



VA_Ports: FDF / Controlling FCF Protocols

Claudio DeSanti

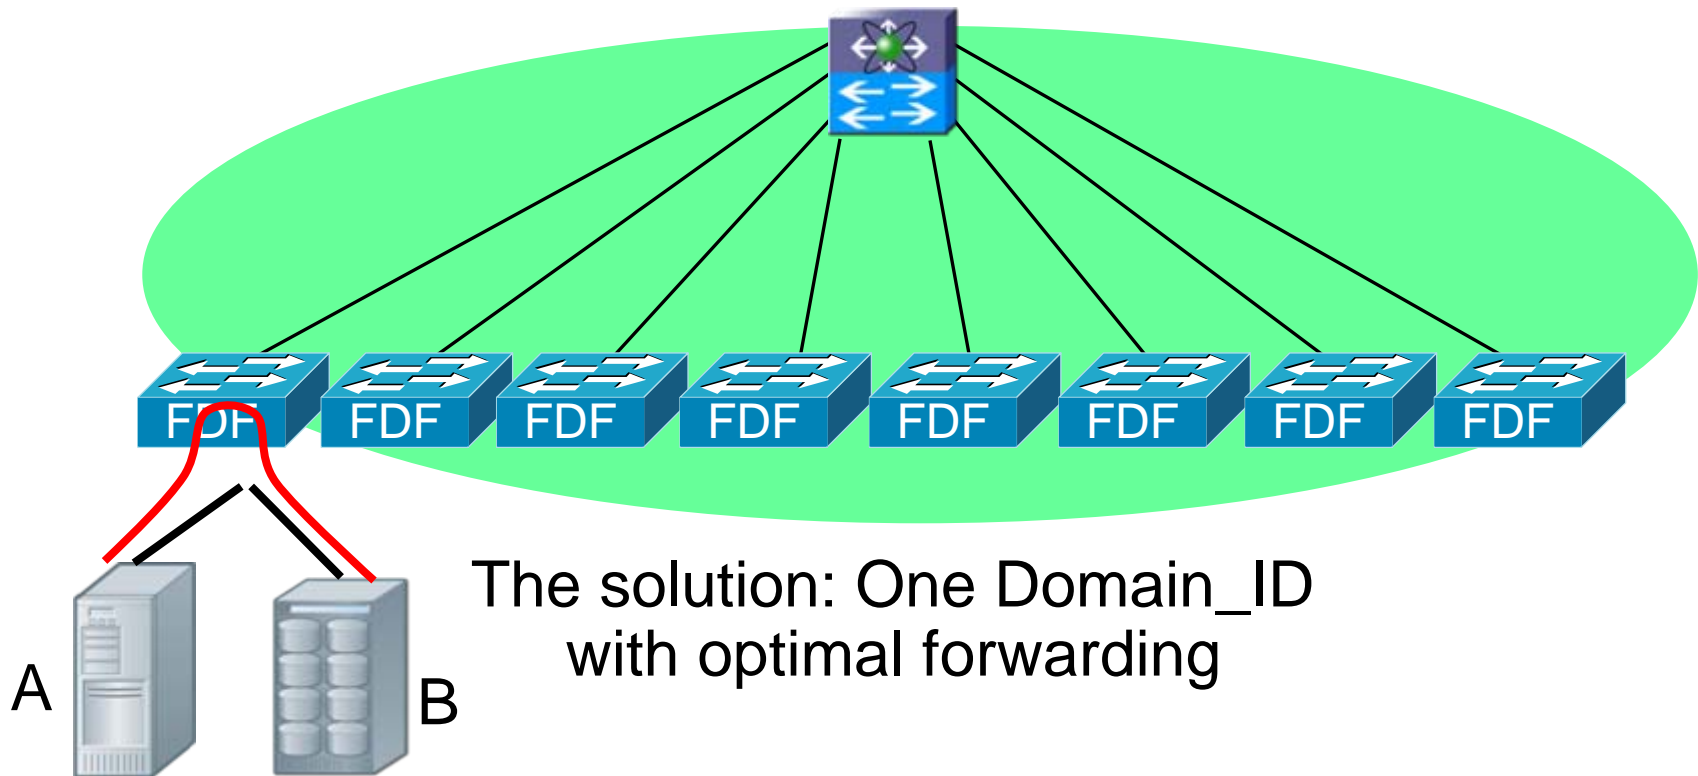
T11/10-271v0, June 2010

Image Credit:
Landon Curt Noll

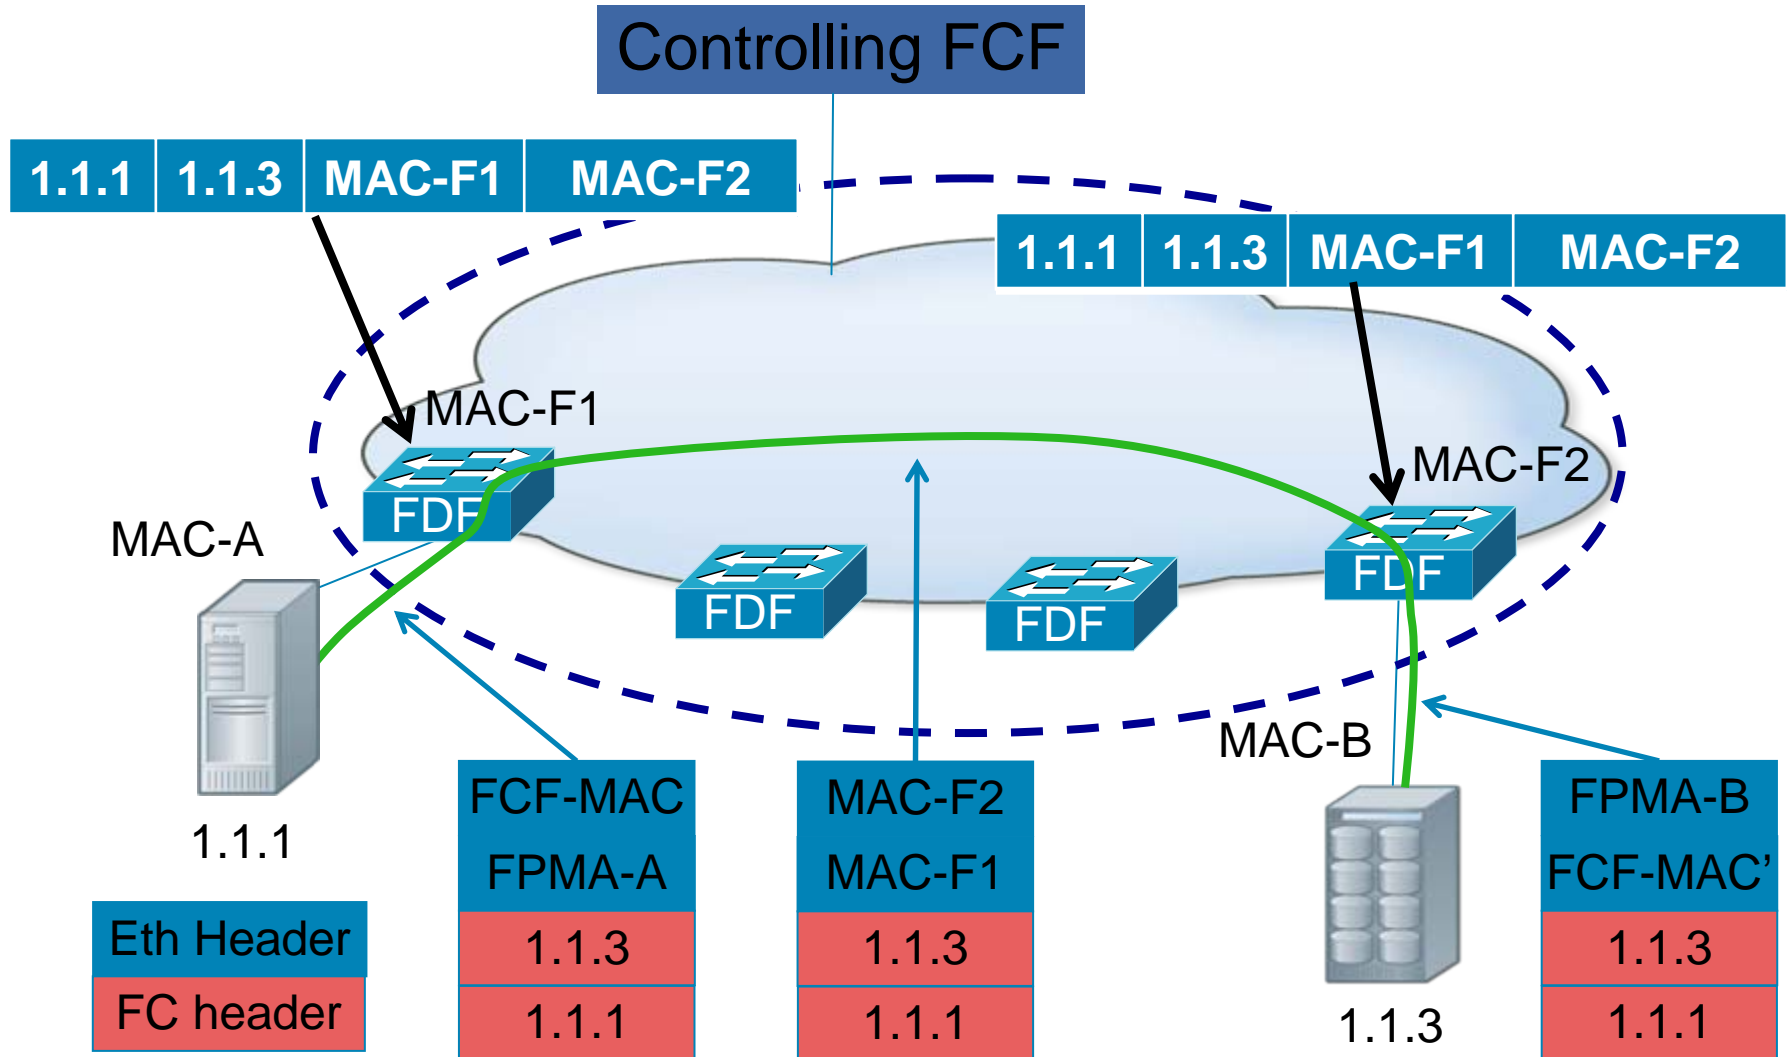


Introducing the Controlling FCF / FDFs

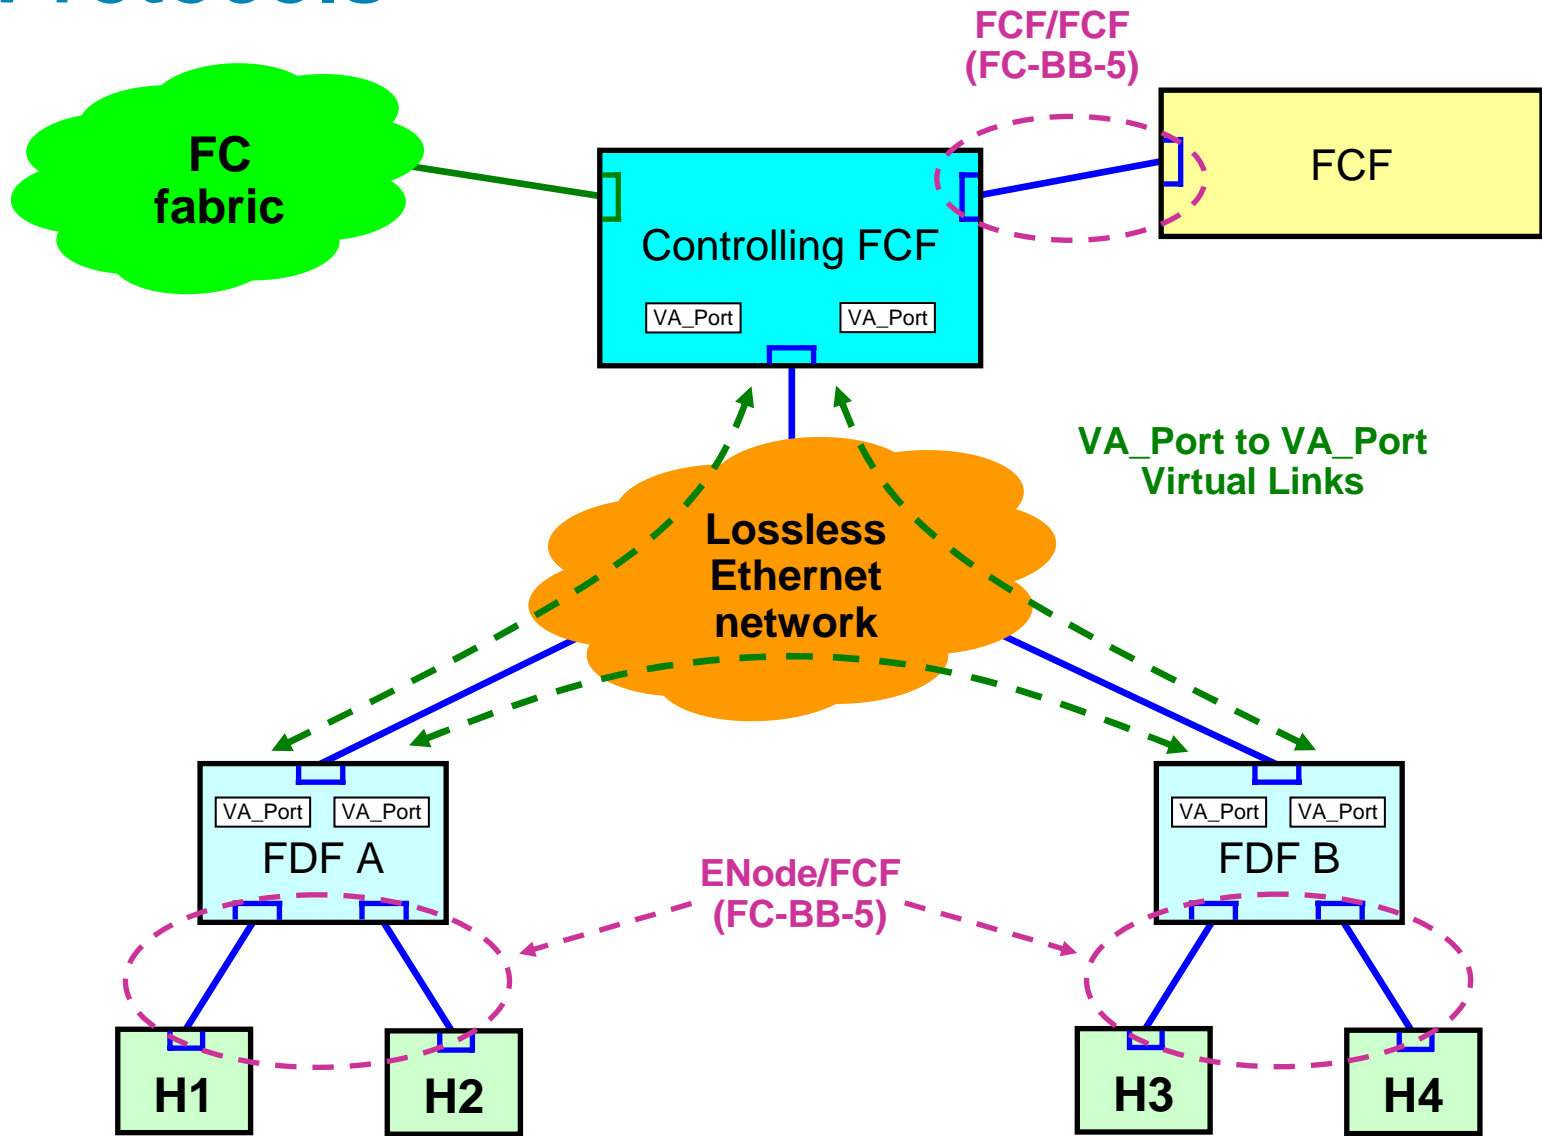
Controlling FCF




Controlling FCF / FDFs



Protocols



VA_Port to VA_Port Virtual Links



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 FDFs to discover possible controlling FCFs**
- **FDFs select the controlling FCF from which they wish to be controlled**

VA_Port to VA_Port Virtual Links

- Upon selecting its controlling FCF, an FDF initiates a FIP ELP Exchange with that FCF with a VA_Port flag set in ELP

An FDF/FCF VA_Port to VA_Port Virtual Link is established

- When a VA_Port to VA_Port Virtual Link is established, an FDF requests to its controlling FCF a range of N_Port_IDs

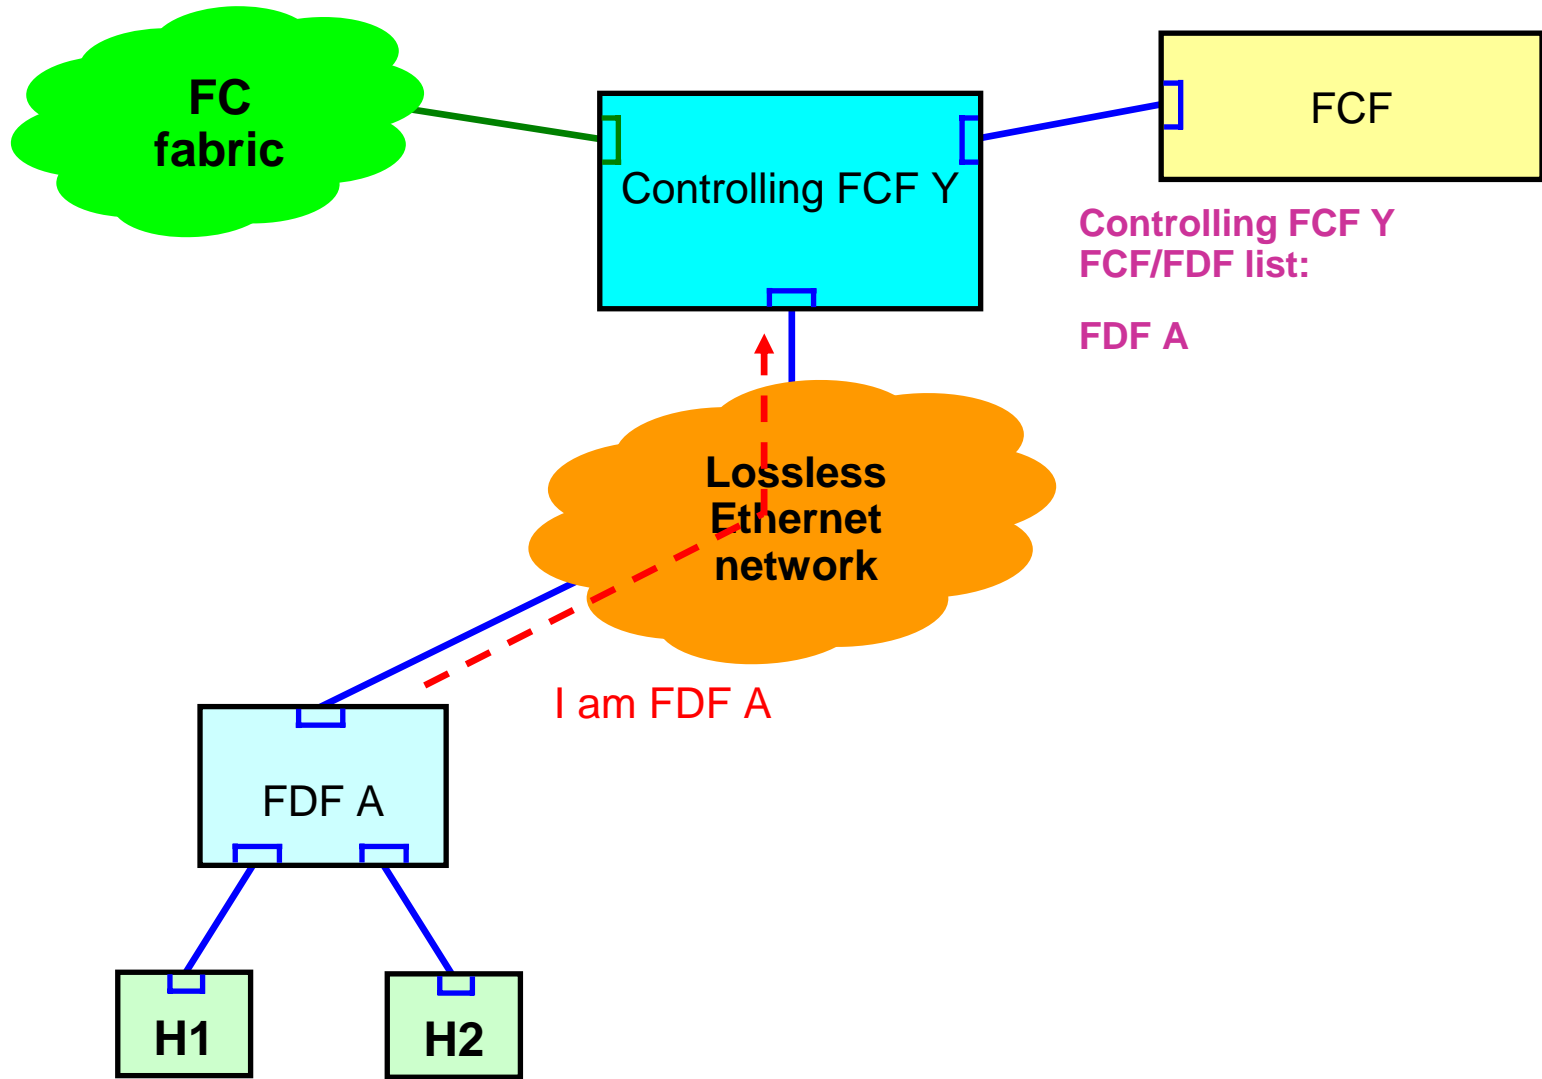
A new unidirectional SW_ILS indicating how many N_Port_IDs

