




VA_Ports: FDF / Controlling FCF Protocols



Claudio DeSanti

T11/10-271v4, December 2010



Agenda

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

Distributed FCF

- **Many Controlling FCFs and FDFs can be connected by a Lossless Ethernet network**
- **A Primary/Secondary pair of Controlling FCF plus a defined set of FDFs constitutes a Distributed FCF**
- **A Distributed FCF is defined by configuration of the Controlling FCFs, listing:**
 - which pair of Controlling FCFs is the Primary/Secondary pair for that Distributed FCF**
 - which FDFs are part of that Distributed FCF**
- **Controlling FCFs should operate in Primary/Secondary pairs for high availability**
 - The Secondary replaces the Primary in case of failure**



Image Credit:
Flickr user Muffet
Creative Commons License

Primary/Secondary Controlling FCFs

- A Primary/Secondary pair of Controlling FCFs instantiates a VE_Port to VE_Port Virtual Link between themselves
- The VE_Port to VE_Port Virtual Link is used to keep synchronized between Primary and Secondary:
 - The FDF topology
 - The N_Port_ID allocation state
- Controlling FCFs instantiate VA_Port to VA_Port Virtual Links with the controlled FDFs
- The Primary Controlling FCF define the routes for the FDF topology and performs N_Port_ID allocations for all its controlled FDFs



Image Credit:
Flickr user Muffet
Creative Commons License



The FDF (1)

- **FDFs tell the controlling FCF(s) about the Virtual Link they establish**
- **An FDF first has to try instantiate a VA_Port to VA_Port Virtual Link with the Controlling FCF**

A connection with the “brain” is needed

- **This enables these FDFs to establish FDF to FDF Virtual Links**
- **Upon instantiating a VA_Port to VA_Port Virtual link with another FDF, the FDF sends a “FDF reachable” message to the Controlling FCF**

In this way the Controlling FCFs maintain an updated topology of the network and compute updated routes

The Controlling FCFs distribute the computed routes to the FDFs

- **An FDF “connected” with the Controlling FCF relays the messages from/to other FDFs**



The FDF (2)

- Upon de-instantiating a VA_Port to VA_Port Virtual link with another FDF, the FDF sends a “FDF unreachable” message to the Controlling FCF

In this way the Controlling FCFs maintain an updated topology of the network and compute updated routes

The Controlling FCFs distribute the computed routes to the FDFs

- An FDF “connected” with the Controlling FCF relays the messages from/to other FDFs



The FDF (3)

- On establishing a VN_Port to VF_Port Virtual Link, an FDF encapsulates the parameters of the received FIP FLOGI or FIP NPIV FDISC Request in a “VN_Port Reachable” 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 (4)

- **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

- 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 Protocol Code = 0001h		Reserved	SubCode = 01h									
Descriptor List Length = 8		F P	S P	Flags			F D	F C				F
Type = 2	Len = 2	FCF-MAC Address										
Type = 3	Len = 2	Reserved										
Reserved	FC-MAP											
Type = 4	Len = 3	Reserved										
Switch_Name												
Type = 6	Len = 1	Max FCoE Size										

FIP Advertisement

FIP Protocol Code = 0001h		Reserved	SubCode = 02h
Descriptor List Length = 12		F P	S P
		Flags	F D
			F C
			A S F
Type = 1	Len = 1	Reserved	Priority
Type = 2	Len = 2		
FCF-MAC Address			
Type = 4	Len = 3	Reserved	
Switch_Name			
Type = 5	Len = 4	Rsvd	VF_ID
Reserved	FC-MAP		
Fabric_Name			
Type = 12	Len = 2	Reserved	D
FKA_ADV_PERIOD			
FIP Pad to Max FCoE Size of soliciting entity, if solicited (i.e., if S=1b), otherwise no FIP Pad			

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**
- **VA_Port to VA_Port Virtual Link:**
 - No EFP, FSPF, Zone Merge**
 - Just FDF protocols (4 messages)**
- **When a VA_Port to VA_Port Virtual Link is established, the primary controlling FCF provides to the FDFs updated route and N_Port_ID allocation information**

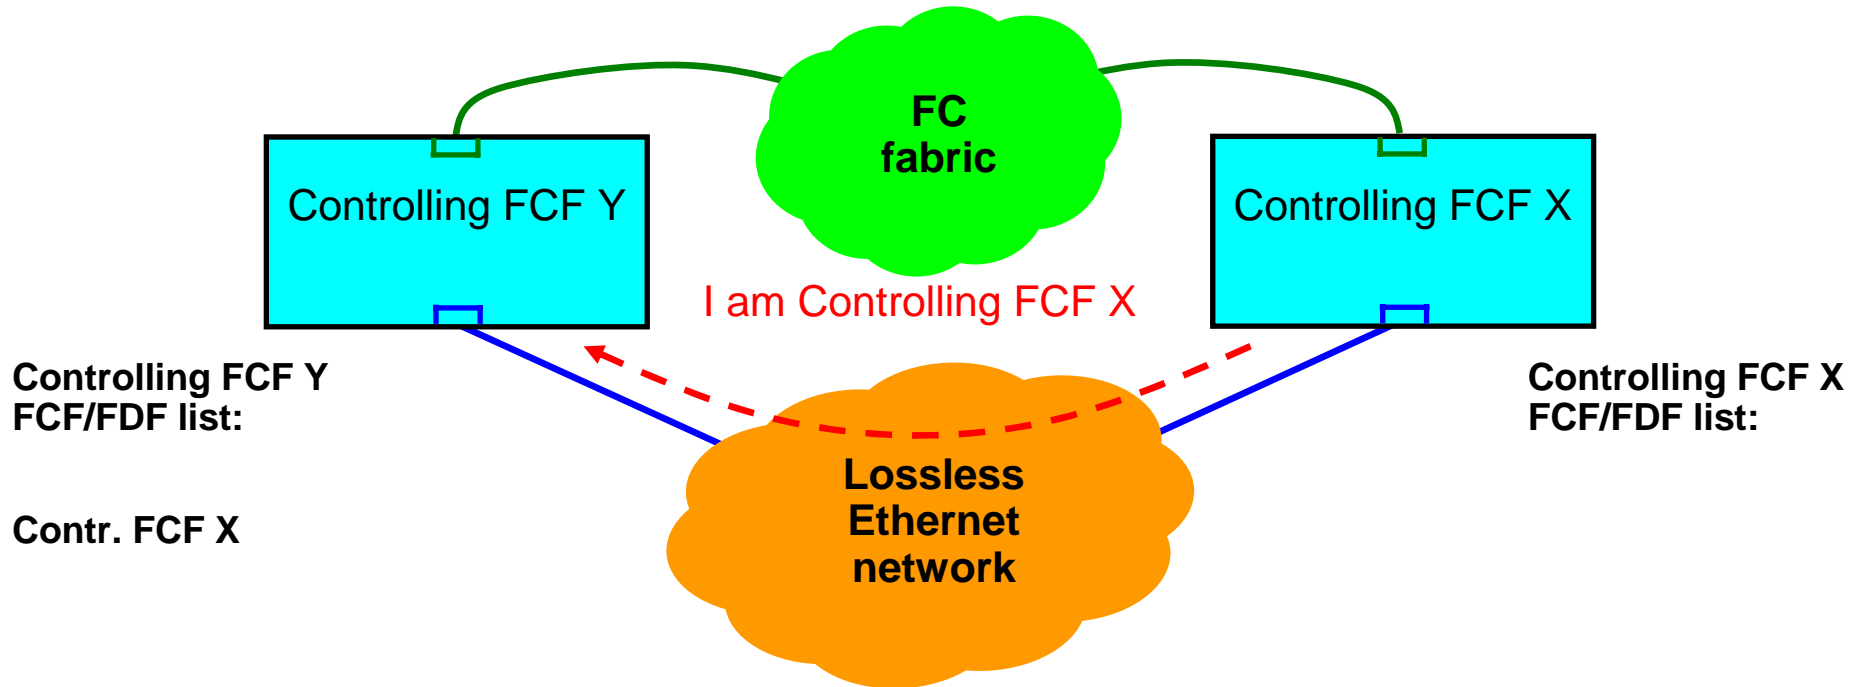
Reaching other Domain_IDs

- **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)

