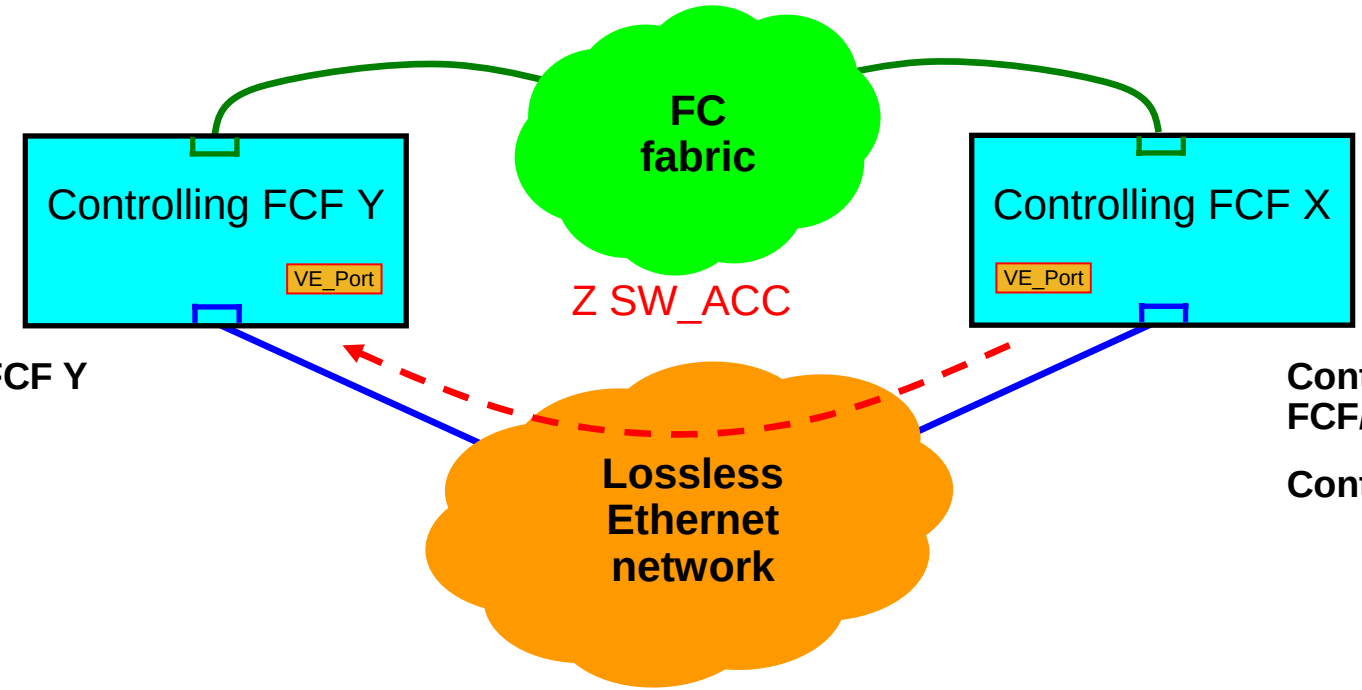

Controlling FCF Y
FCF/FDF list:

Contr. FCF X

Controlling FCF X
FCF/FDF list:

Contr. FCF Y

Advertised FSPF Topology

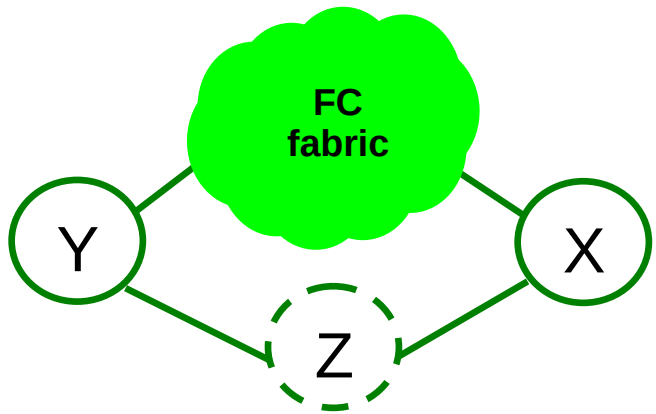

Virtual Domain Announcement SW_ILS

Item	Size
SW_ILS Code	4
Primary Controlling FCF Switch_Name	8
Secondary Controlling FCF Switch_Name	8
Virtual Domain Switch_Name	8
Virtual Domain_ID	4

Virtual Domain Announcement SW_ACC

Item	Size
SW_ACC Code	4
Secondary Controlling FCF Switch_Name	8
Primary Controlling FCF Switch_Name	8
Virtual Domain Switch_Name	8
Virtual Domain_ID	4

Agenda

- **Protocols**
- **Primary to Secondary Controlling FCF**
- **Controlling FCF to FDF**

FDF A

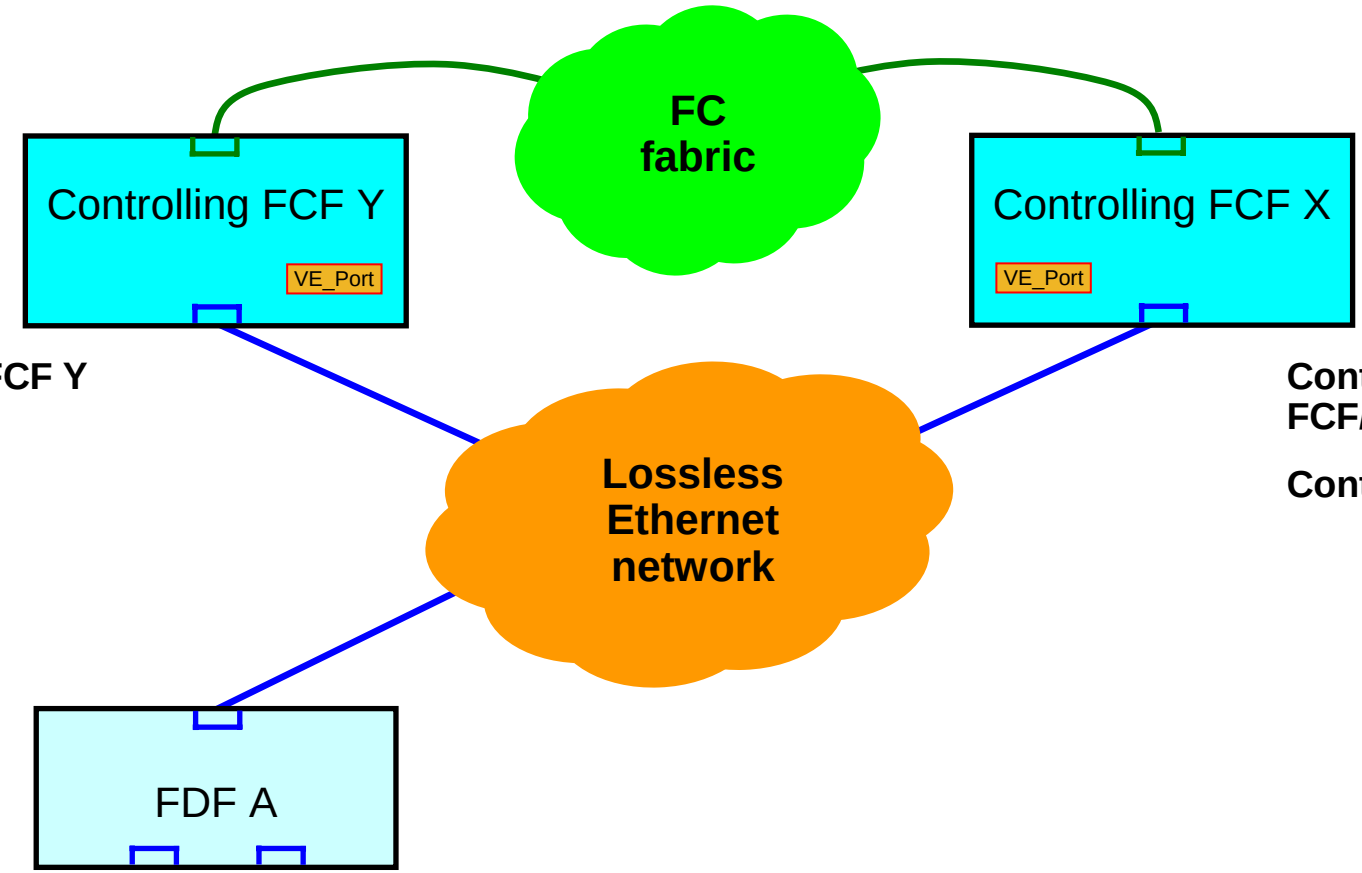

Controlling FCF Y
FCF/FDF list:

Contr. FCF X

Controlling FCF X
FCF/FDF list:

Contr. FCF Y

FDF A FIP Solicitation

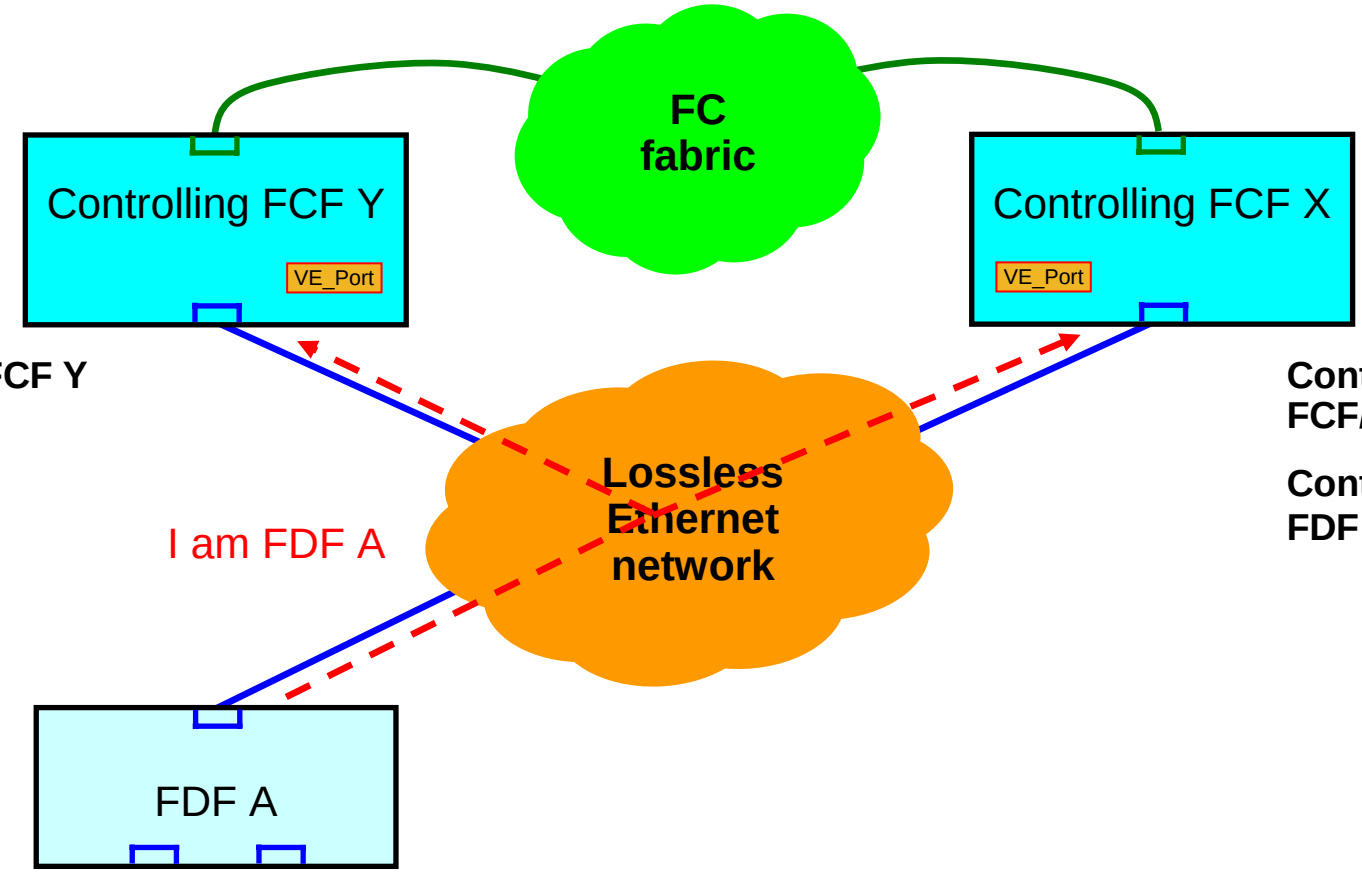

FIP Solicited Advertisements

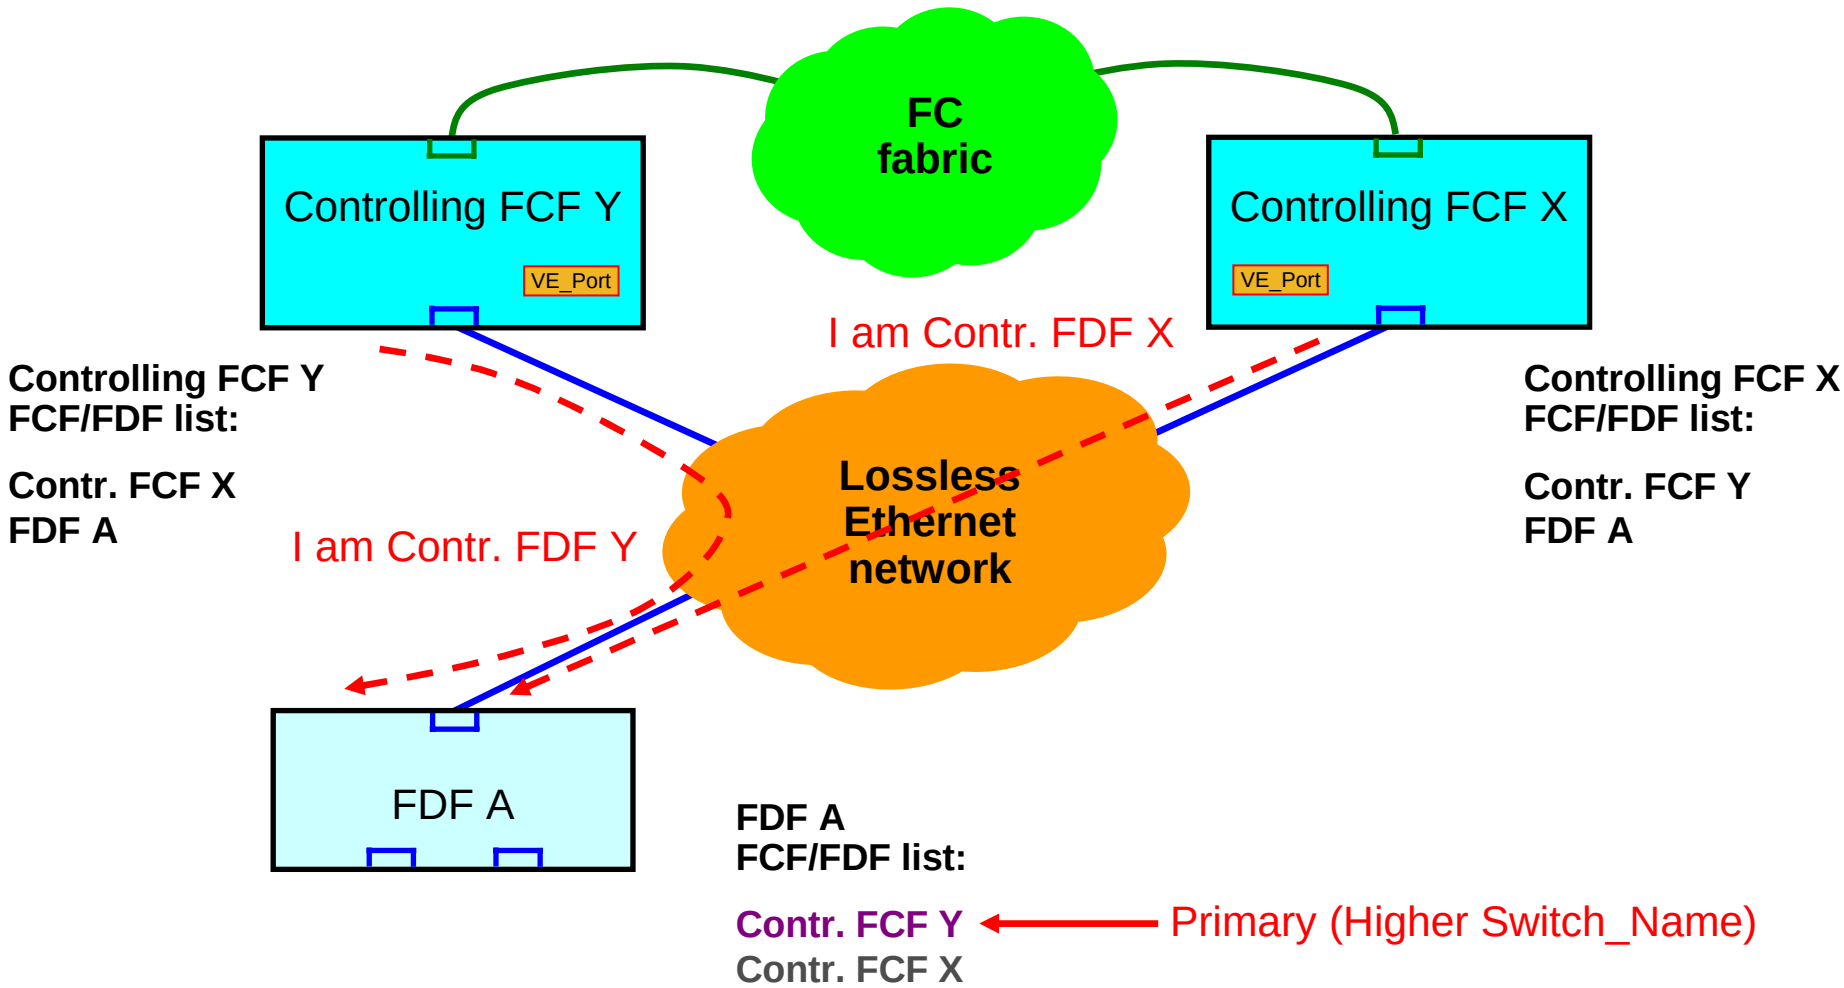

FIP ELP Requests

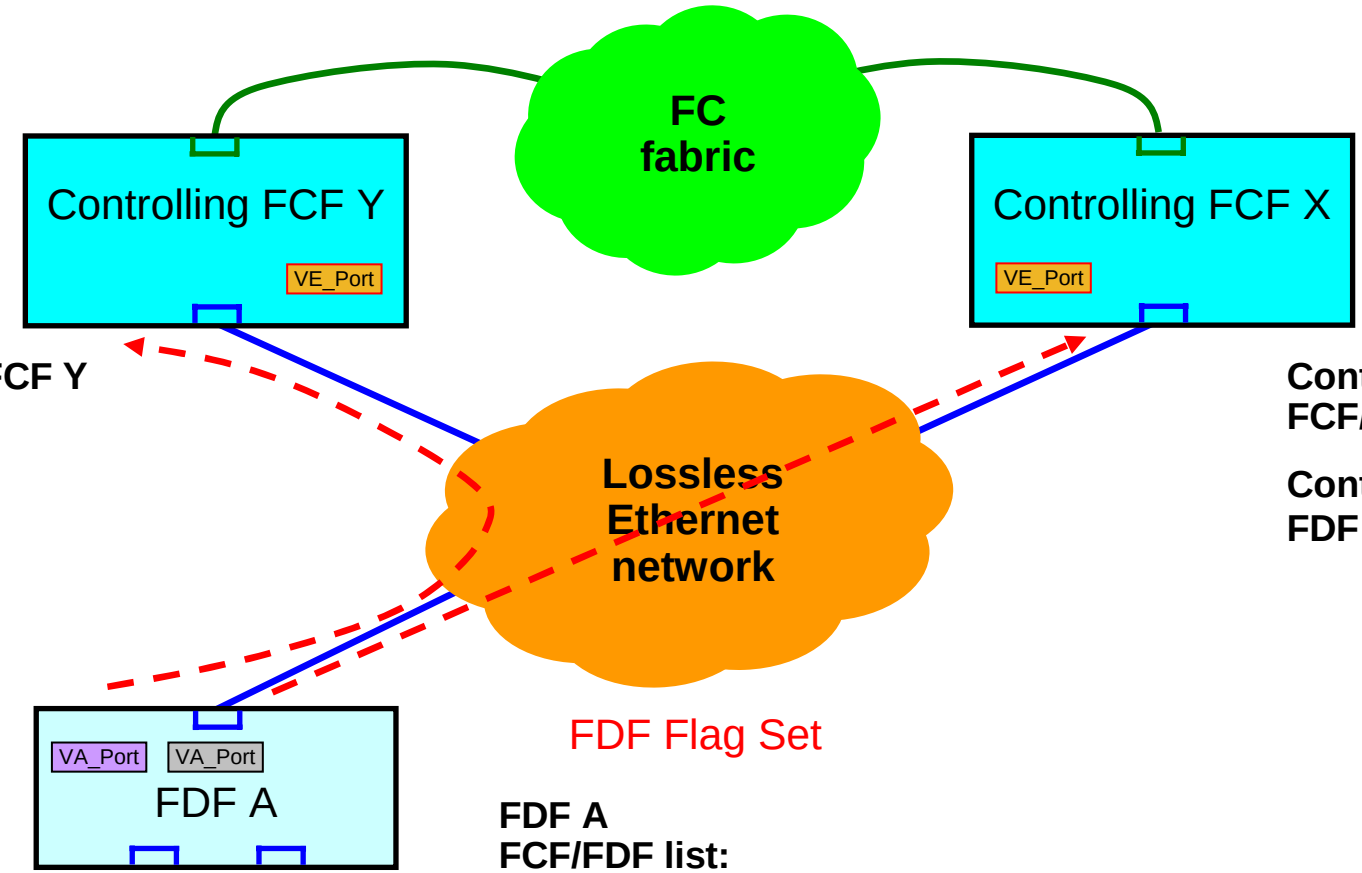

FIP ELP SW_ACC