- The controlling FCF allocates that range and provides it to ALL its FDFs on the same lossless Ethernet network

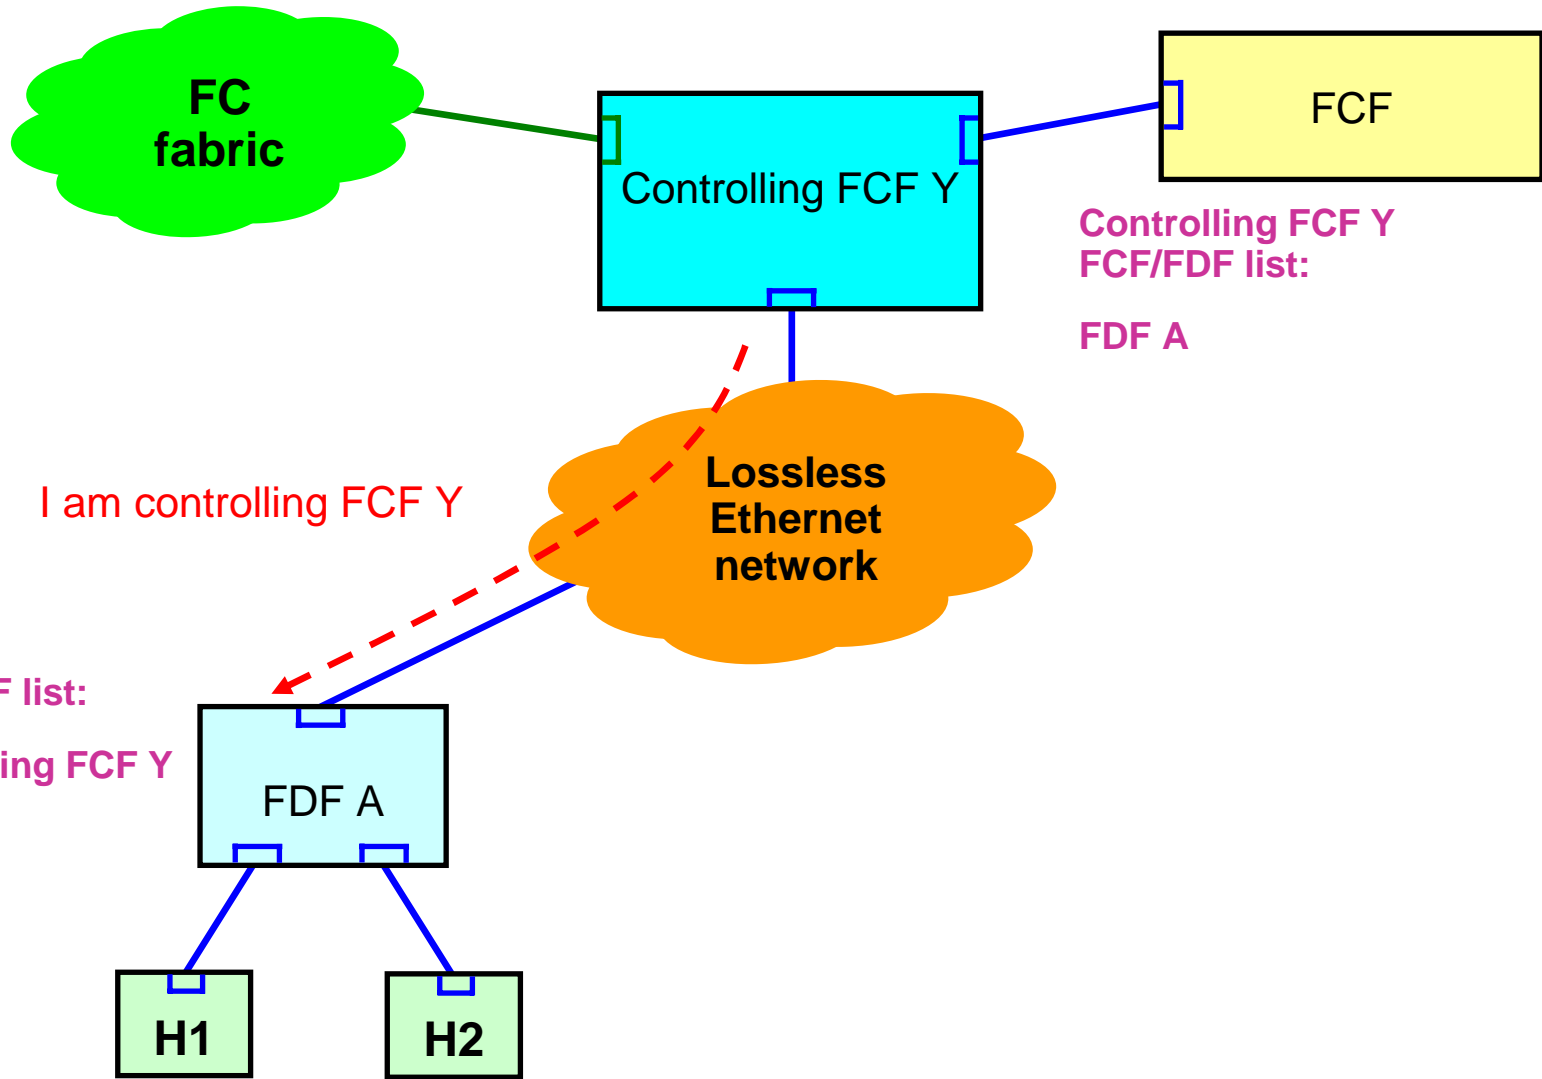
A unidirectional SW_ILS per each VA_Port to VA_Port Virtual Link

- This enables FDFs controlled by the same controlling FCF to discover each other and to establish additional VA_Port to VA_Port Virtual Links among themselves