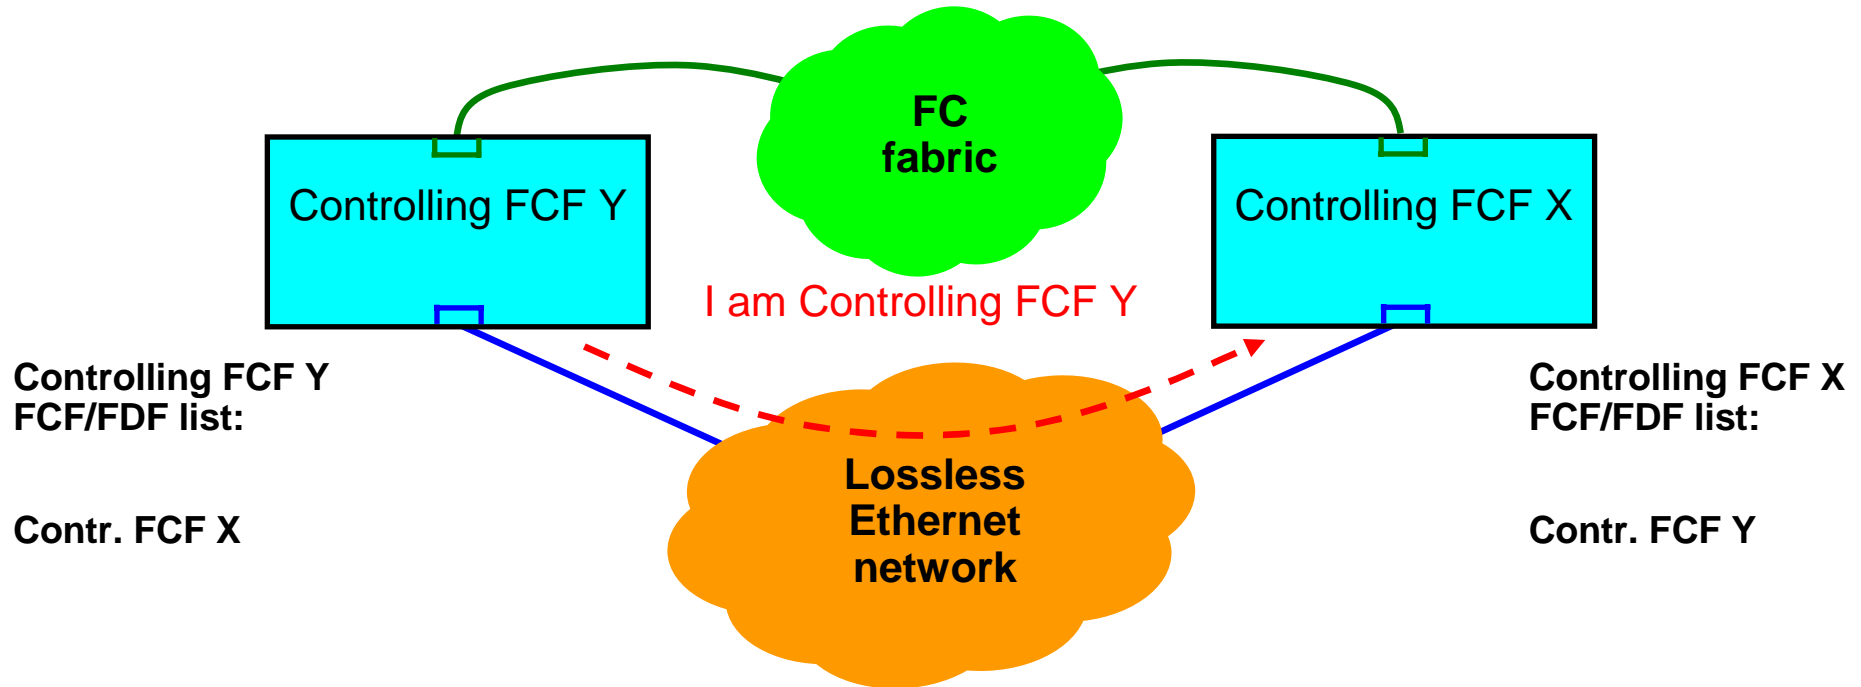
Agenda

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

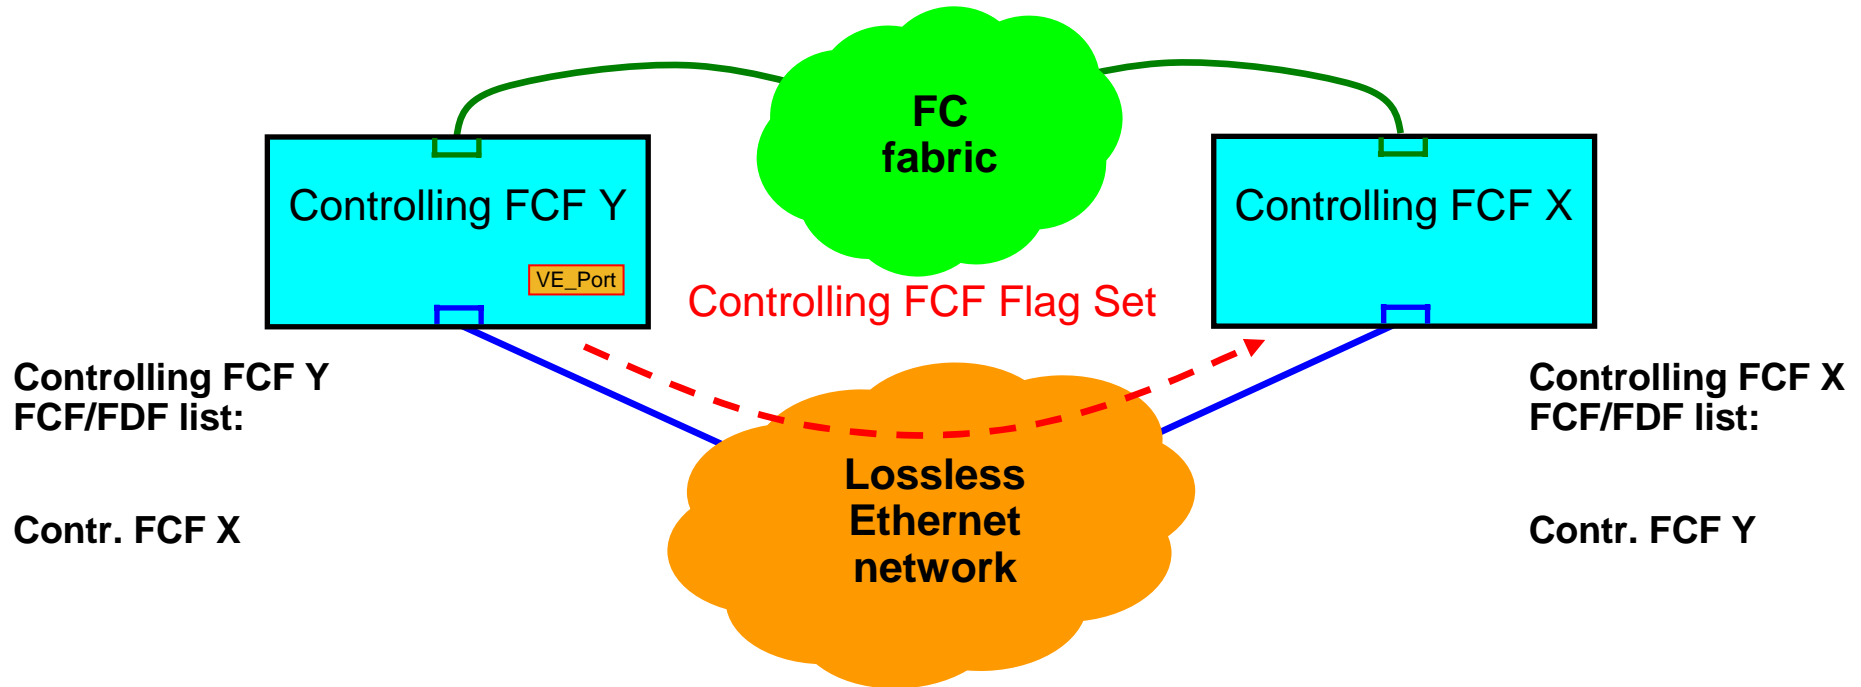
FIP Solicitation



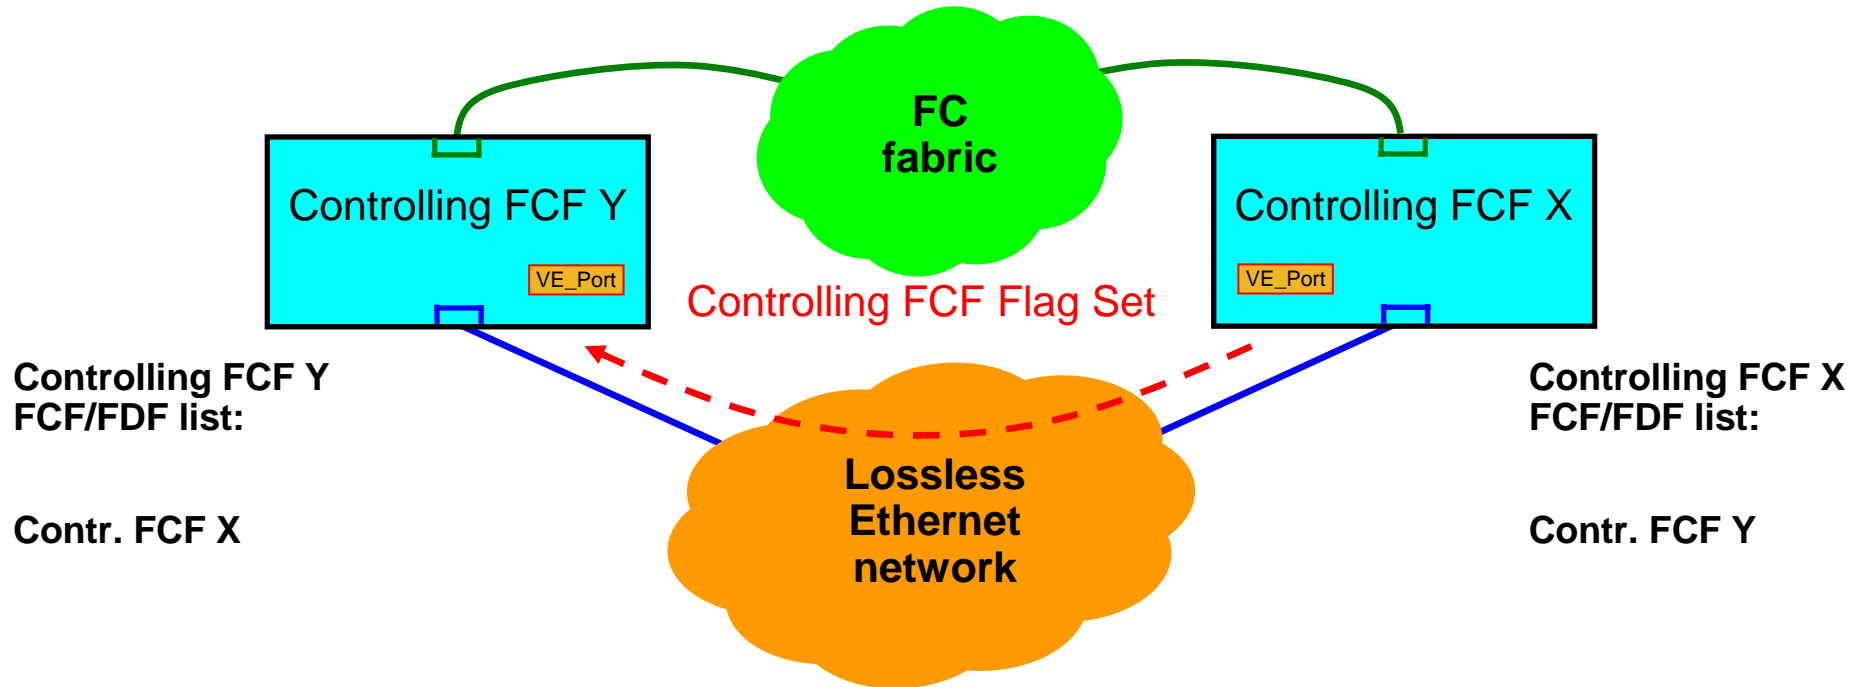
FIP Solicited Advertisement



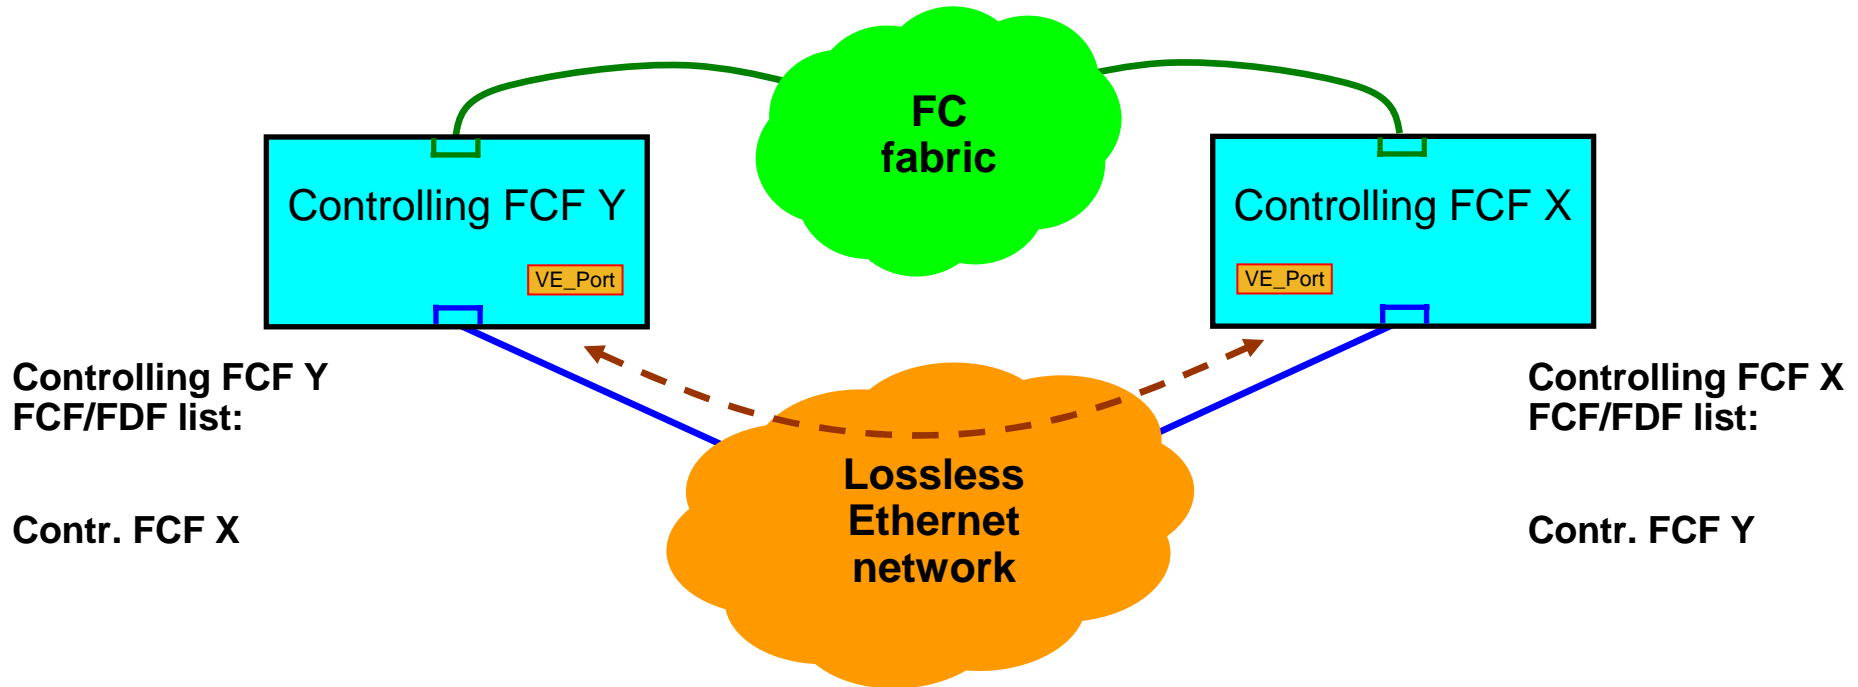
FIP ELP Request



FIP ELP SW_ACC



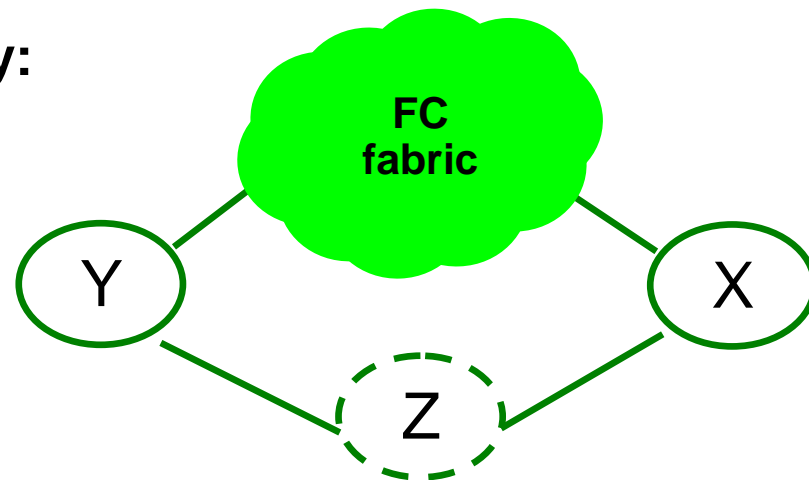
VE_Port to VE_Port Virtual Link



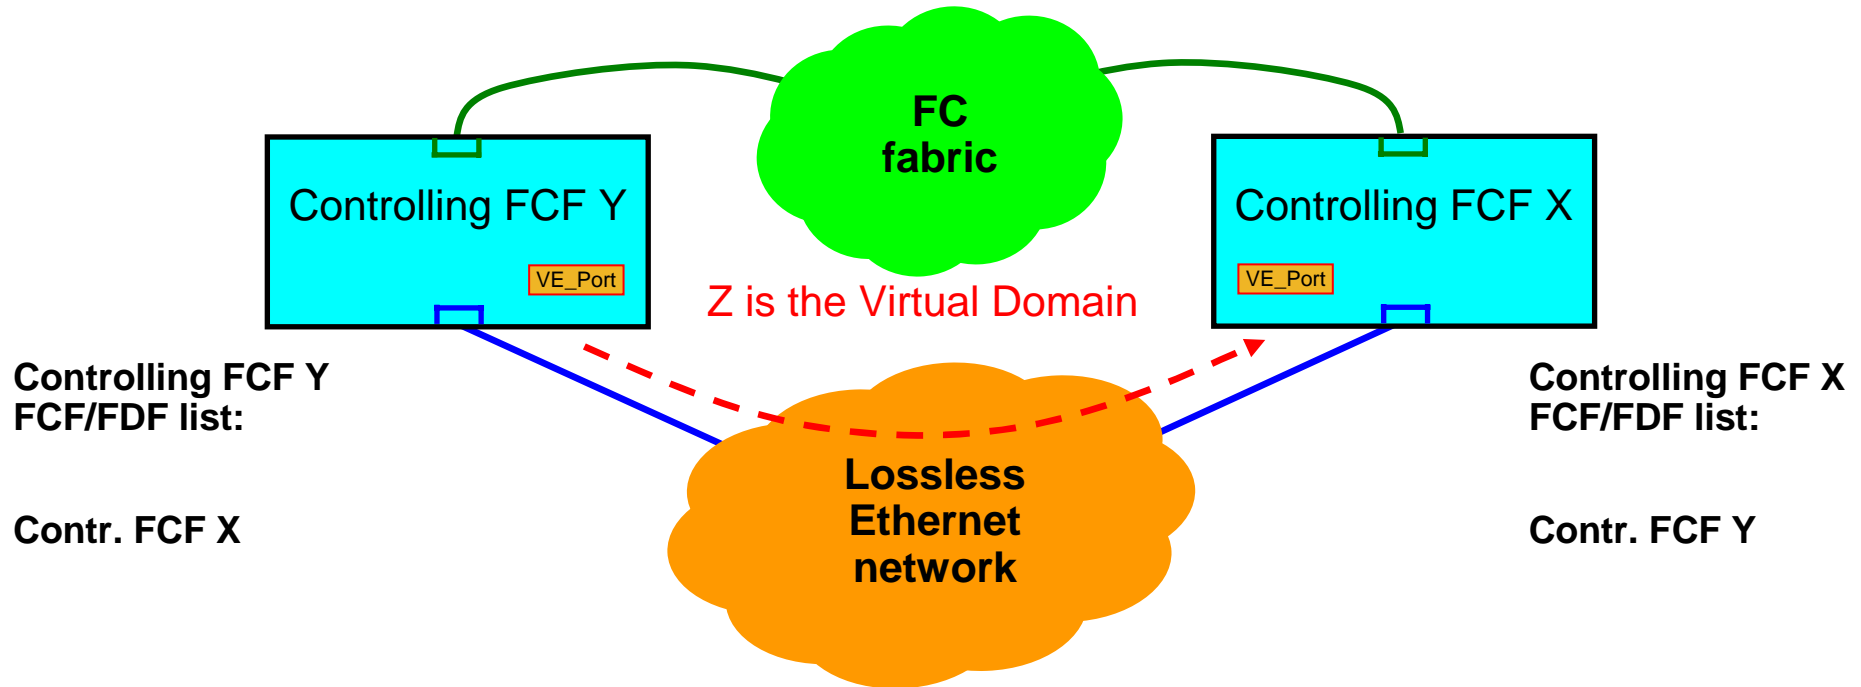
- The VE_Port to VE_Port Virtual Link between two Controlling FCFs is used to establish a high available configuration
- The Controlling FCF with the higher Switch_Name is the Primary, the other one is the Secondary