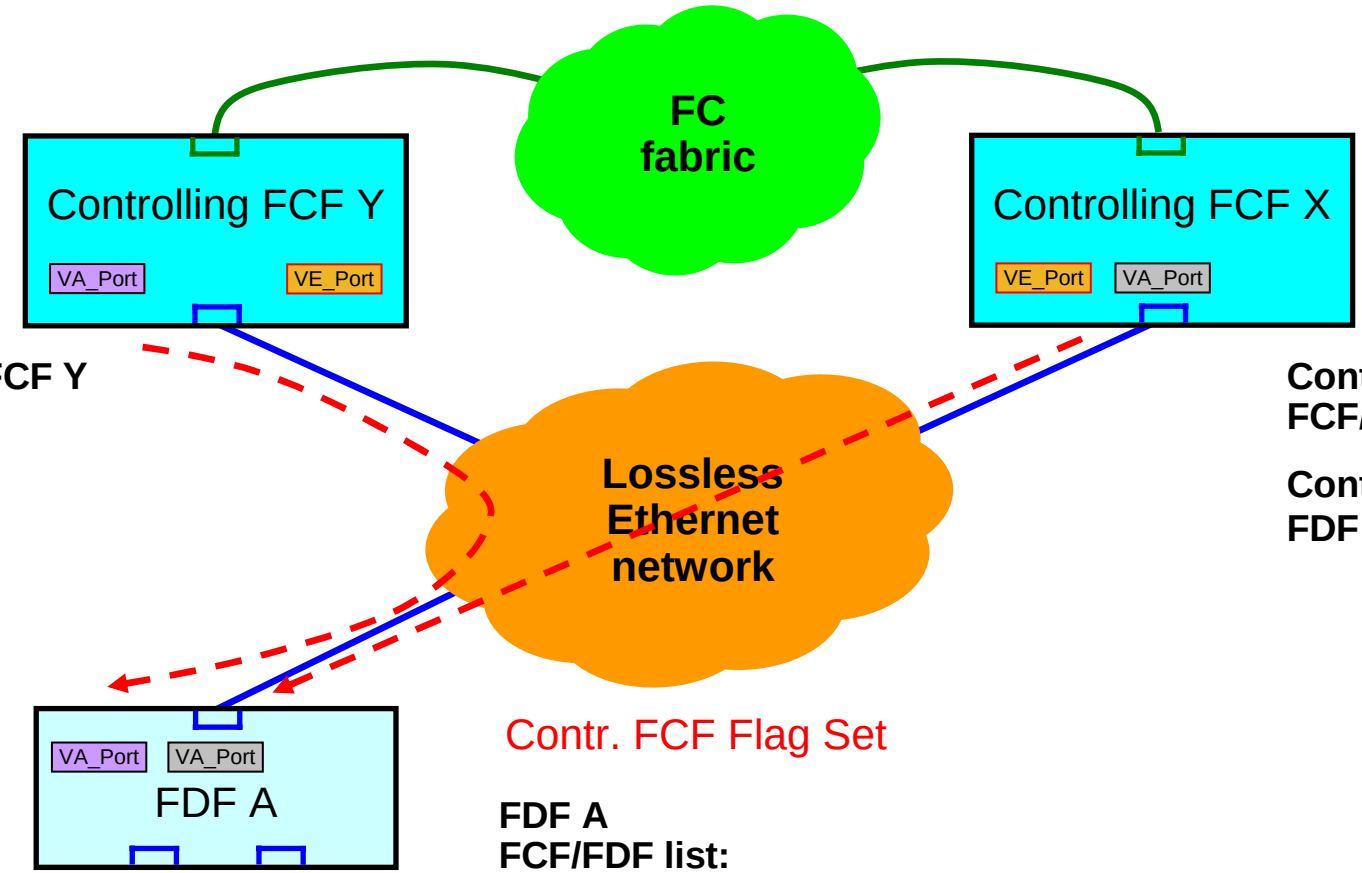

VA_Port to VA_Port Virtual Links

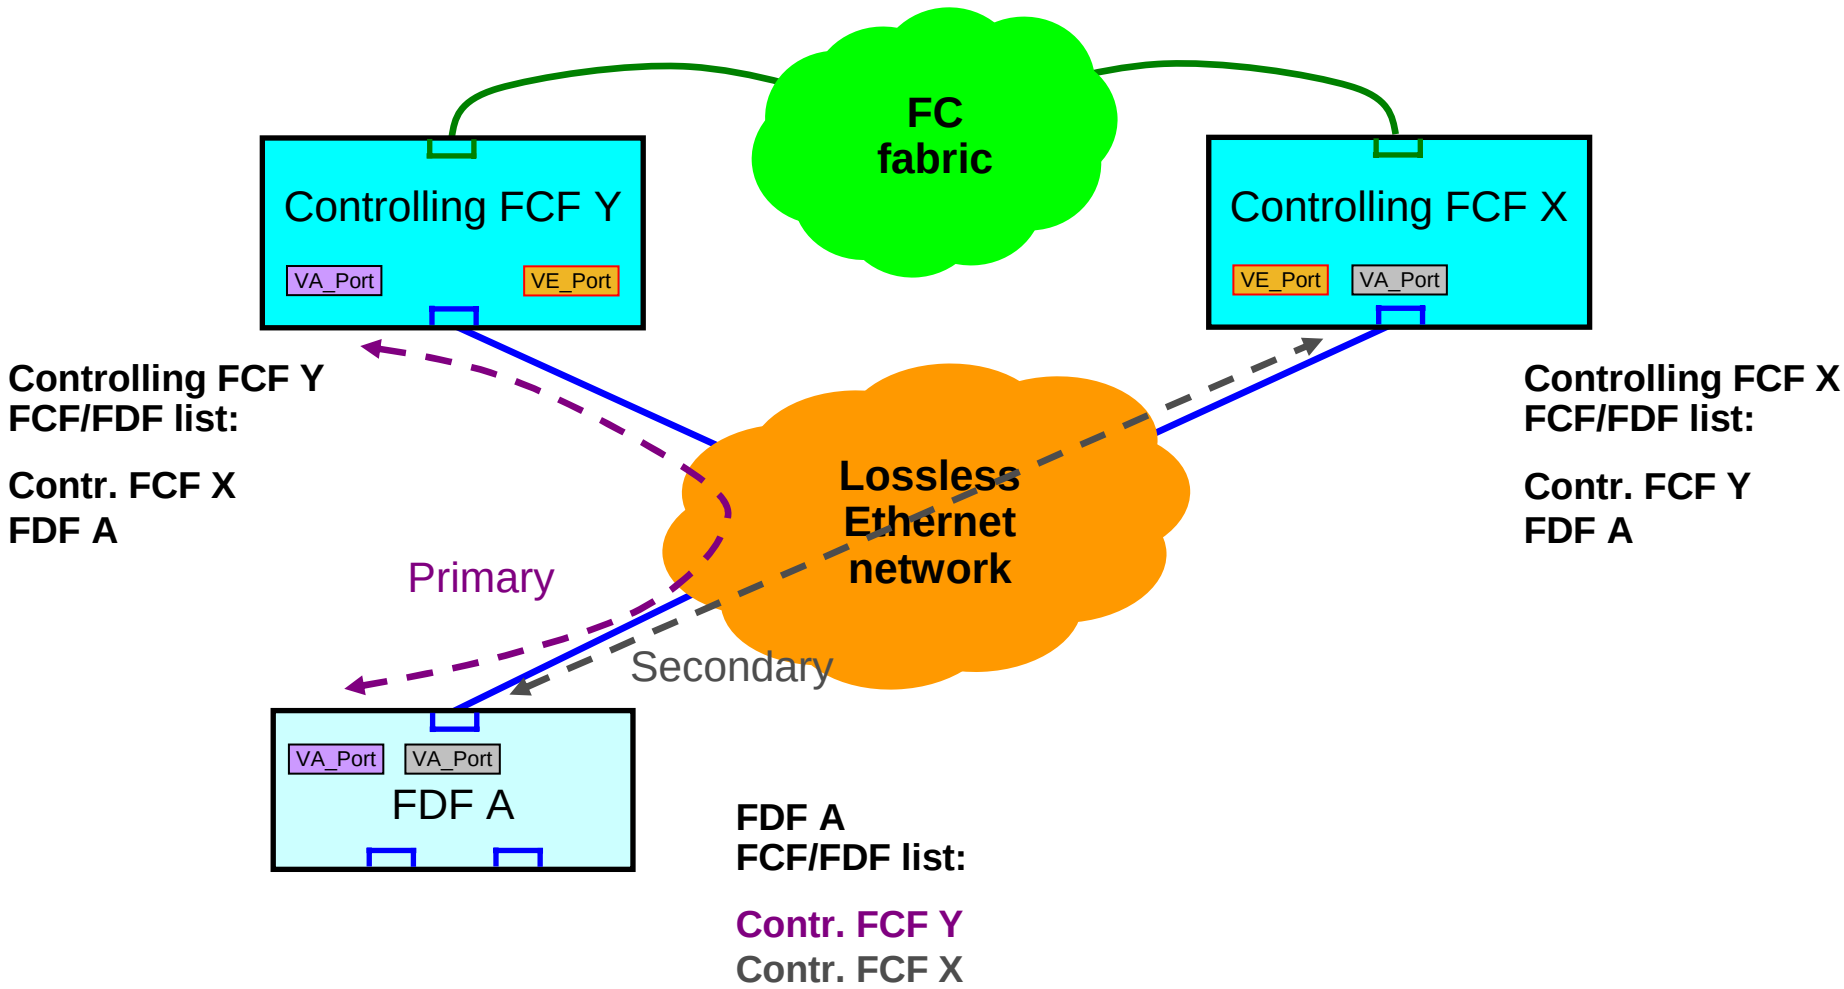

FDF B

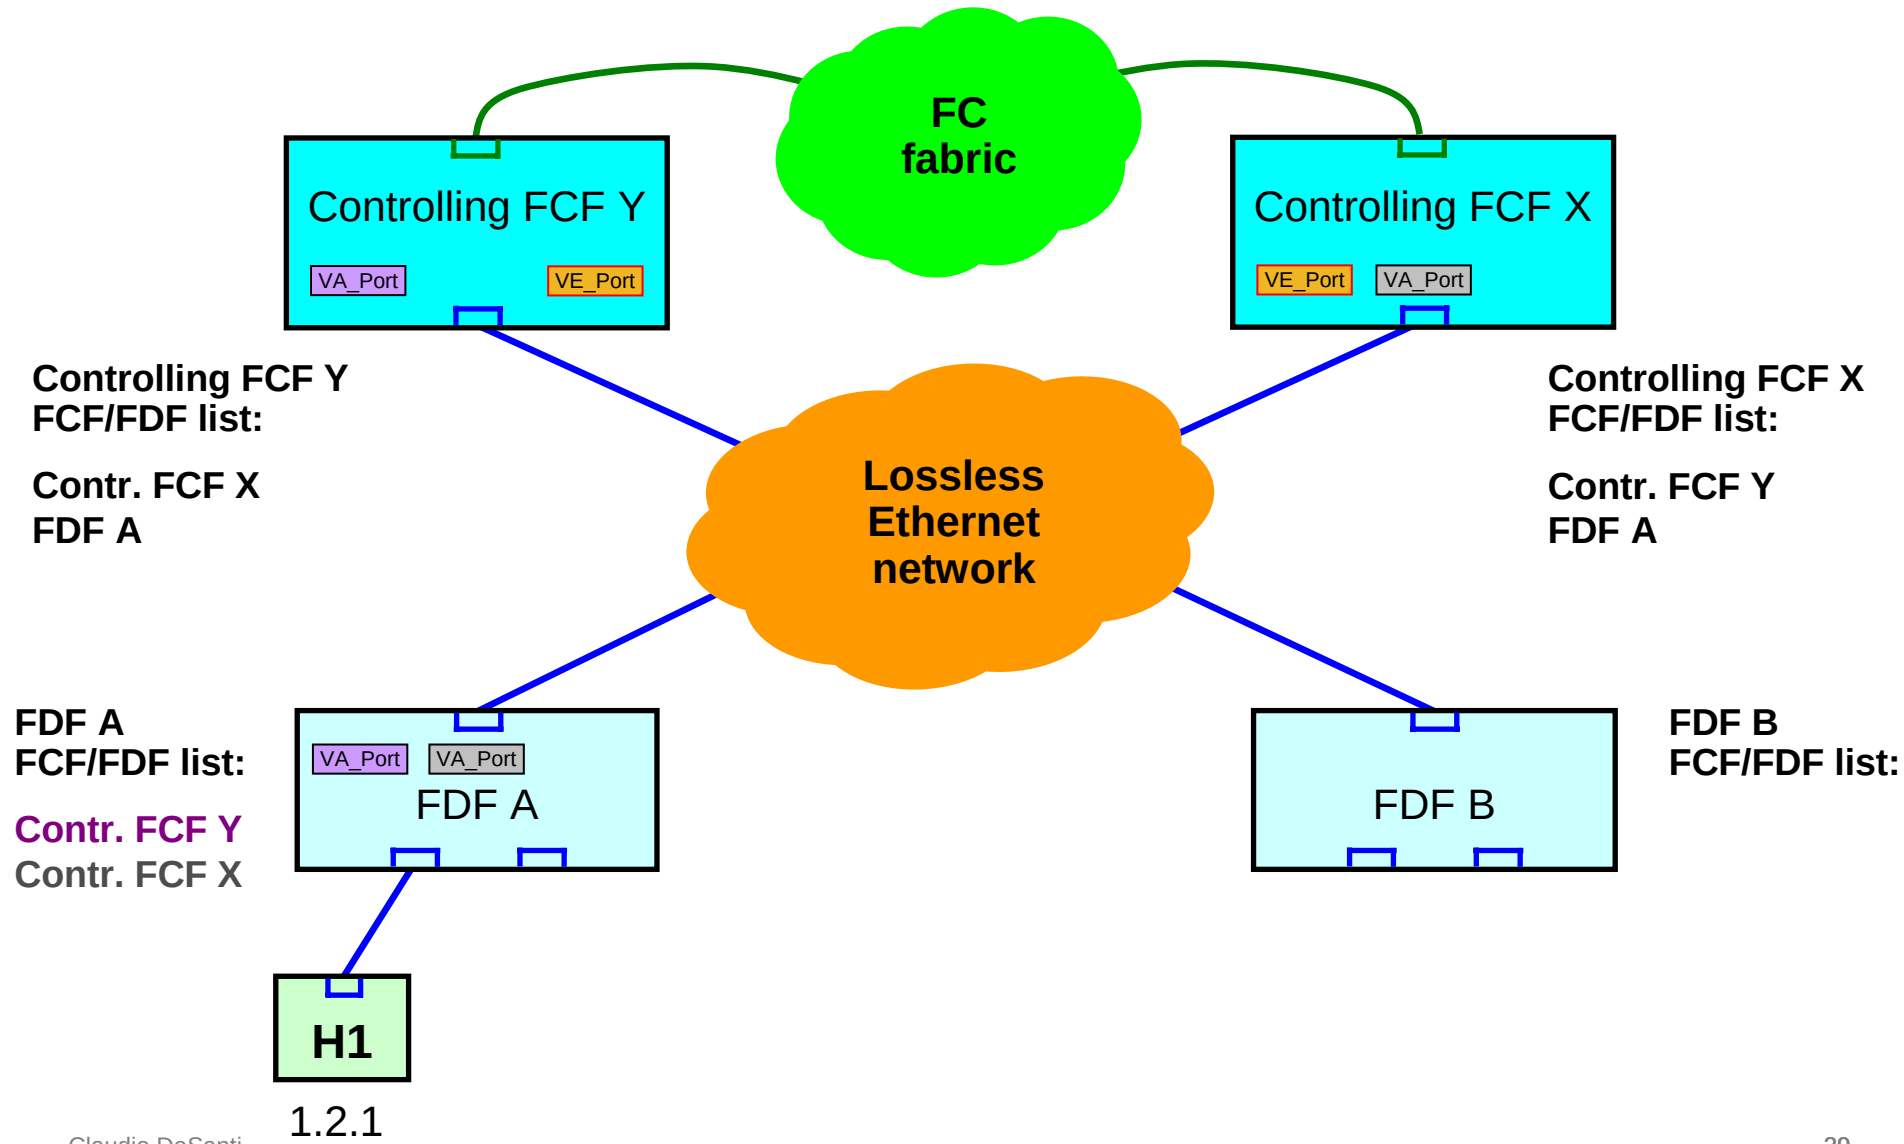
FC
fabric
Controlling FCF Y
VA_Port
VE_Port
Controlling FCF X
VE_Port
VA_Port
Controlling FCF Y
FCF/FDF list:
Contr. FCF X
FDF A
Lossless
Ethernet
network
Controlling FCF X
FCF/FDF list:
Contr. FCF Y
FDF A
FDF A
FCF/FDF list:
Contr. FCF Y
Contr. FCF X
VA_Port
VA_Port
FDF A
FDF B
FDF B
FCF/FDF list:
H1
1.2.1