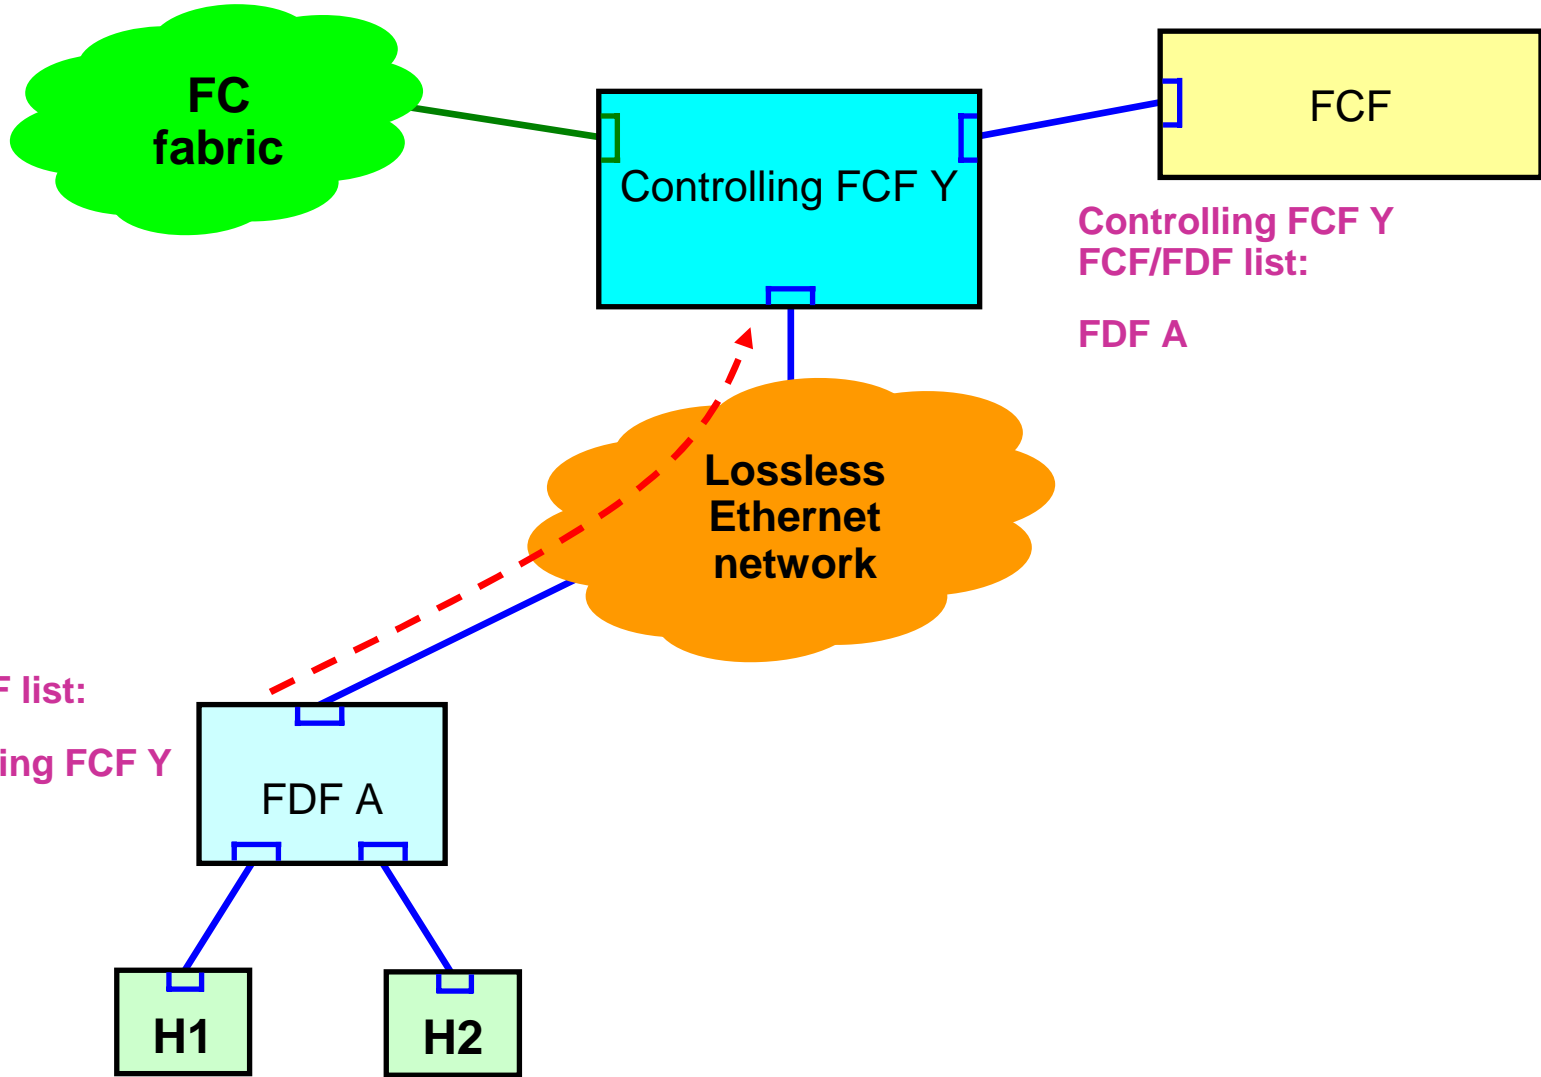
FDF A FIP Solicitation



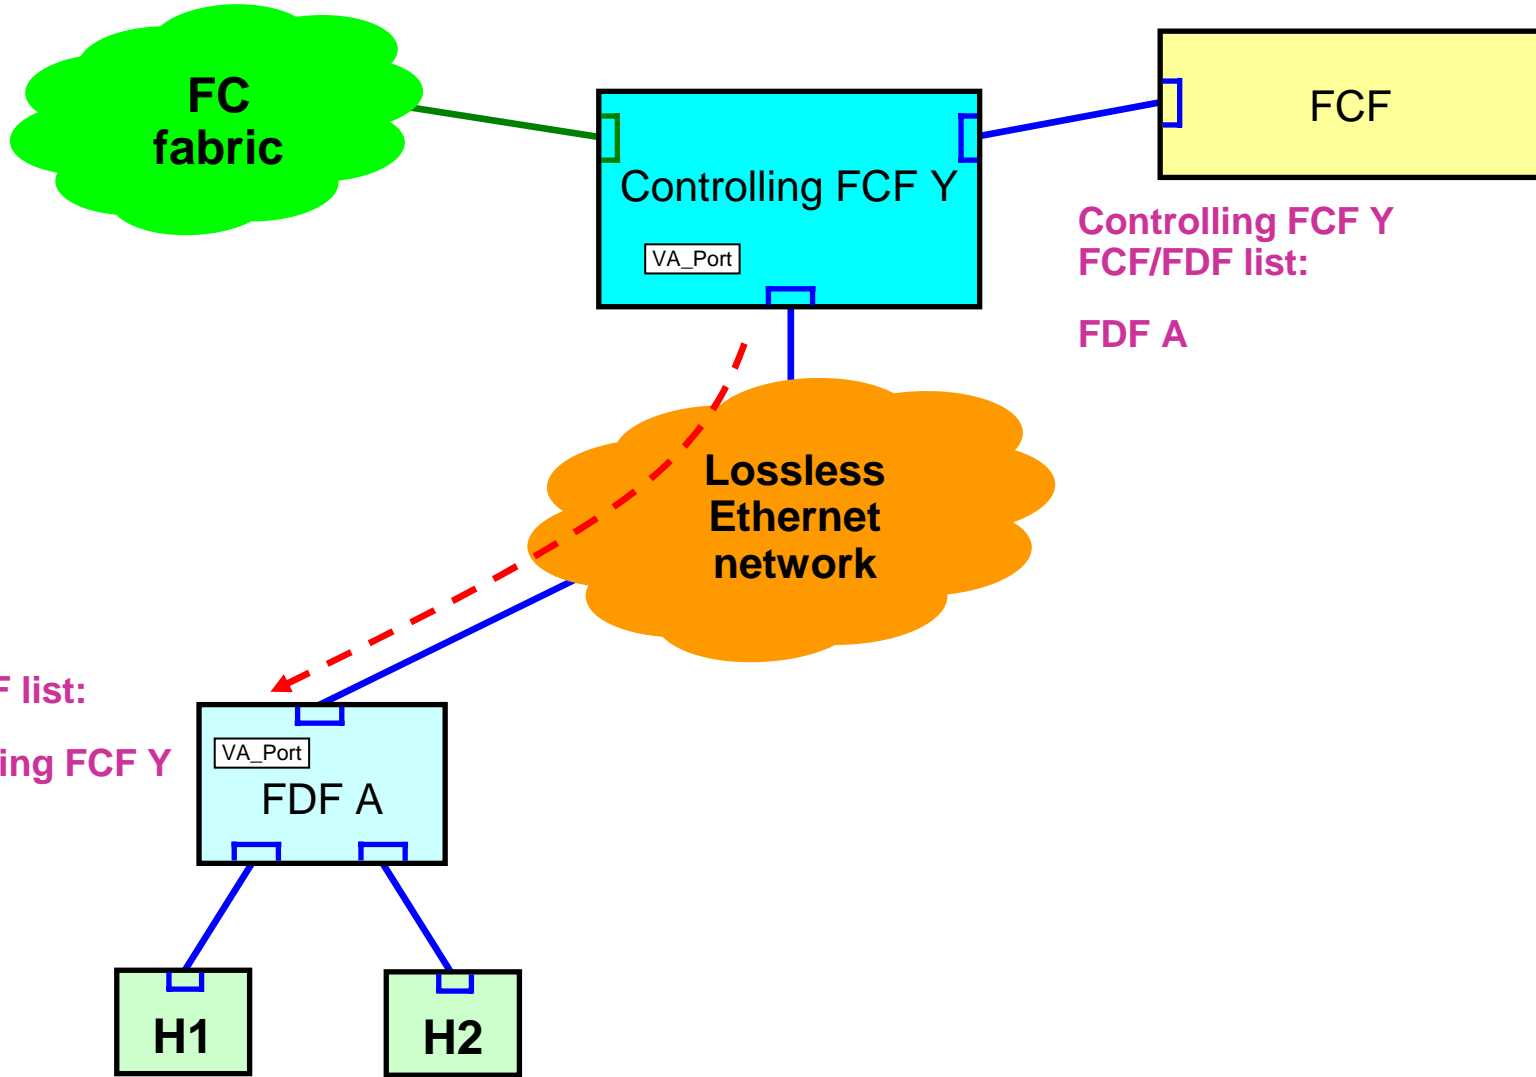
FIP Solicited Advertisement



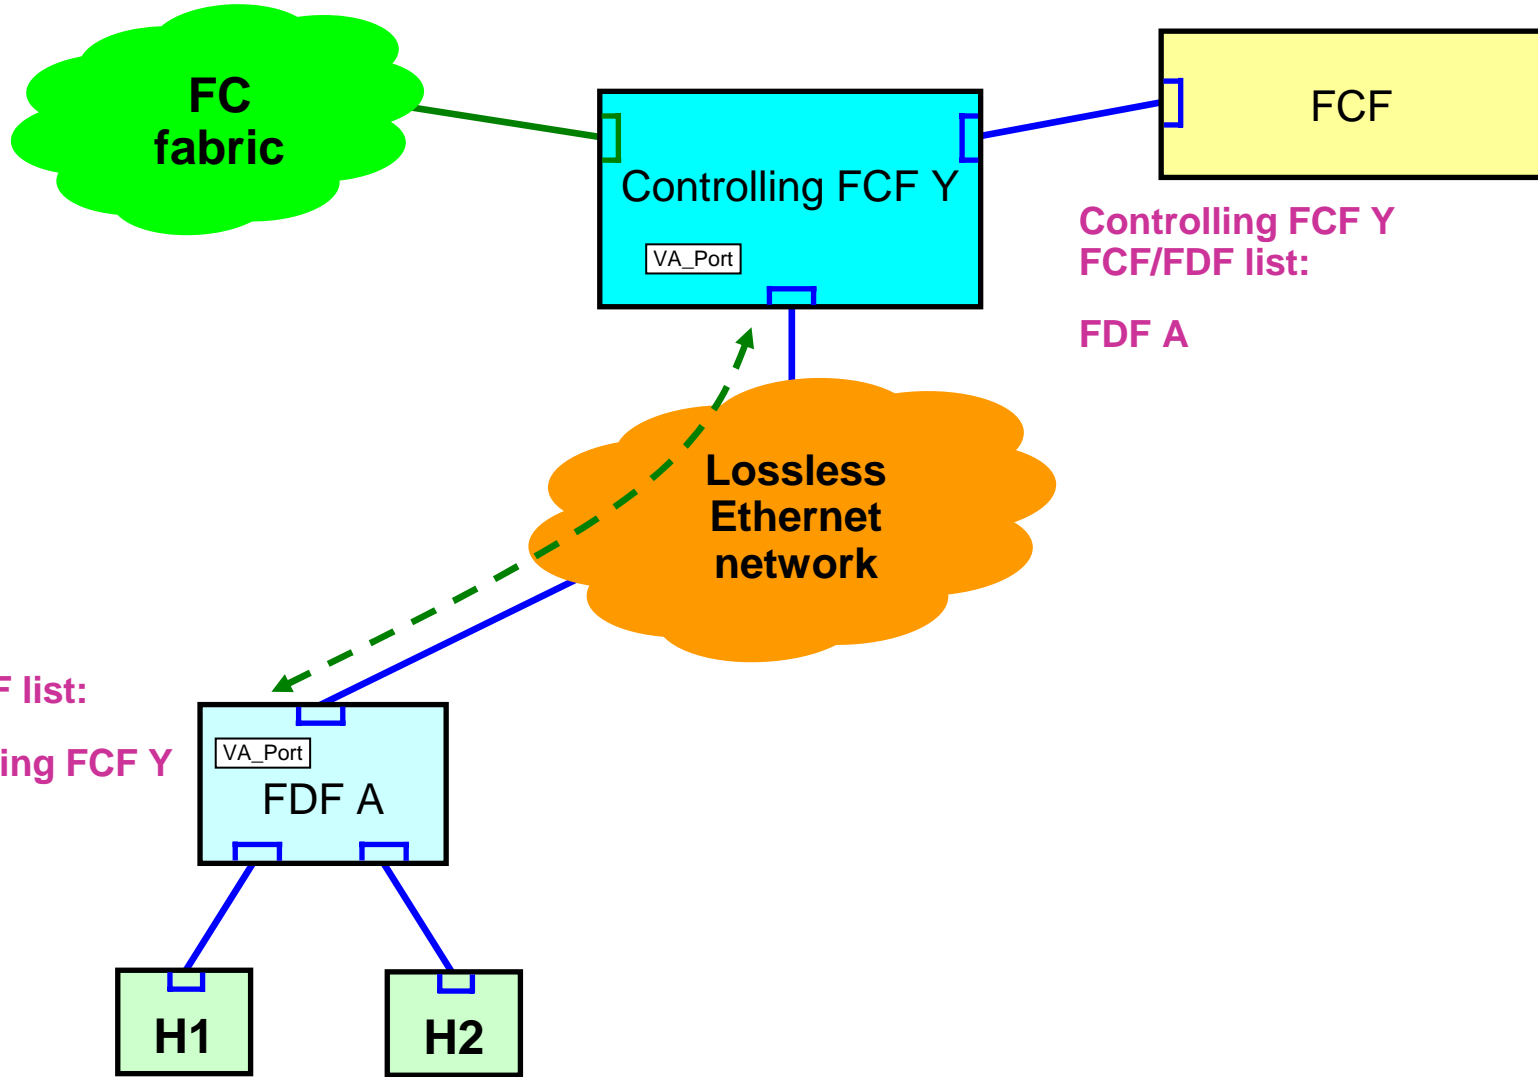
VA_Port FIP ELP Request



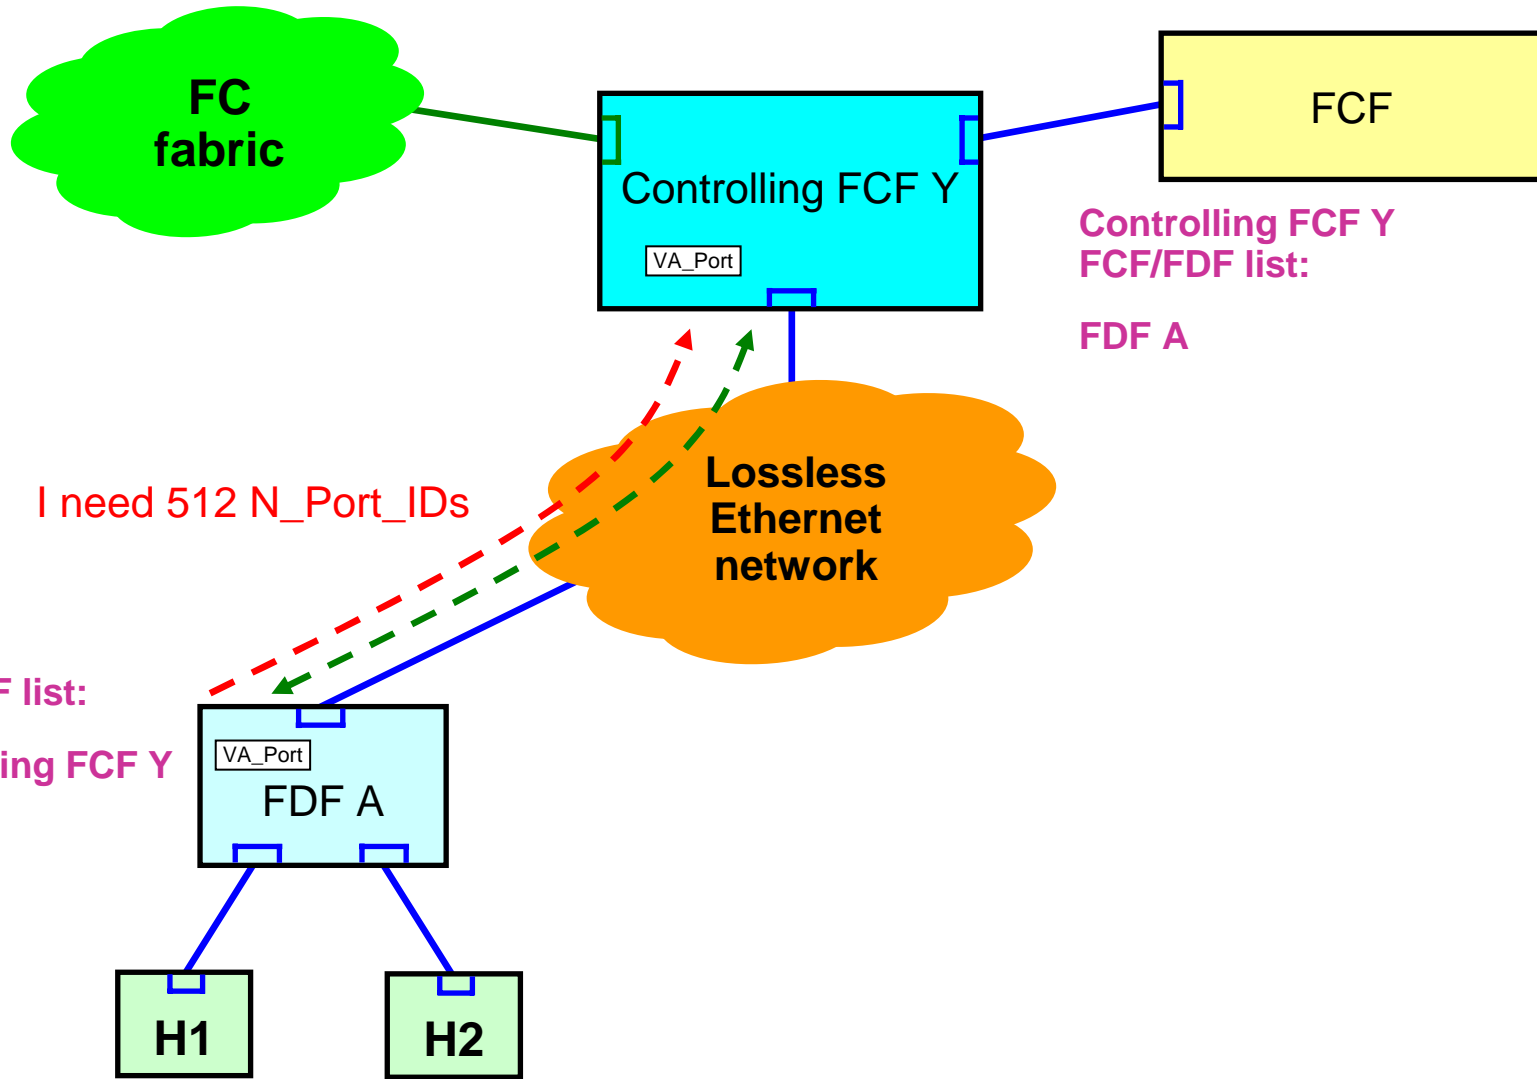
VA_Port FIP ELP SW_ACC



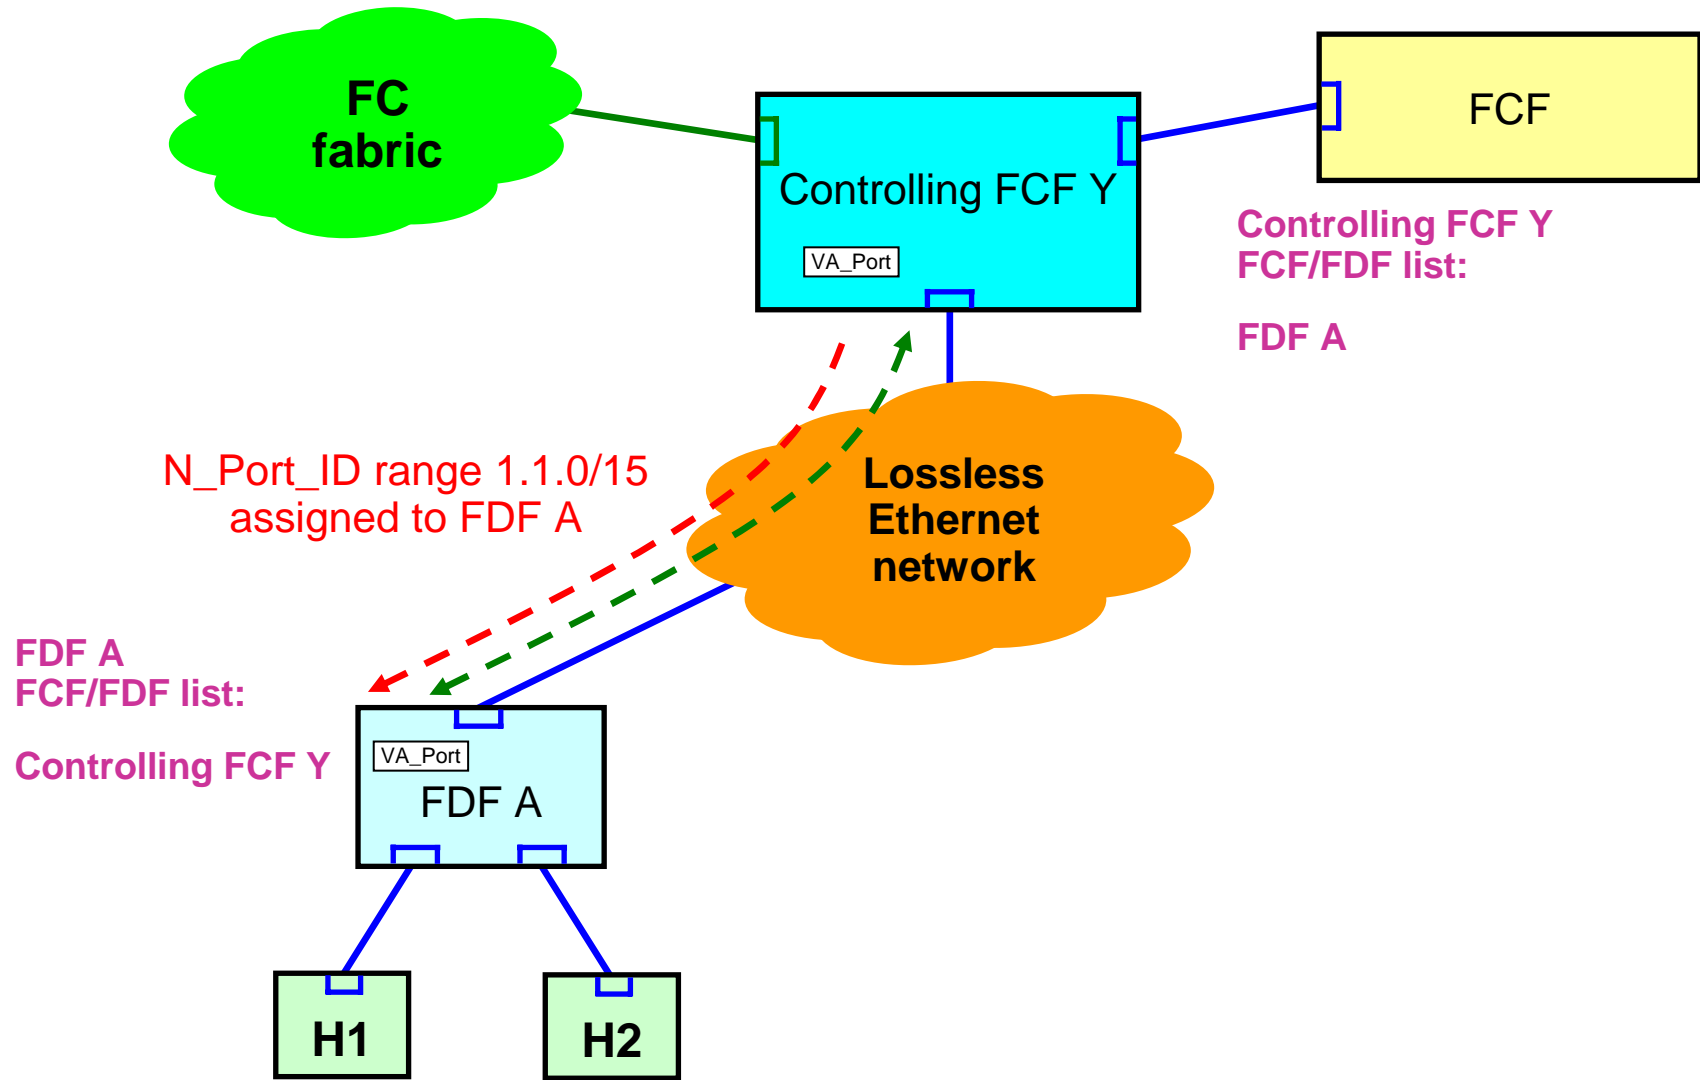
FDF/FCF VA_Port to VA_Port Virtual Link



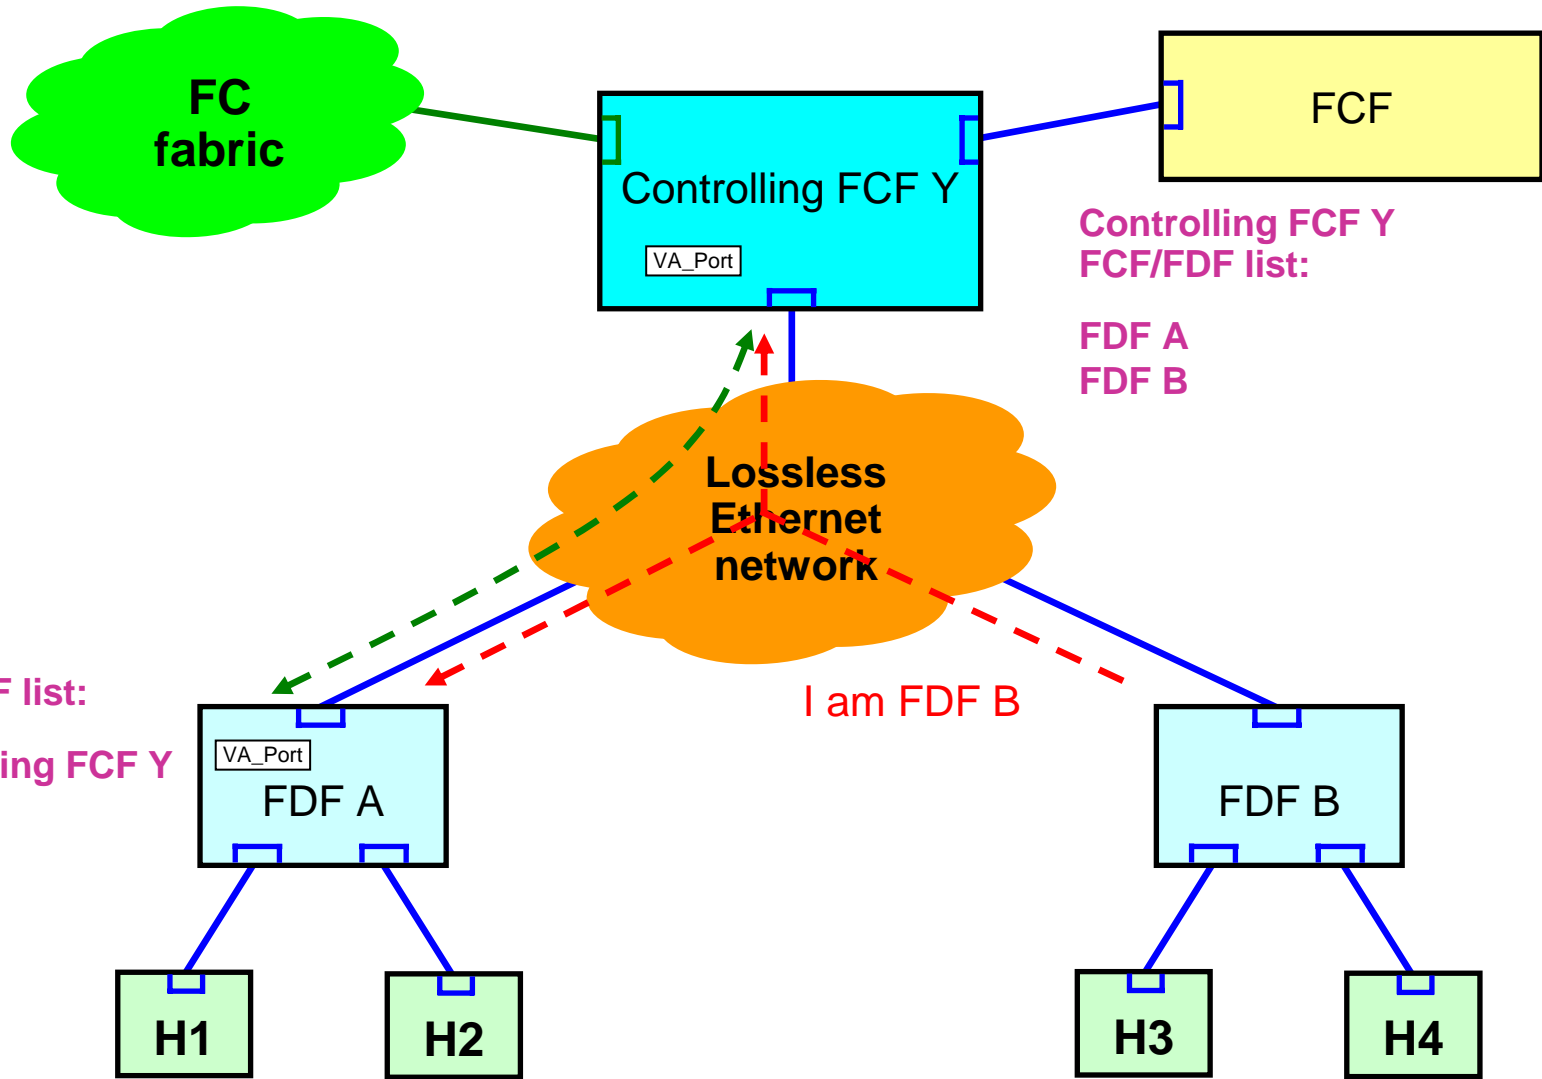
FCoE N_Port_ID Range Request SW_ILS



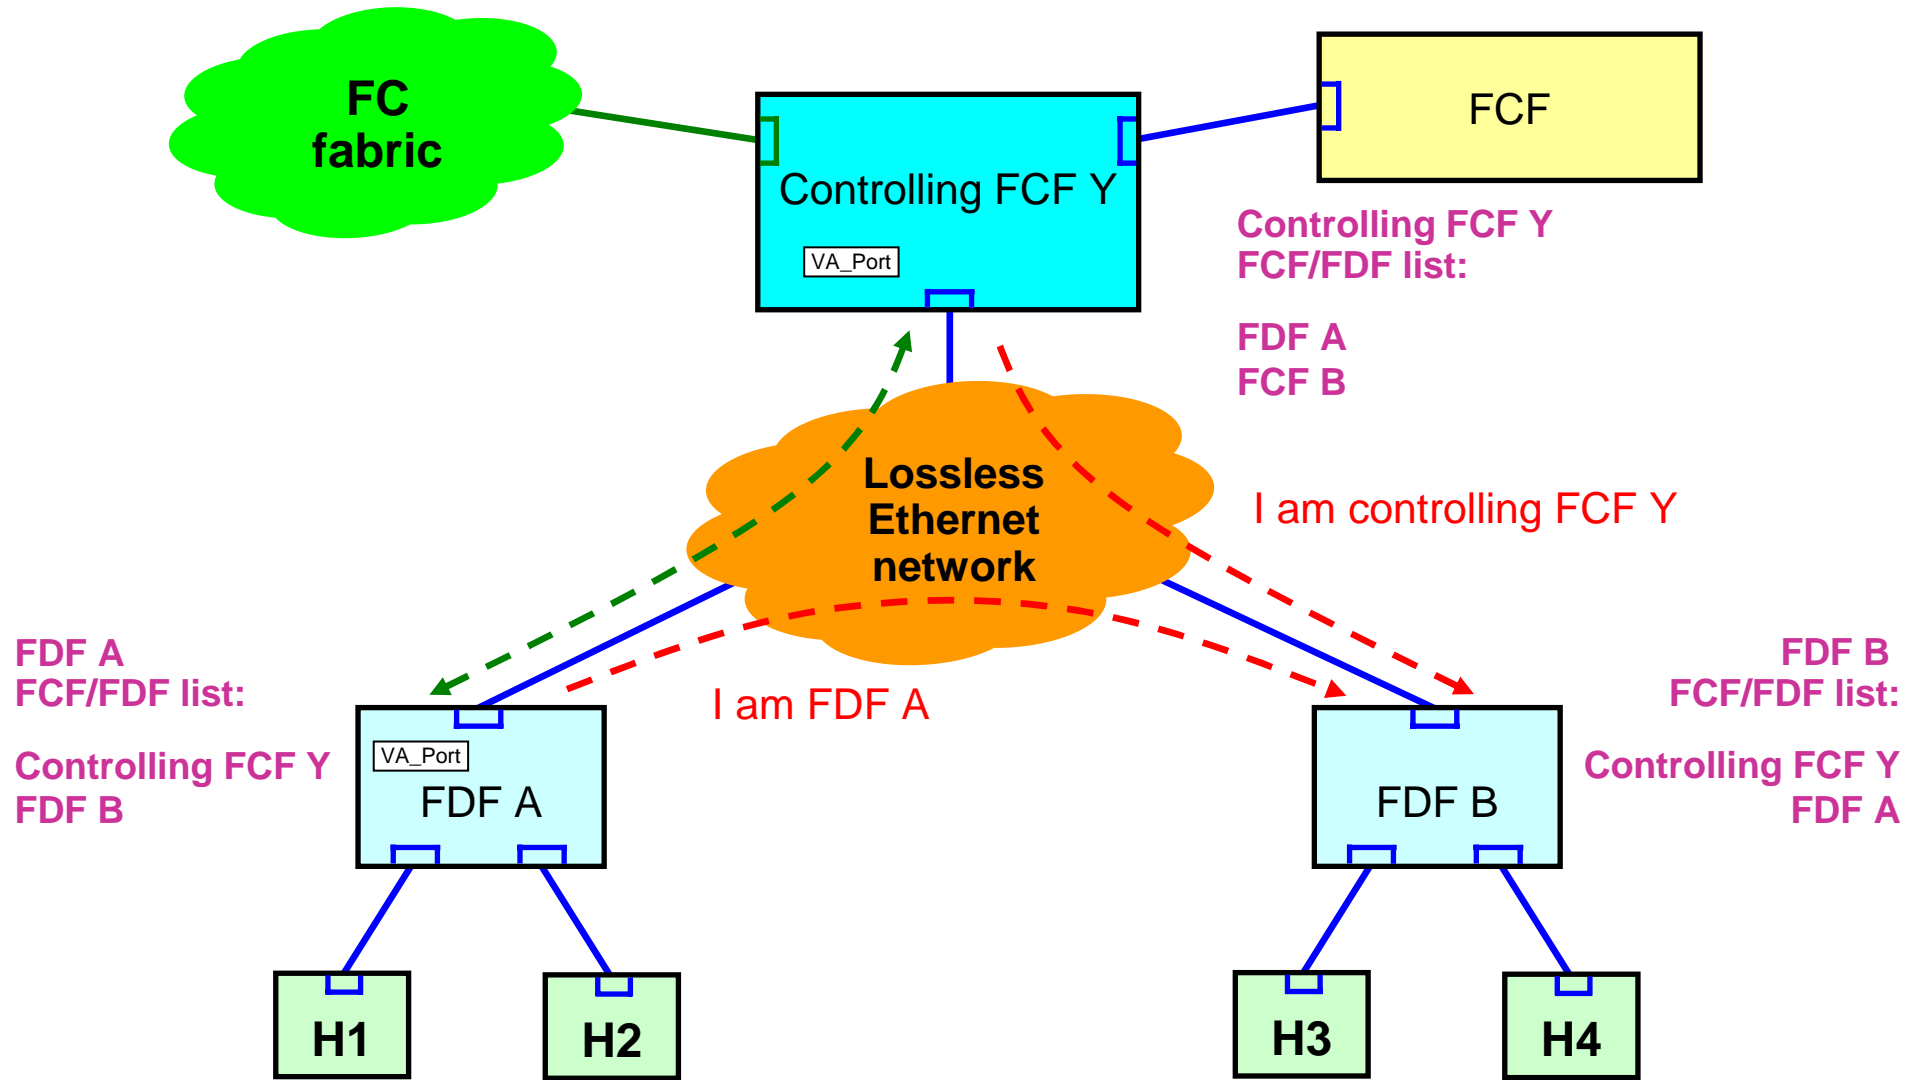
FCoE N_Port_ID Range Assignment SW_ILS



