



VA_Ports: FDF / Controlling FCF Protocols

Claudio DeSanti

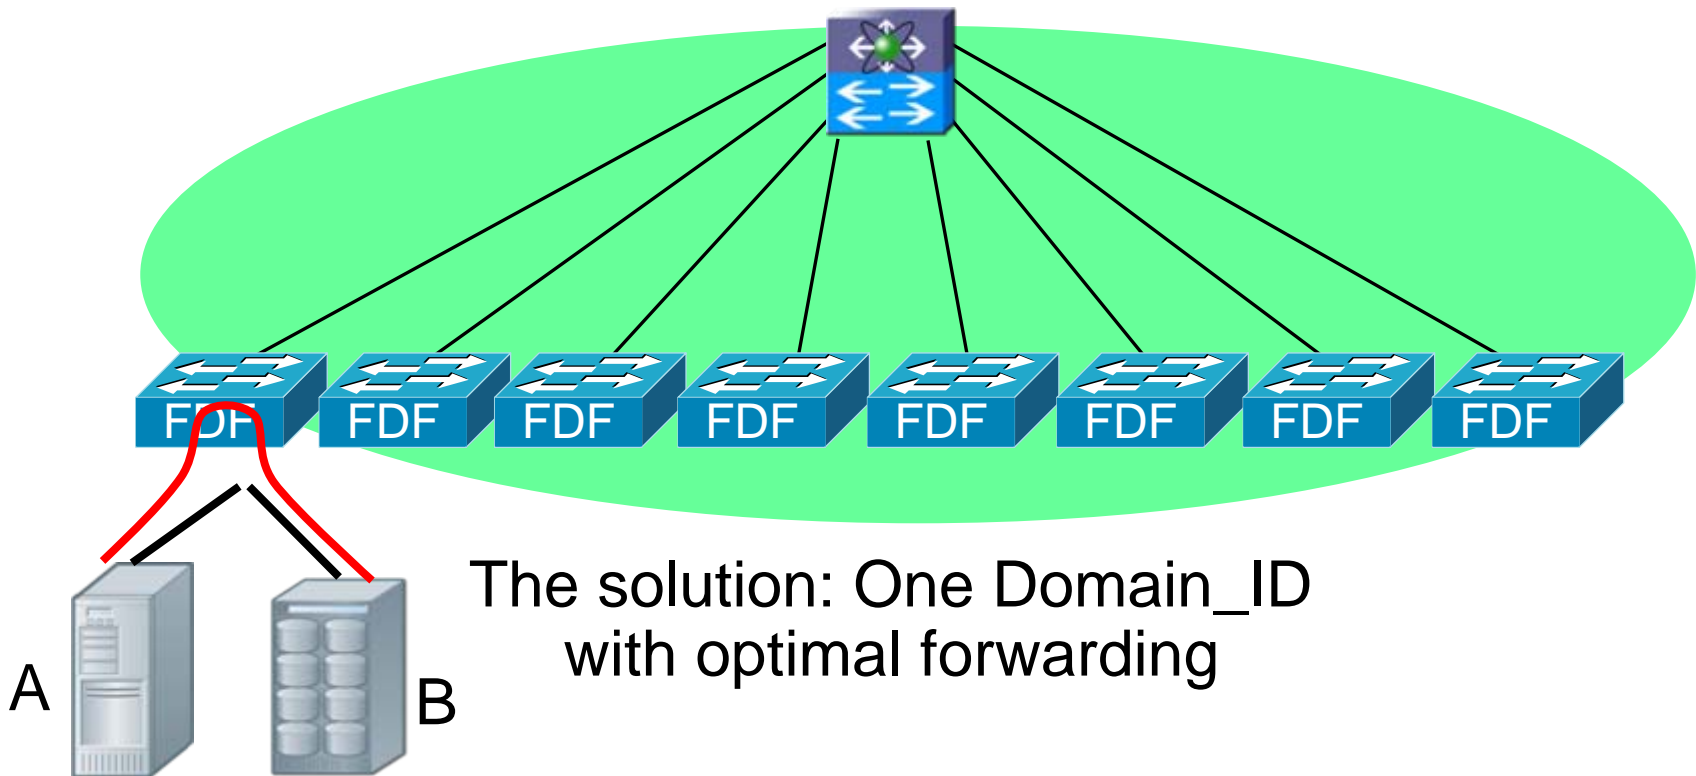
T11/10-271v2, August 2010



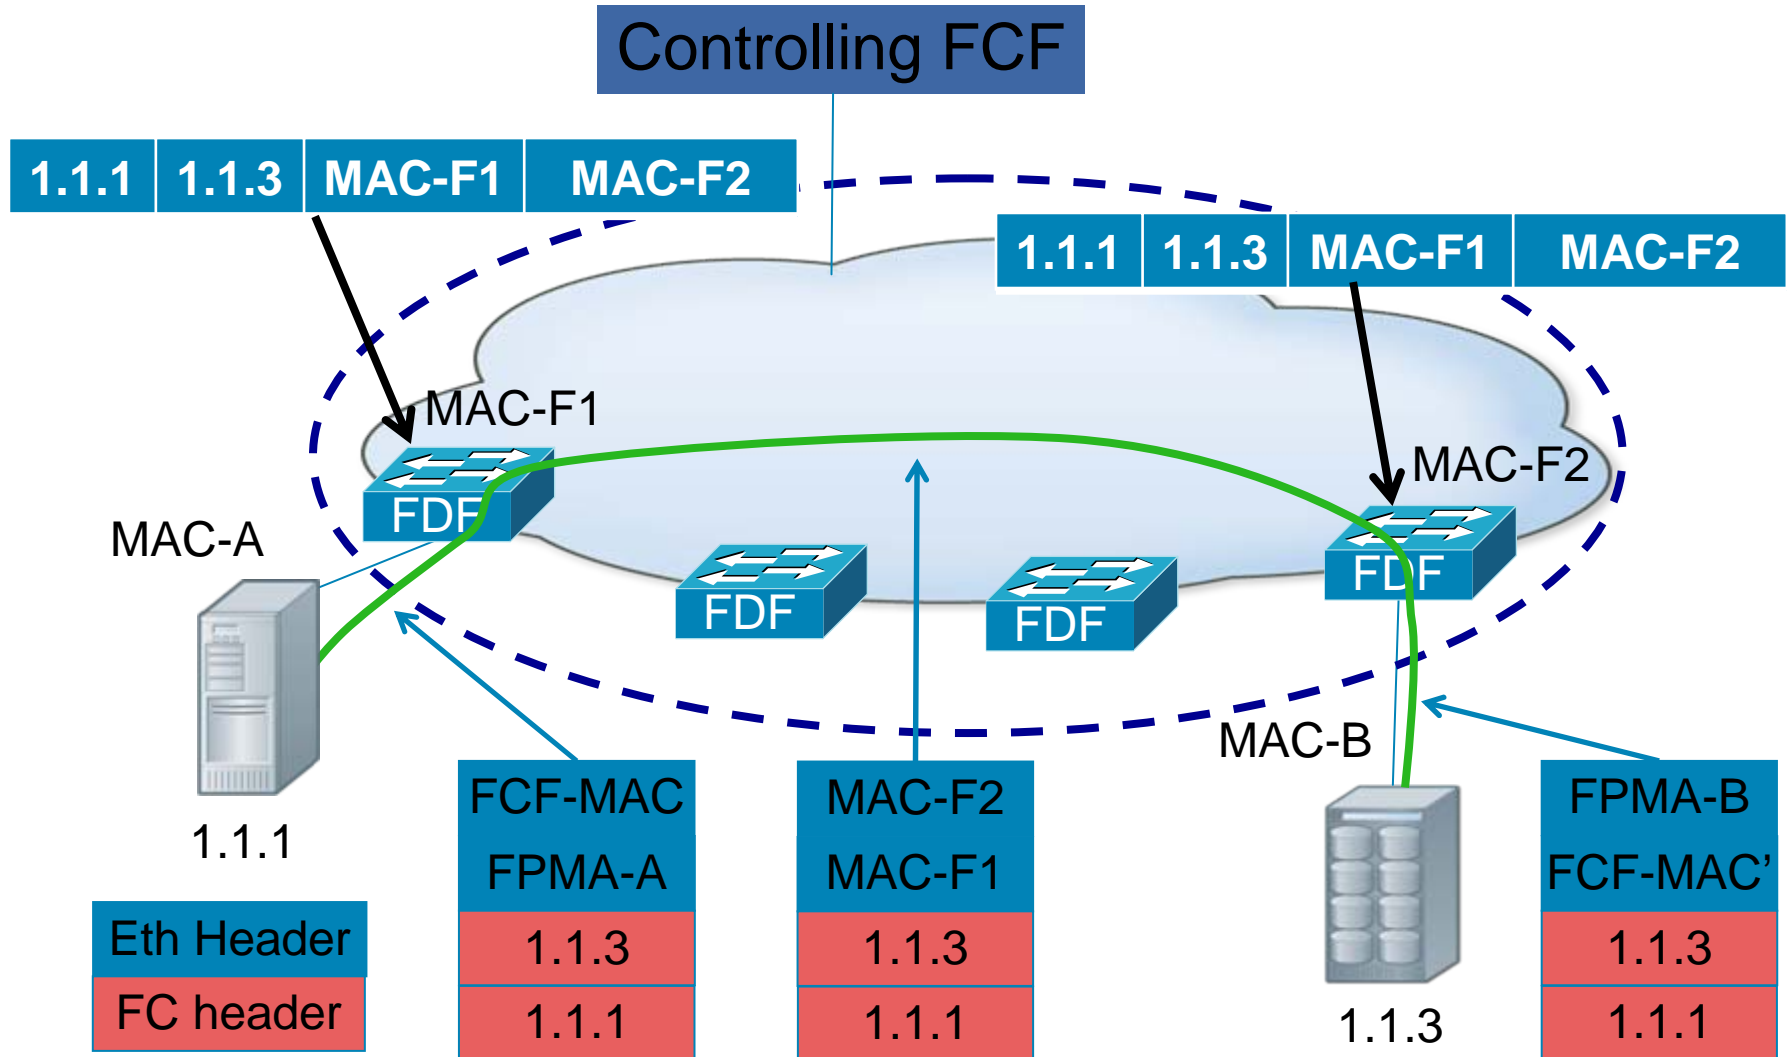
VA Port Authority
Flickr user:conner395
Creative Commons License

Introducing the Controlling FCF / FDFs

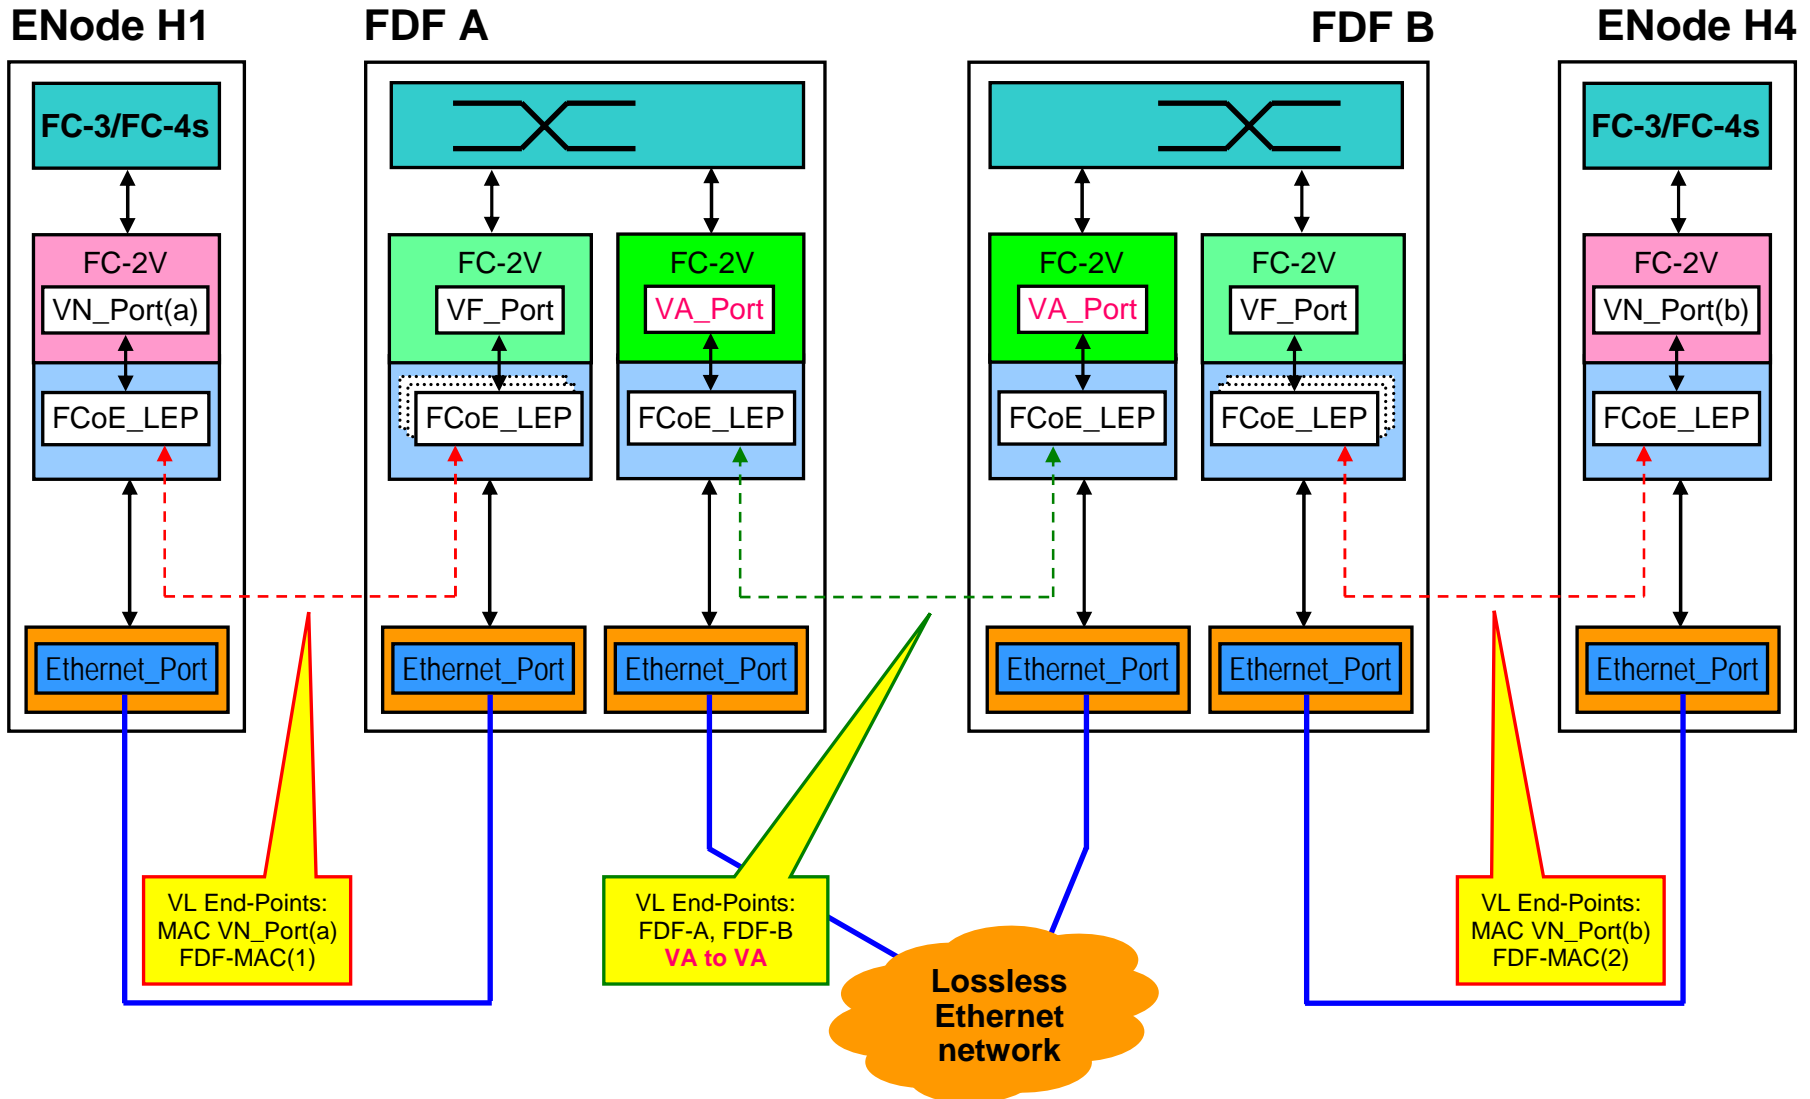
Controlling FCF



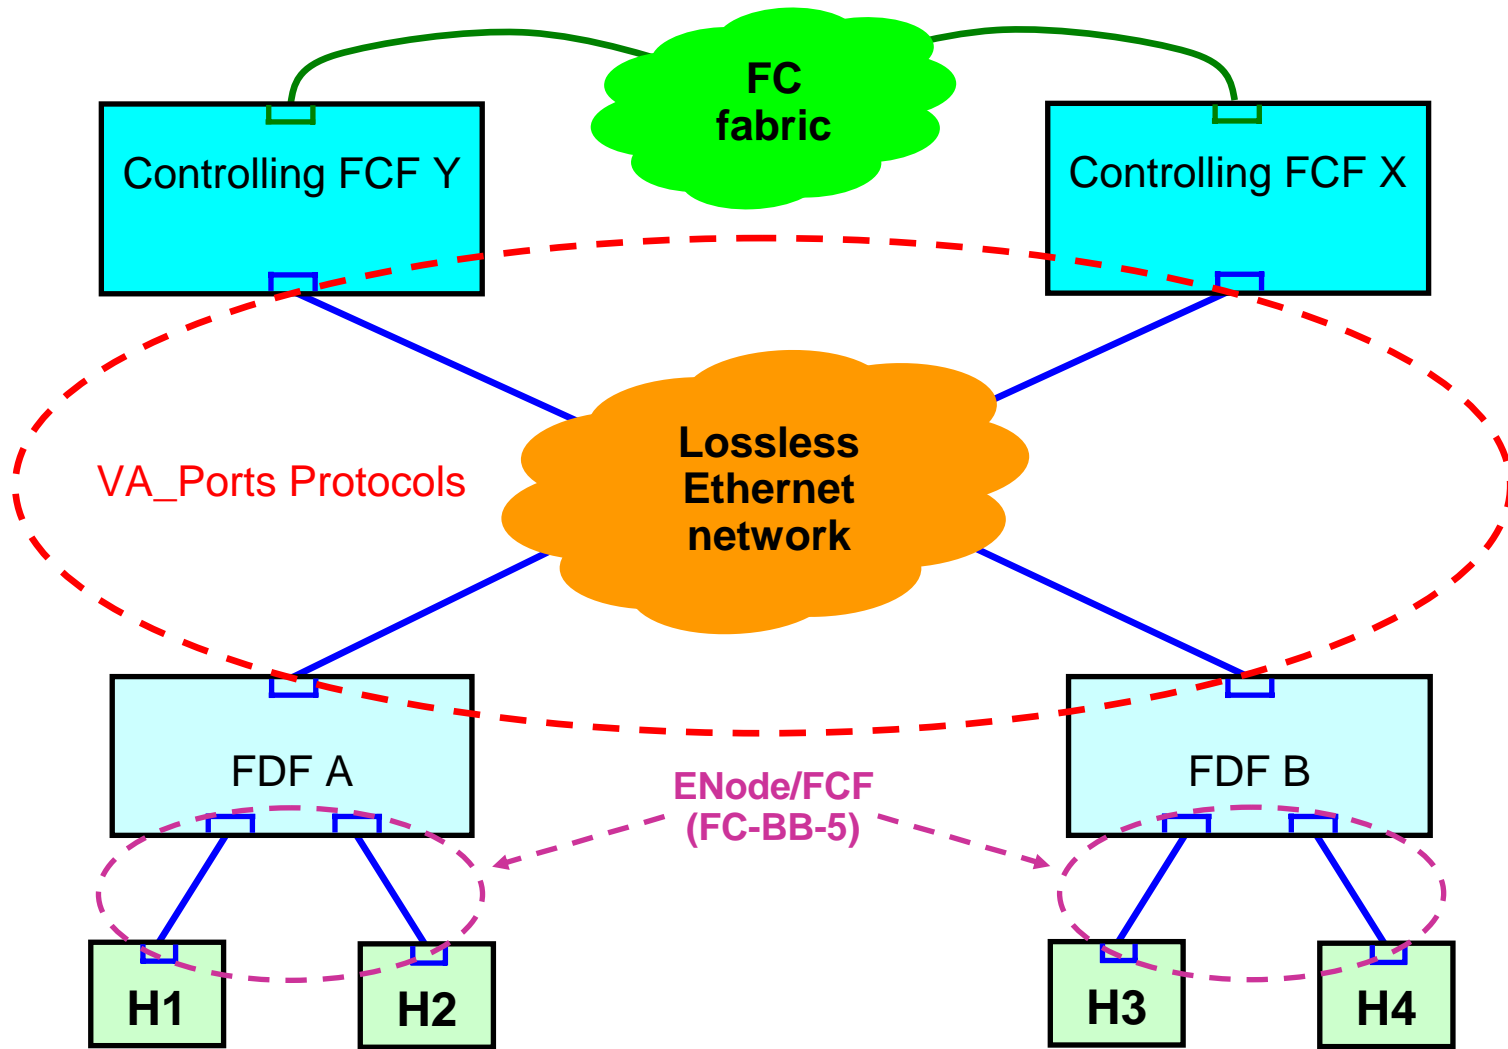
Controlling FCF / FDFs



VA_Port to VA_Port Virtual Links



Protocols



The Controlling FCF

- **Controlling FCFs should operate in primary/secondary pairs for high availability**
- **Controlling FCFs instantiate VE_Port to VE_Port Virtual Links between them**
 - 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**