FDF B FIP Solicitation

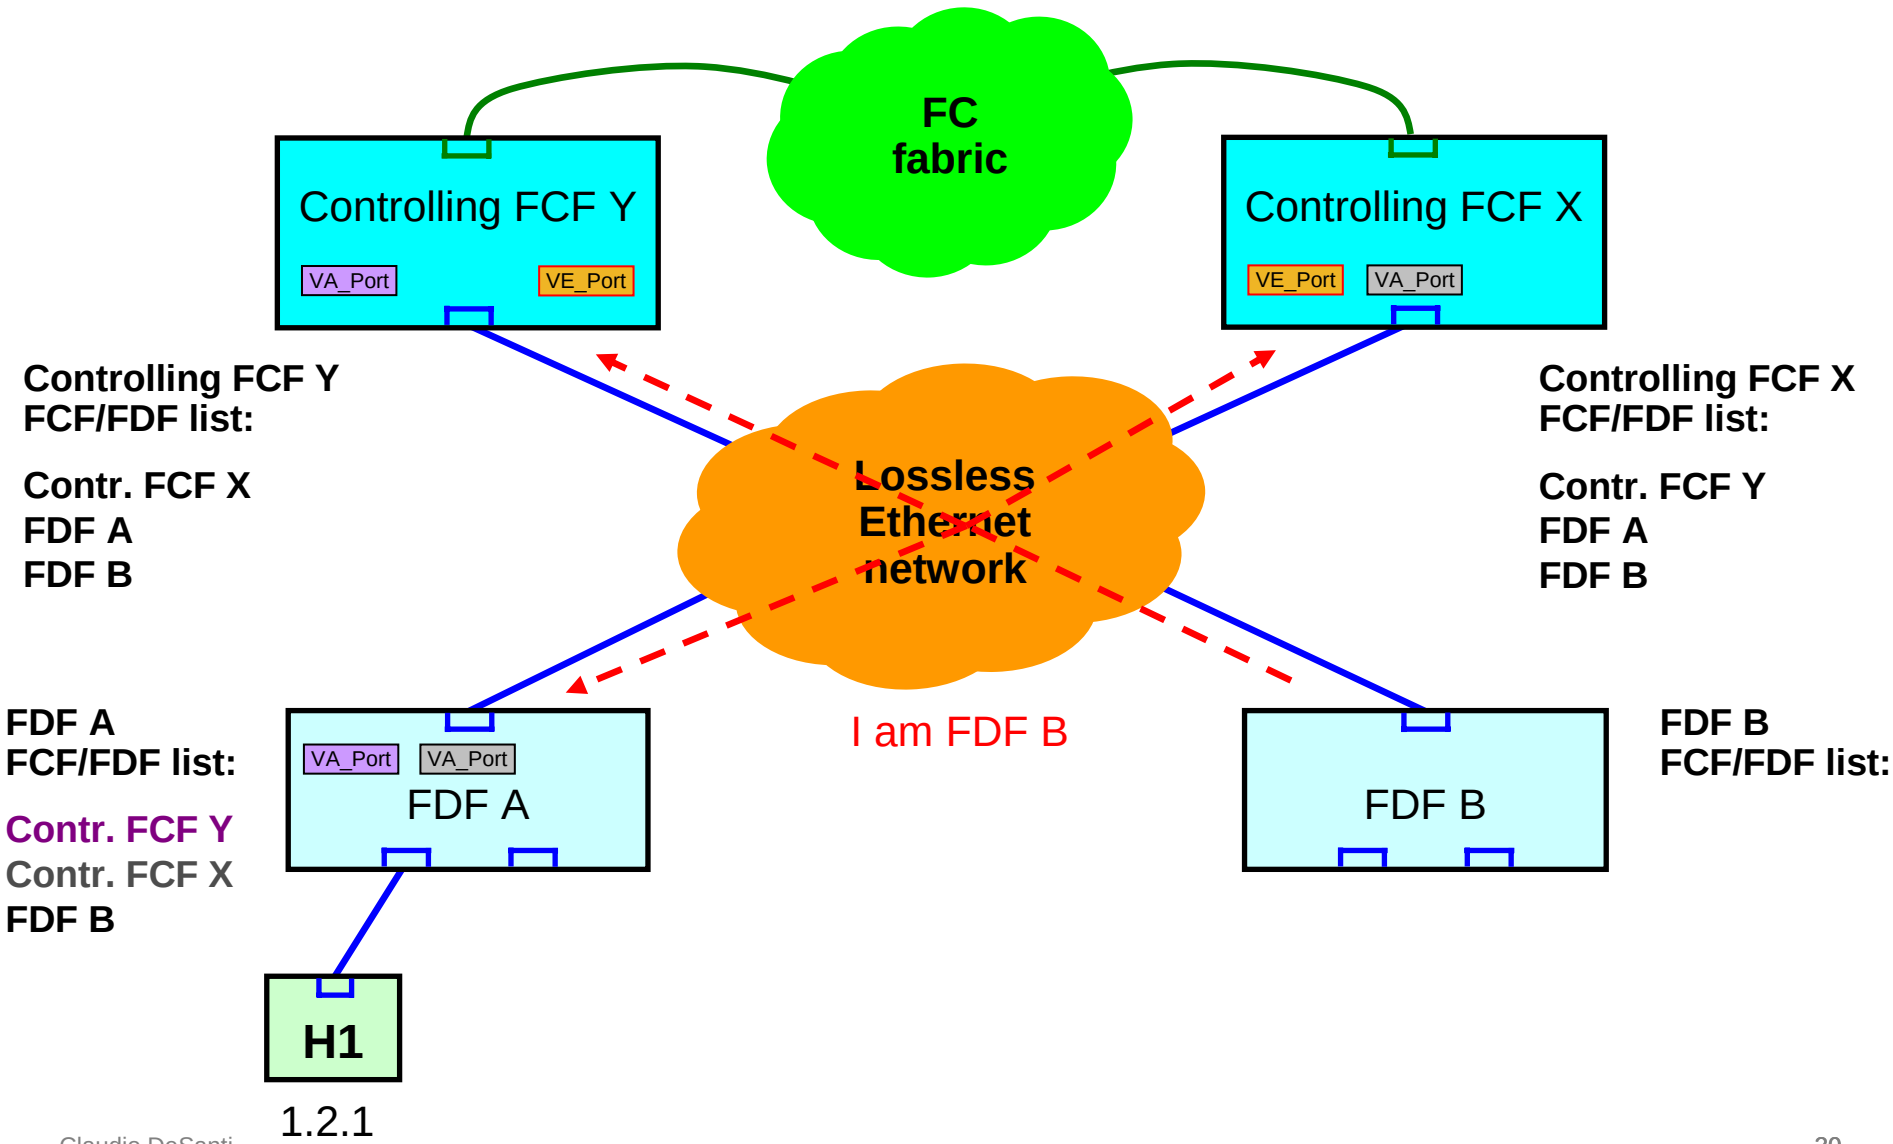

FIP Solicited Advertisements