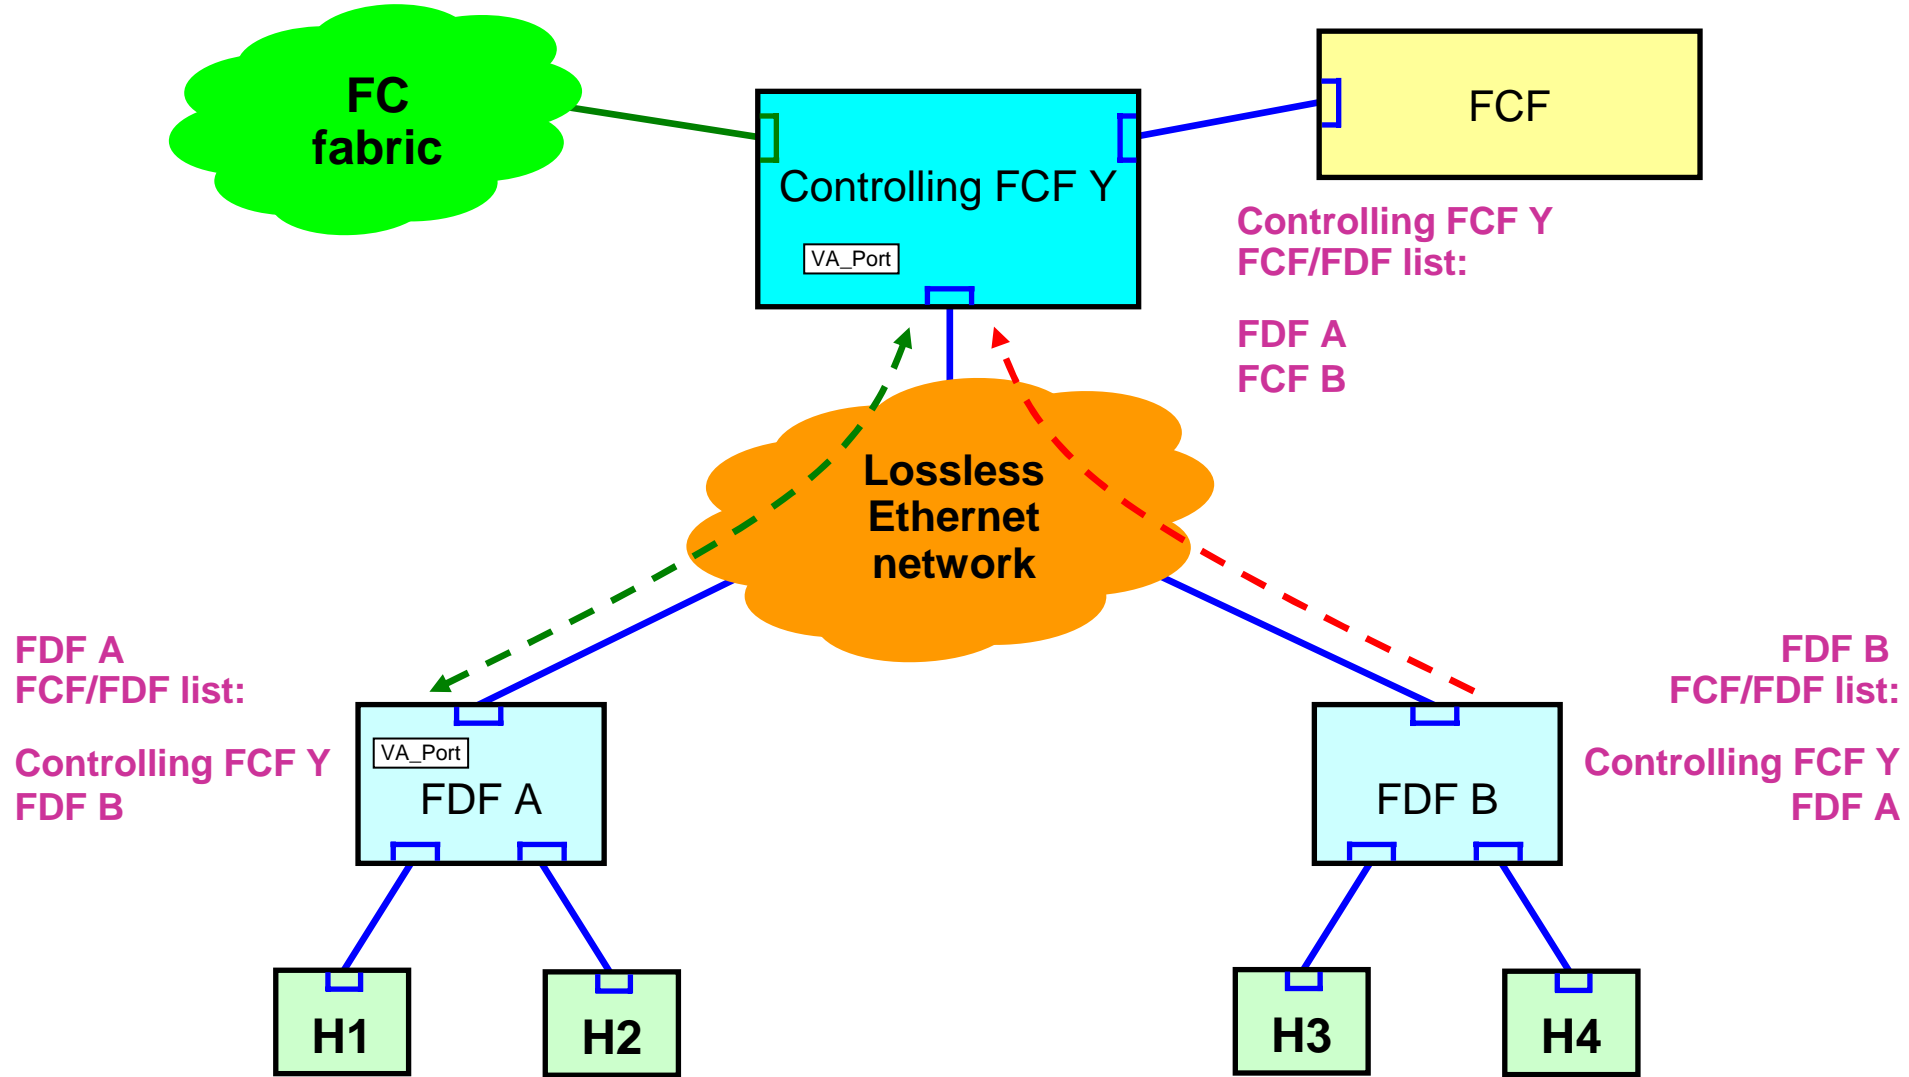
FDF B FIP Solicitation



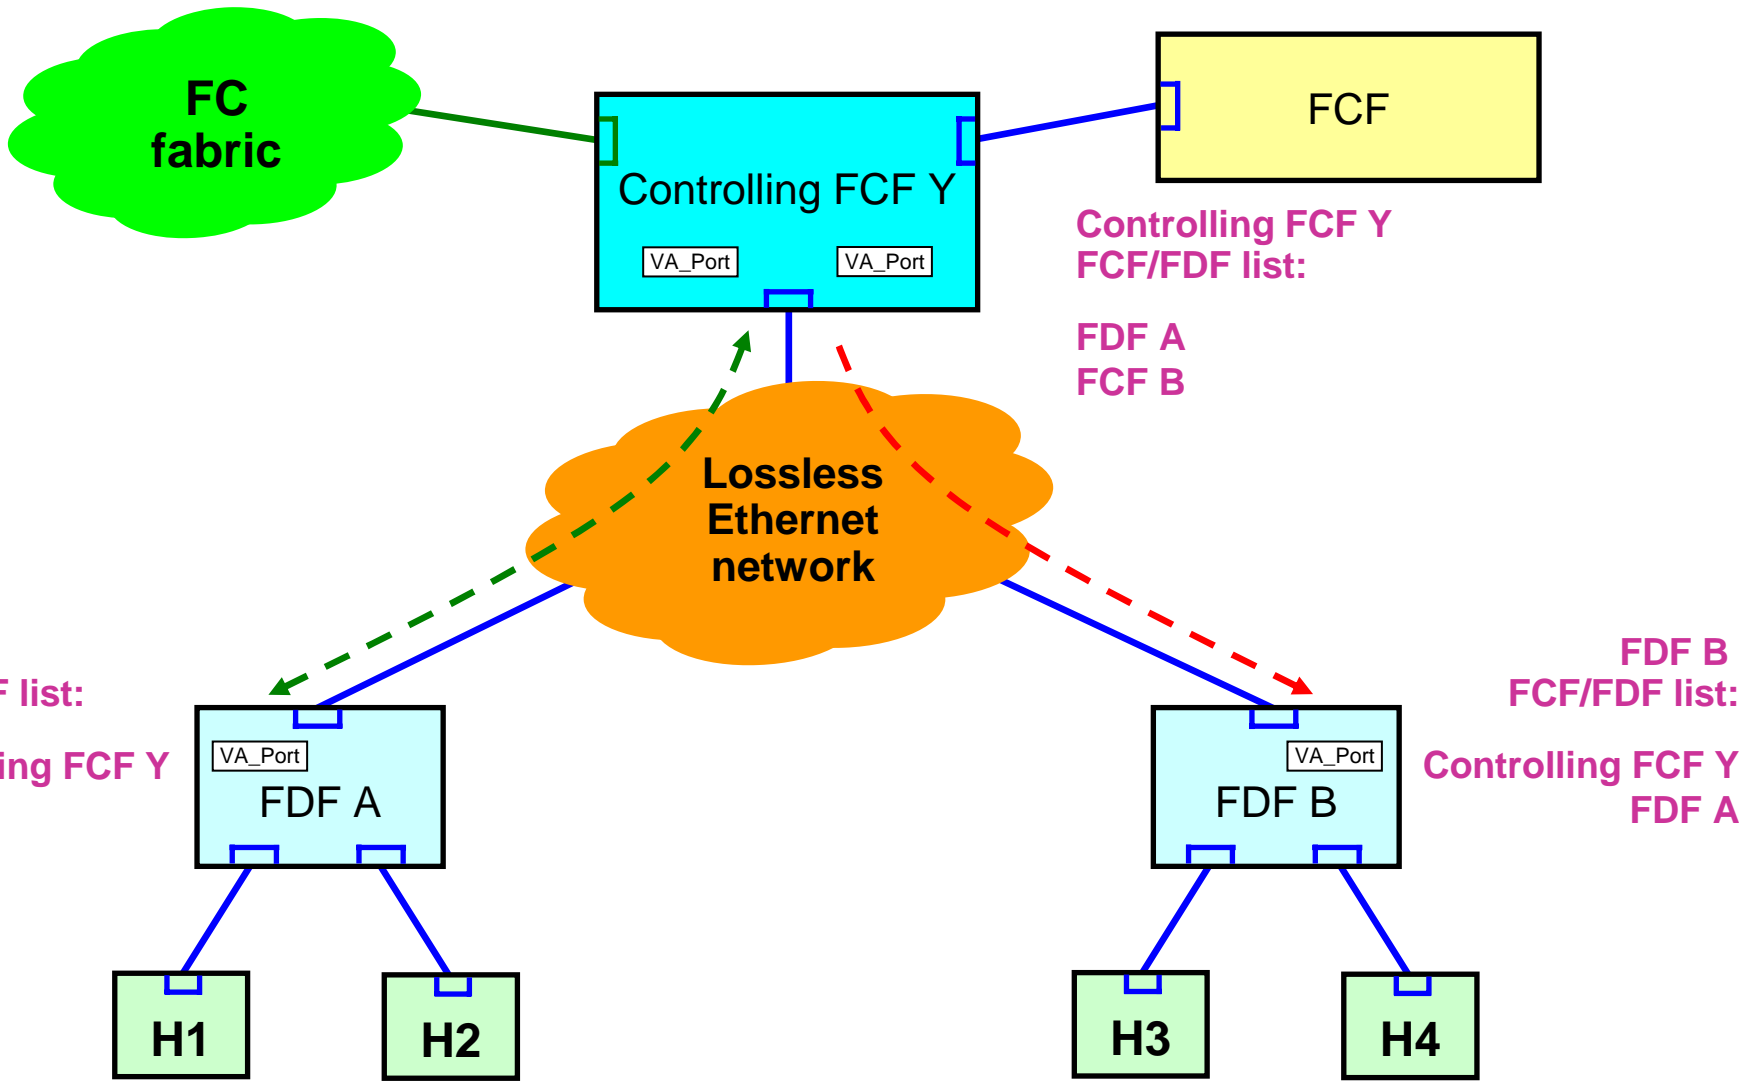
FIP Solicited Advertisements



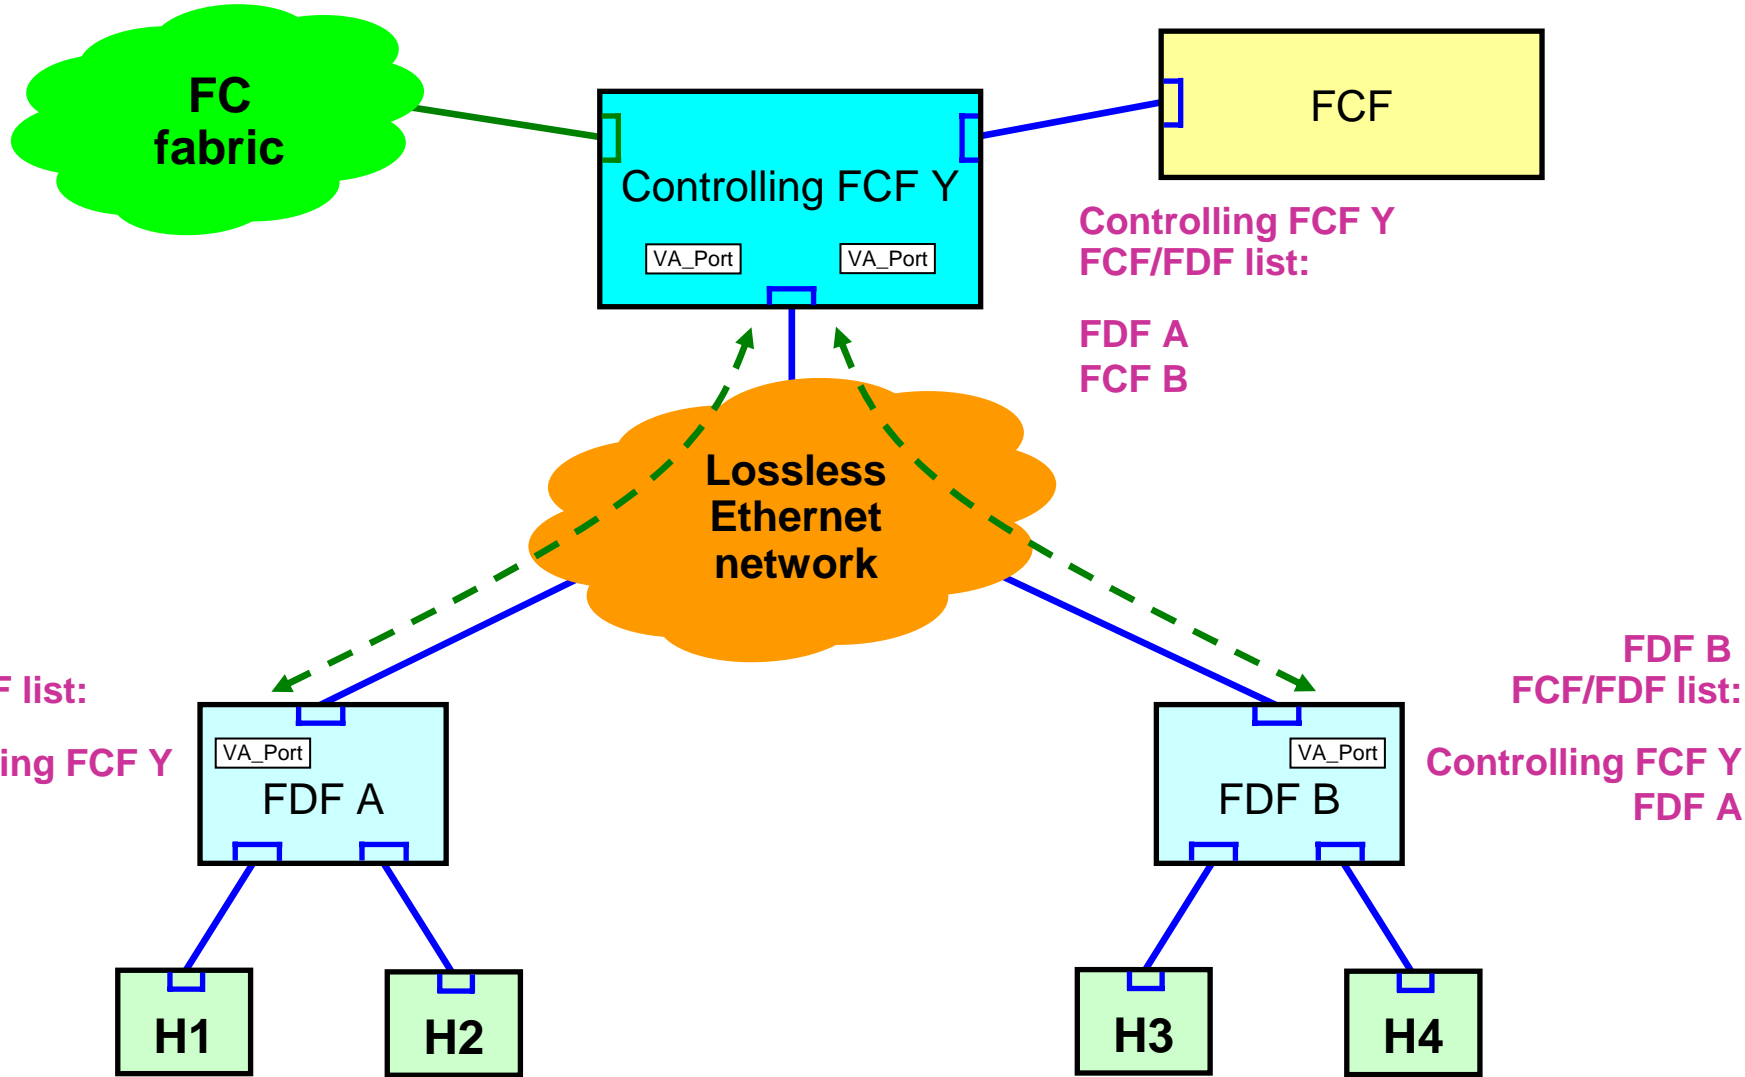
VA_Port FIP ELP Request



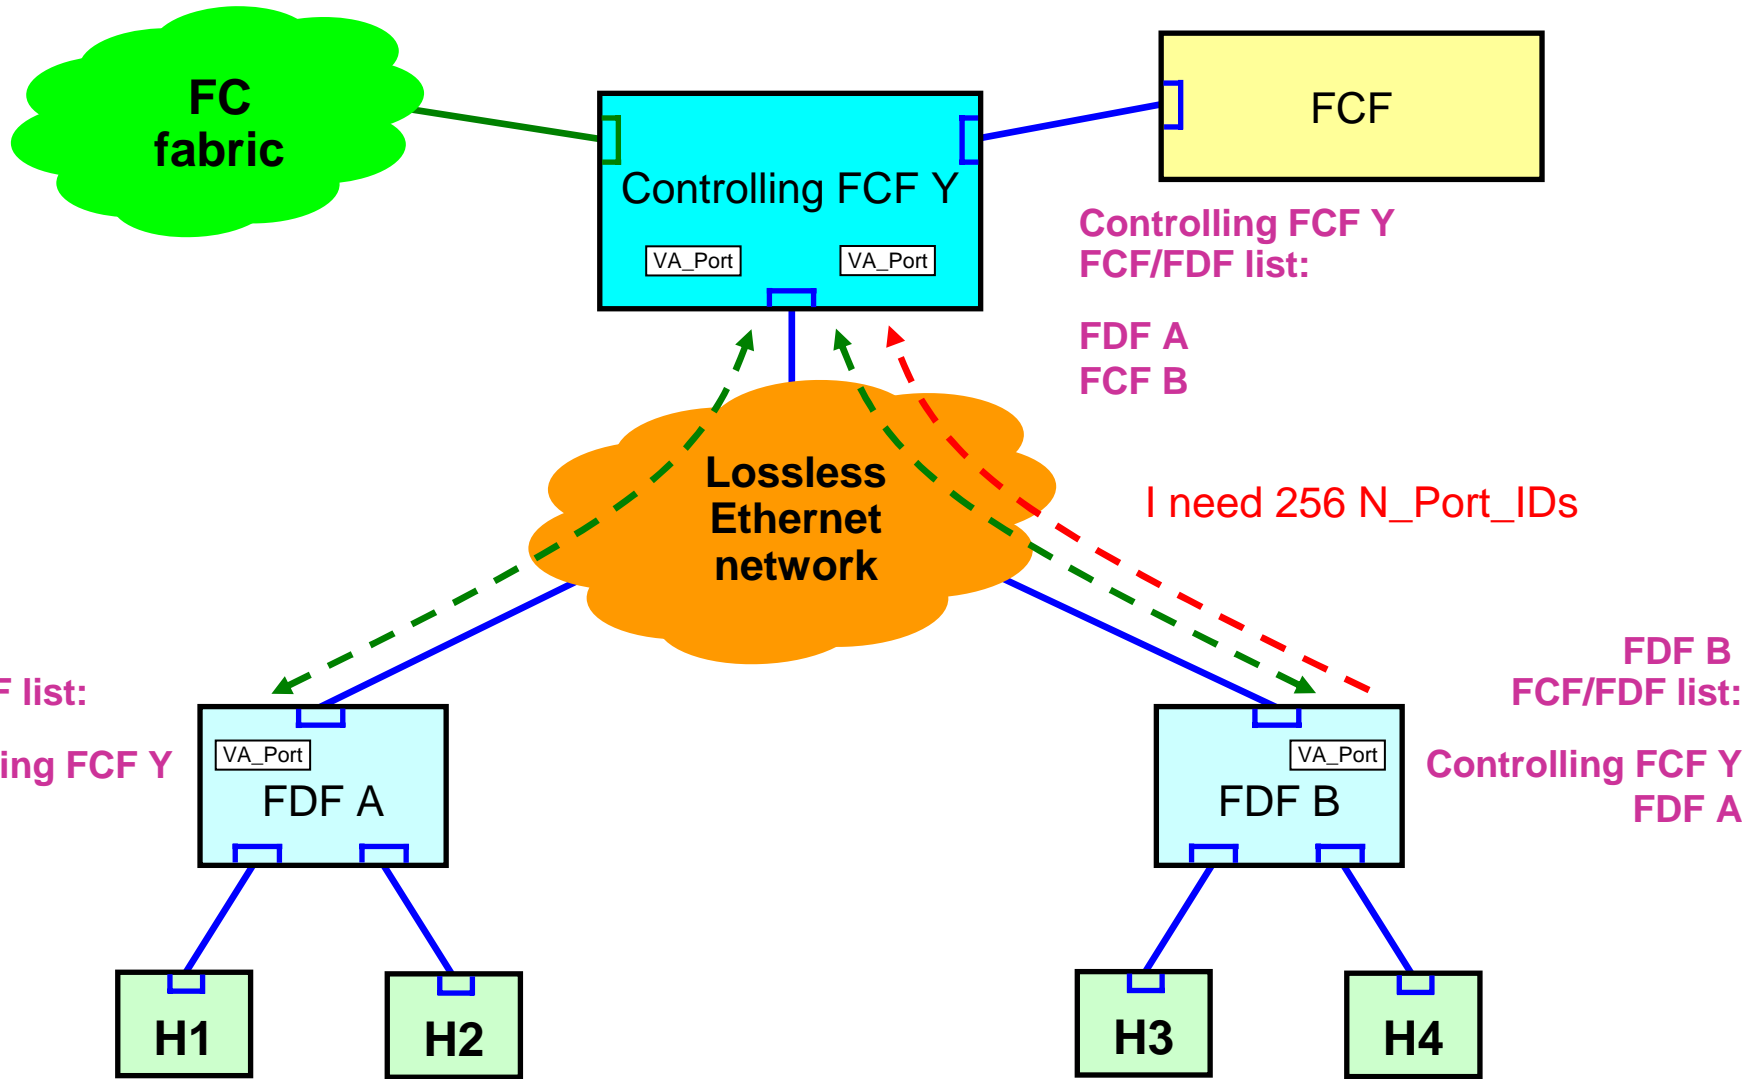
VA_Port FIP ELP SW_ACC



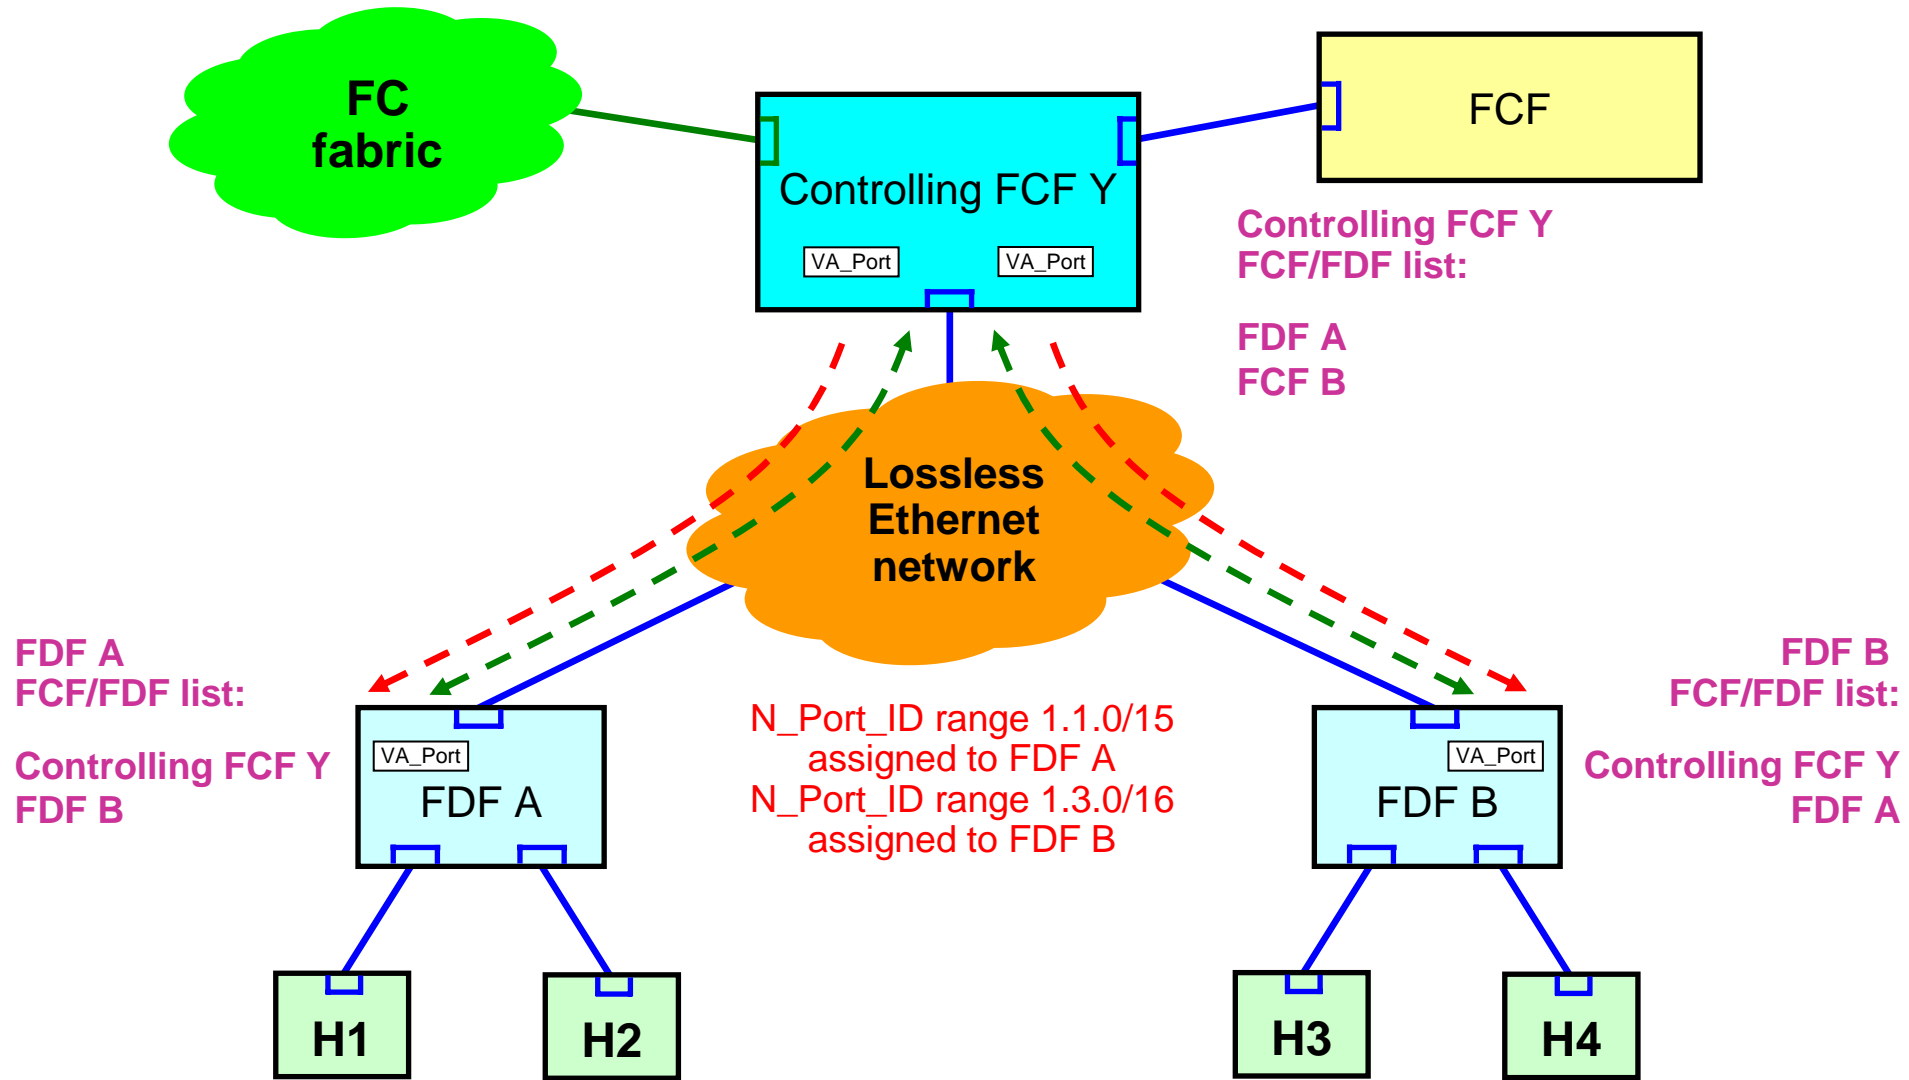
FDF/FCF VA_Port to VA_Port Virtual Link



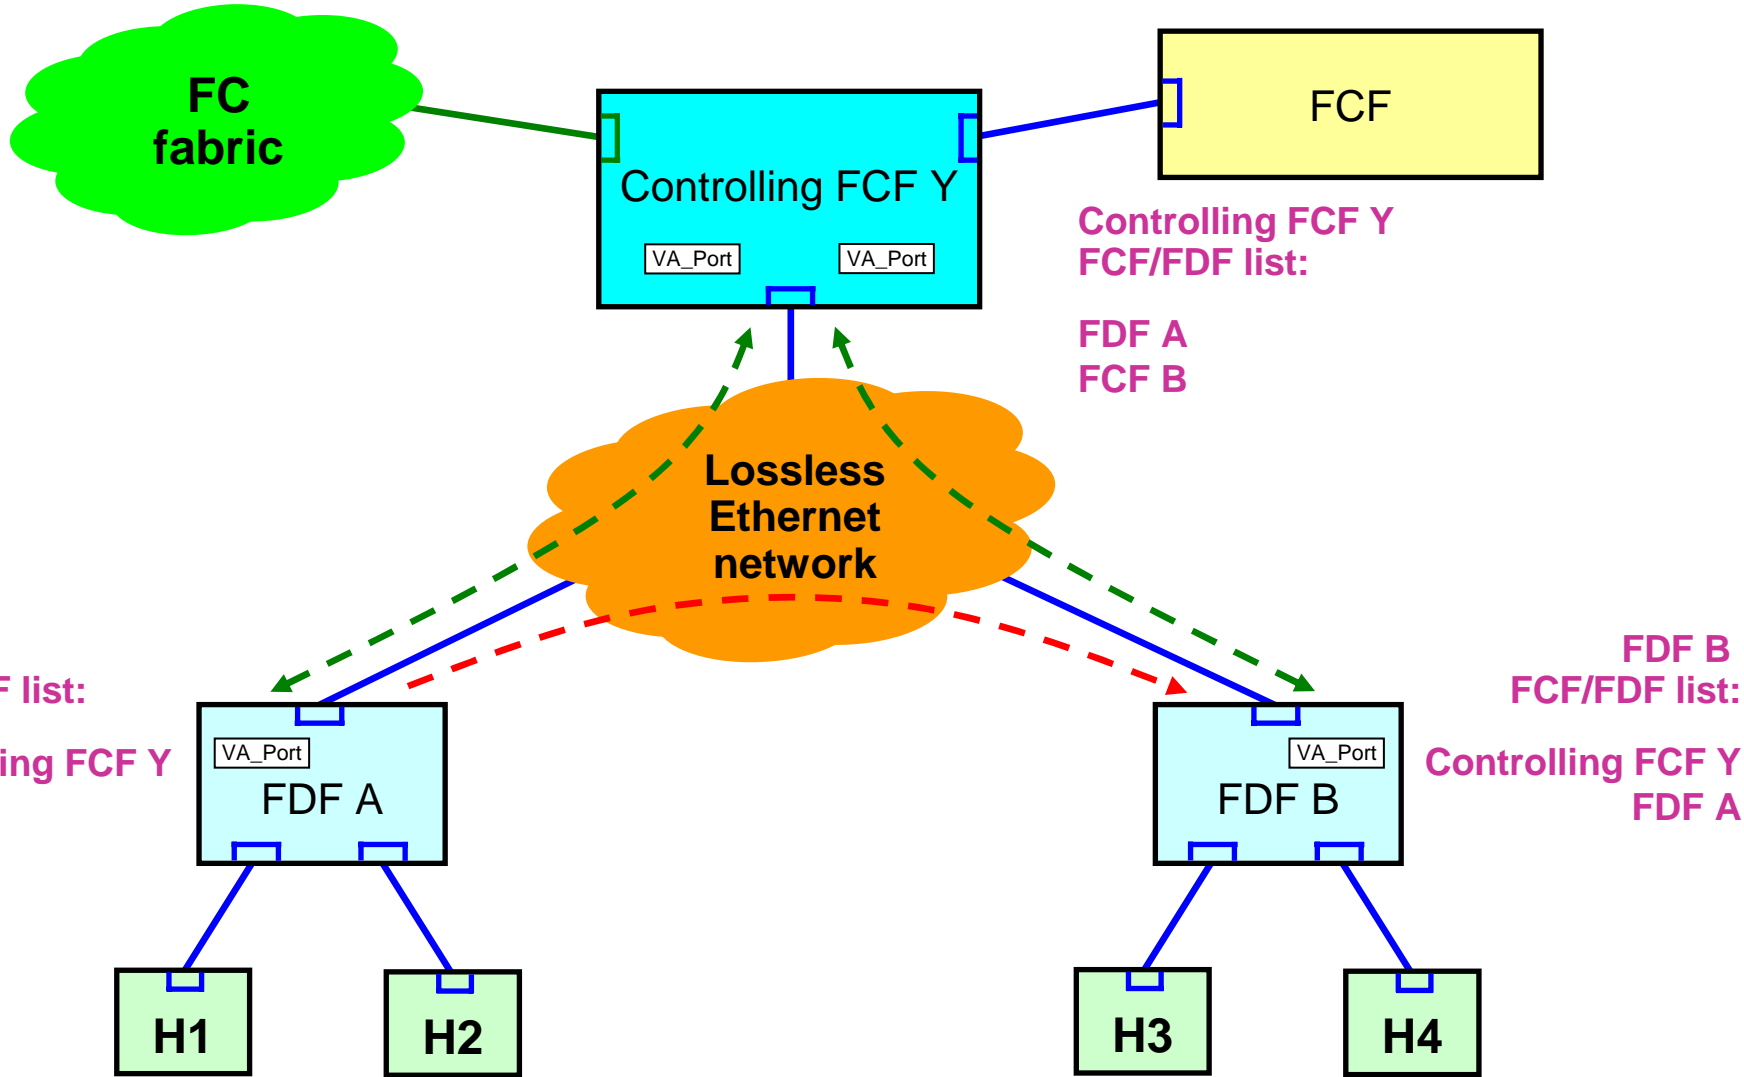
FCoE N_Port_ID Range Request SW_ILS



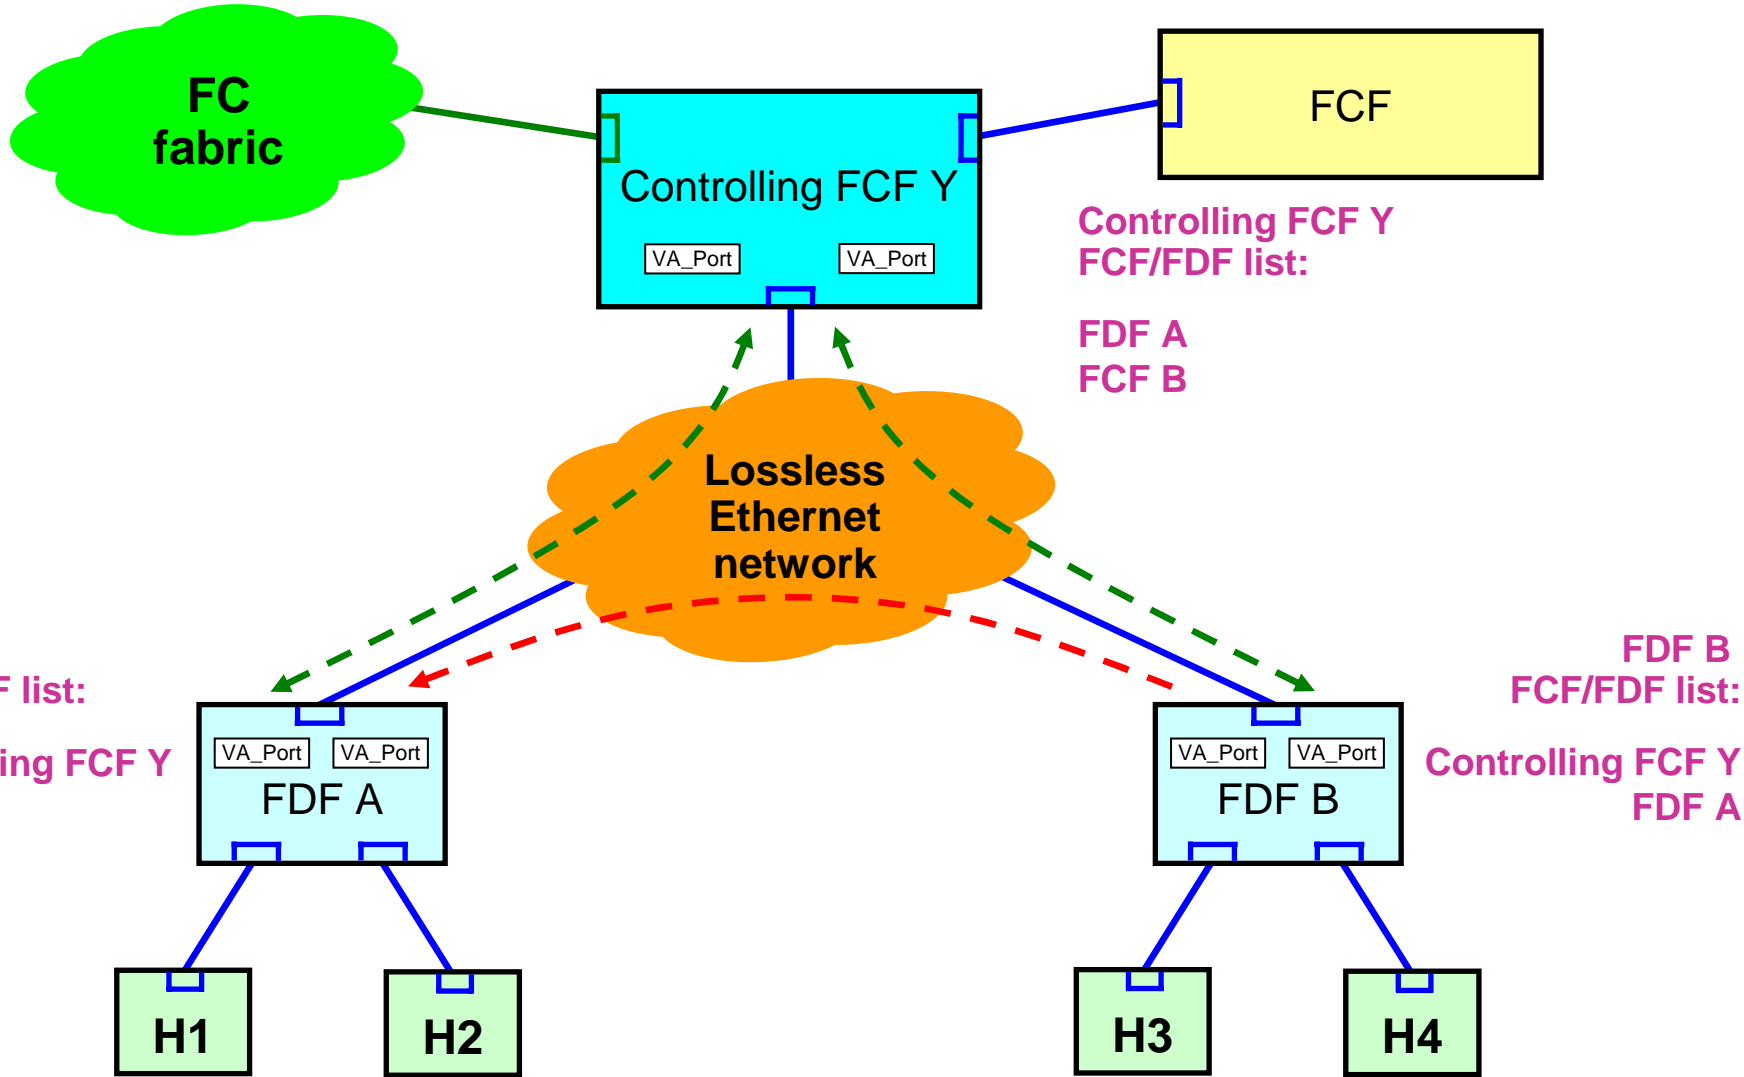
FCoE N_Port_ID Range Assignment SW_ILSs