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

- 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

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**
 - FCF 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)
- **Enabling FDFs to reach other Domain_IDs would require a change in FSPF**

Summary information (as in OSPF) is needed to achieve optimal forwarding if arbitrary connections are possible between FDFs and other FCFs

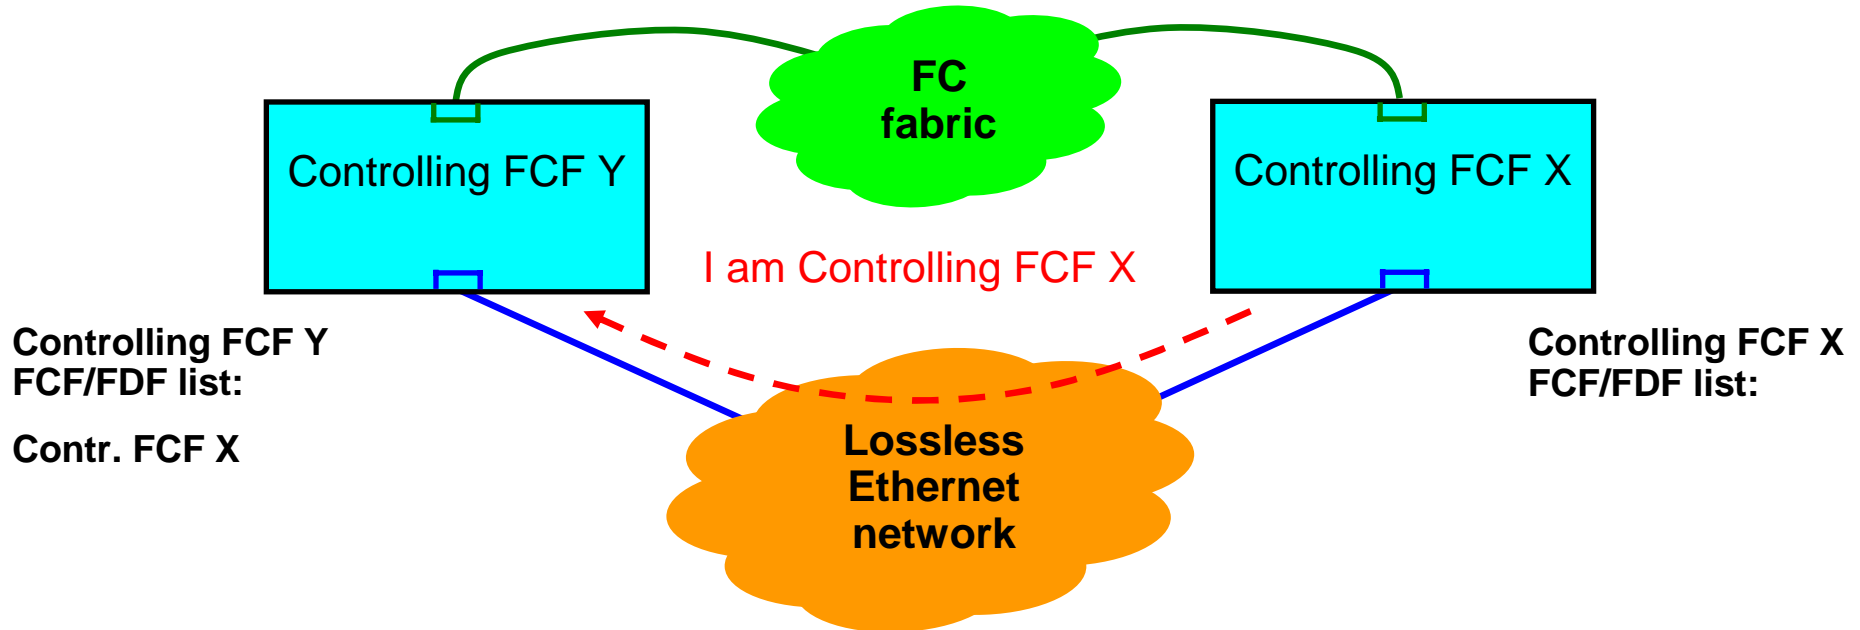
An FCF of FC switch would need to know over which port an N_Port_ID range can be reached at the lowest cost

Currently FSPF provides only reachability information for Domain_IDs, not for N_Port_ID ranges

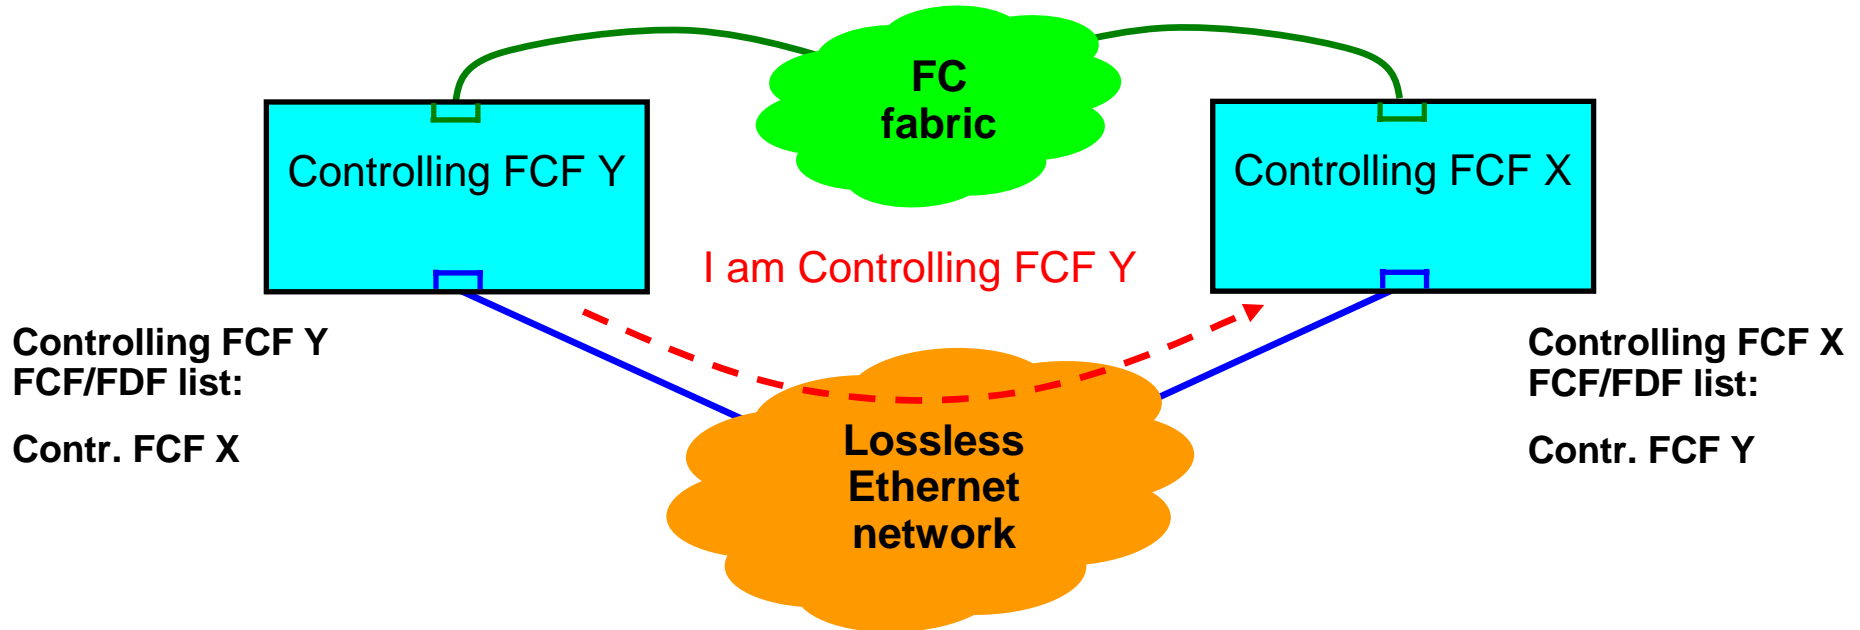
- **This seems too big of a change**

It would require an upgrade of the FC Fabric to deploy FDFs

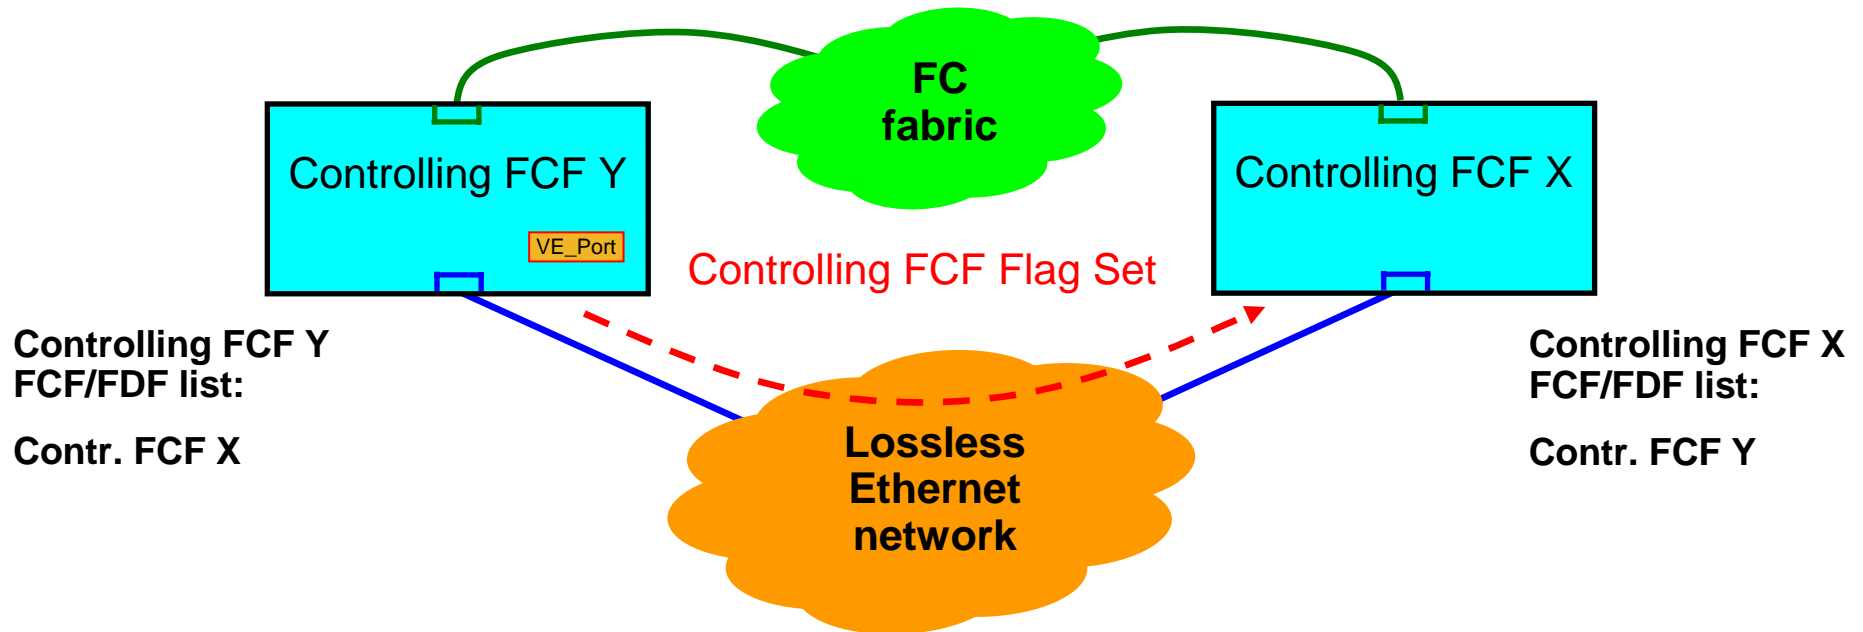
FIP Solicitation



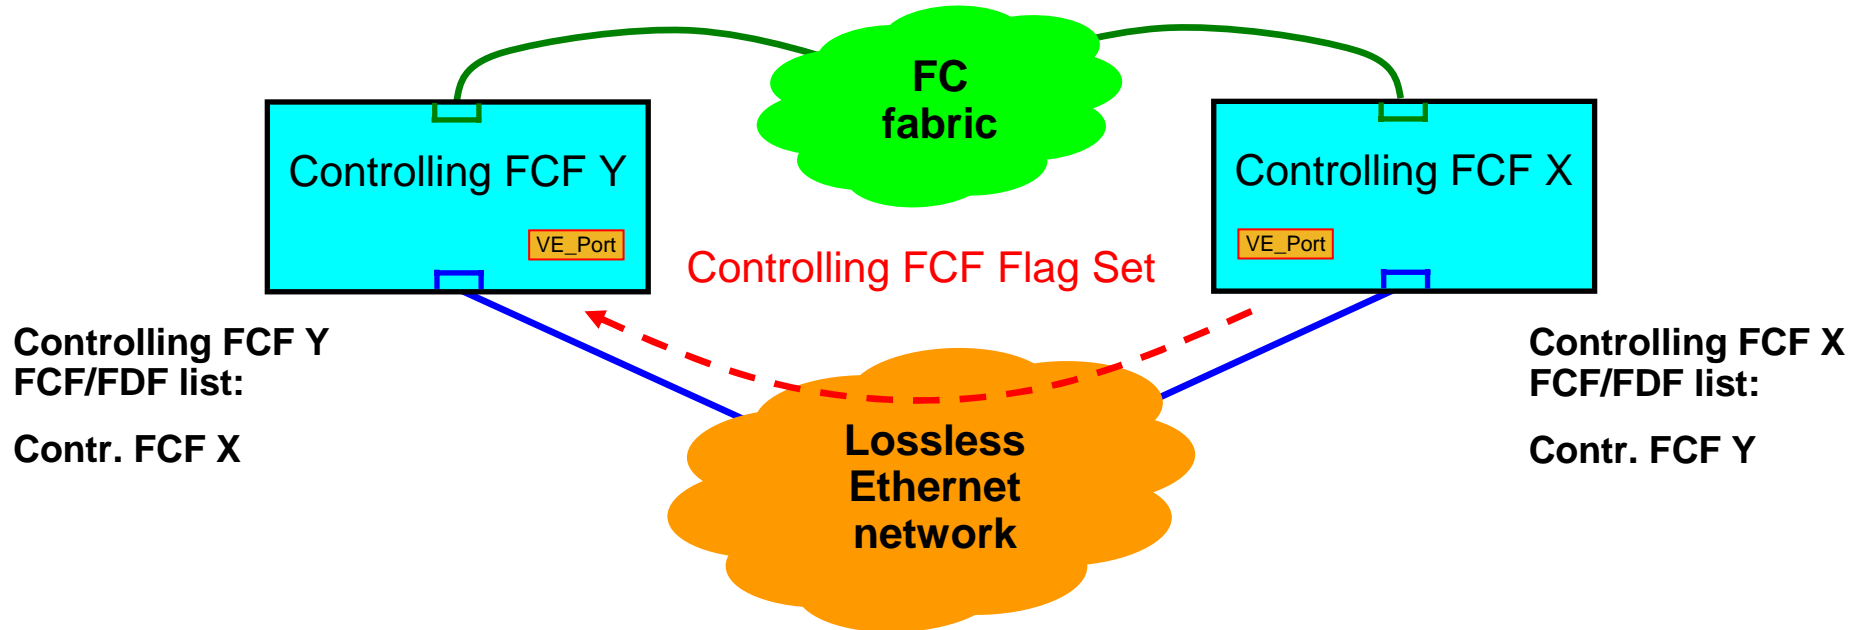
FIP Solicited Advertisement




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



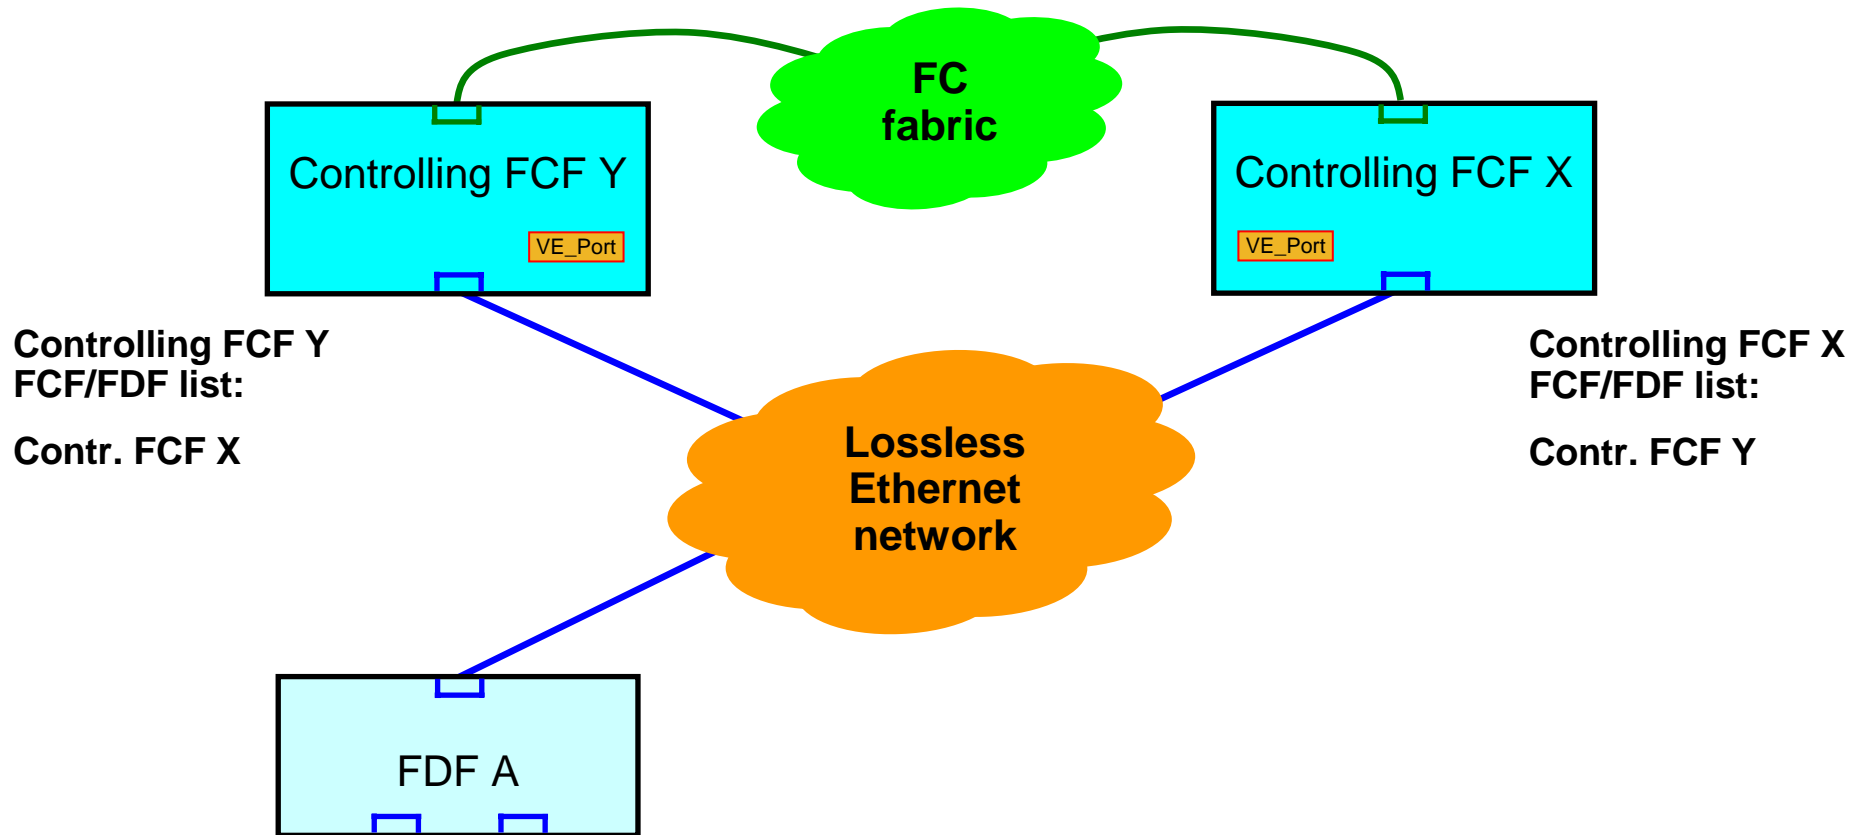
Controlling FCF Y
FCF/FDF list:
Contr. FCF X