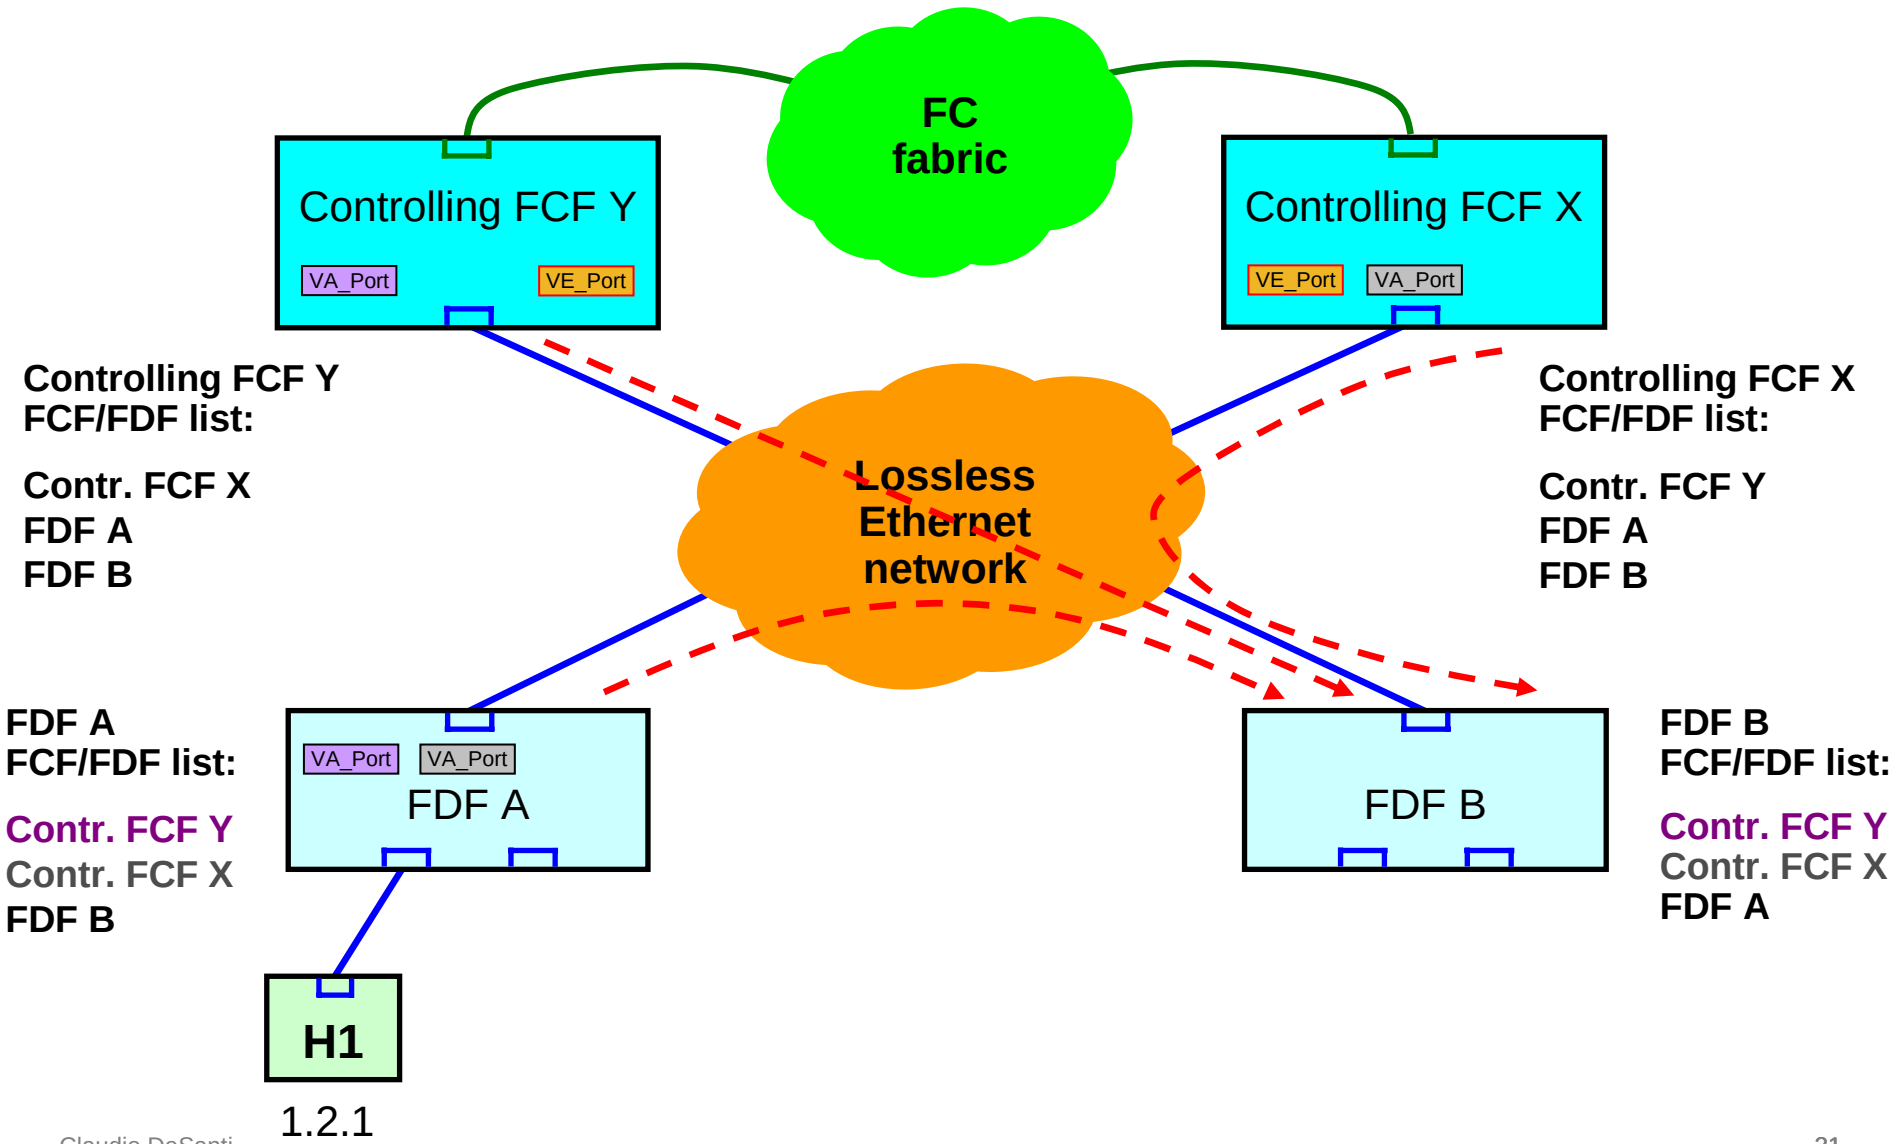

FDF B FIP ELP Requests

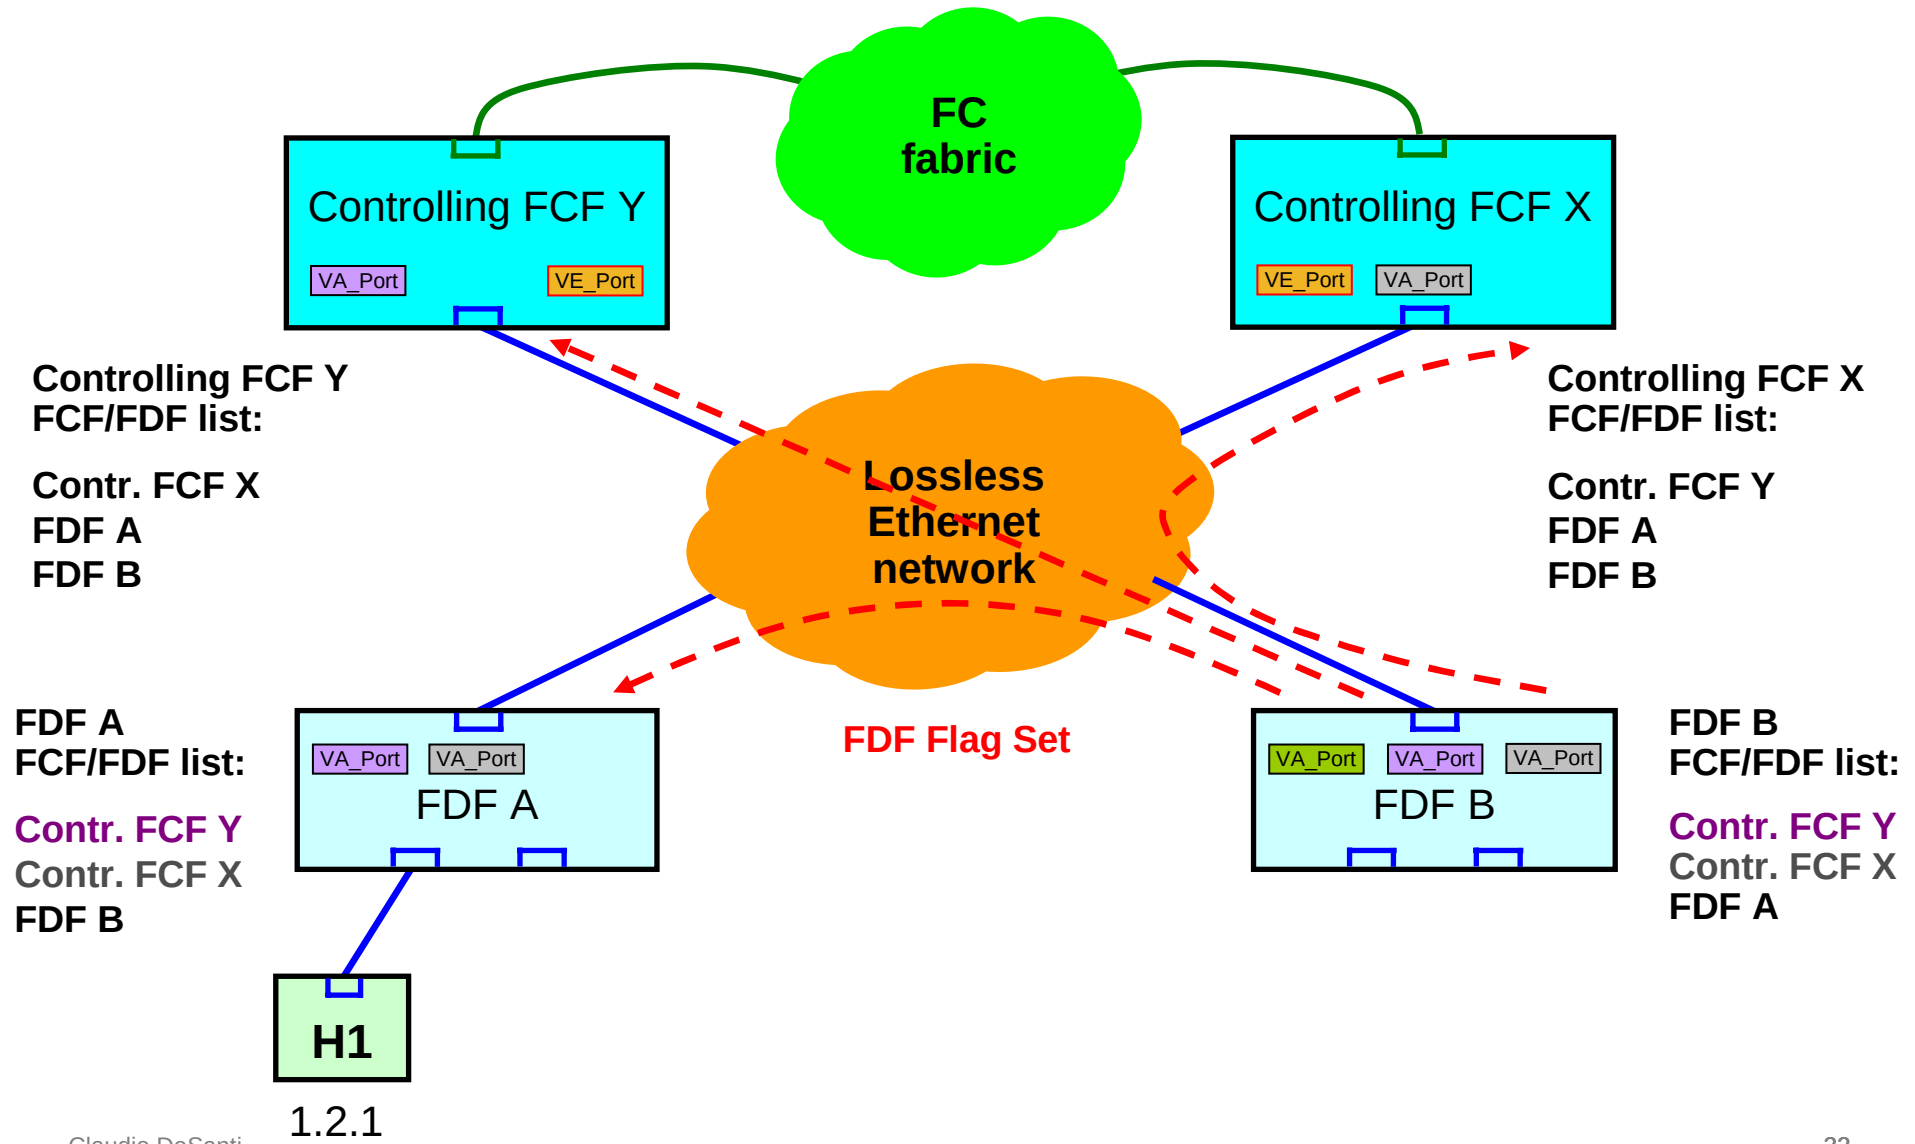