High Availability

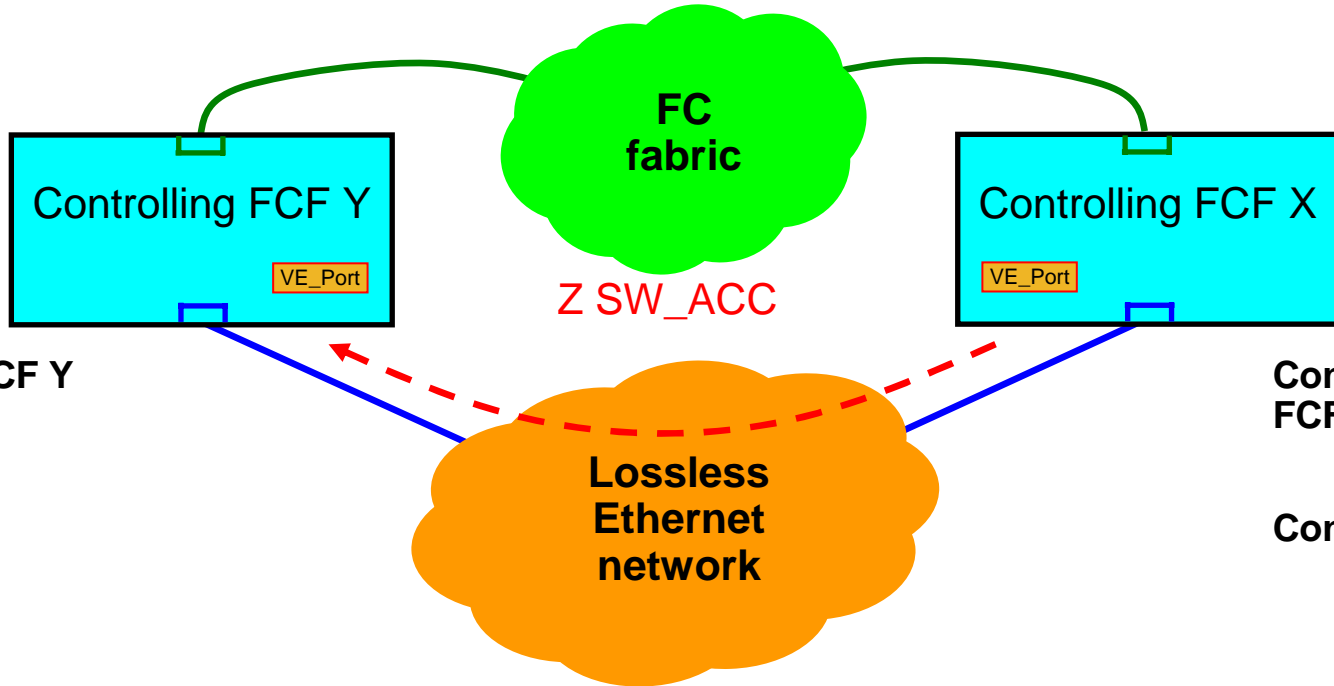
- The Primary Controlling FCF obtains an additional Domain_ID from the Principal Switch
 - Used as Virtual Domain_ID to allocate N_Port_IDs to FDFs
- The Primary Controlling FCF communicates the Virtual Domain_ID to the Secondary Controlling FCF
 - Through a FCoE encapsulated SW_ILS
- Then the Primary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



Virtual Domain_ID Z Announcement



Virtual Domain Z Announcement SW_ACC



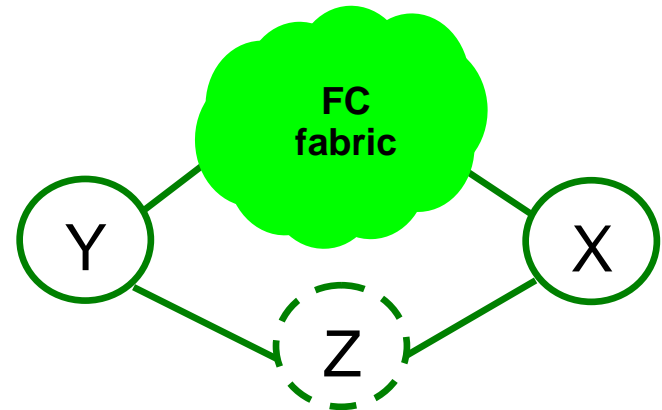
Controlling FCF Y
FCF/FDF list:

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

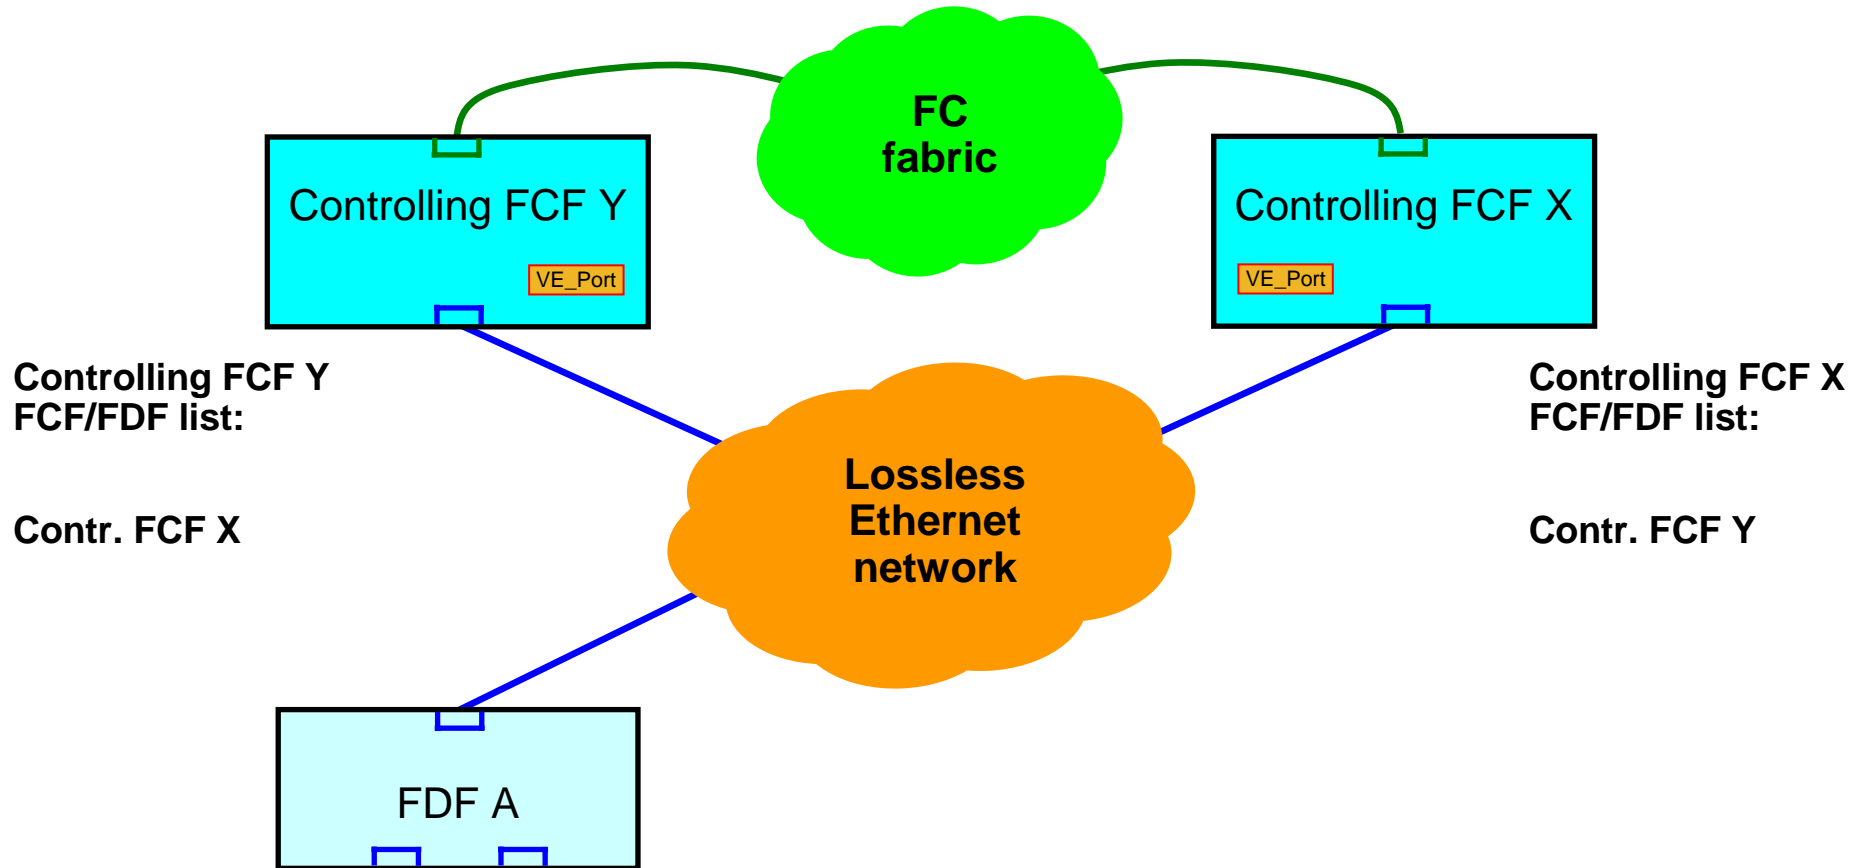
Virtual Domain Announcement SW_ACC

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

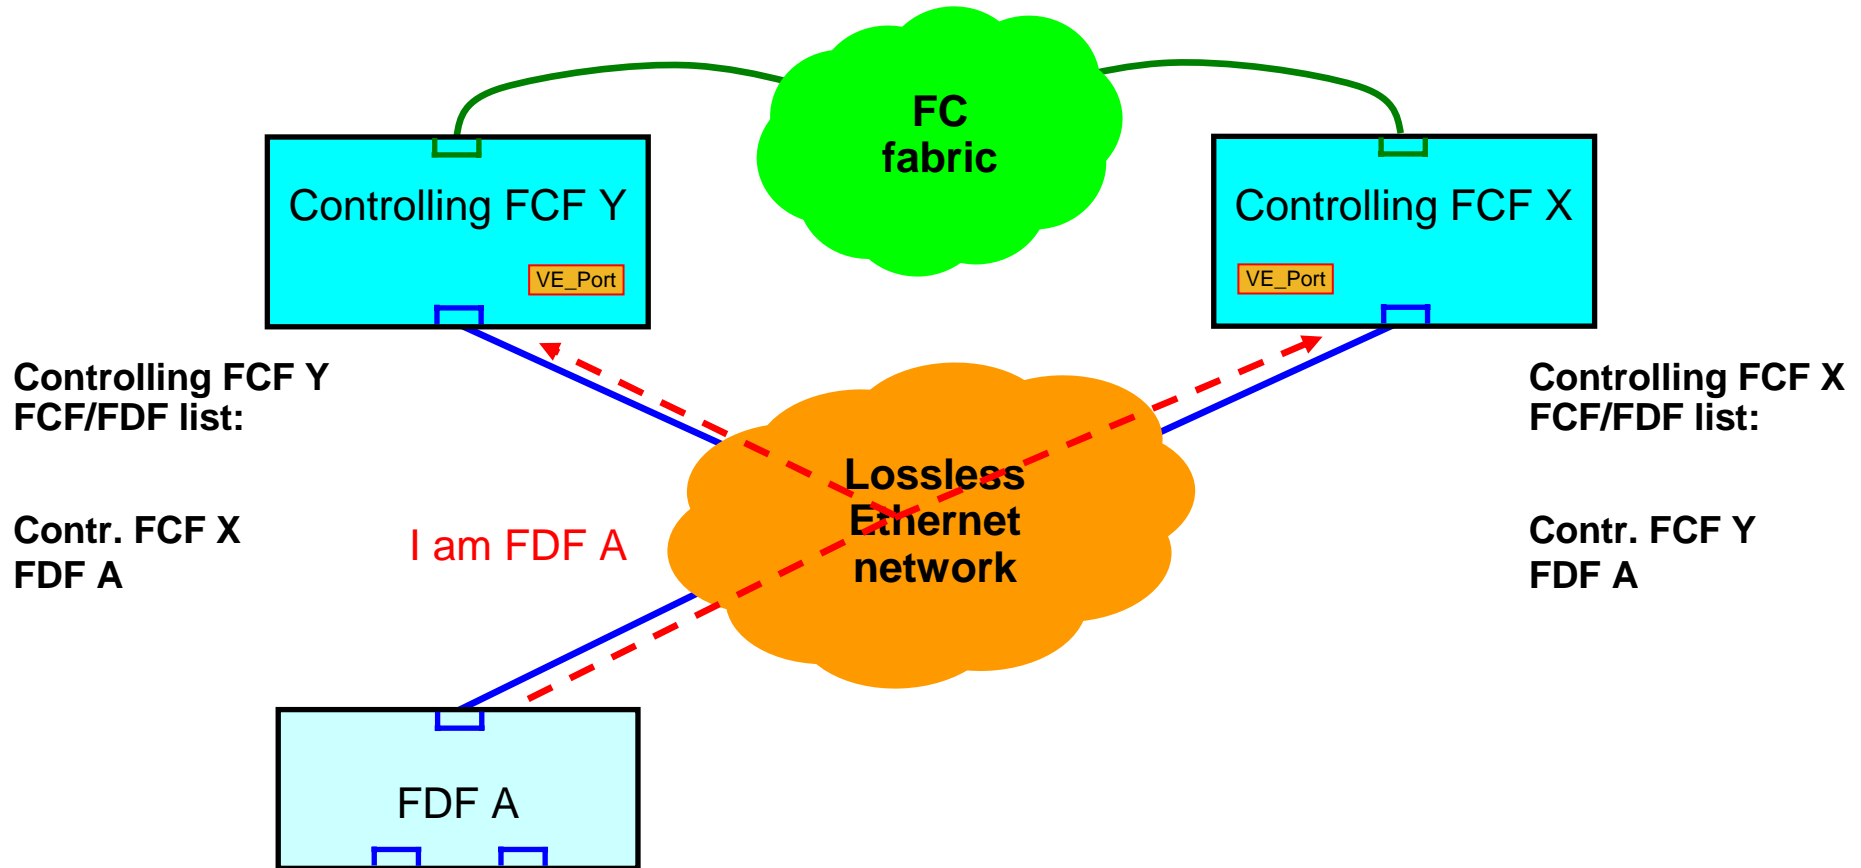