Controlling FCF X
FCF/FDF list:
Contr. FCF Y

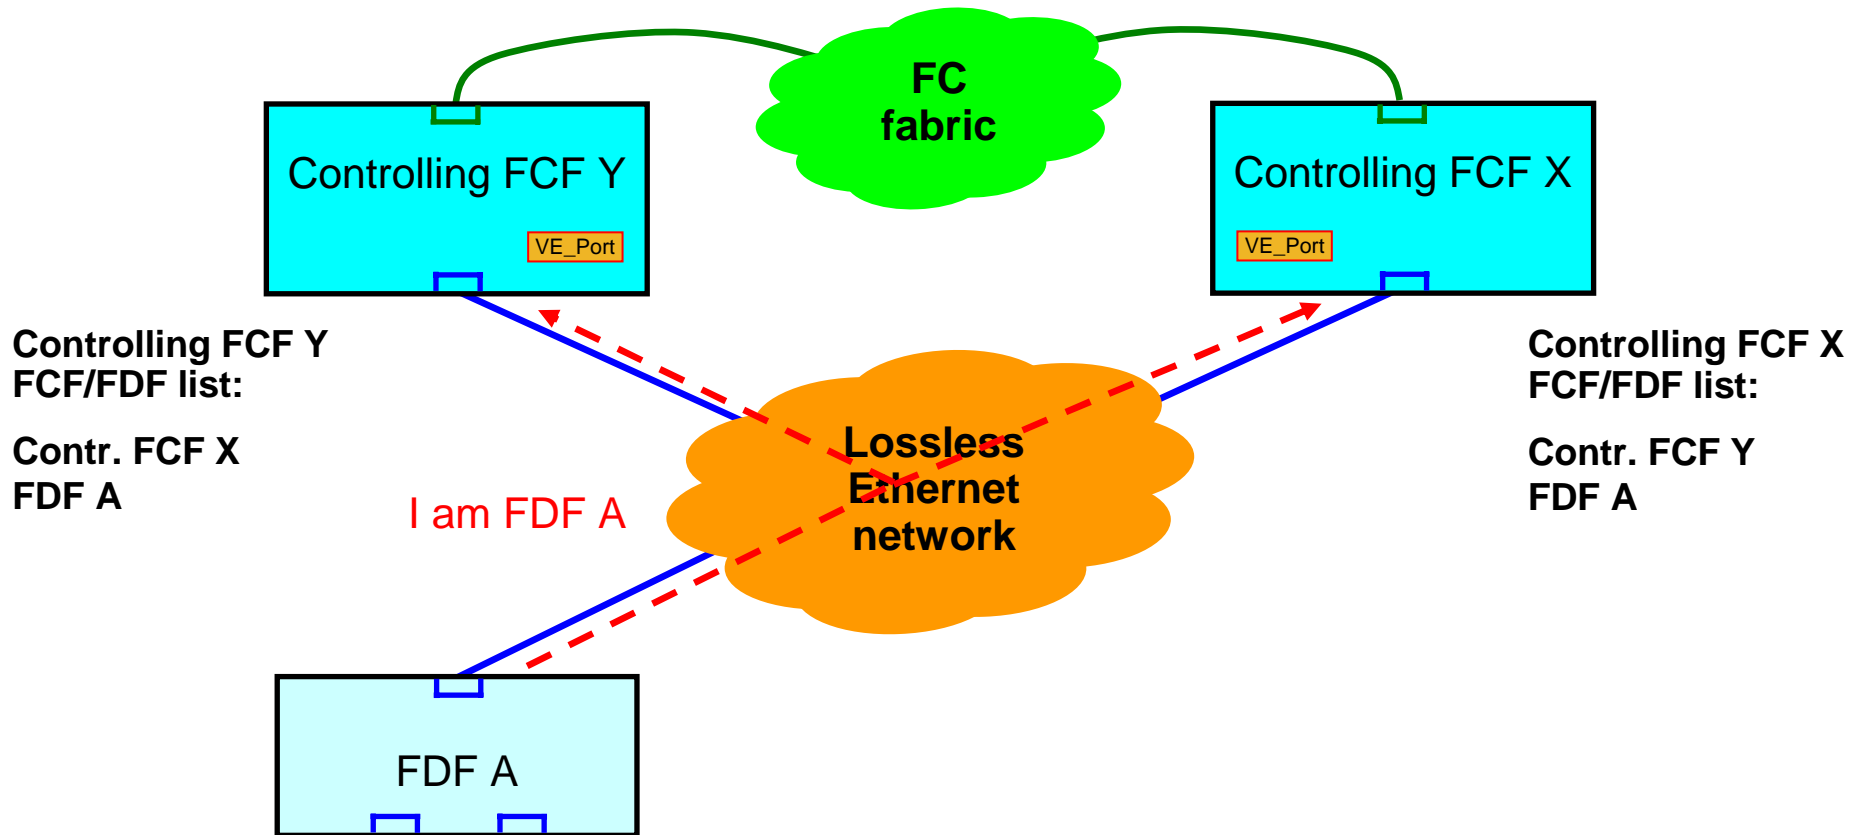
Advertised FSPF Topology



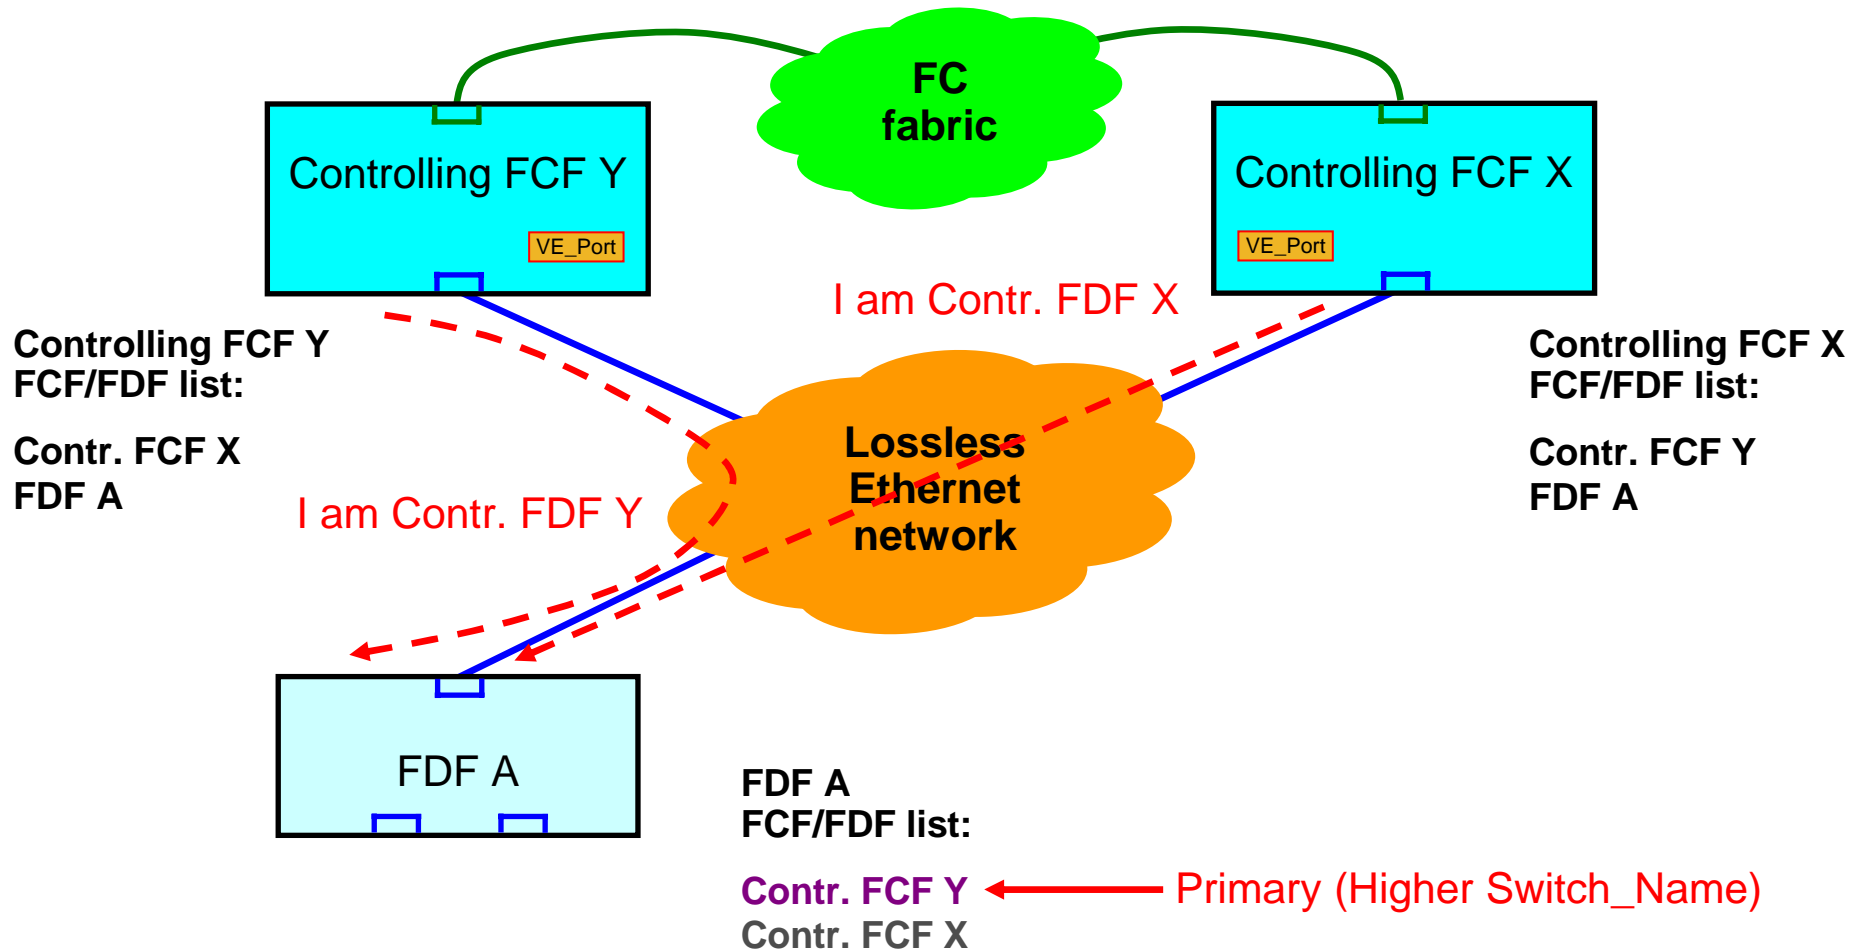
FDF A



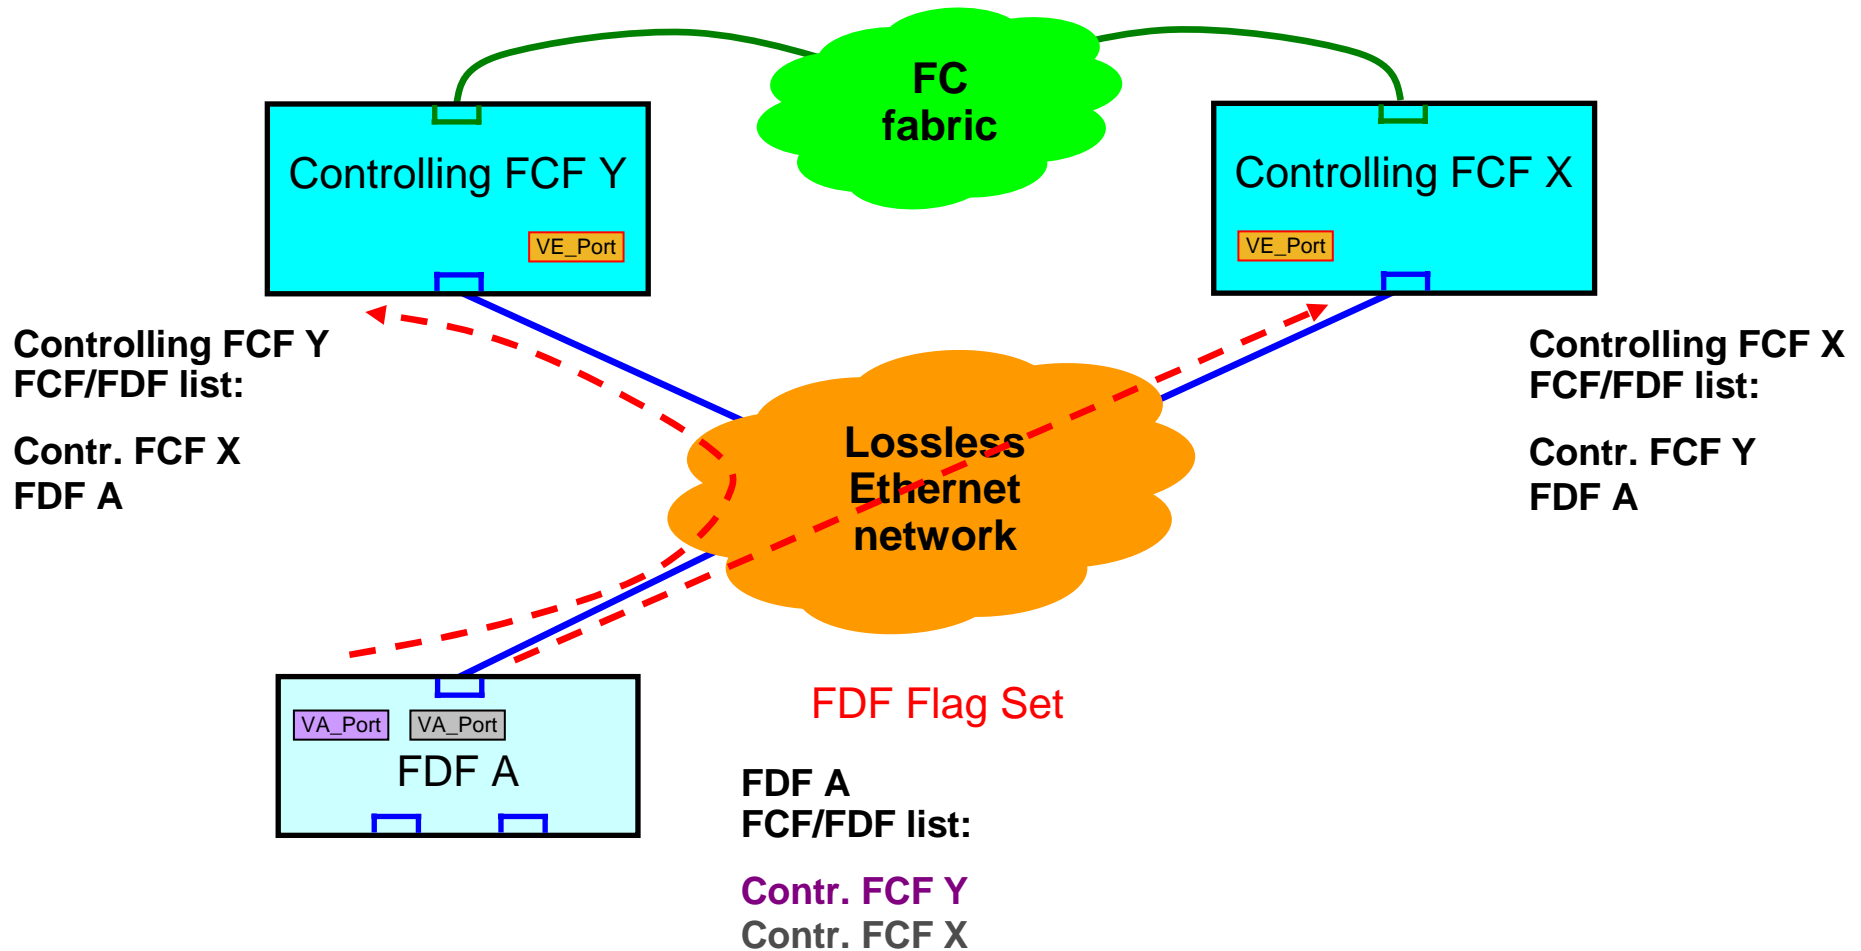
FDF A FIP Solicitation



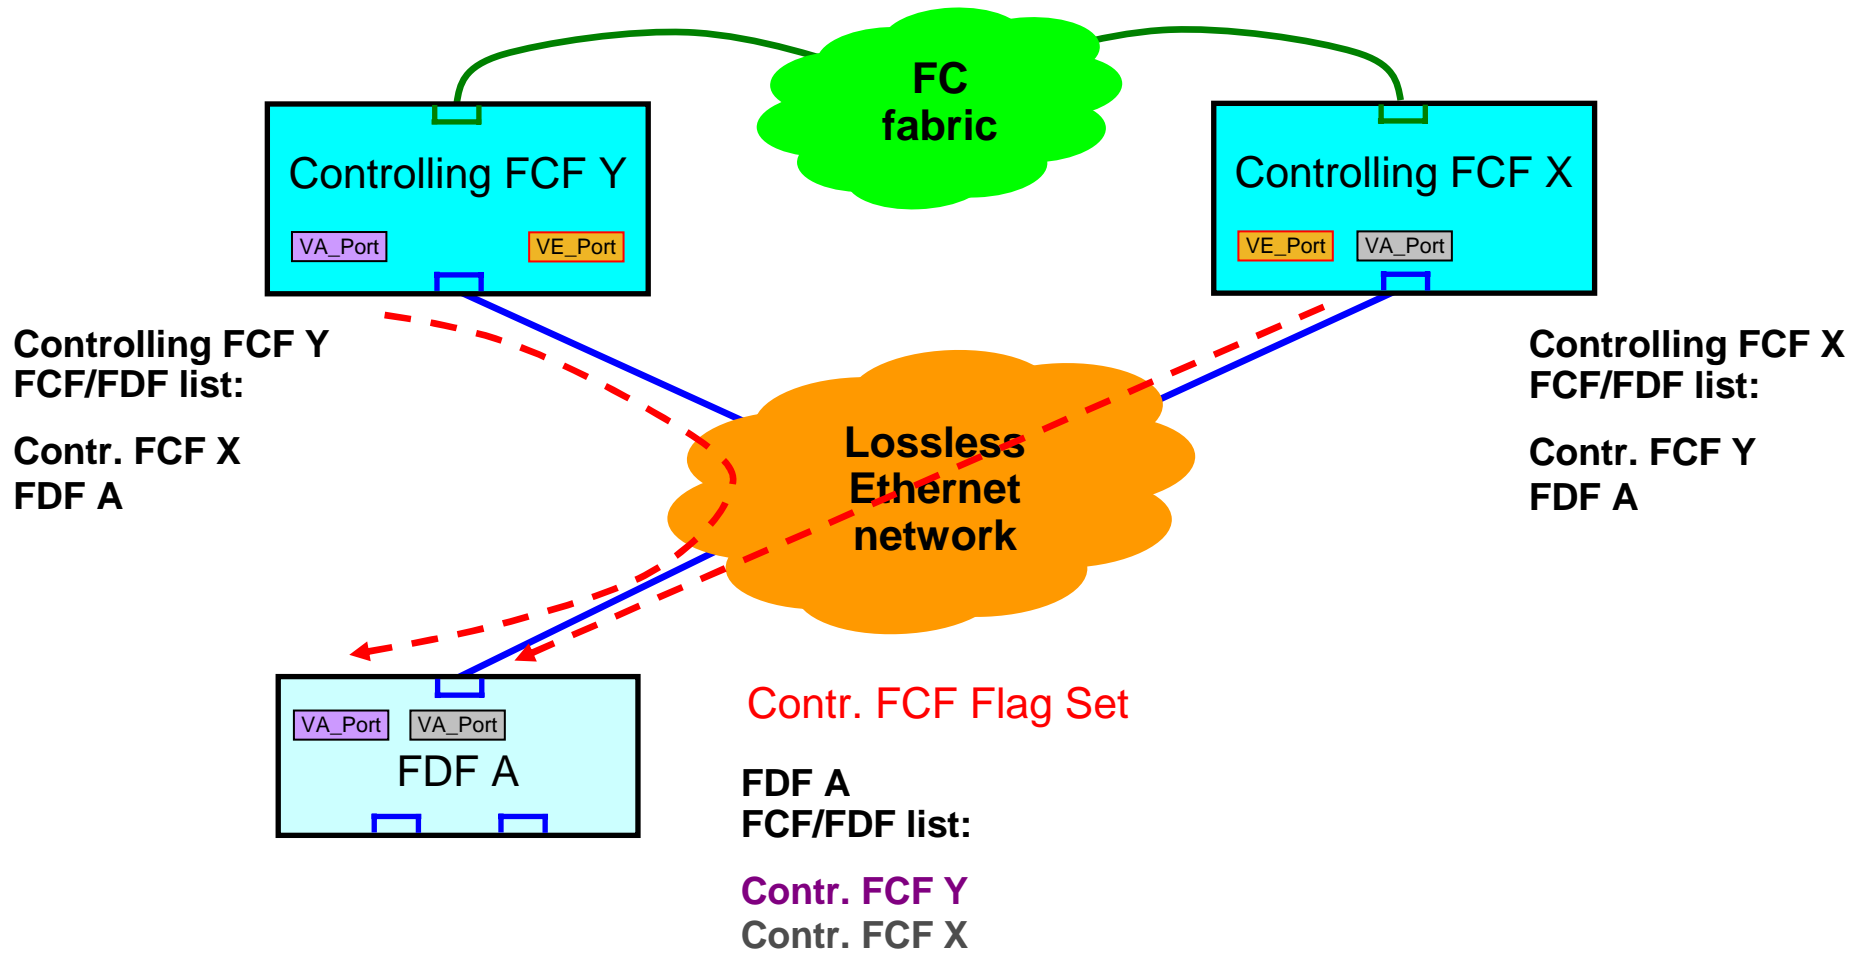
FIP Solicited Advertisements



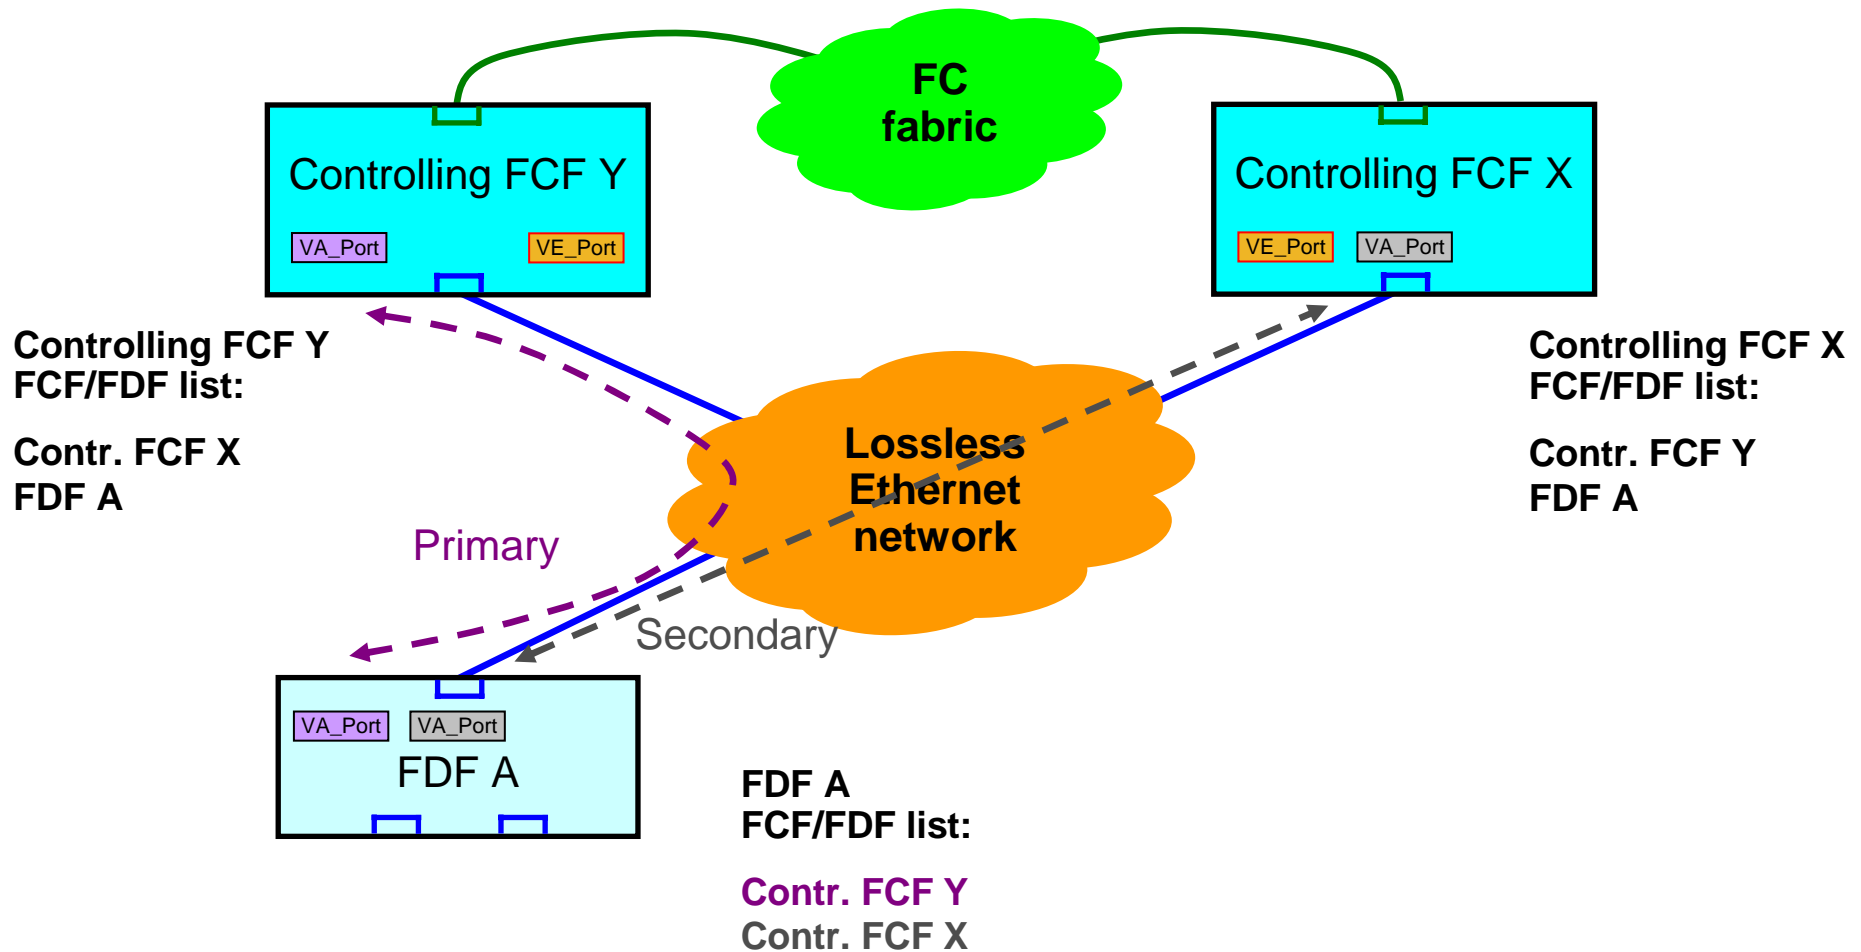
FIP ELP Requests



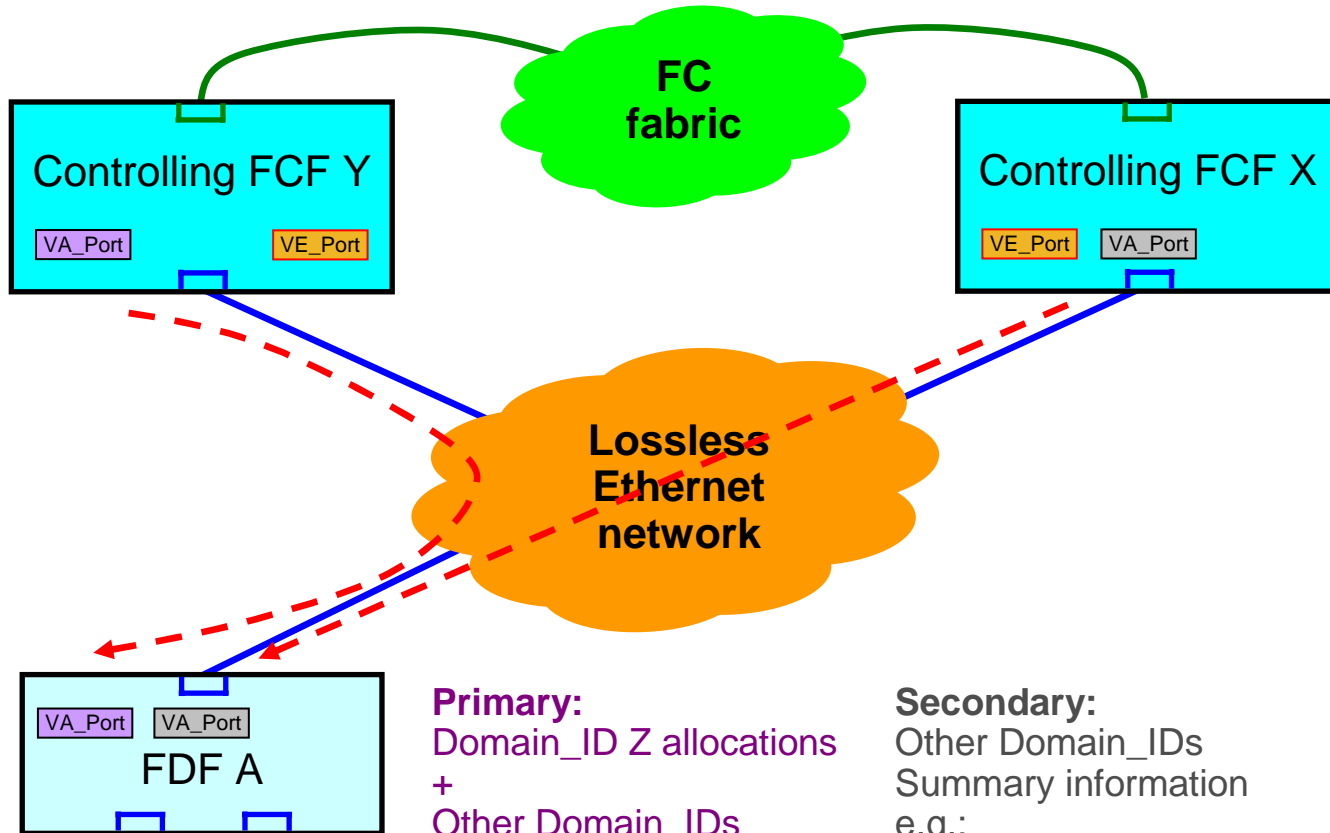
FIP ELP SW_ACC



VA_Port to VA_Port Virtual Links



N_Port_ID Range Allocation Distribution



FDF A Forwarding Table:

- M → Contr. FCF Y
- N → Contr. FCF X
- P → Contr. FCF Y

Primary:
 Domain_ID Z allocations
 +
 Other Domain_IDs
 Summary information
 e.g.:
 Domain_ID M, cost my
 Domain_ID N, cost ny
 Domain_ID P, cost py