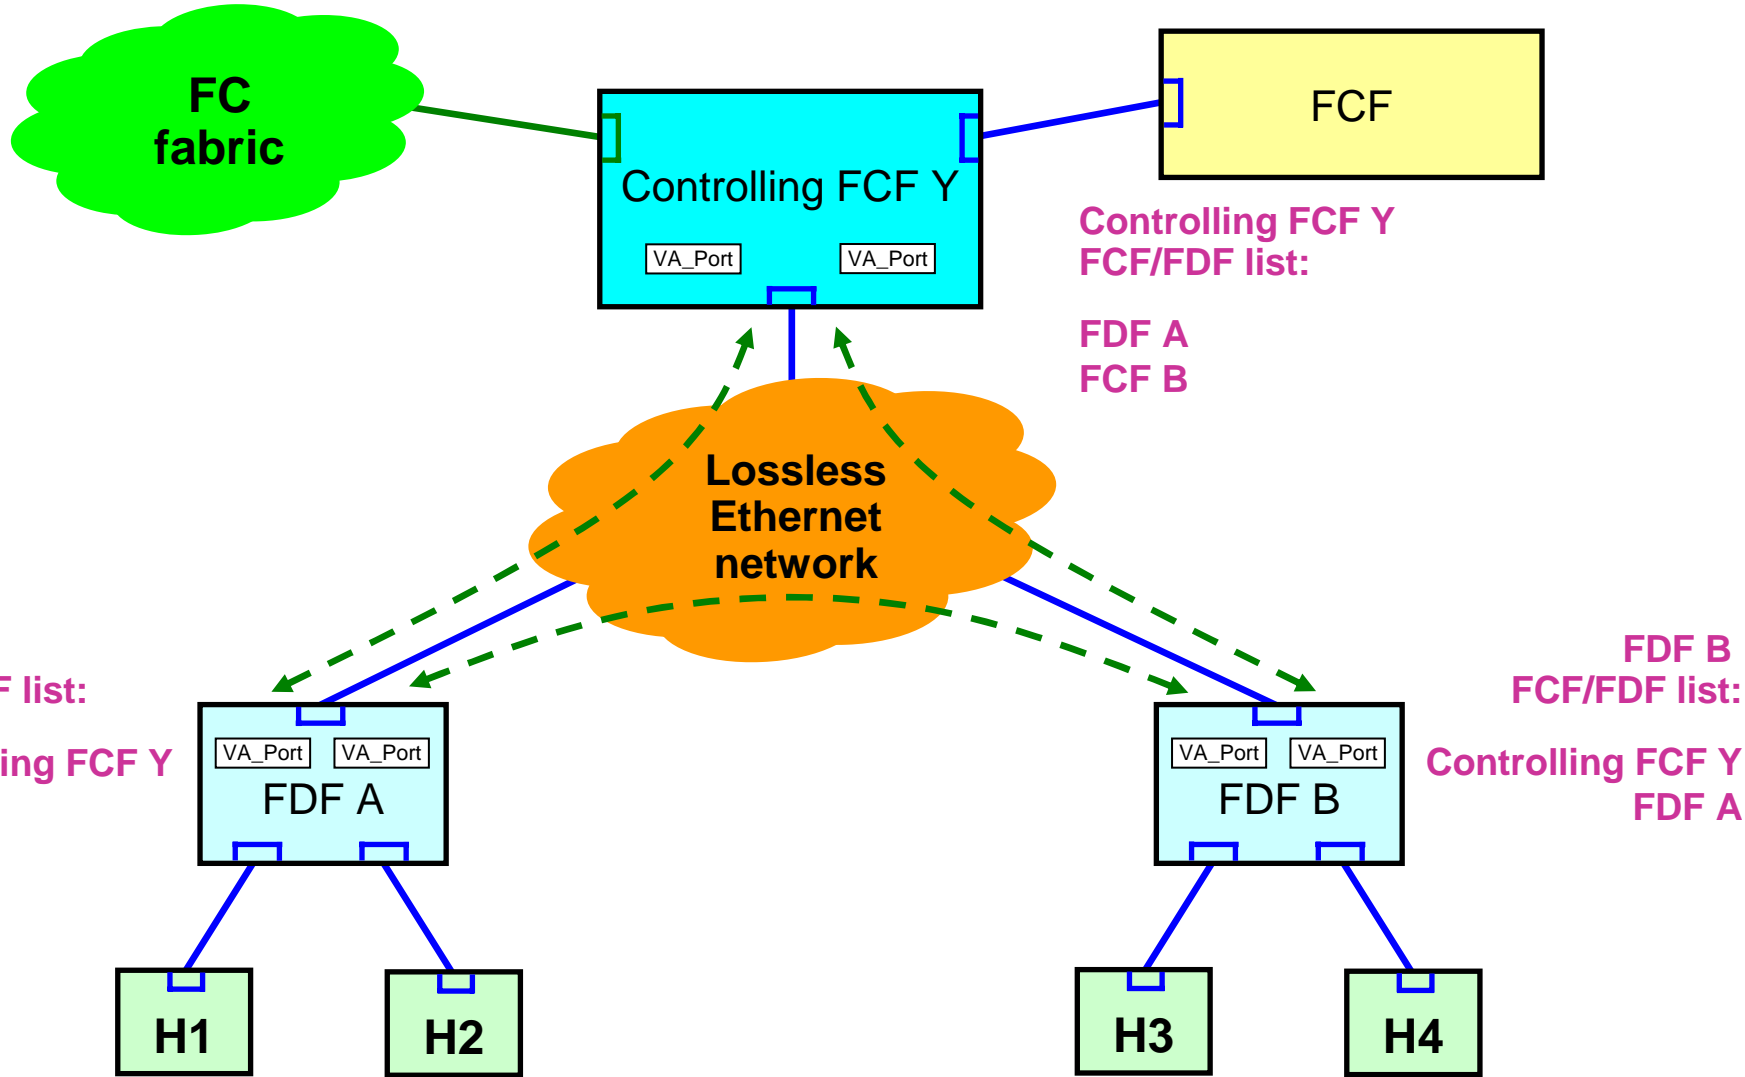
VA_Port FIP ELP Request



VA_Port FIP ELP SW_ACC



FDF/FDF VA_Port to VA_Port Virtual Link



Forwarding Tables

- **The N_Port_ID ranges allocation provided by the controlling FCF to all its FDFs enable the controlling FCF and all FDFs to set up proper forwarding tables for these N_Port_ID ranges**
- **Over time, if an FDF exhausts its N_Port_ID range allocation, it may request additional N_Port_ID ranges**
- **The additional N_Port_ID ranges are communicated by the controlling FCF to all its FDFs, to keep the forwarding tables updated**