Agenda

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

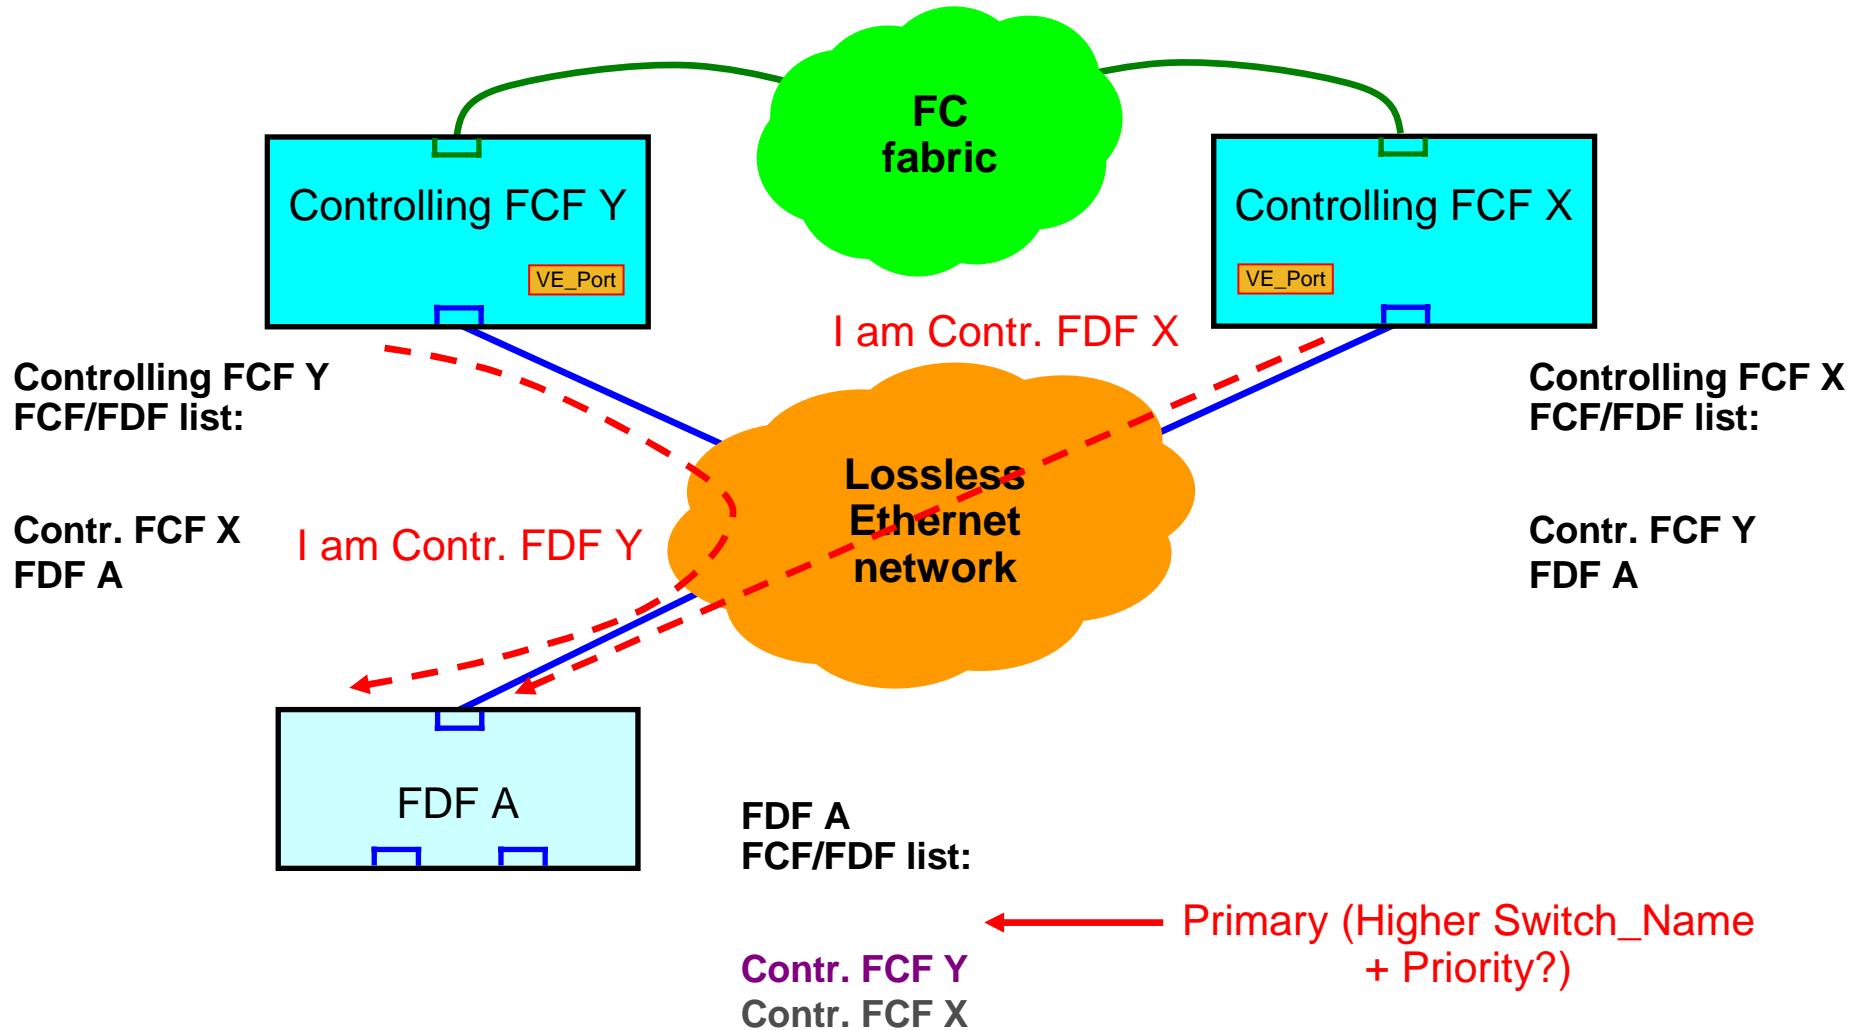
FDF A



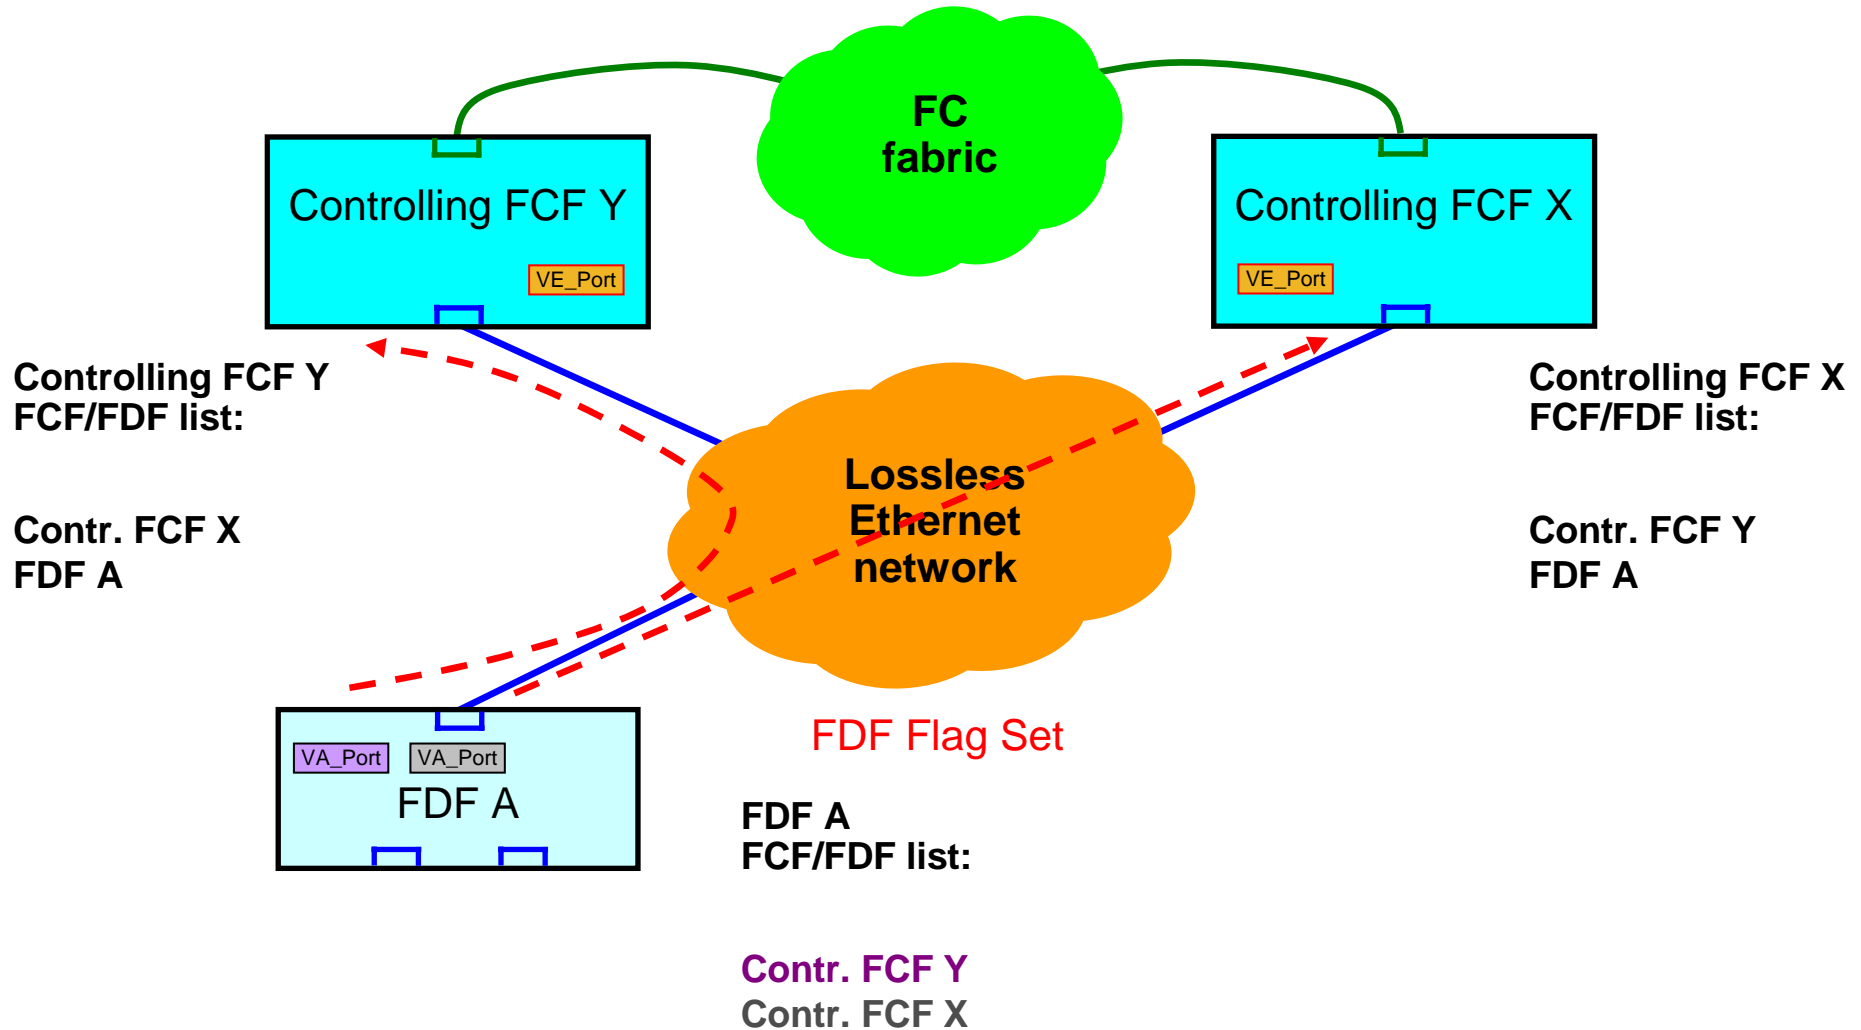
FDF A FIP Solicitation