Secondary:
 Other Domain_IDs
 Summary information
 e.g.:
 Domain_ID M, cost mx
 Domain_ID N, cost nx
 Domain_ID P, cost px

H1 N_Port_ID Allocation

Allocate 1.2.1 to H1

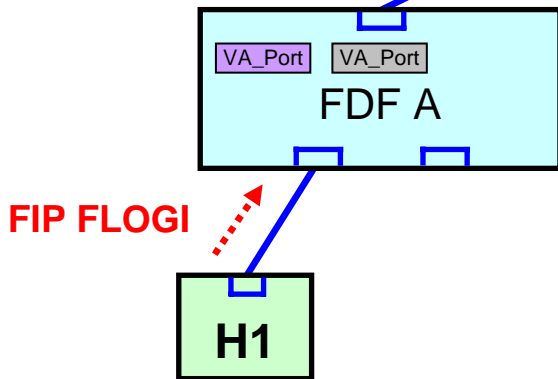


H1 requested FLOGI



Active Zone Set:
Zone P: {H1, H2}
Zone Q: {H1, H4}
Zone R: {H2, H3}

New Zoning ACL:




N_Port_ID and Zoning ACL Distribution

Allocate 1.2.1 to H1




H1: 1.2.1, FDF A, alloc



H1: 1.2.1, FDF A, alloc


Active Zone Set:
 Zone P: {H1, H2}
 Zone Q: {H1, H4}
 Zone R: {H2, H3}



FDF A Forwarding Table:

1.2.1 → self
 M → Contr. FCF Y
 N → Contr. FCF X
 P → Contr. FCF Y

FIP FLOGI ACC



1.2.1

FDF B



Controlling FCF Y
FCF/FDF list:



Contr. FCF X
FDF A

Controlling FCF X
FCF/FDF list:


Contr. FCF Y
FDF A



**Lossless
Ethernet
network**

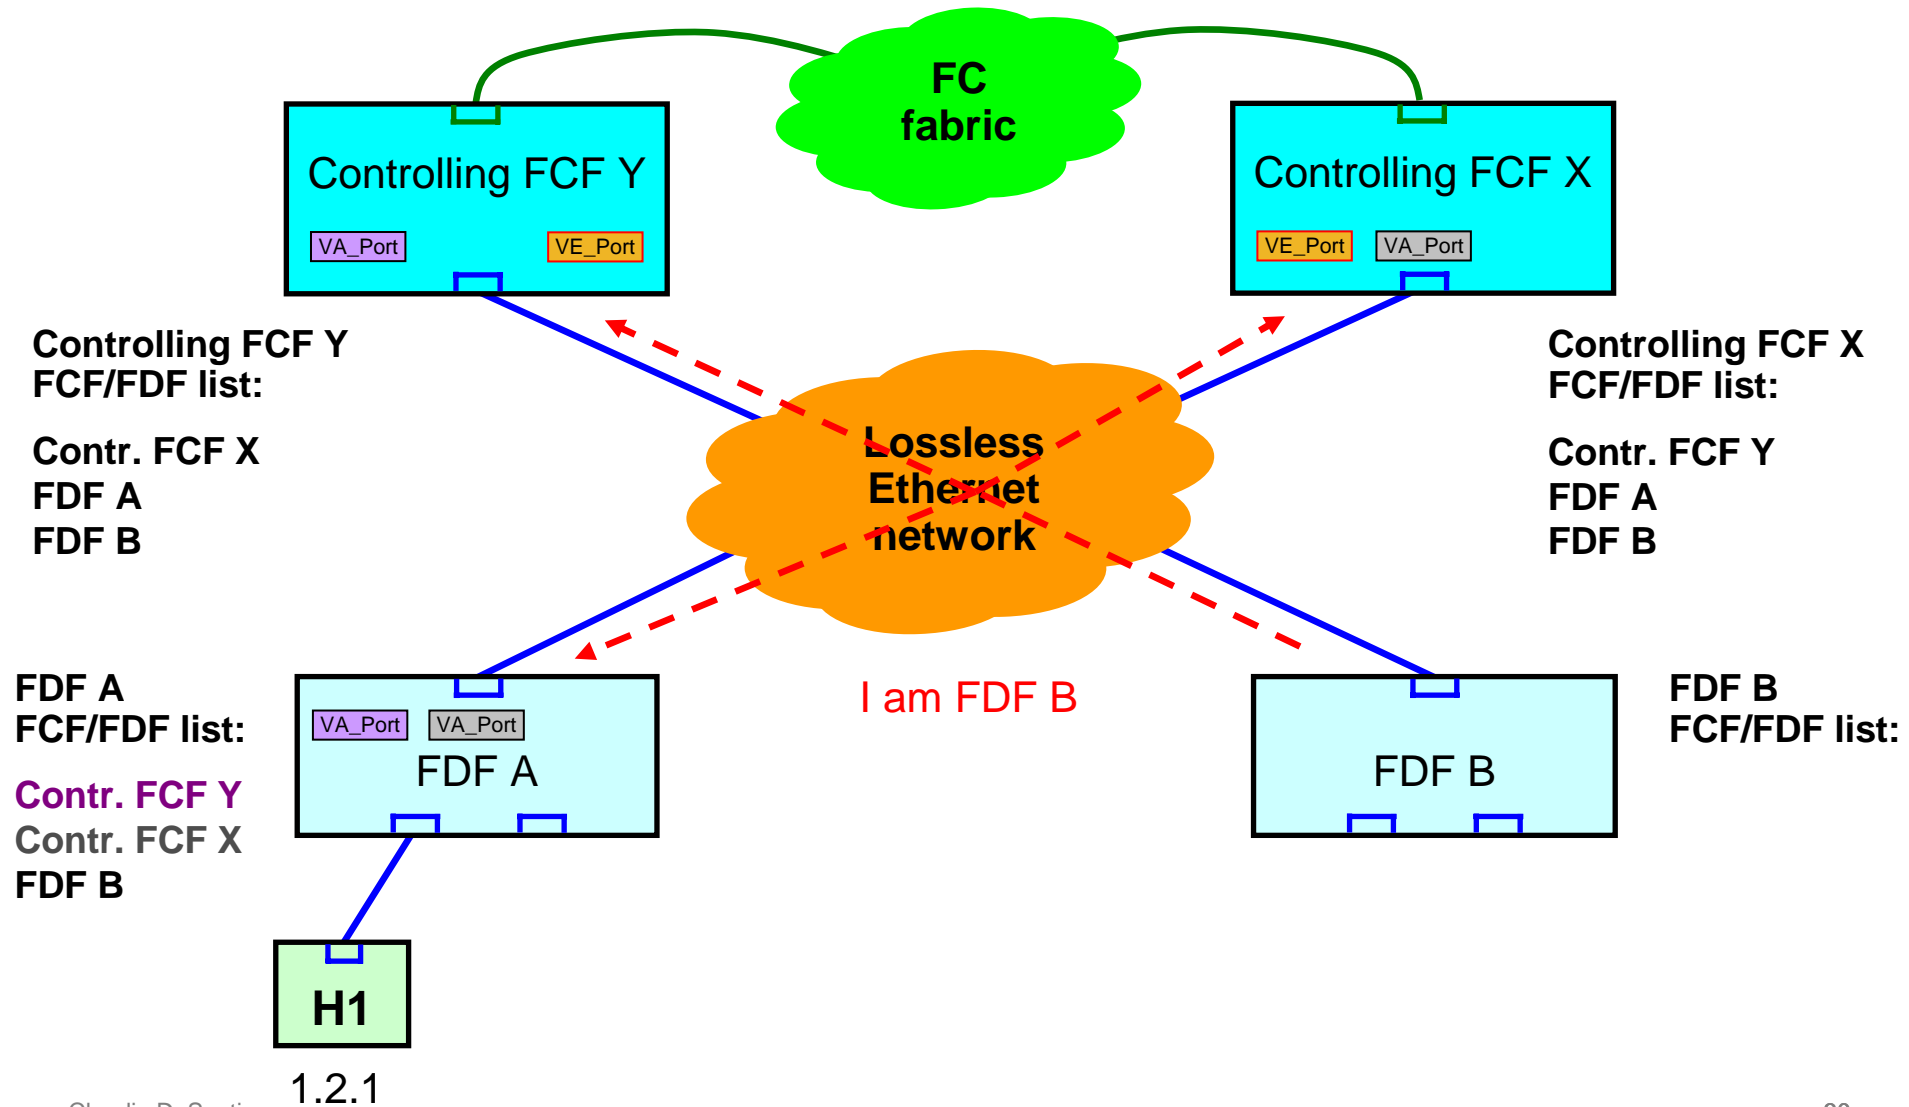


1.2.1

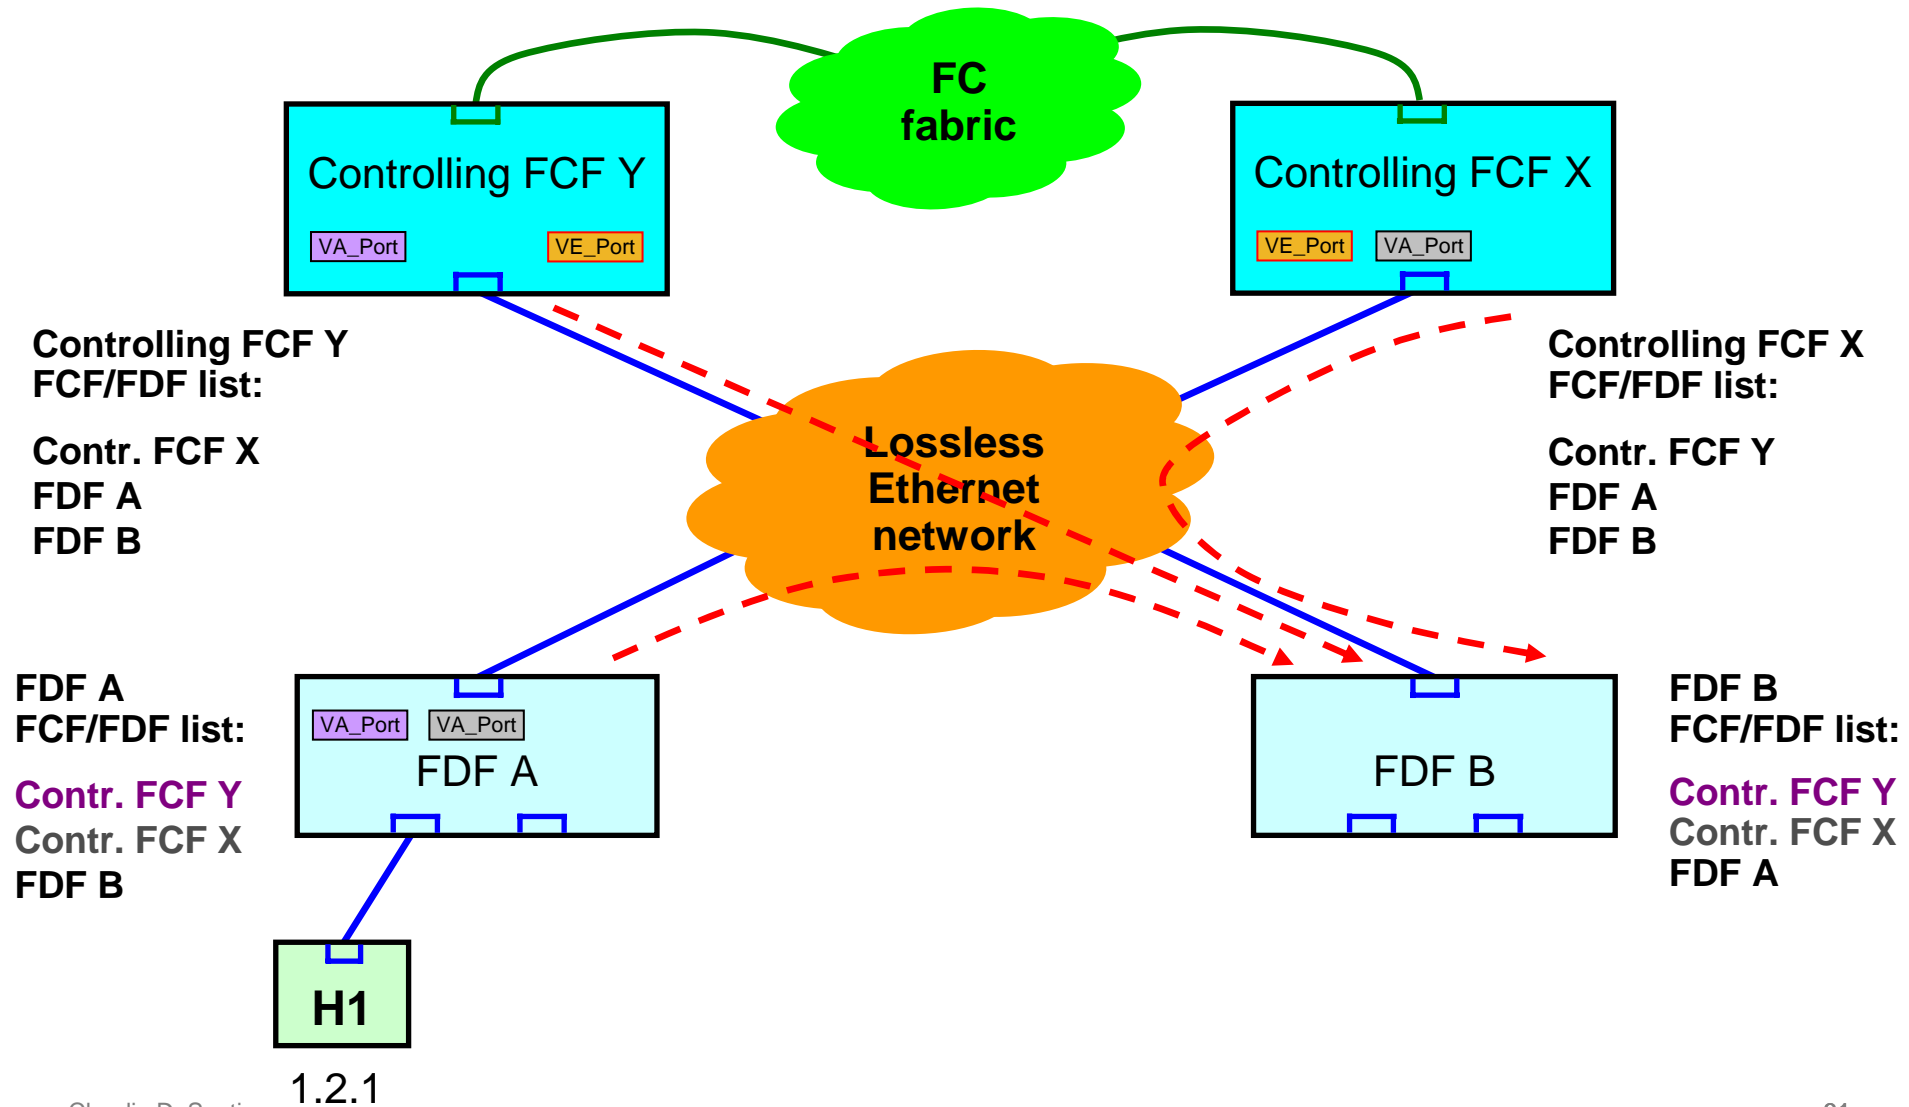


FDF B
FCF/FDF list:

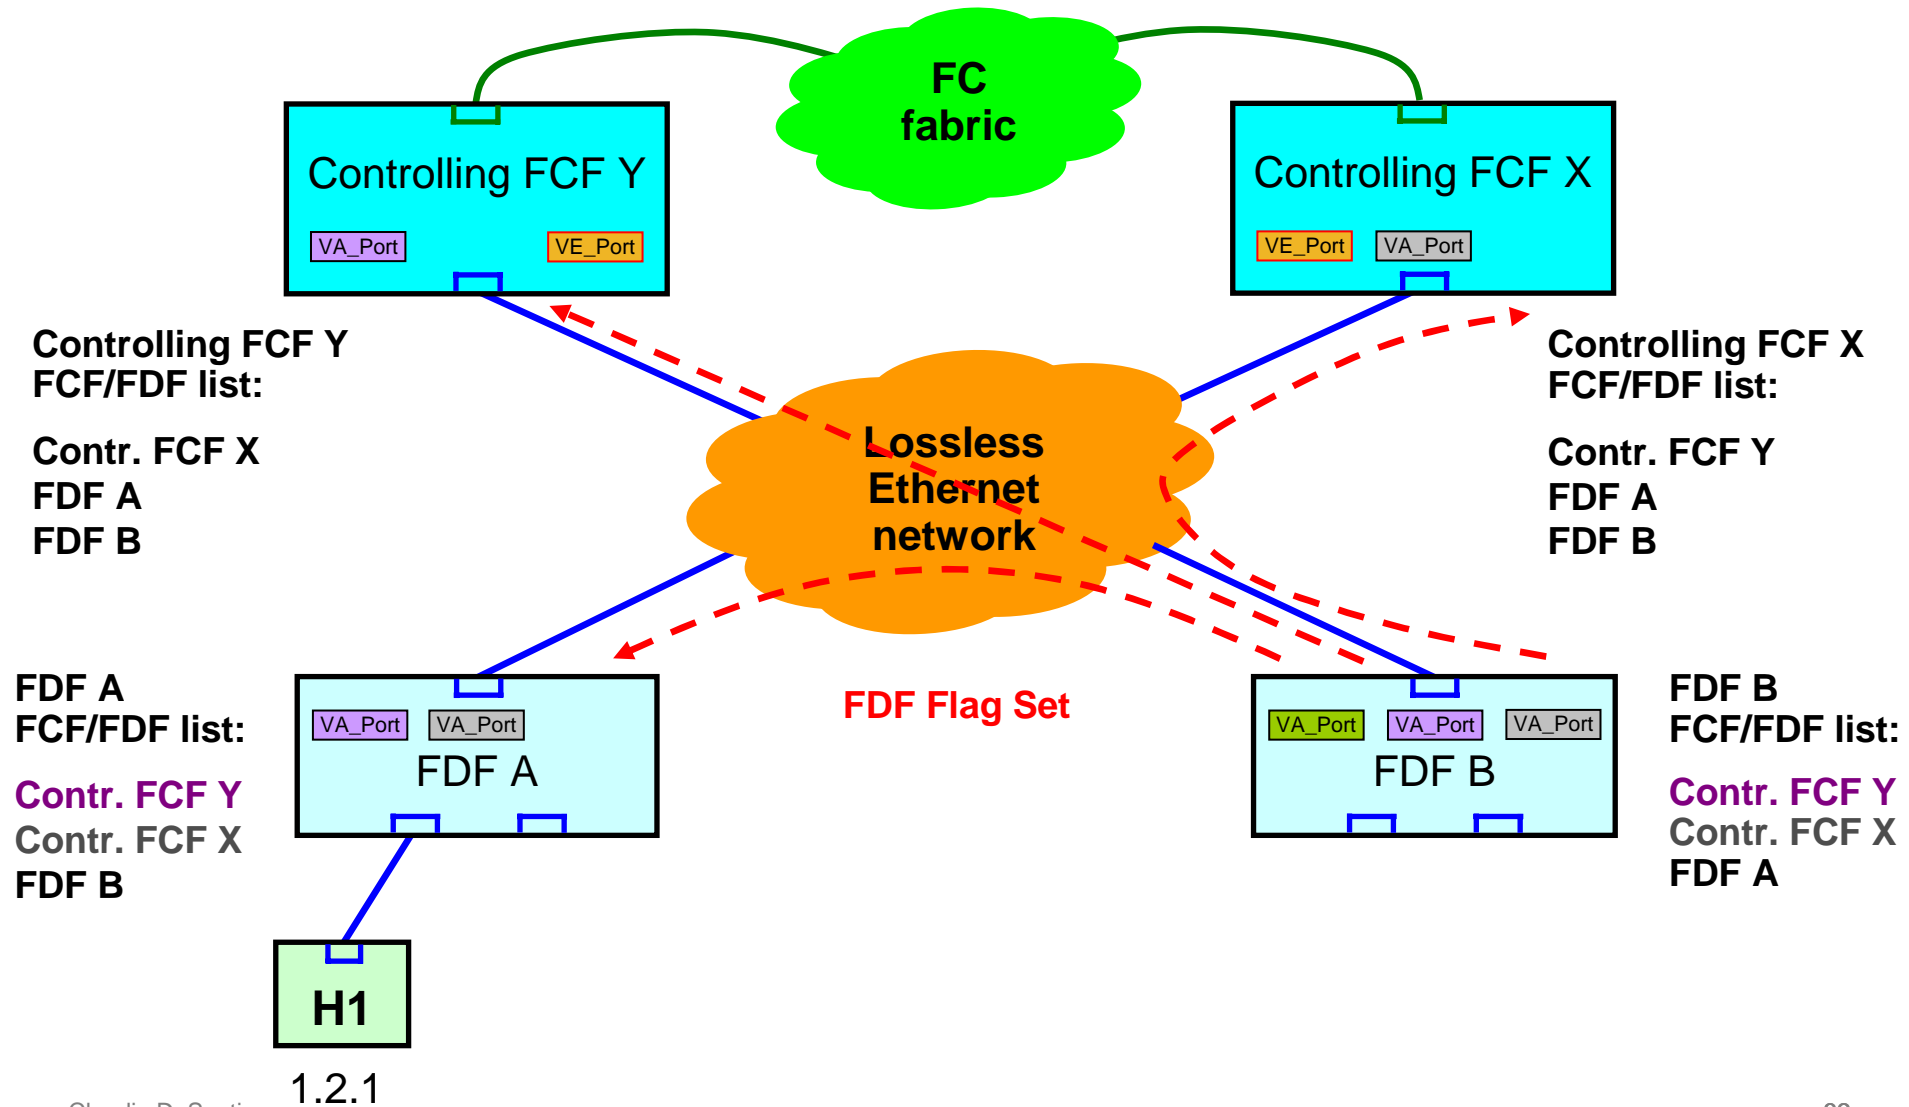
FDF B FIP Solicitation



FIP Solicited Advertisements



FDF B FIP ELP Requests



FDF B FIP ELP SW_ACC




Contr. FCF Flag Set

Controlling FCF Y
FCF/FDF list:




Contr. FCF X
FDF A
FDF B

Controlling FCF X
FCF/FDF list:

Contr. FCF Y
FDF A
FDF B



FDF Flag Set



1.2.1

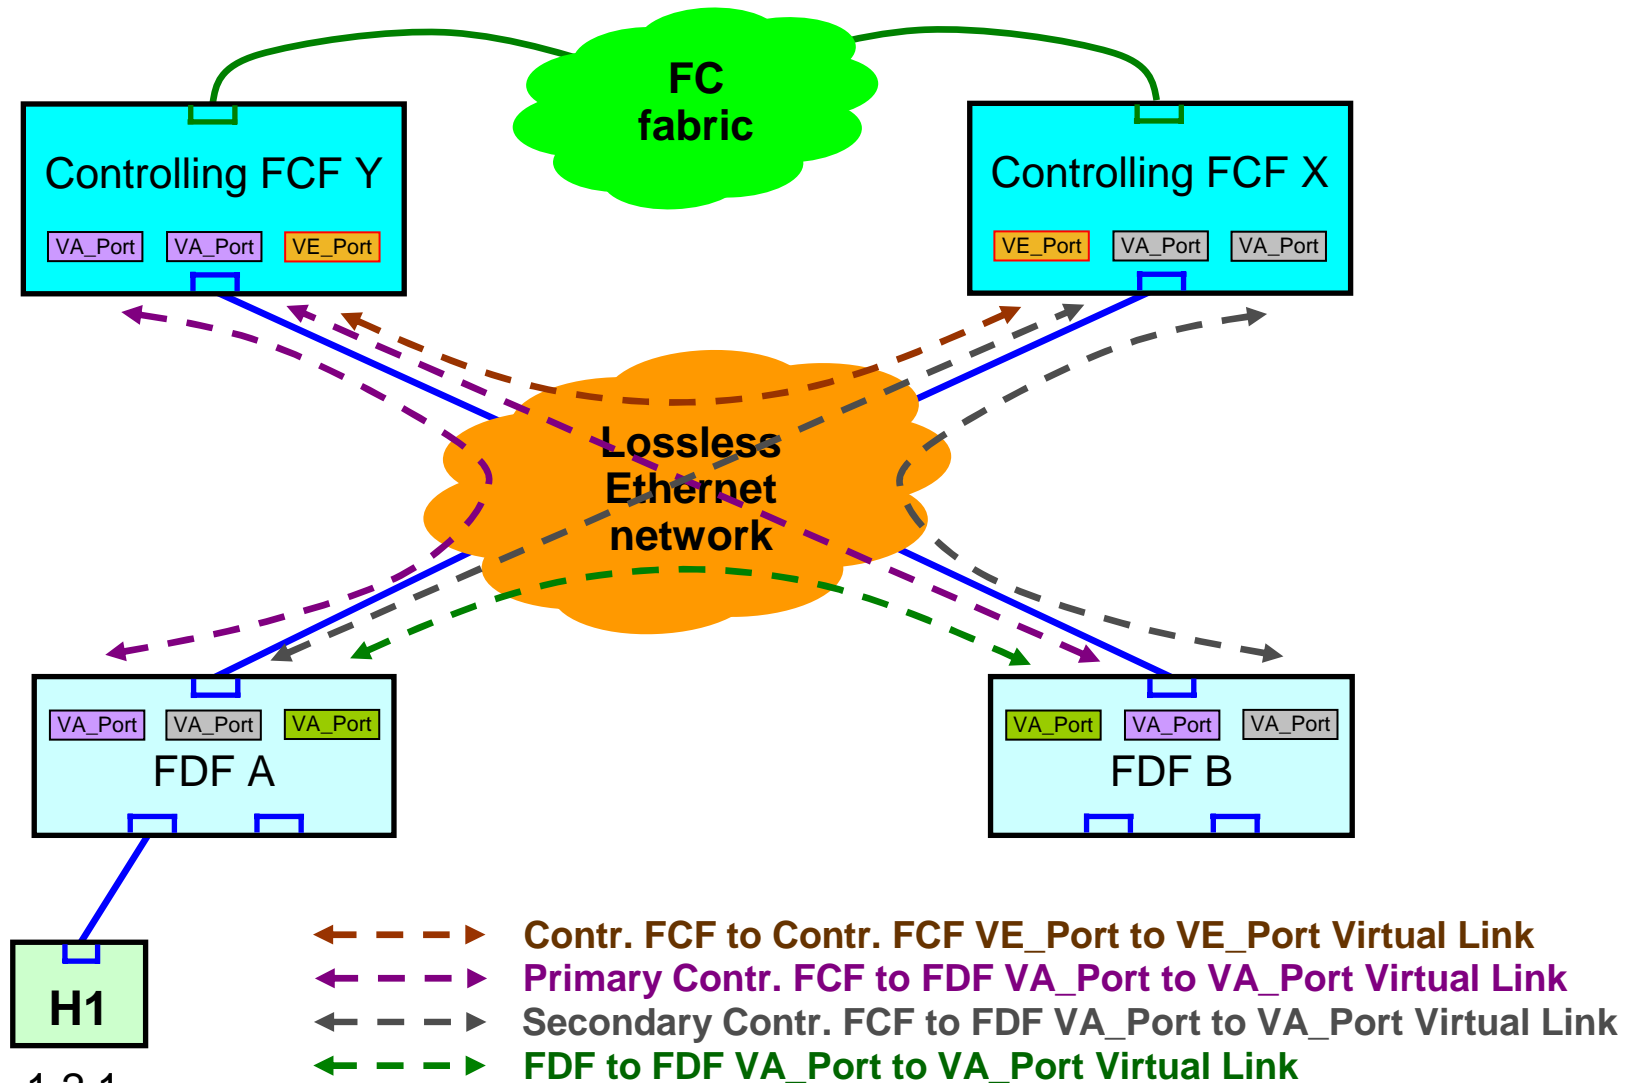
FDF A
FCF/FDF list:

Contr. FCF Y
Contr. FCF X
FDF B

FDF B
FCF/FDF list:

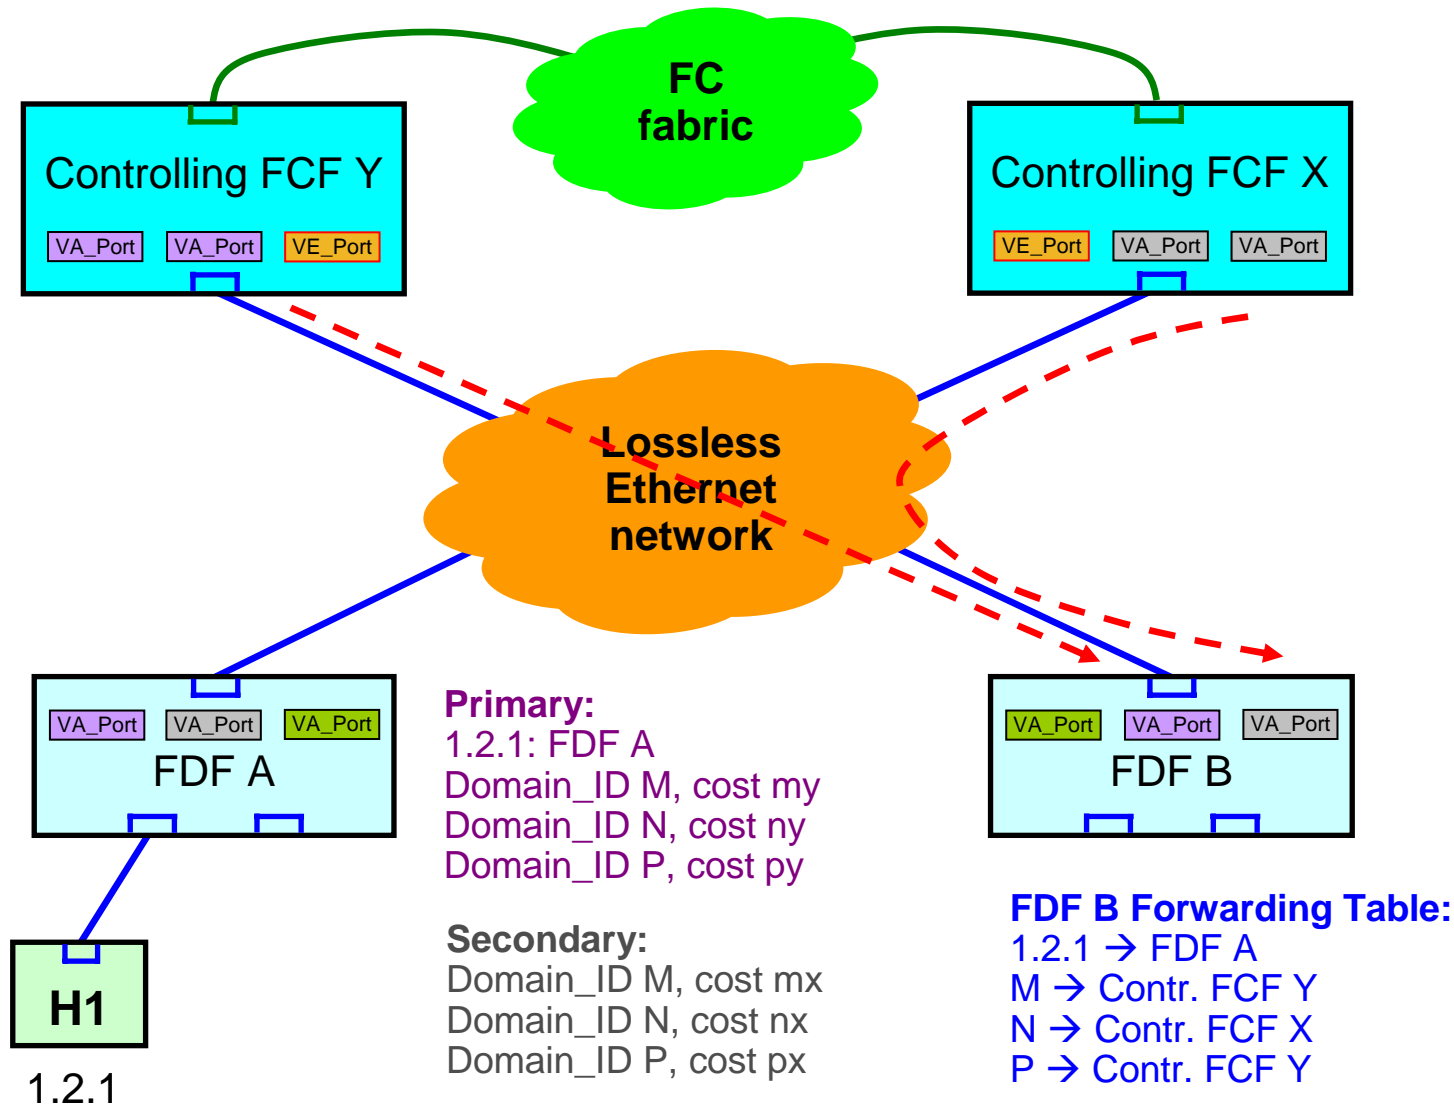
Contr. FCF Y
Contr. FCF X
FDF A

Virtual Links

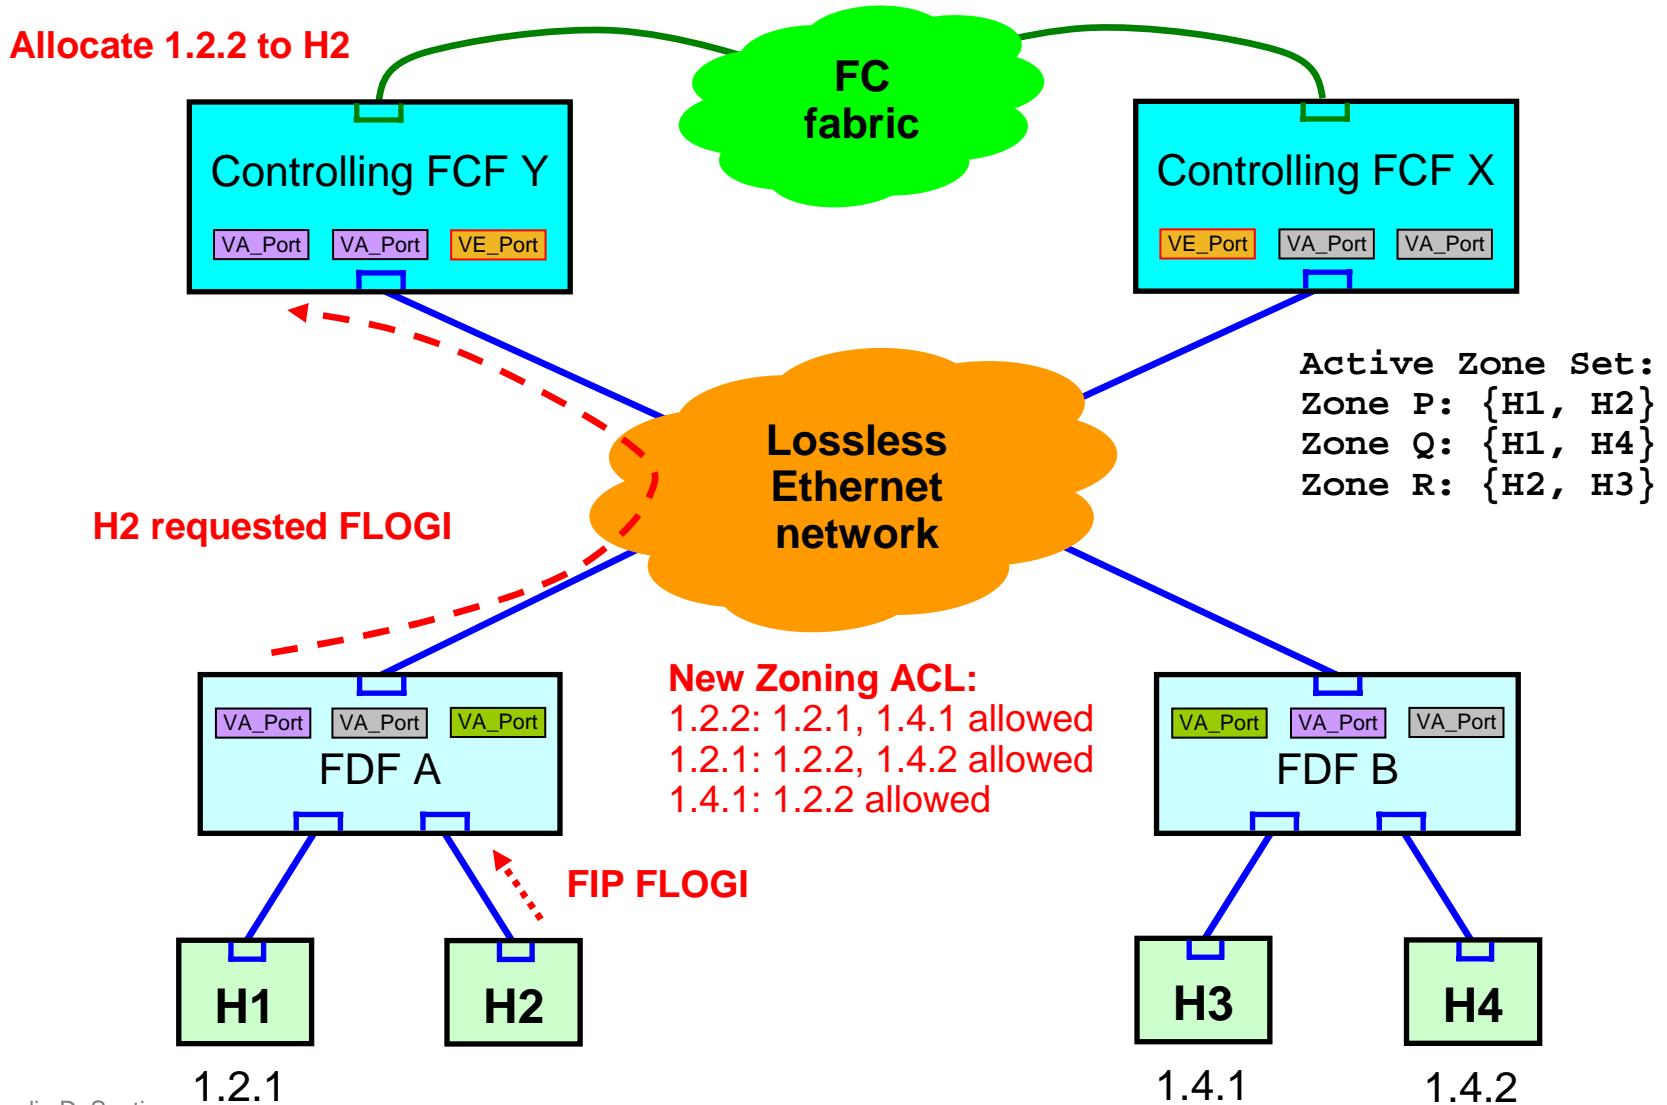


1.2.1

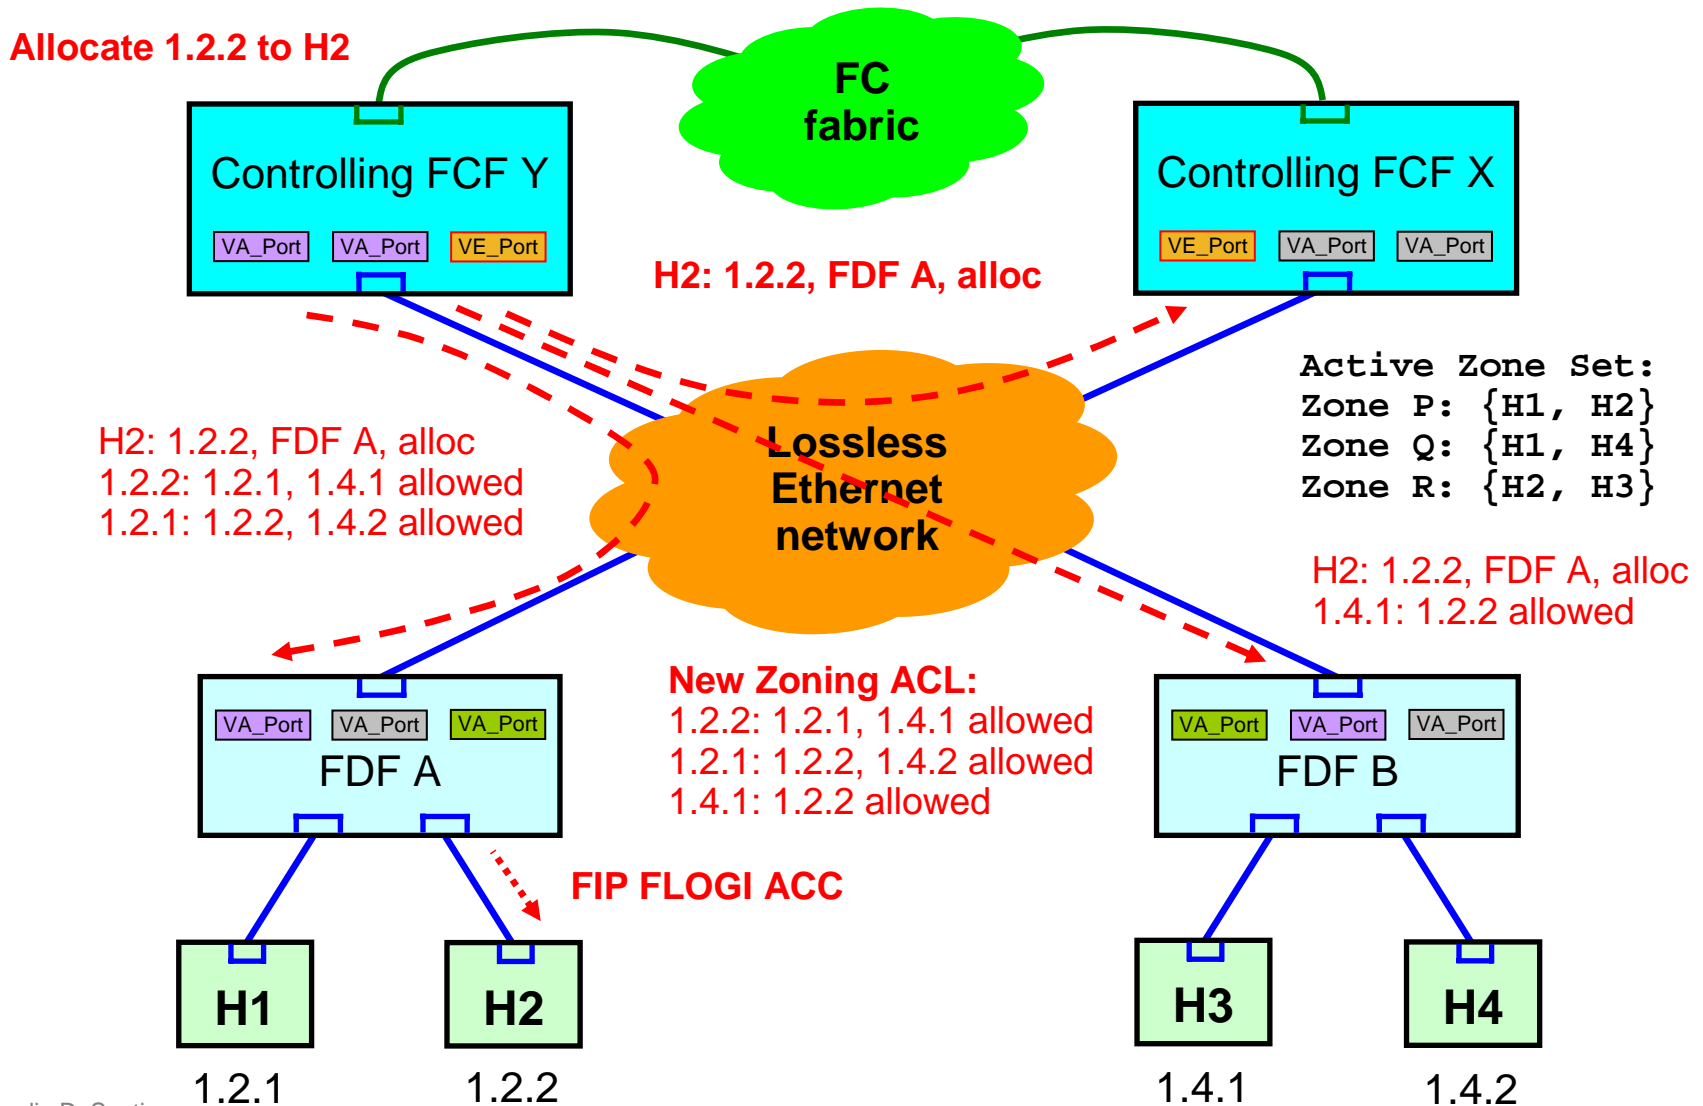
N_Port_ID Range Allocation Distribution



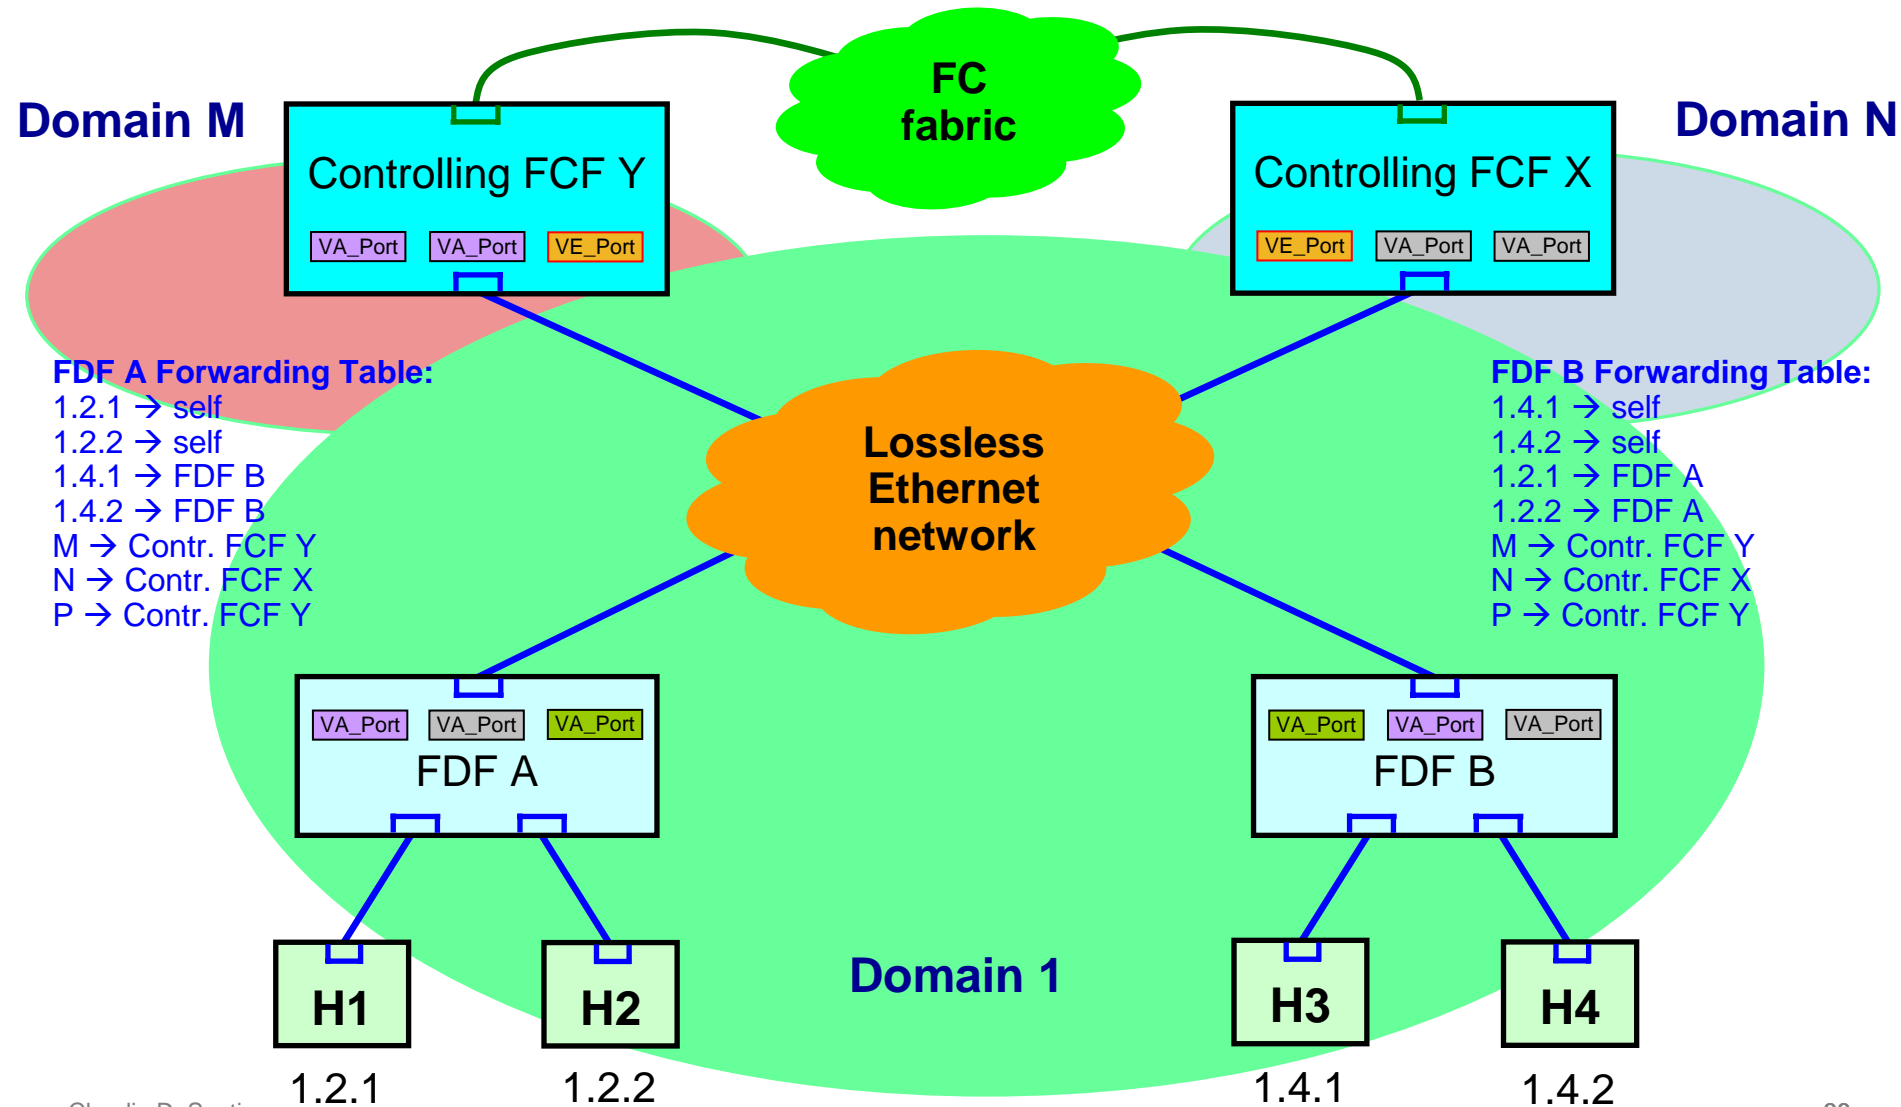
N_Port_ID Allocation (2)



N_Port_ID and Zoning ACL Distribution (2)



Forwarding Tables



Generic Services

- Fabric based Generic Services (e.g., Name Server, Management Server) are provided by the controlling FCF