- **FDFs rely on the controlling FCF to reach other Domain_IDs**
i.e., only the controlling FCF can be connected to other Domain_IDs (other FCFs or native FC Fabric)

FDFs cannot

Considerations on Forwarding

- **Why FDFs cannot be directly connected to other Domain_IDs?**
- **Because this 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

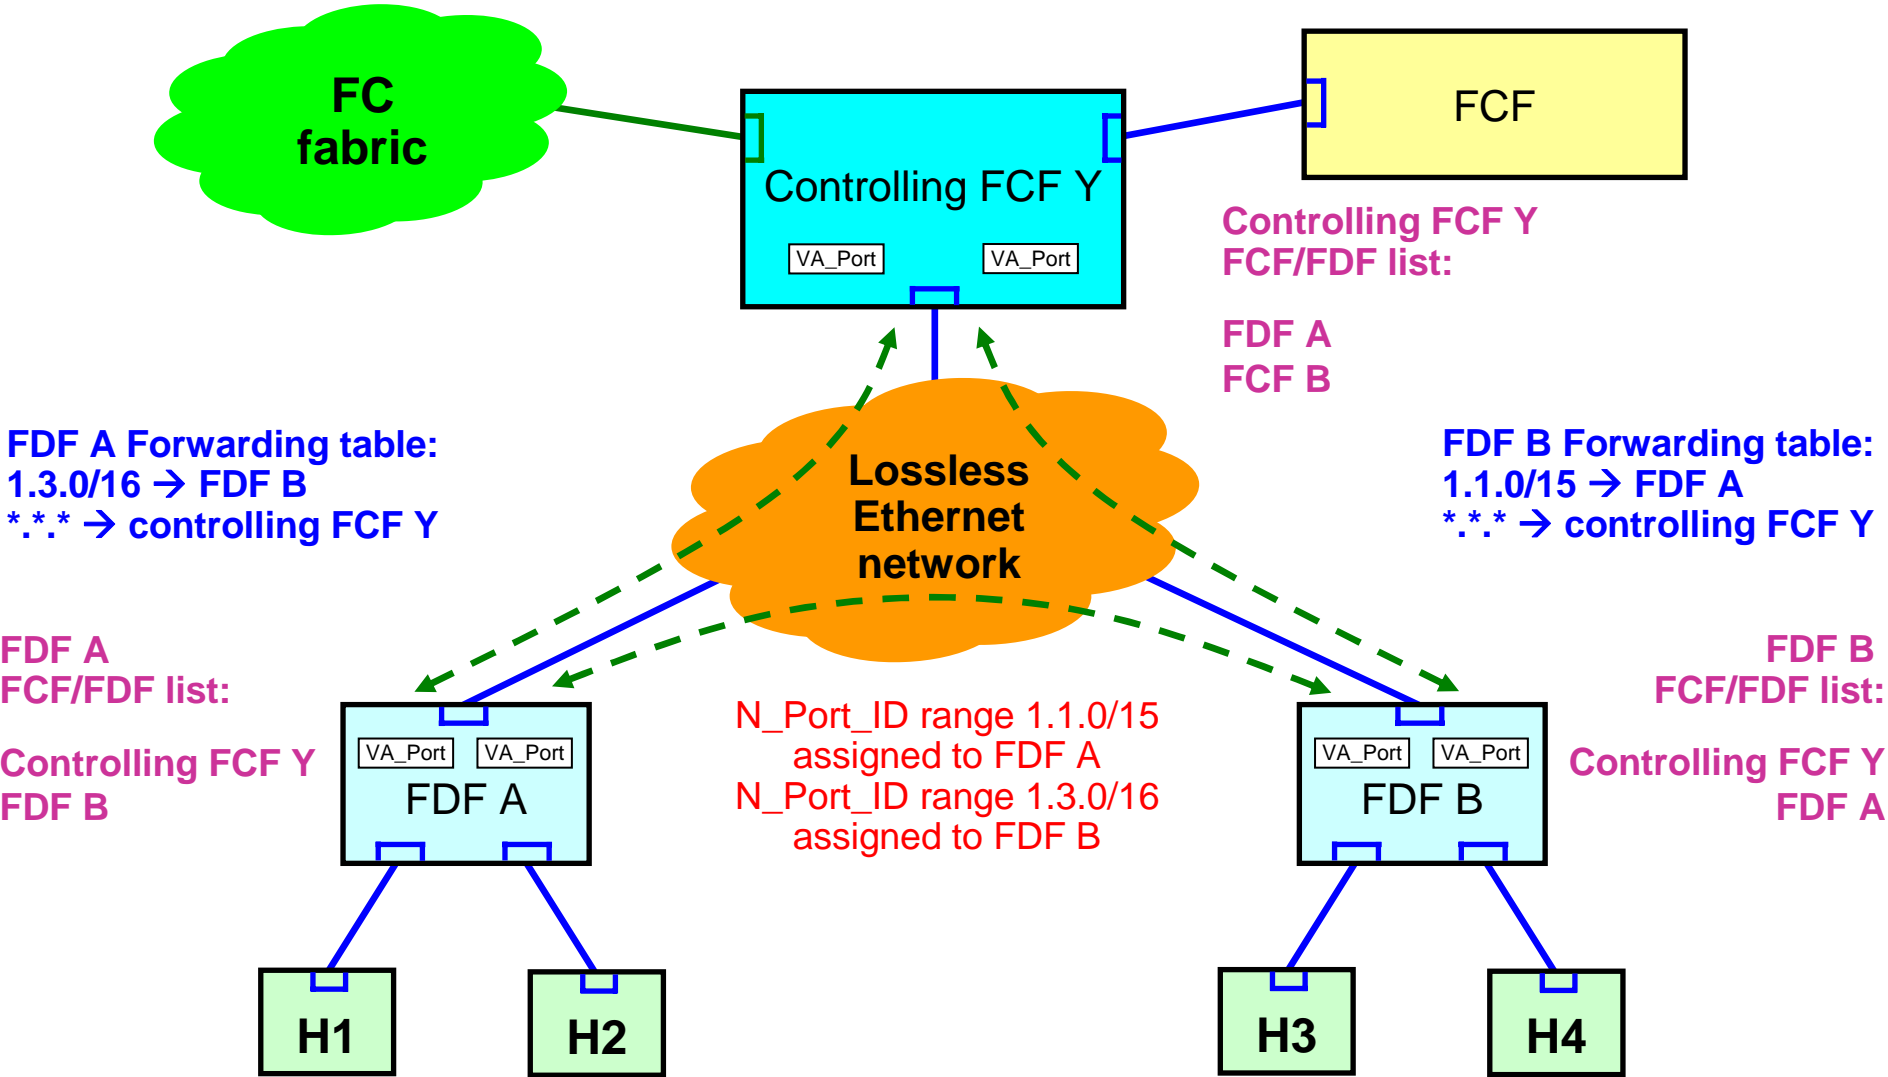
An FCF of FC switch would need to know over which port an N_Port_ID range can be reached at the lowest cost