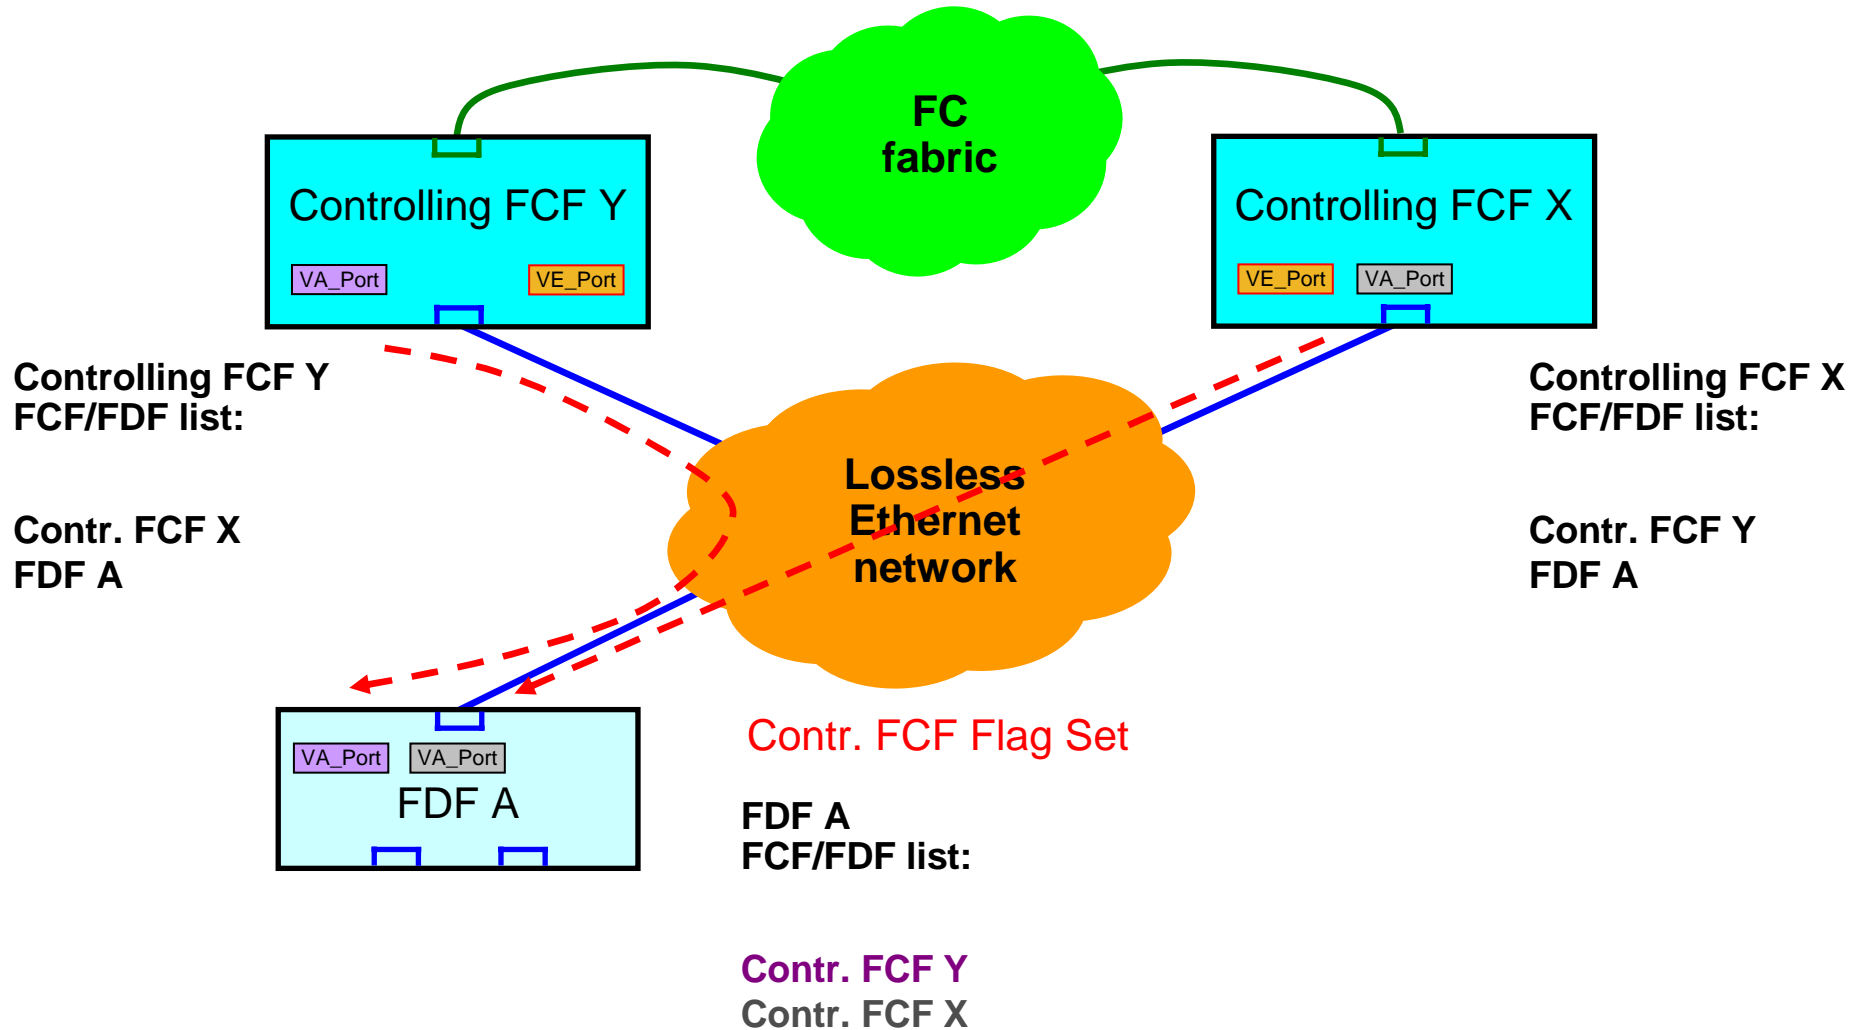
FIP Solicited Advertisements



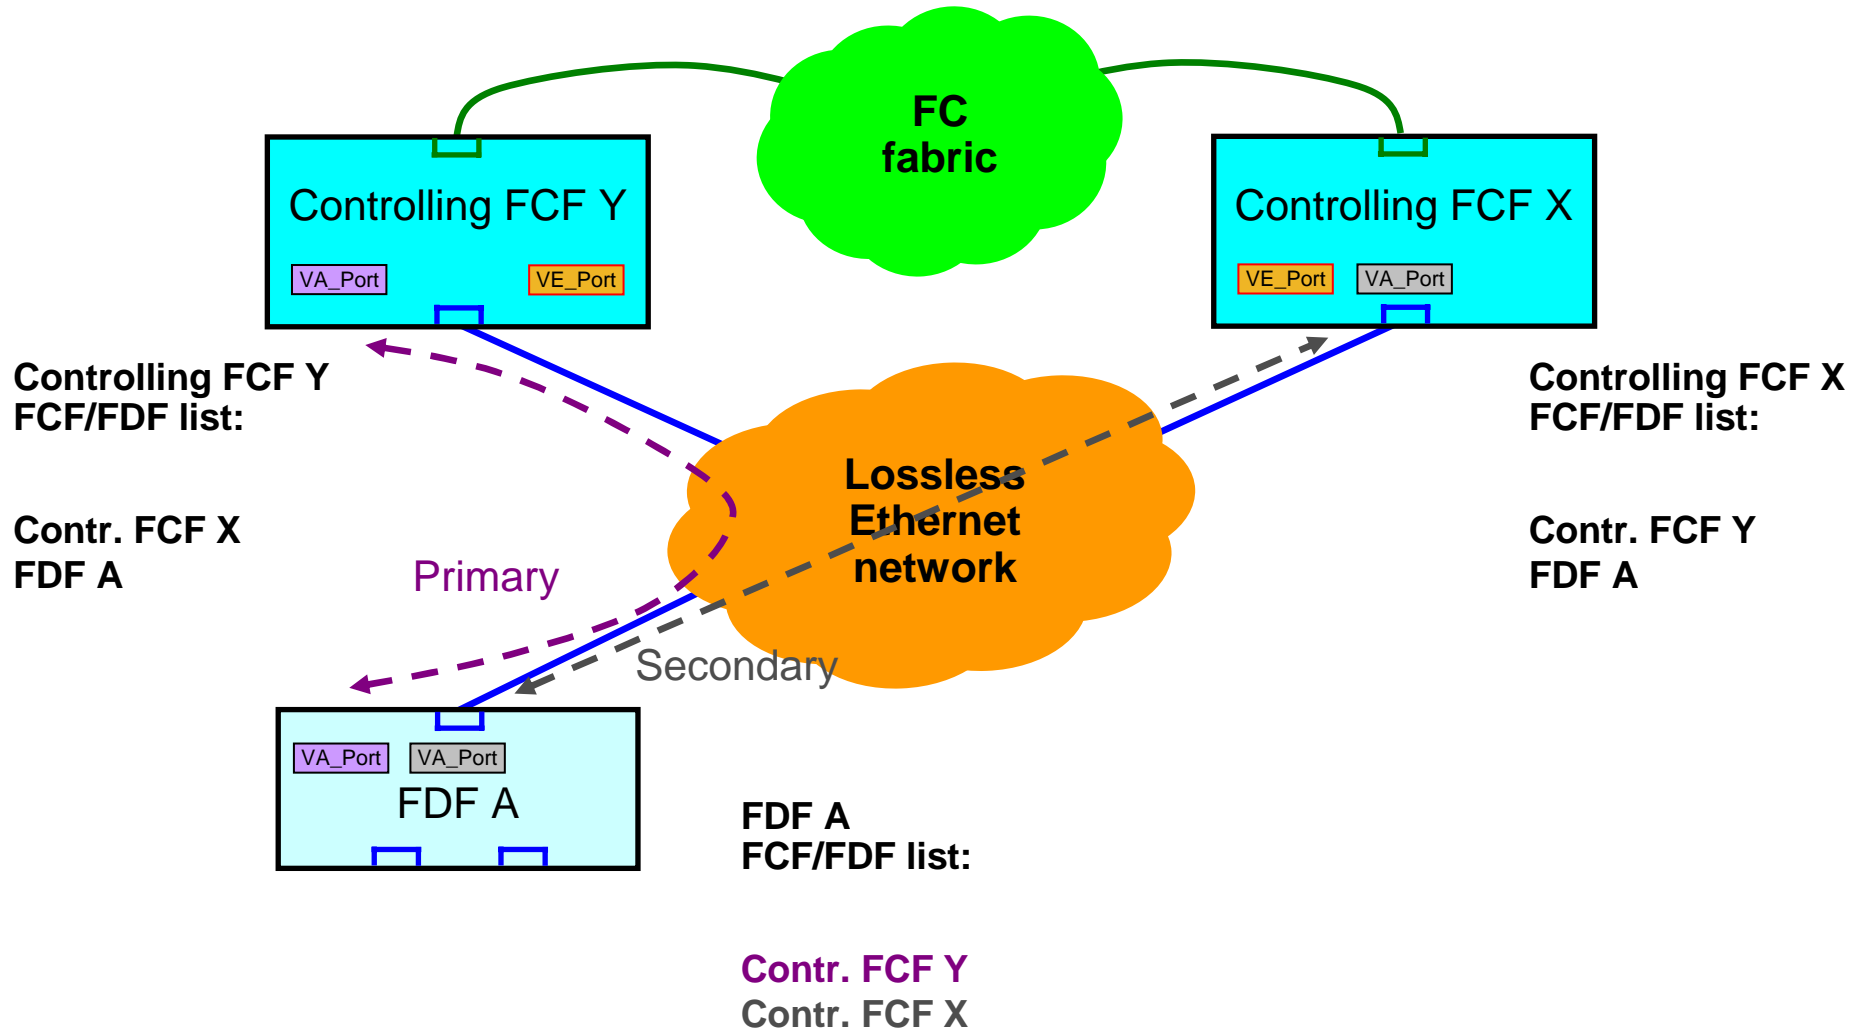
FIP ELP Requests



FIP ELP SW_ACC



VA_Port to VA_Port Virtual Links



FDF B



Controlling FCF Y
FCF/FDF list:

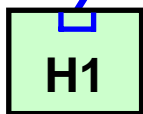
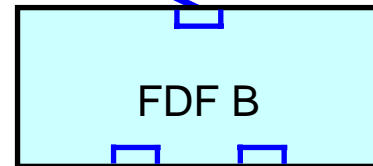
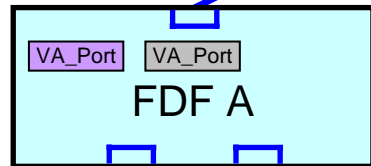
Controlling FCF X
FCF/FDF list:

Contr. FCF X
FDF A

Contr. FCF Y
FDF A



**Lossless
Ethernet
network**




1.2.1

FDF A
FCF/FDF list:

FDF B
FCF/FDF list:

Contr. FCF Y
Contr. FCF X

FDF B FIP Solicitation



Controlling FCF Y
FCF/FDF list:

Controlling FCF X
FCF/FDF list:

Contr. FCF X
FDF A
FDF B

Contr. FCF Y
FDF A
FDF B

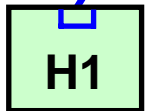
Lossless Ethernet network

I am FDF B

FDF A
FCF/FDF list:

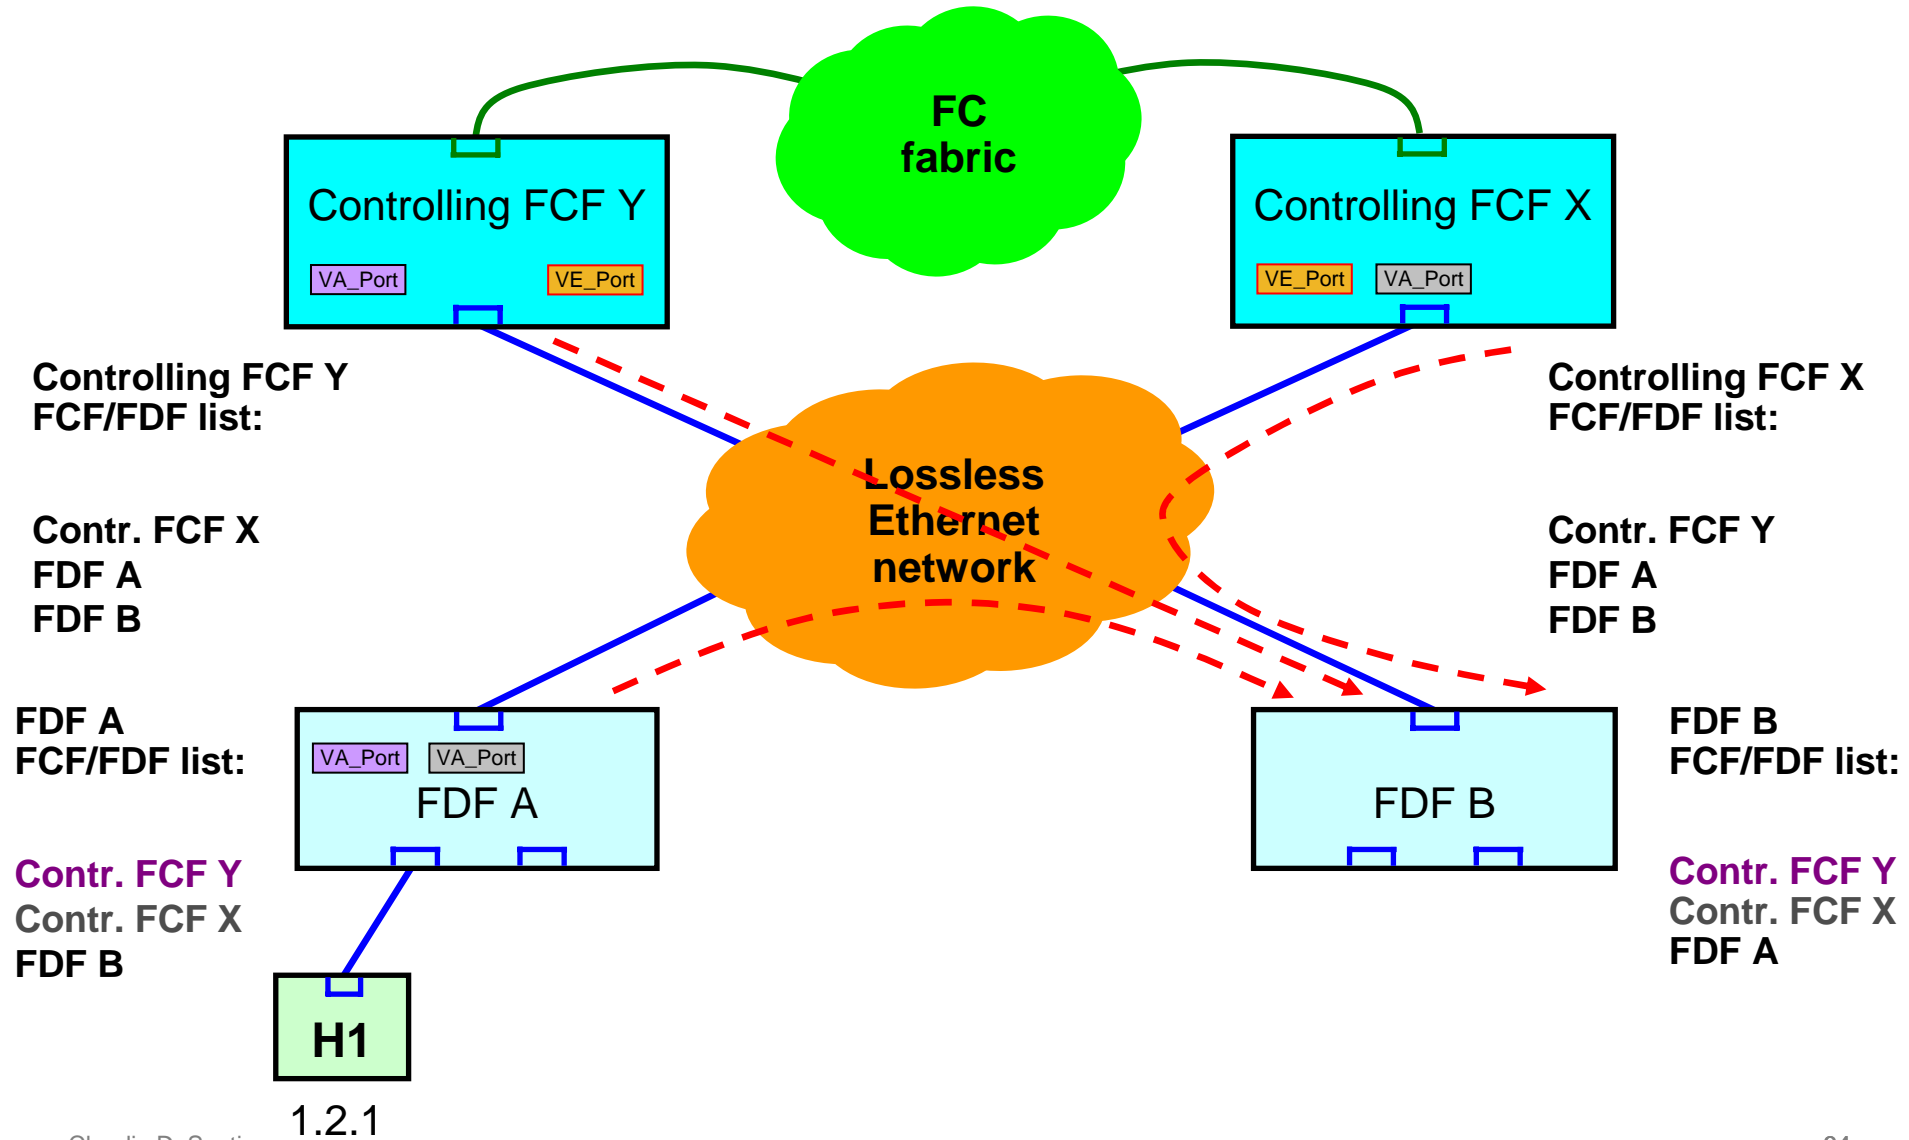
FDF B
FCF/FDF list:

Contr. FCF Y
Contr. FCF X
FDF B

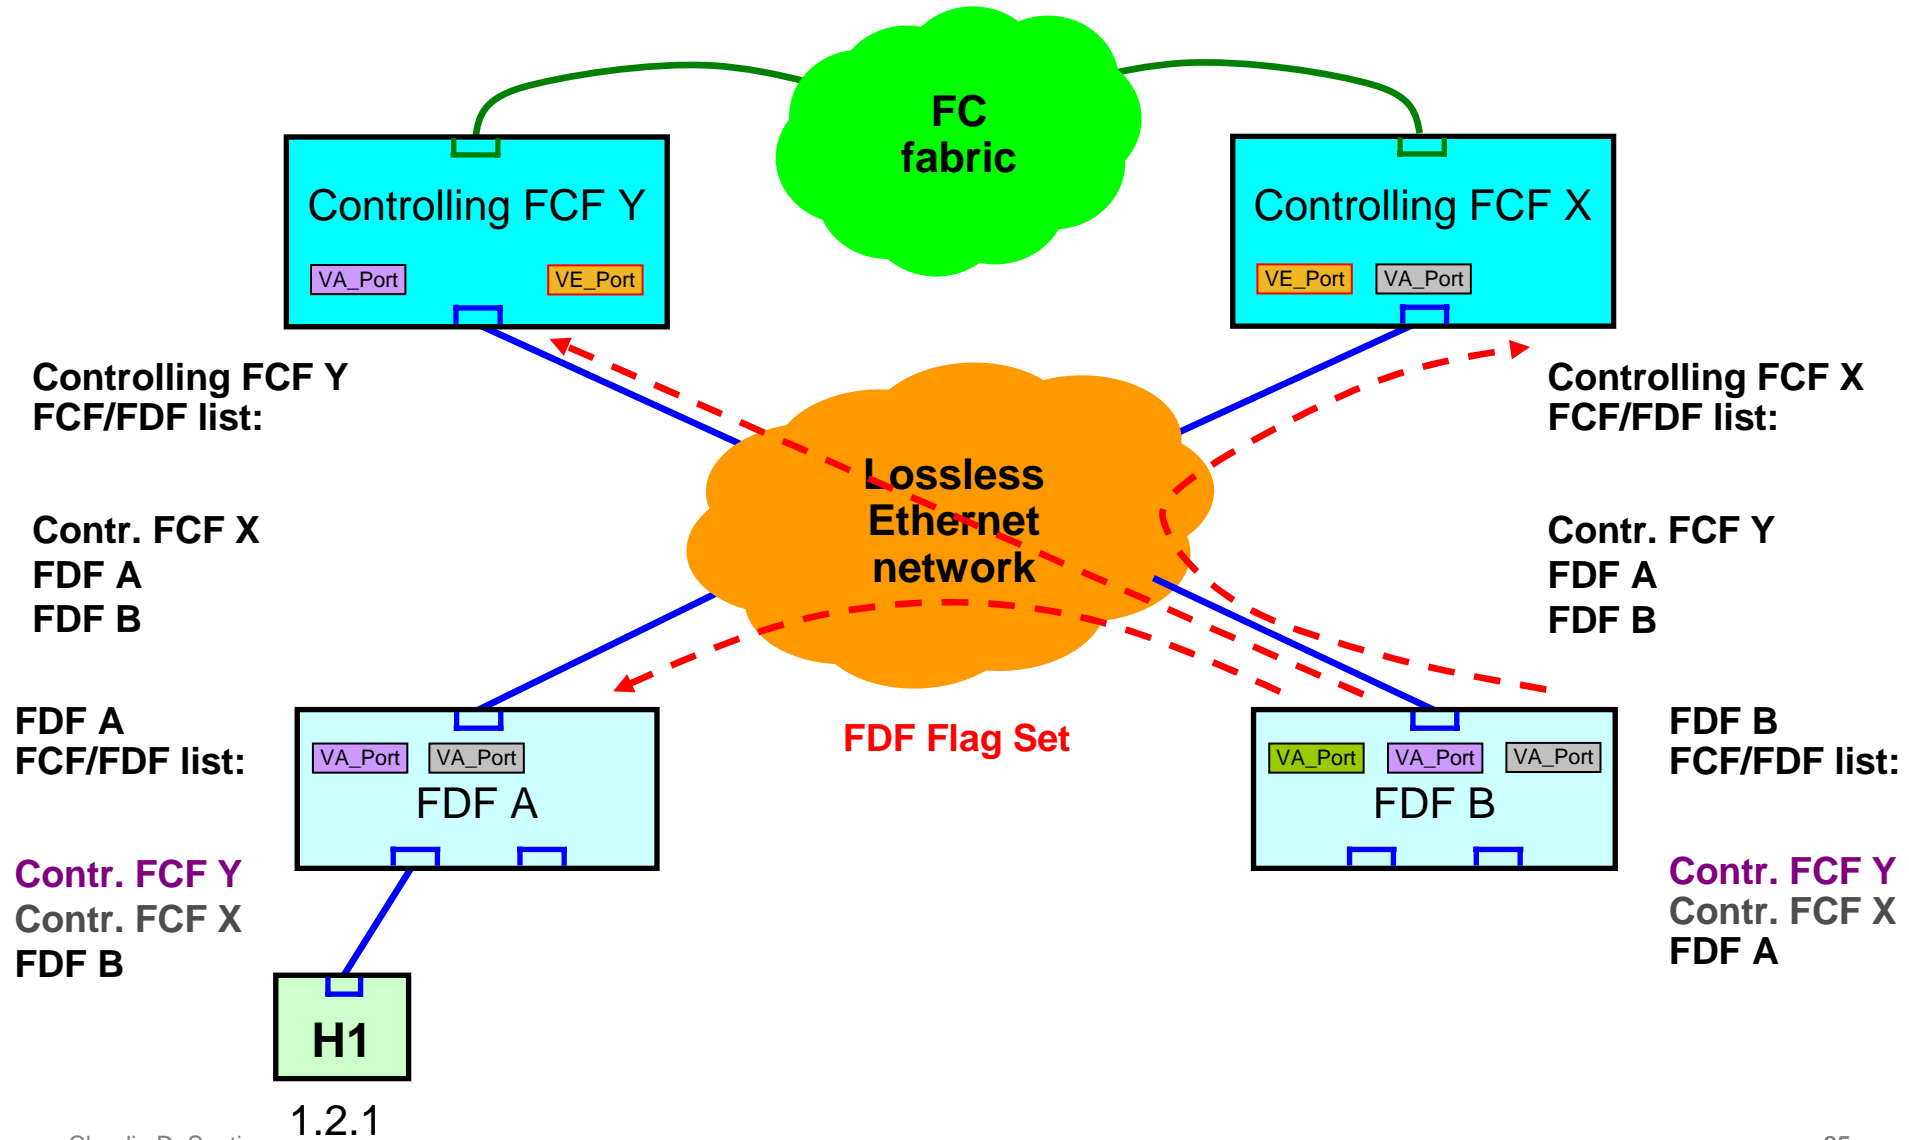


1.2.1

FIP Solicited Advertisements



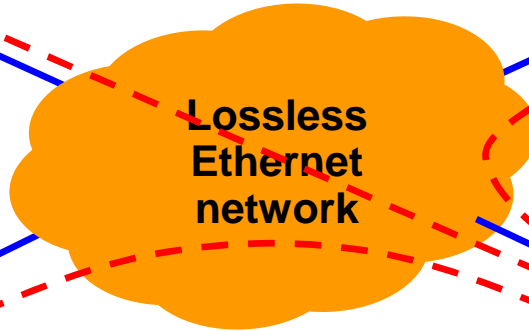
FDF B FIP ELP Requests



FDF B FIP ELP SW_ACC



Contr. FCF Flag Set



FDF Flag Set

Controlling FCF Y
FCF/FDF list:

Controlling FCF X
FCF/FDF list:

Contr. FCF X
FDF A
FDF B

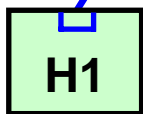
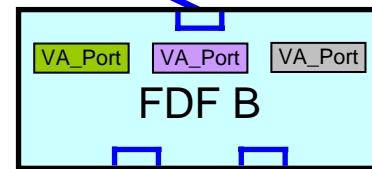
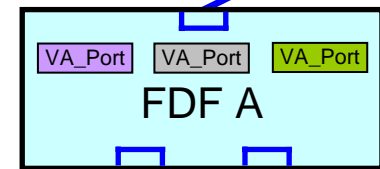
Contr. FCF Y
FDF A
FDF B

FDF A
FCF/FDF list:

FDF B
FCF/FDF list:

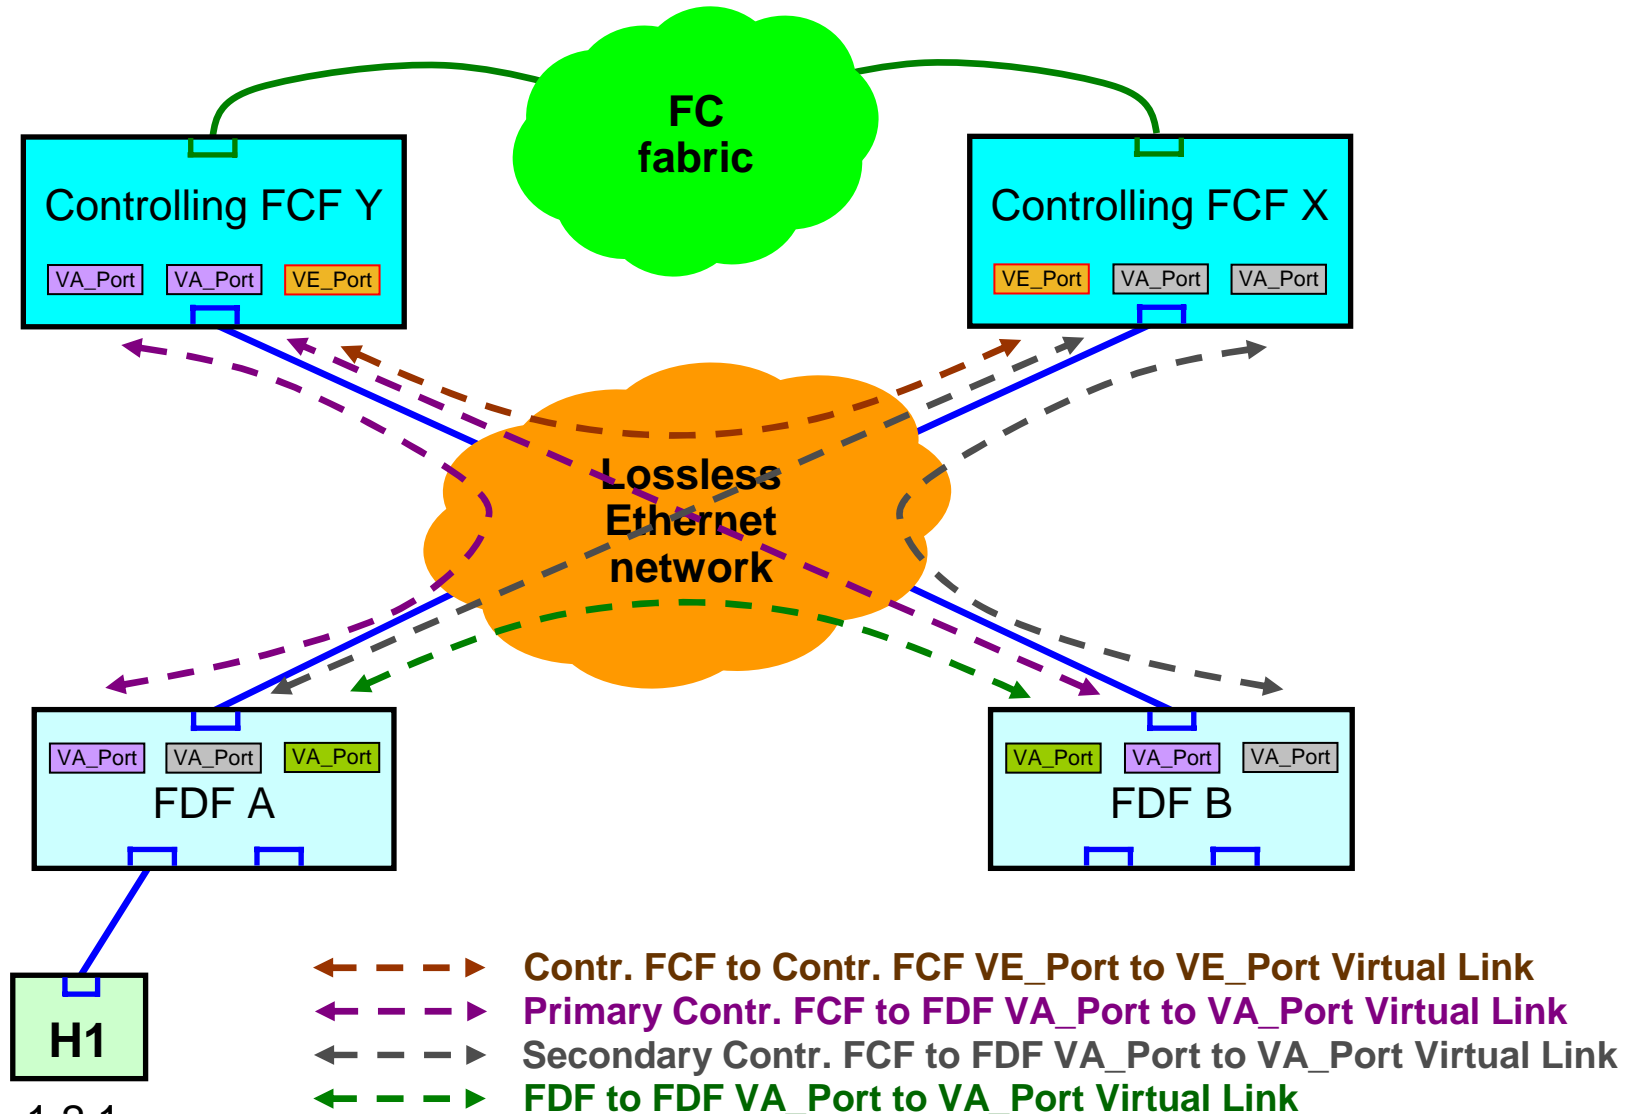
Contr. FCF Y
Contr. FCF X
FDF B

Contr. FCF Y
Contr. FCF X
FDF A



1.2.1

Virtual Links



1.2.1

Thank You



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