- On receiving a Common Transport frame, an FDF forwards it to its controlling FCF over the VA_Port to VA_Port virtual link

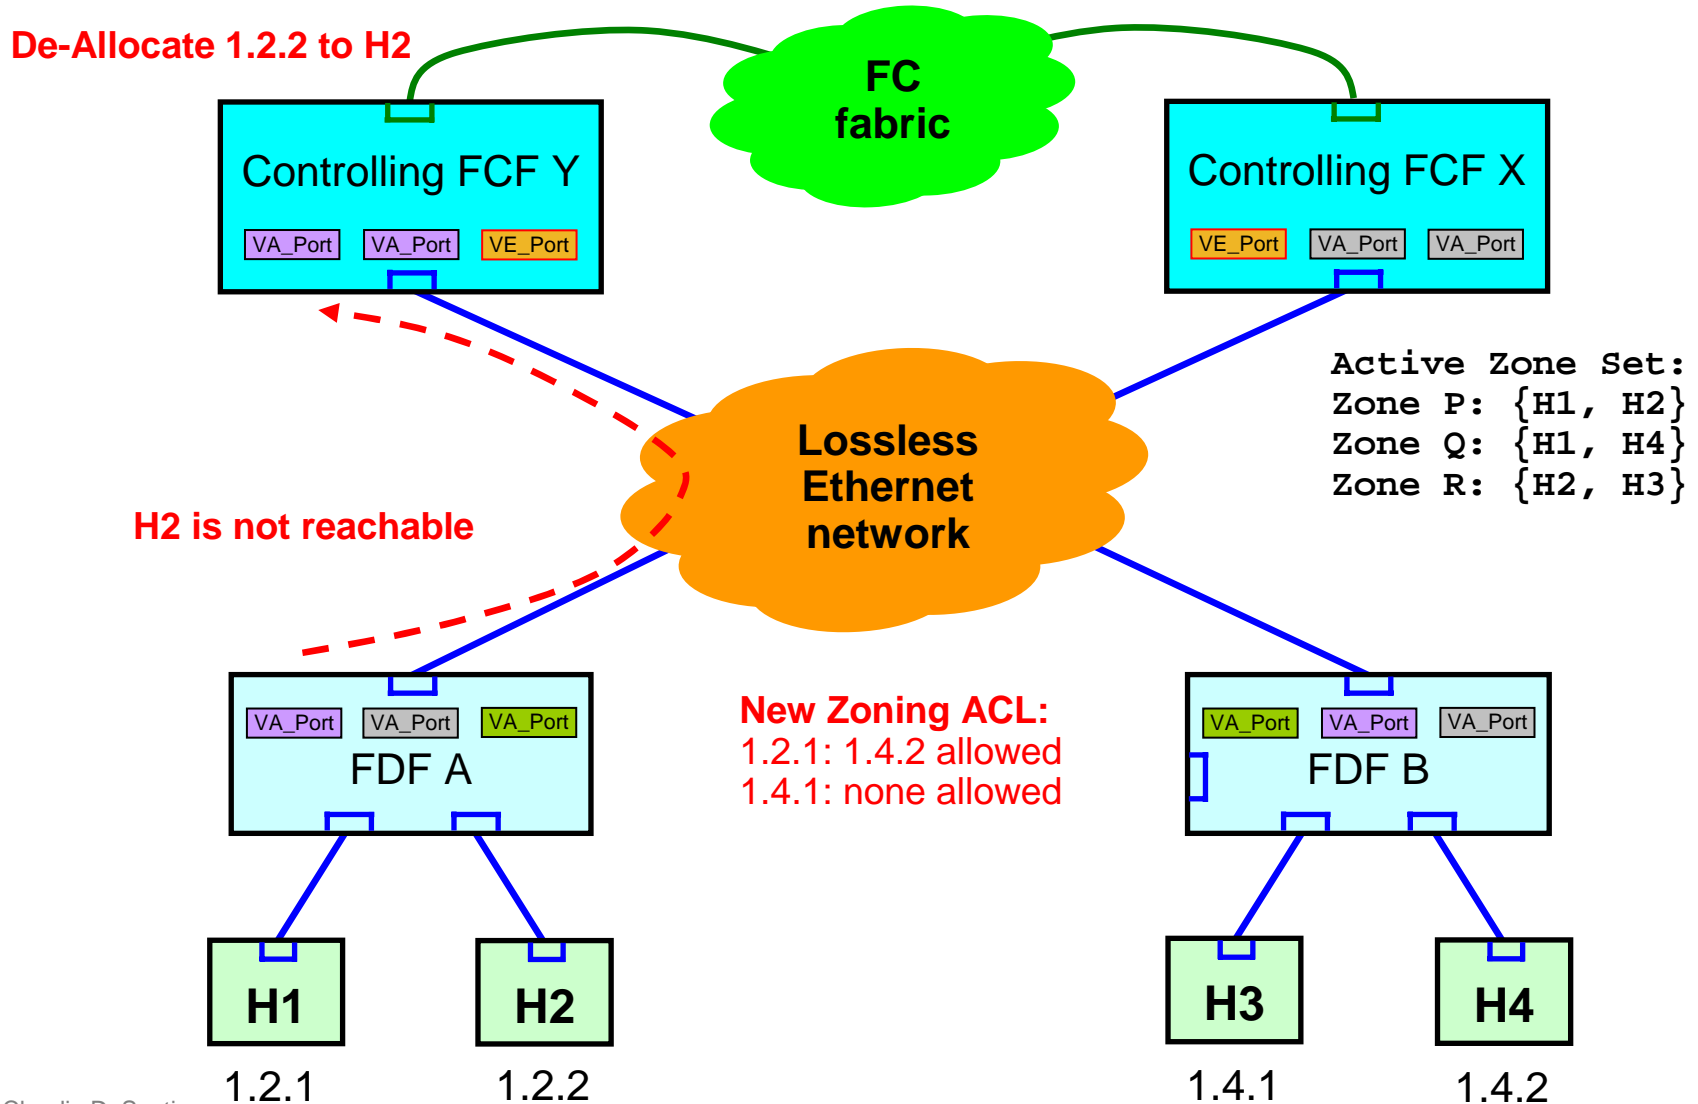


Image Credit:
Flickr user shazam791
Creative Commons License

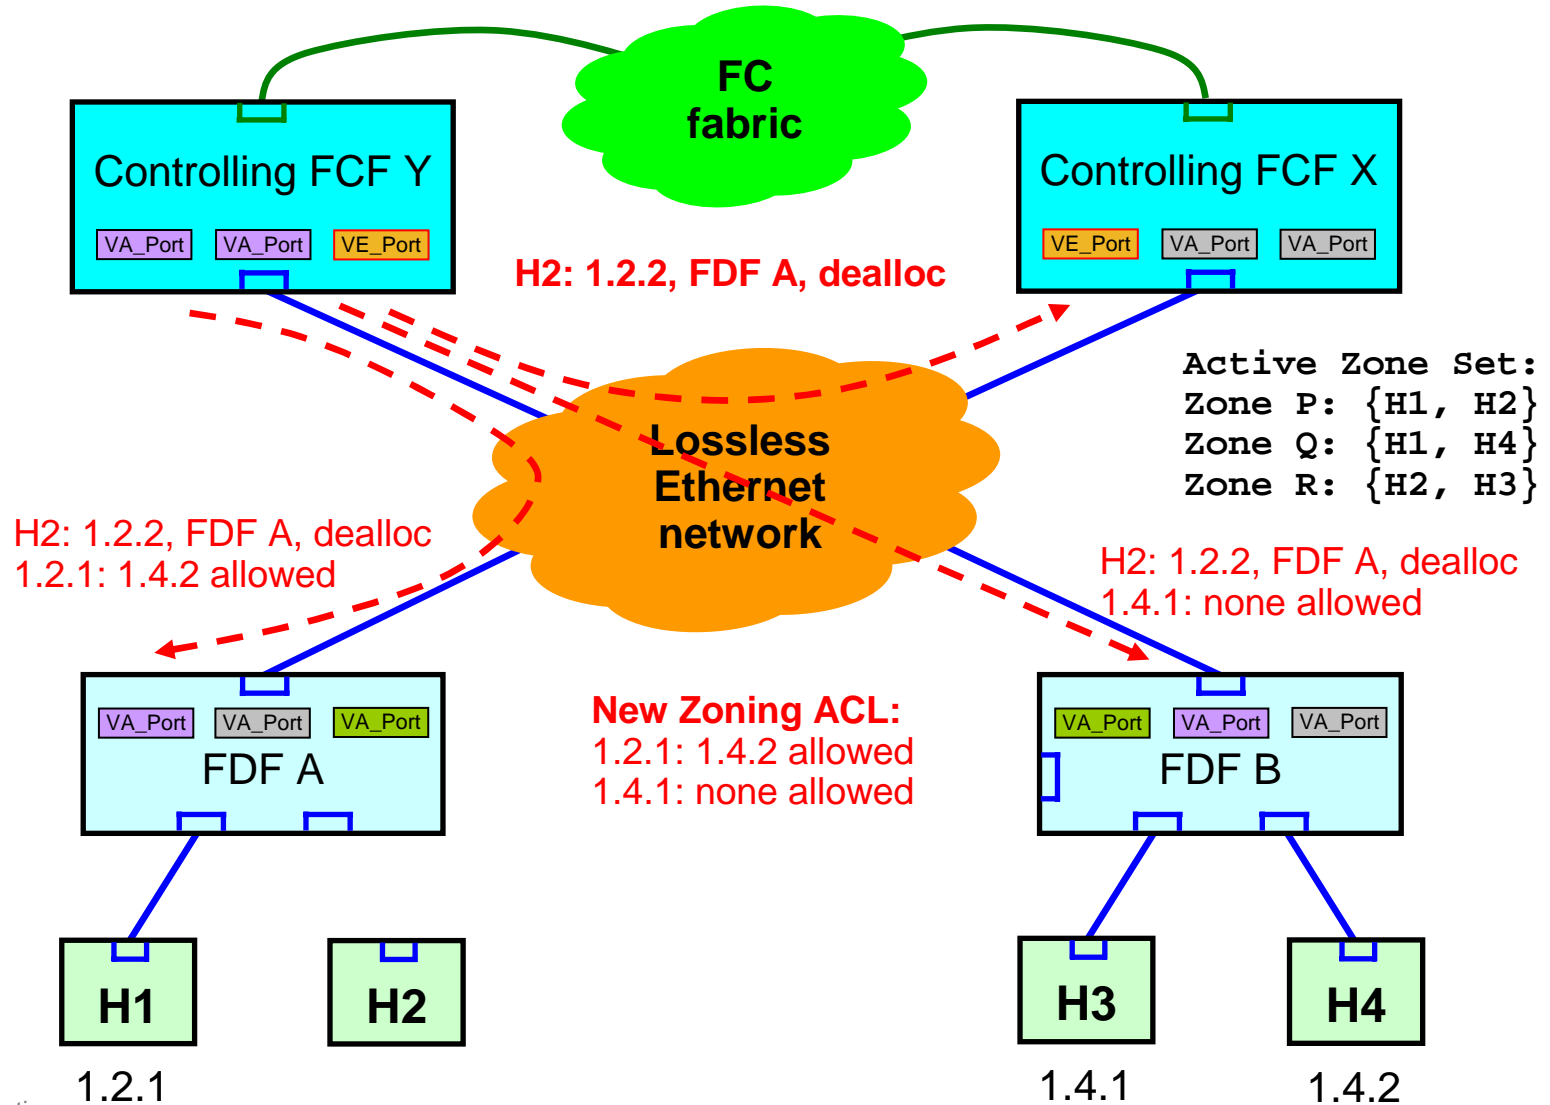
N_Port_ID Move

- **Centralizing the N_Port_ID allocation function on the controlling FCF enables N_Port_ID moves**

N_Port_ID Move (1)



N_Port_ID Move (2)




N_Port_ID Move (3)

Allocate 1.2.2 to H2

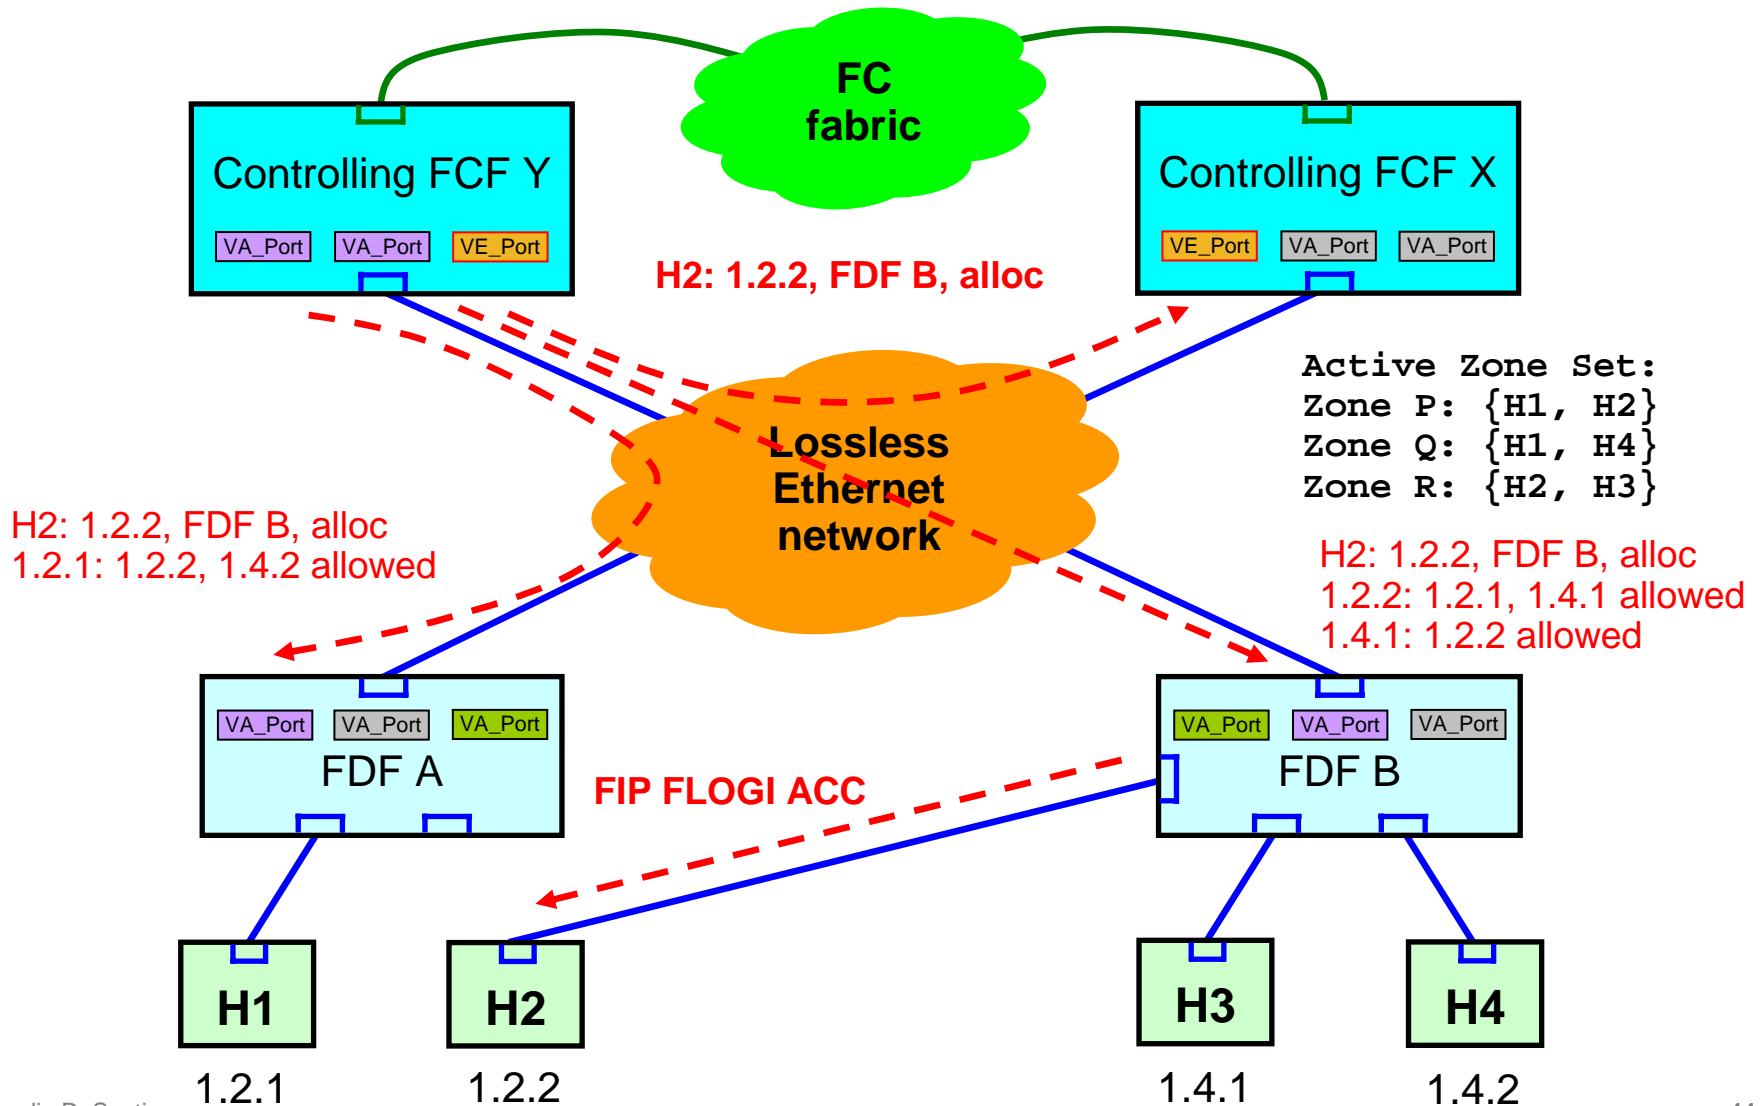


New Zoning ACL:
1.2.2: 1.2.1, 1.4.1 allowed
1.2.1: 1.2.2, 1.4.2 allowed
1.4.1: 1.2.2 allowed