FDF B FIP ELP SW_ACC

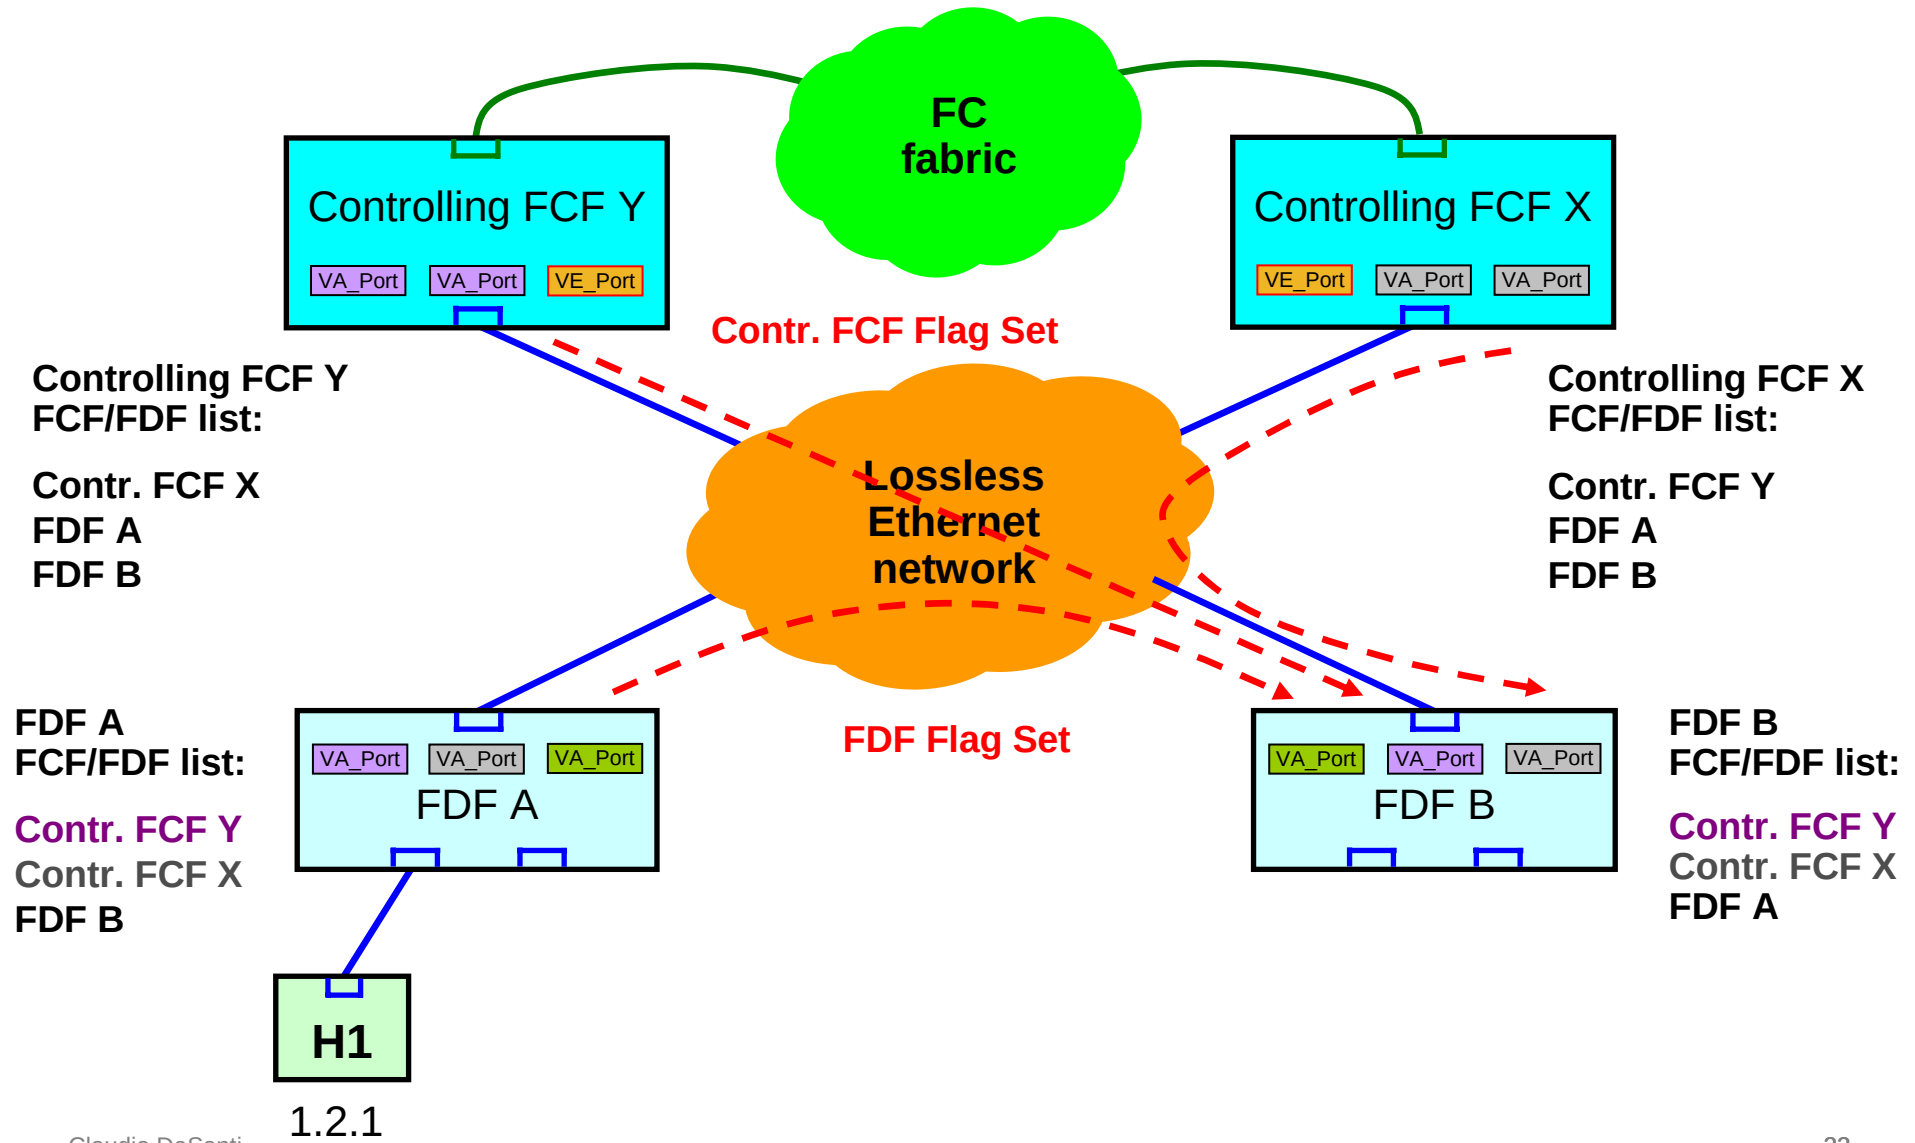

Virtual Links

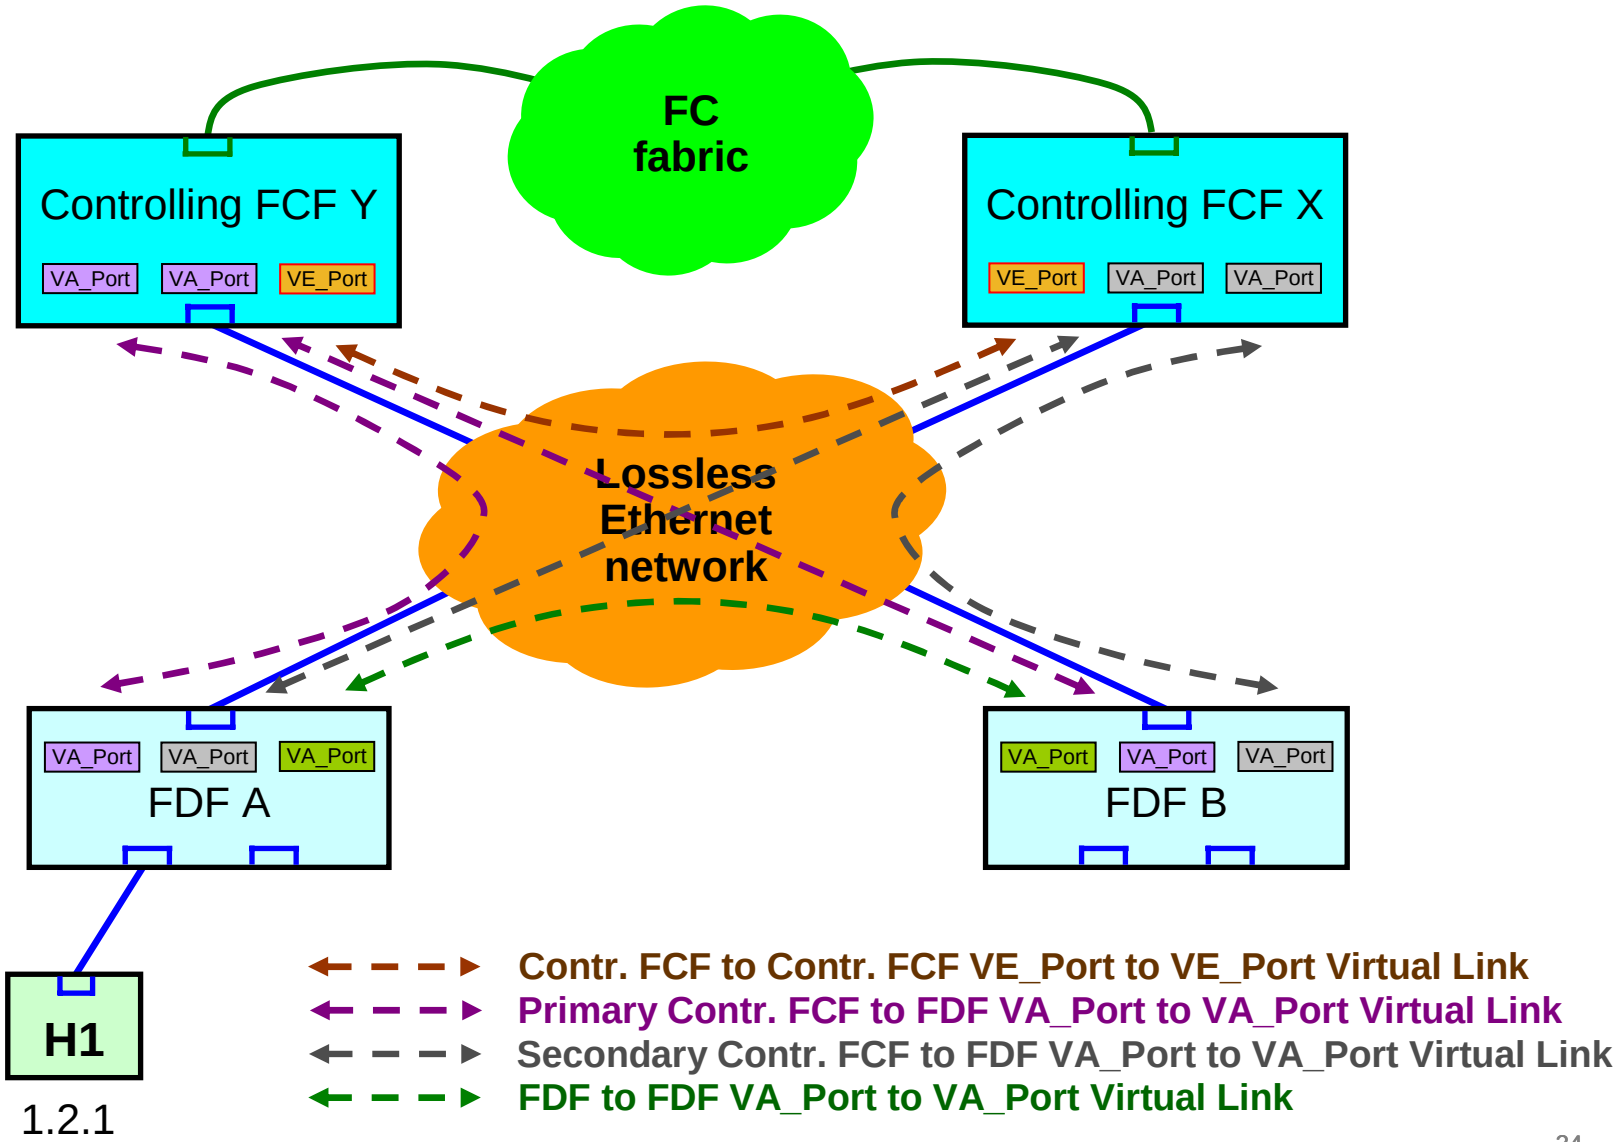

Thank You

Total Solar Eclipse July 11, 2010
As seen by
Landon Curt Noll and Claudio DeSanti