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

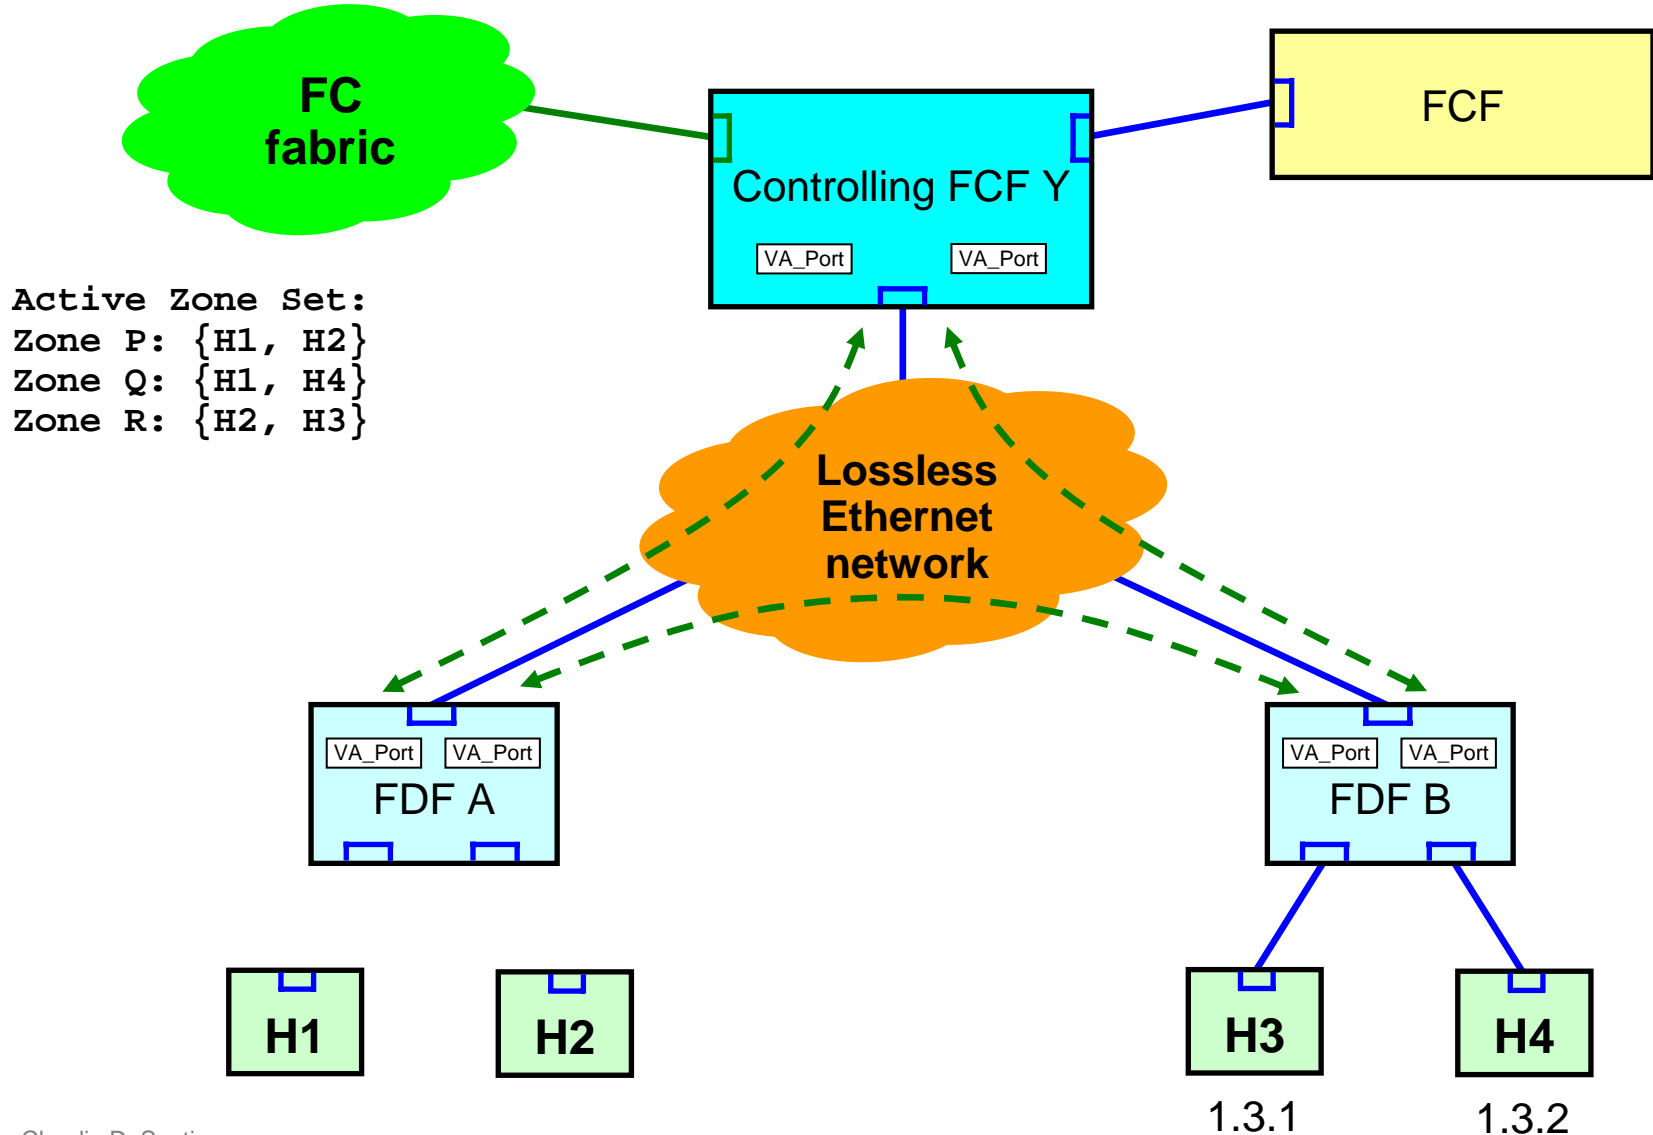
Forwarding Tables (2)



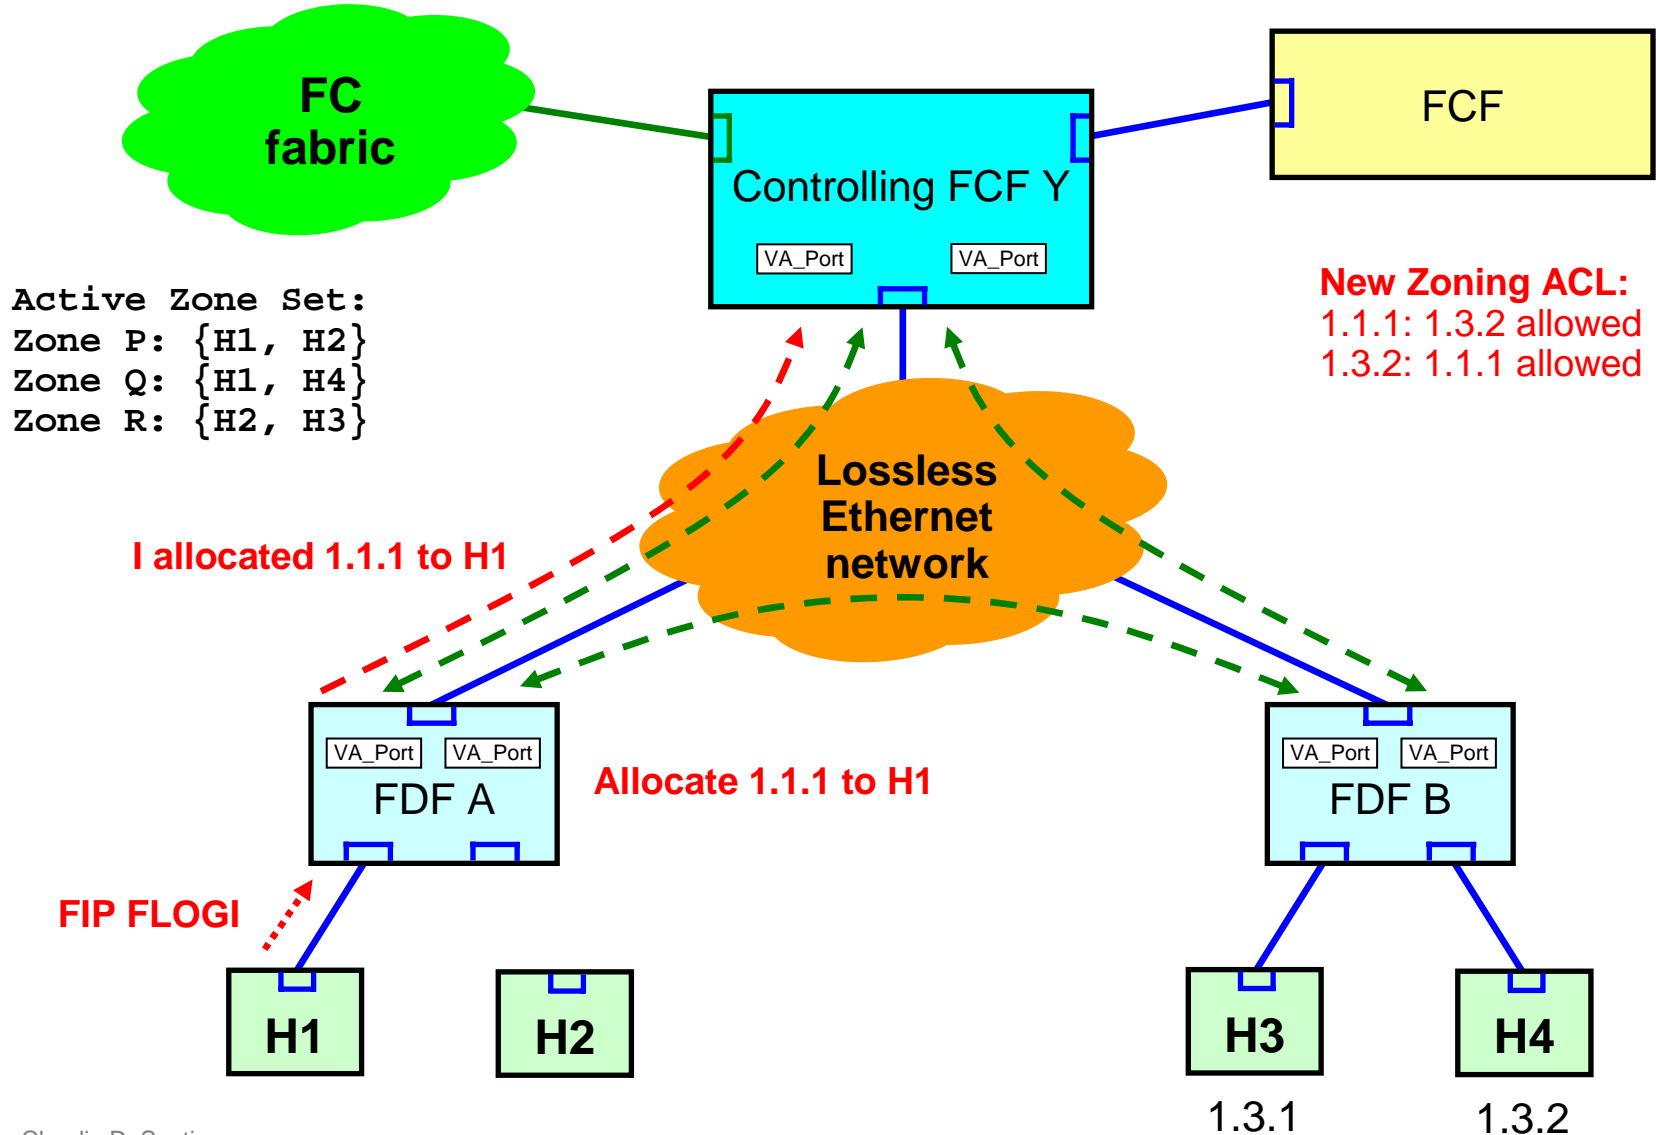
The FDF (1)

- **The FDF is identified by an FDF-MAC and an FDF-Name**
- **The FDF instantiates VF_Ports**
- **On receiving a FIP FLOGI or FIP NPIV FDISC Request an FDF allocates an N_Port_ID for the requester**
- **Then the FDF notifies the controlling FCF of this allocation**
 - A unidirectional SW_ILS
- **The controlling FCF updates the FC Name Server and generates the appropriate RSCN(s)**
- **The controlling FCF also re-computes the zoning “ACLs”**
 - The list of N_Port_IDs able to communicate with the newly allocated one
 - This may change also some other VN_Ports
- **The controlling FCF distributes the updated zoning ACLs to its affected FDFs**
- **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**

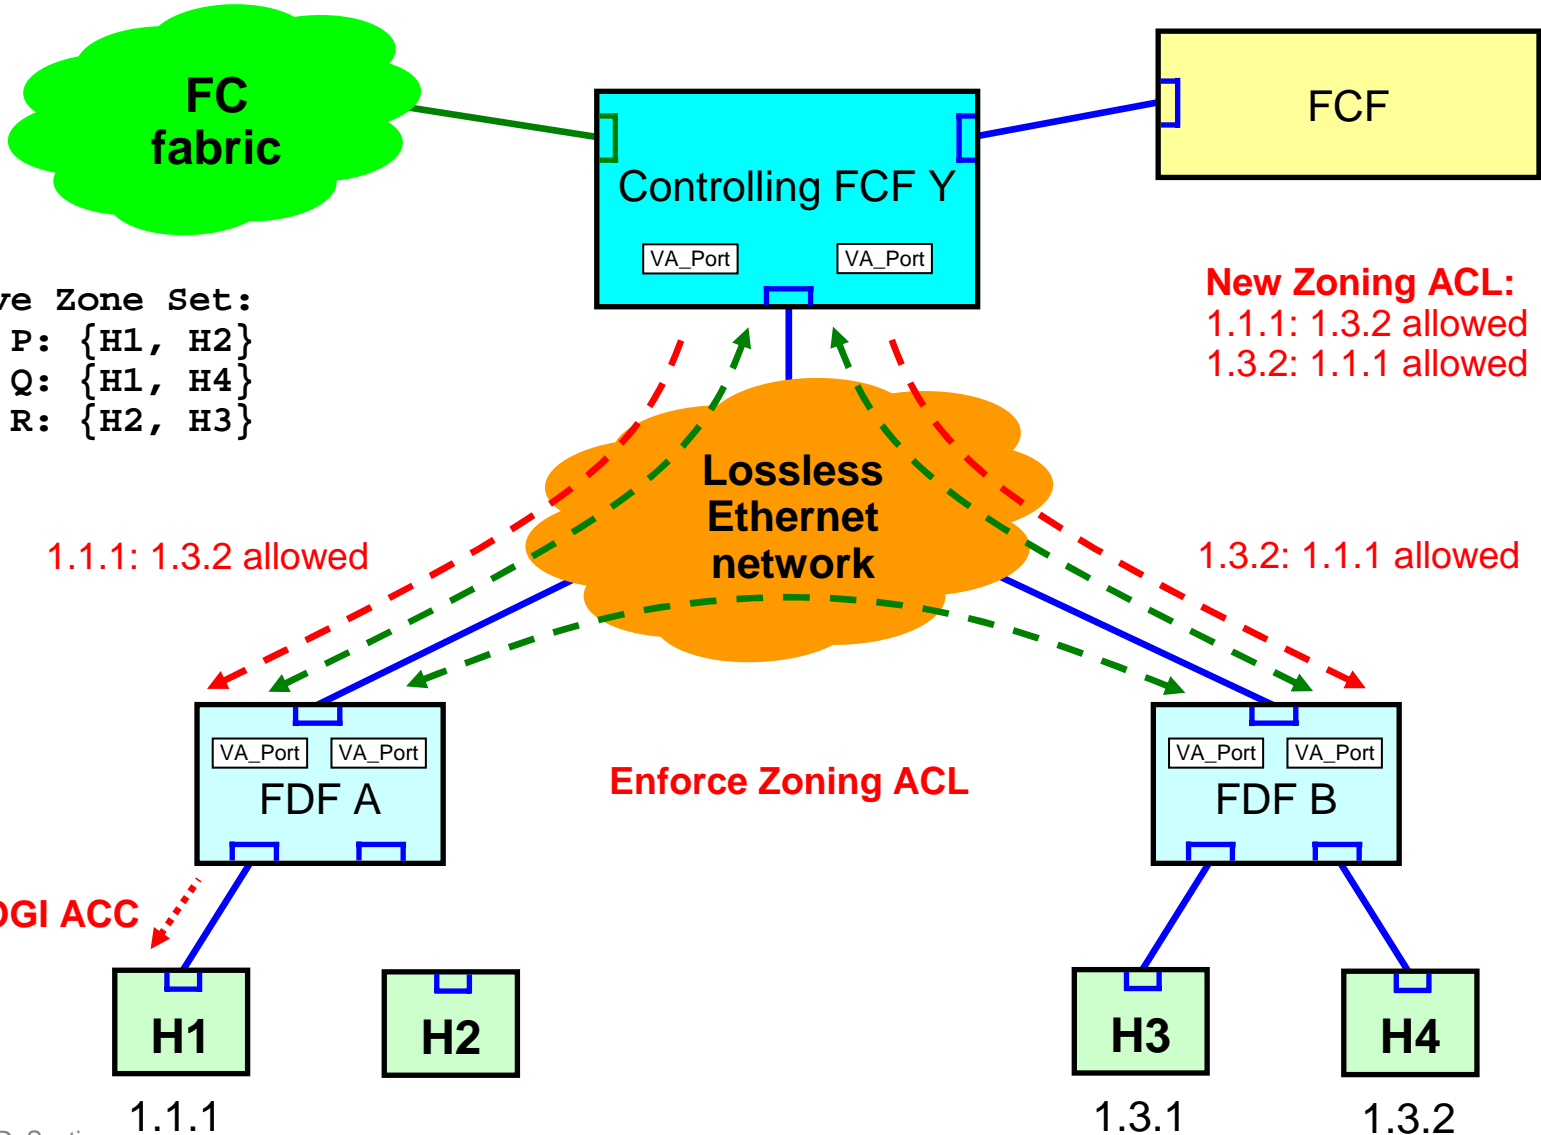
Zoning Configuration Example



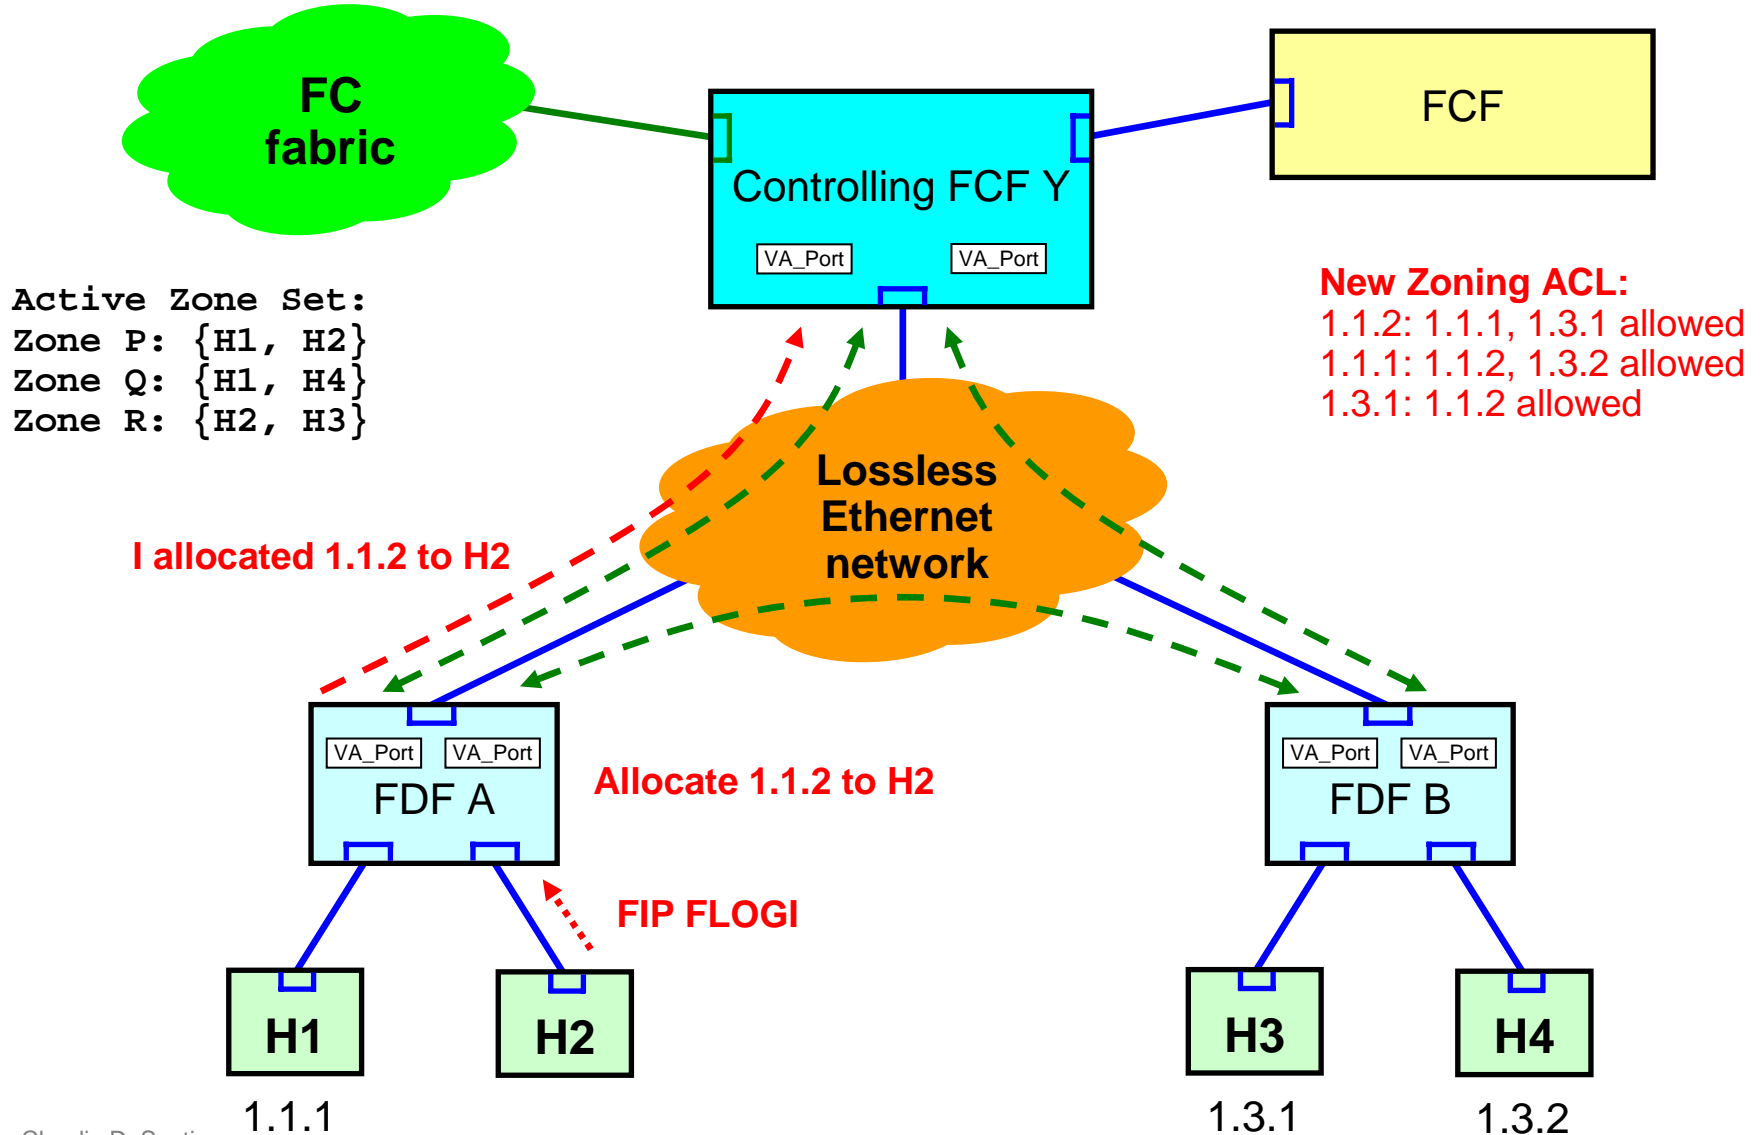
N_Port_ID Allocation (1)