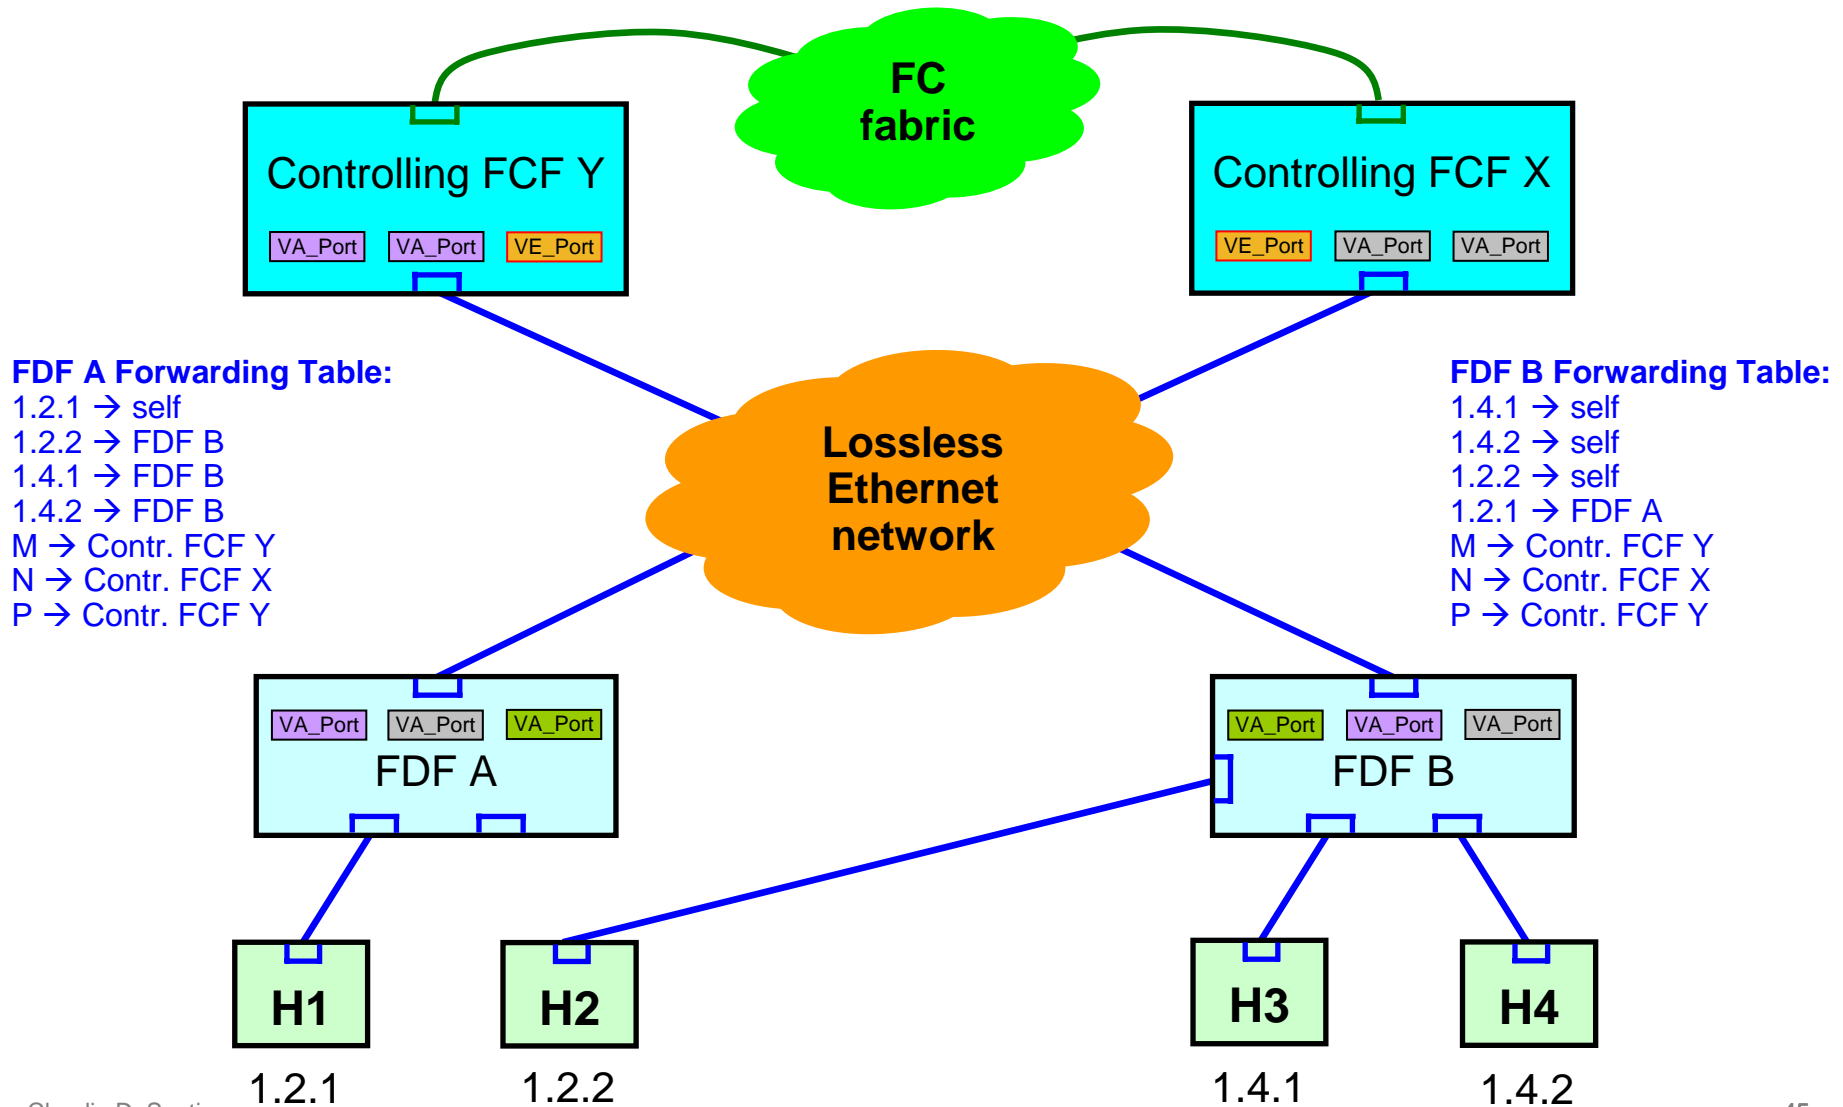
Active Zone Set:
Zone P: {H1, H2}
Zone Q: {H1, H4}
Zone R: {H2, H3}



N_Port_ID Move (4)



Forwarding Tables



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 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

N_Port_ID Range Allocation SW_ILS

Item	Size
SW_ILS Code	4
Controlling FCF Switch_Name	8
Recipient FDF_Name	8
Number of N_Port_ID Range Entries (p)	4
N_Port_ID Range Entry #1	
N_Port_ID Range Entry #2	
...	
N_Port_ID Range Entry #p	
Number of Reachable Domain_ID Entries (r)	4
Reachable Domain_ID Entry #1	
Reachable Domain_ID Entry #2	
...	
Reachable Domain_ID Entry #r	

N_Port_ID Range and Reachable Domain_ID Entries

Item	Size
FDF_Name	8
Number of N_Port_ID Ranges (q)	4
N_Port_ID Range #1	4
N_Port_ID Range #2	4
...	
N_Port_ID Range #q	4

The N_Port_ID Range is defined by an N_Port_ID in the least significant three bytes, and by the number of bits defining the range in the most significant byte (e.g., 1.2.0/15)

Item	Size
Reachable Domain_ID, cost	4

The Reachable Domain_ID is defined by a Domain_ID in the most significant byte, and by the cost to reach it in the least significant three bytes (e.g., 24, cost 1250)

N_Port_ID Allocation Request SW_ILS

Item	Size
SW_ILS Code	4
Requesting FDF_Name	8
Controlling FCF Switch_Name	8
F_Port_Name	8
FLOGI/NPIV FDISC Parameters	116

N_Port_ID Deallocation Request SW_ILS

Item	Size
SW_ILS Code	4
Requesting FDF_Name	8
Controlling FCF Switch_Name	8
Unreachable N_Port_Name	8
Unreachable N_Port_ID	4

N_Port_ID and Zoning ACL Distribution SW_ILS

Item	Size
SW_ILS Code	4
Controlling FCF Switch_Name	8
Recipient FDF_Name / Controlling FCF Switch_Name	8
Flags	4
N_Port_Name of Allocated/Deallocated N_Port_ID	8
Allocated/Deallocated N_Port_ID	4
FDF_Name of the FDF to which the Allocated/Deallocated N_Port_ID is connected	8
FLOGI Parameters	116
Number of Peering Entries (h)	4
Peering Entry #1	
Peering Entry #2	
...	
Peering Entry #h	

The most significant bit of the most significant byte indicates deallocation vs. allocation

Present only if the appropriate flag is set to one

Peering Entry Format

Item	Size
Peering N_Port_ID	4
Number of Allowed Peers (k)	4
Peer N_Port_ID #1	4
Peer N_Port_ID #2	4
...	
Peer N_Port_ID #k	4

The FDF

- **Instantiates VF_Ports and VA_Ports**

 - No VE_Ports!**

- **VF_Ports operate as defined in FC-BB-5**

 - No change**

- **VA_Ports:**

 - Are not VE_Ports**

 - Are instantiated through FIP ELP with FDF flag**

 - Performs Virtual Link maintenance through FIP advertisements**

 - Use only four new unidirectional SW_ILSs:**

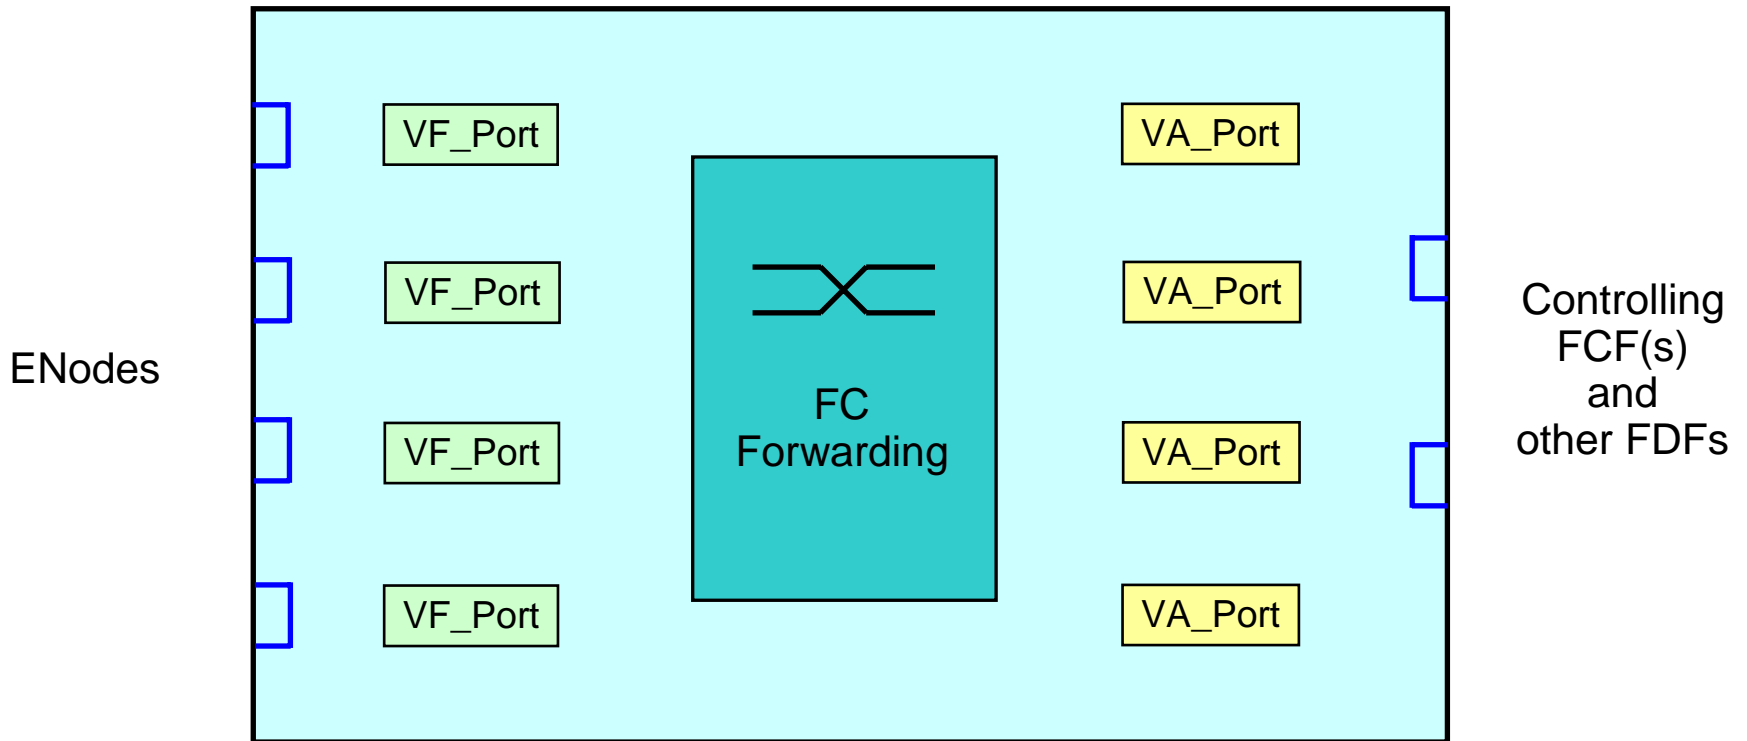
 - N_Port_ID Range Allocation

 - N_Port_ID Allocation Request

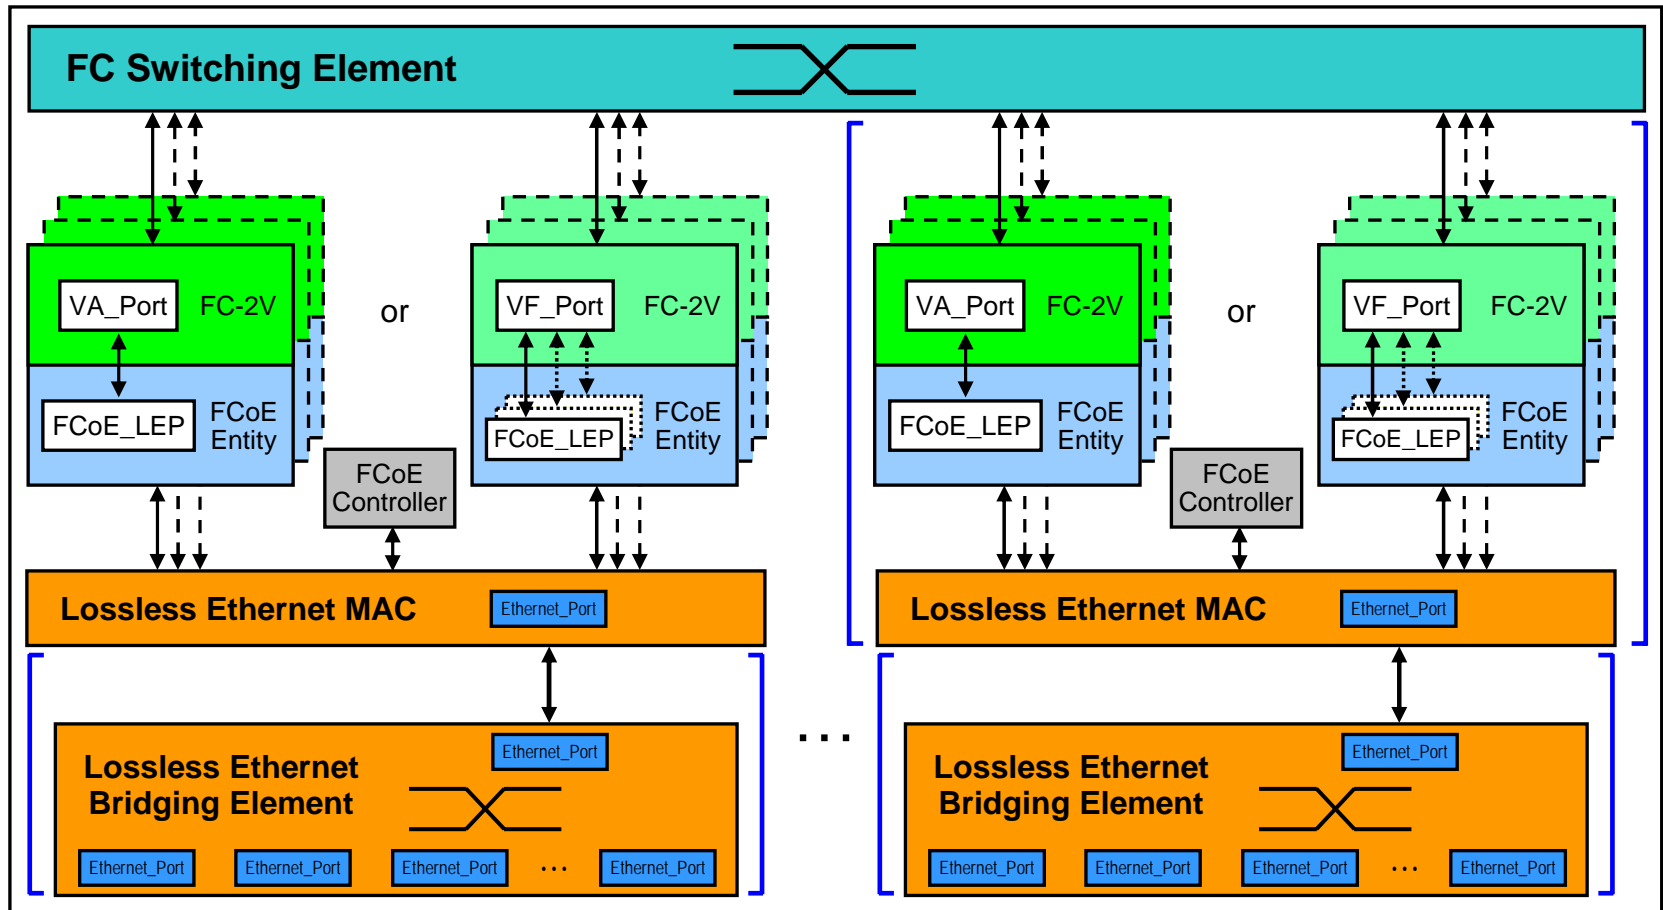
 - N_Port_ID Deallocation Request

 - N_Port_ID and Zoning ACL Distribution


FDF Forwarding



FDF Functional Model



Thank You



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