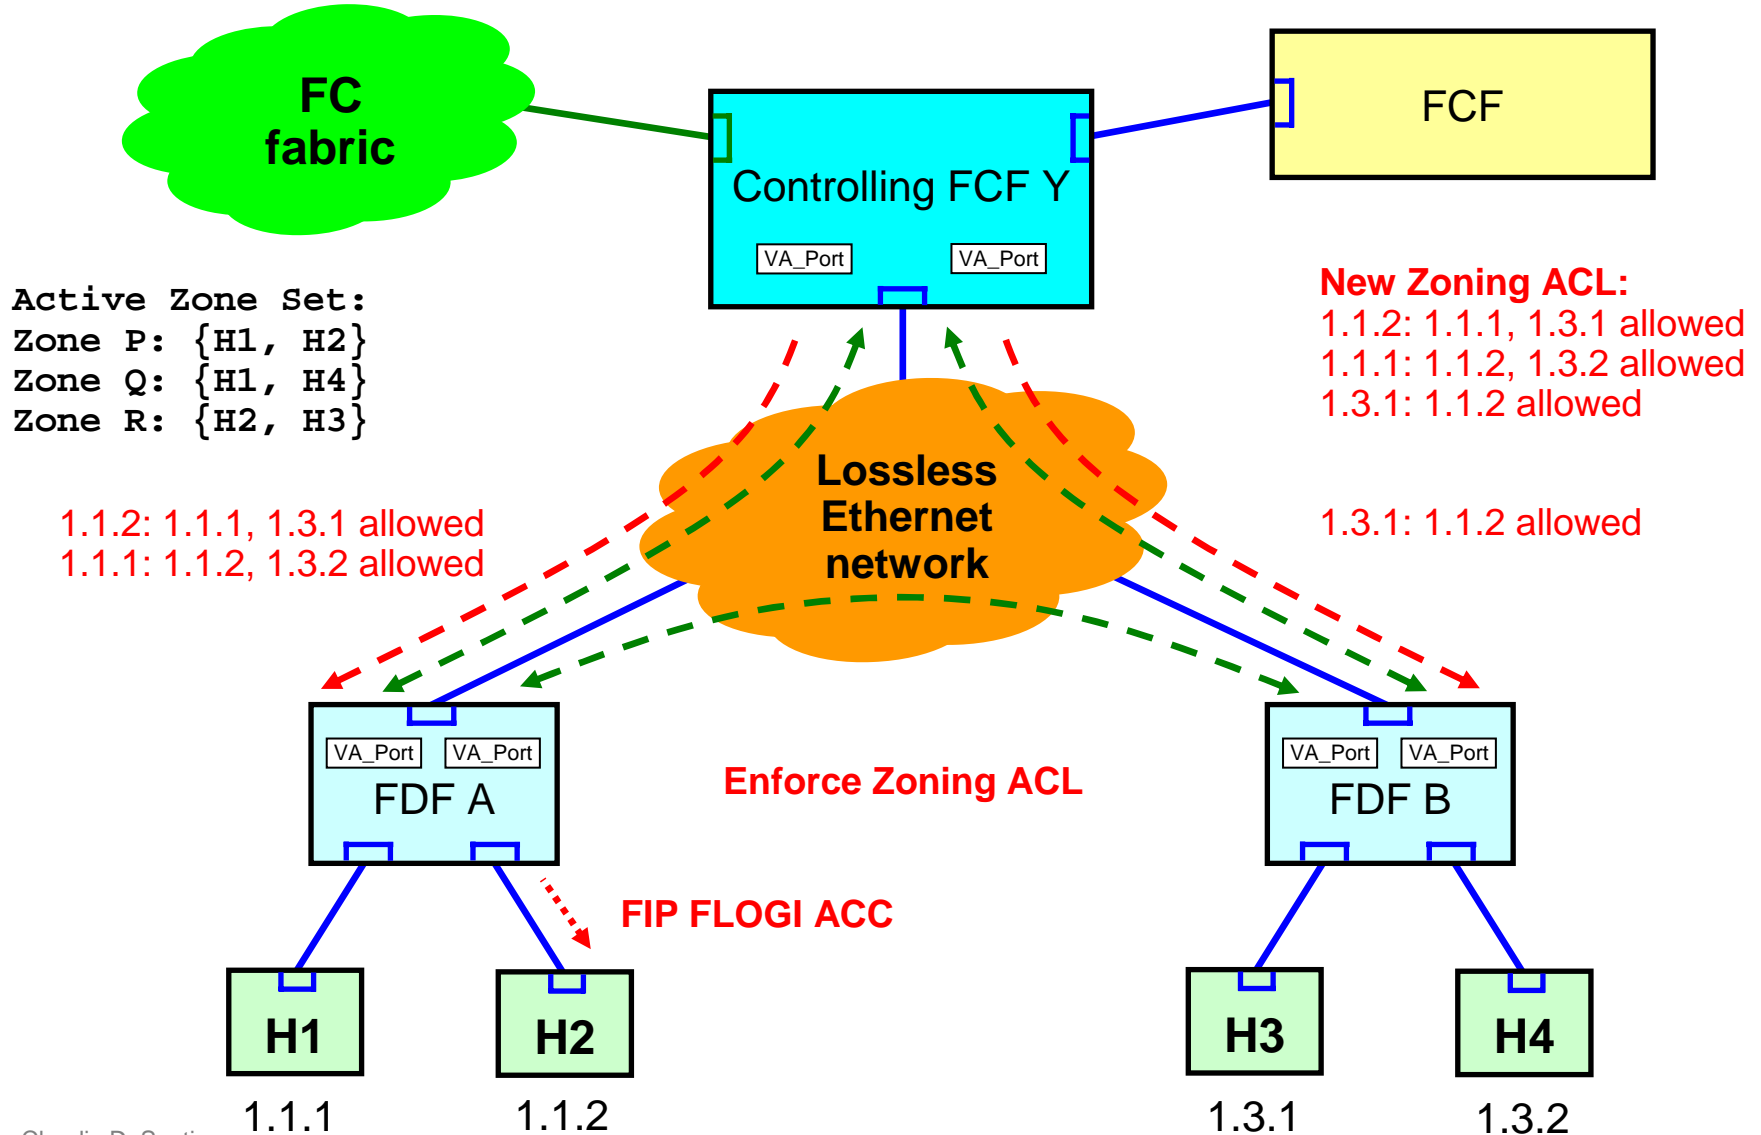
Zoning ACL Distribution (1)



N_Port_ID Allocation (2)



Zoning ACL Distribution (2)



The FDF (2)

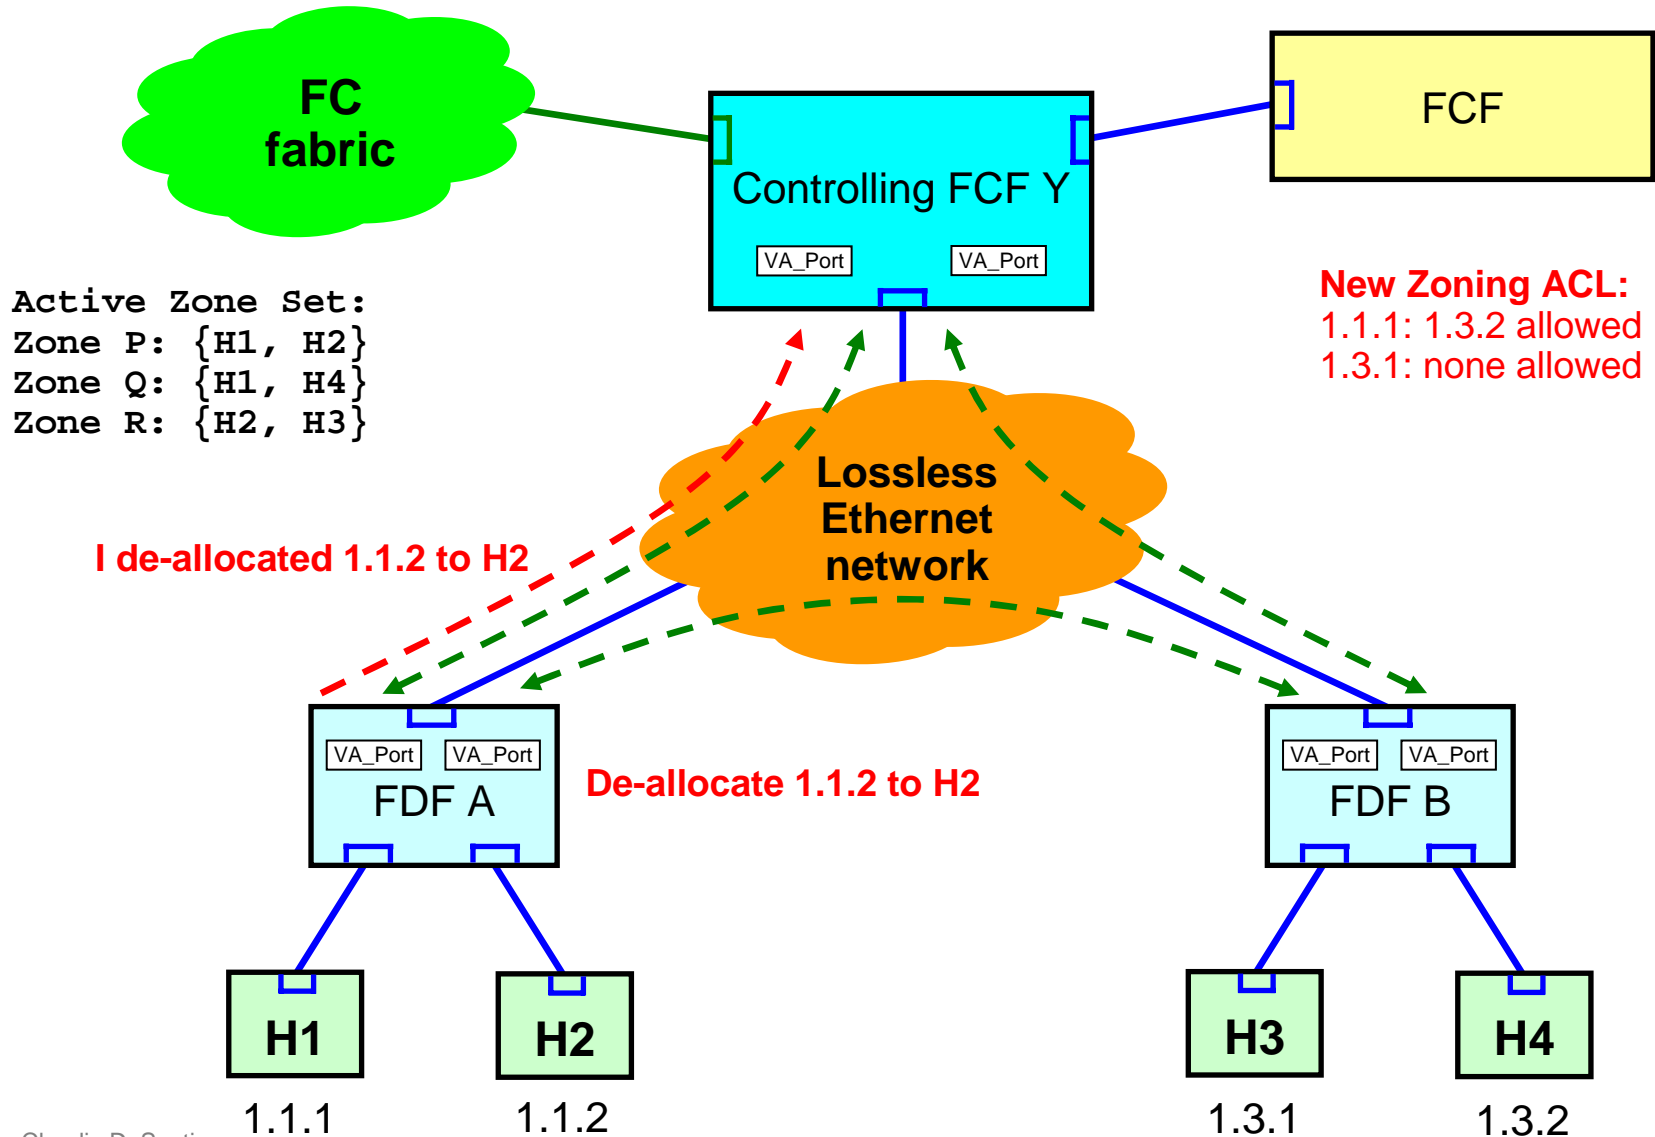
- **The FDF performs all ENode/FCF FIP functions**

It implements the VF_Ports

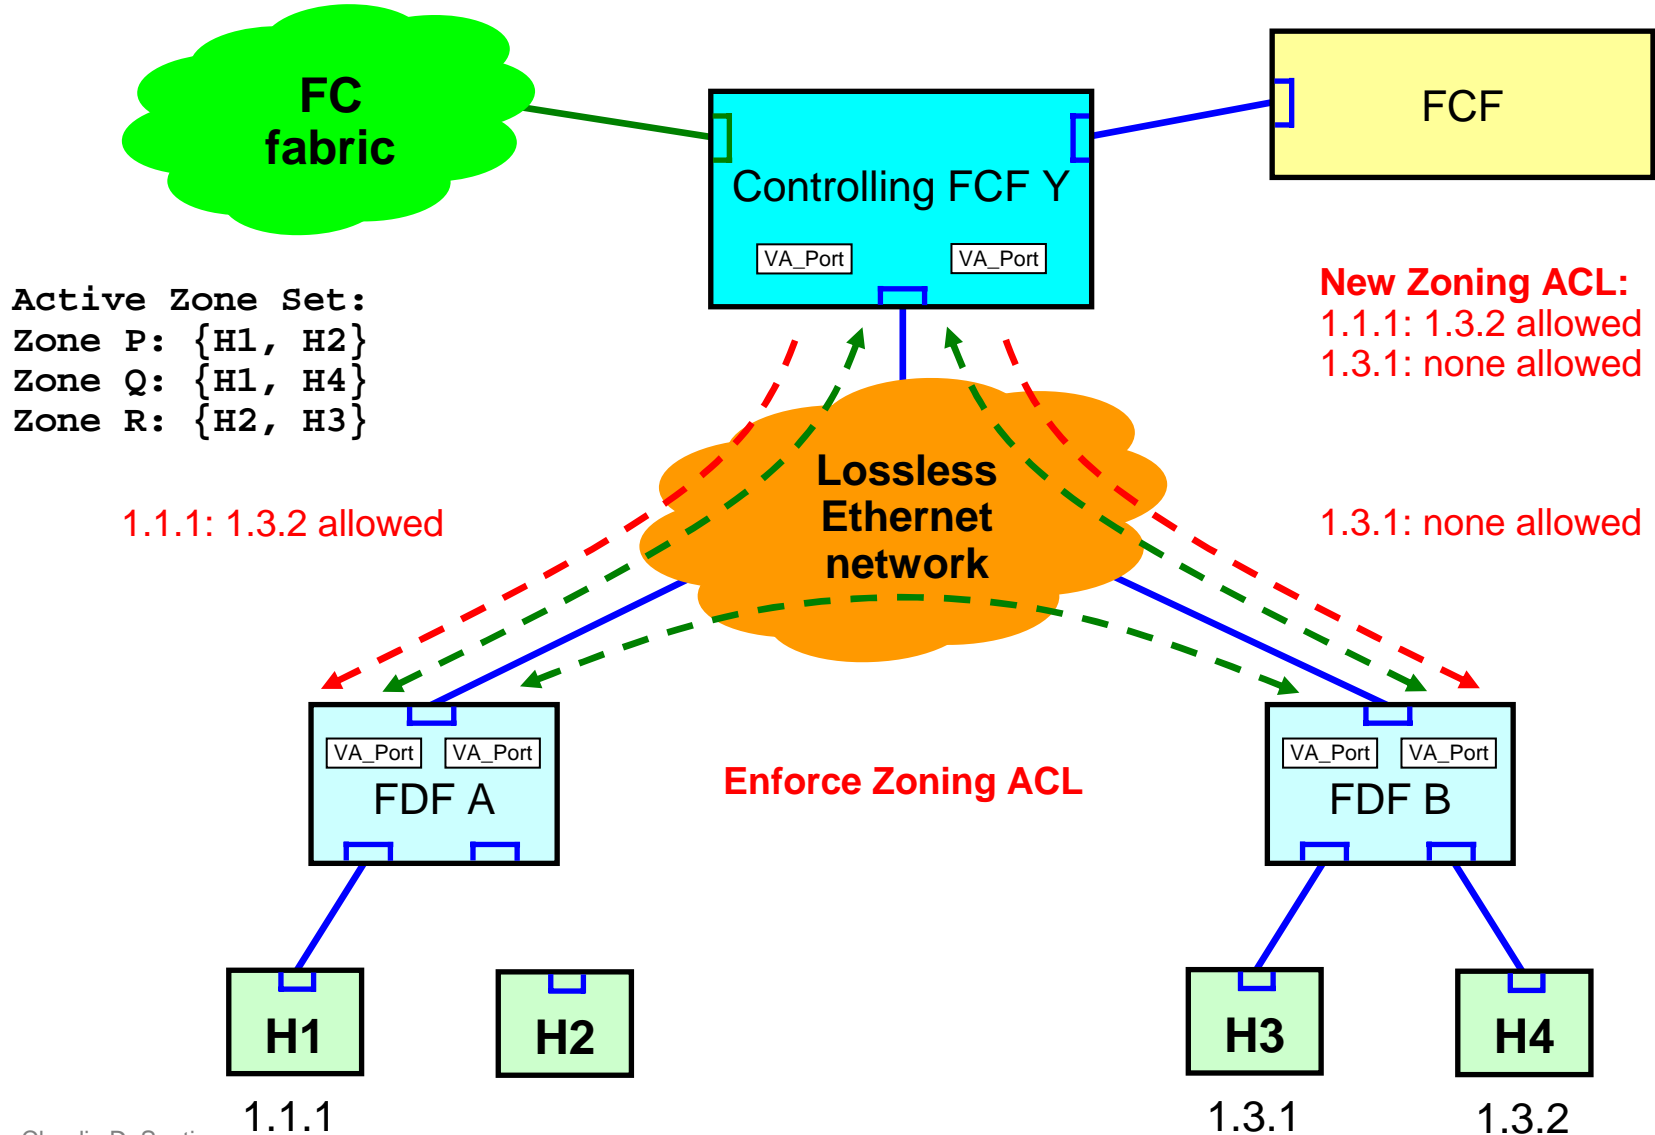
In particular it terminates FIP for Virtual Link maintenance

- **When a VF_Port to VN_Port Virtual Link is deinstantiated, the FDF sends a message to the controlling FCF to communicate that VN_Port is not anymore reachable**
- **The controlling FCF has to update the FC Name Server, recompute the zoning ACL and sends them to the affected FDFs, generate the appropriate RSCN(s)**

N_Port_ID De-allocation



Zoning ACL Distribution (3)



Generic Services

- **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**

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			

N_Port_ID Range Request SW_ILS

Item	Size
SW_ILS Code	4
Requesting FDF Switch_Name	8
Controlling FCF Switch_Name	8
Number of N_Port_IDs requested	4

N_Port_ID Range Assignment SW_ILS

Item	Size
SW_ILS Code	4
Controlling FCF Switch_Name	8
Recipient FDF Switch_Name	8
Initial N_Port_ID allocated	4
Ending N_Port_ID allocated	4

N_Port_ID Allocation/Deallocation SW_ILS

Item	Size
SW_ILS Code	4
Requesting FDF Switch_Name	8
Controlling FCF Switch_Name	8
Number of allocated N_Port_IDs (n)	4
Allocated N_Port_ID #1	4
Allocated N_Port_ID #2	4
...	
Allocated N_Port_ID #n	4
Number of deallocated N_Port_IDs (m)	4
Deallocated N_Port_ID #1	4
Deallocated N_Port_ID #2	4
...	
Deallocated N_Port_ID #m	4

N_Port_ID Peers Distribution SW_ILS

Item	Size
SW_ILS Code	4
Controlling FCF Switch_Name	8
Recipient FDF Switch_Name	8
Number of Peering Entries (h)	4
Peering Entry #1	4
Peering Entry #2	4
...	
Peering Entry #h	4

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

Thank You



Image Credit
Flickr user Kaeru
Creative Commons License