



UCS EHM Data Path Concepts and Troubleshooting



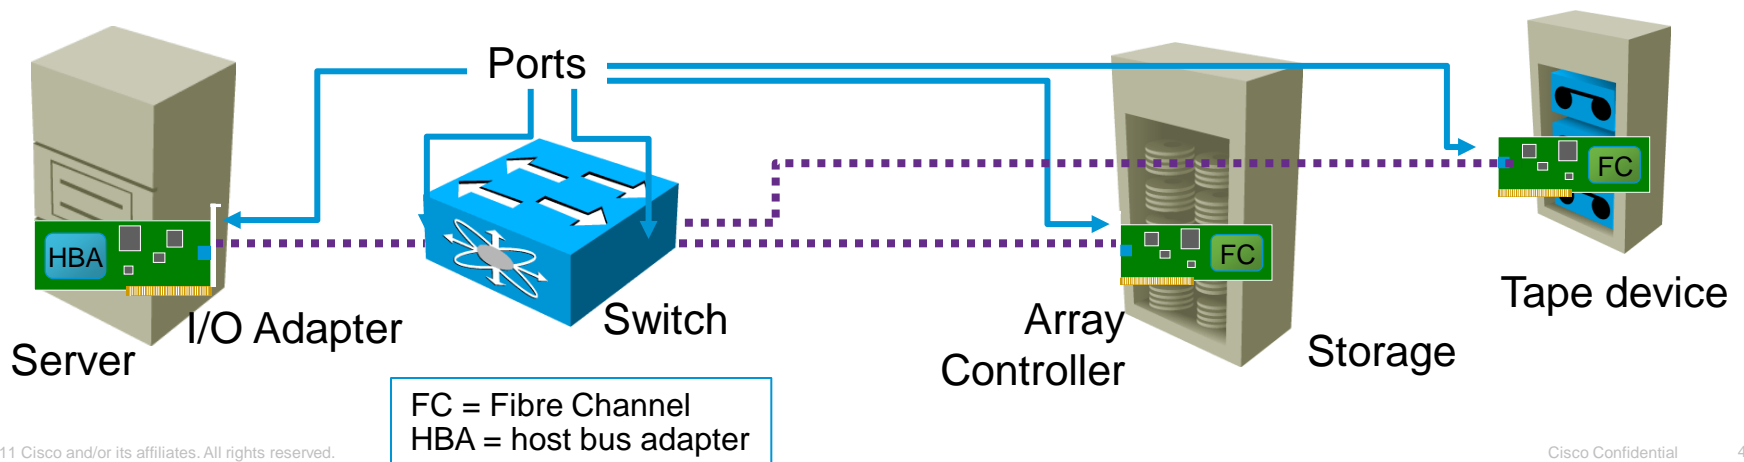
Agenda

- Fabric Channel Protocol Basic
- UCS EHM SAN Data Path & troubleshooting
- UCS EHM Ethernet Data Path
- Pinning
- DLU (Disjoint L2 Upstream)

Fibre Channel

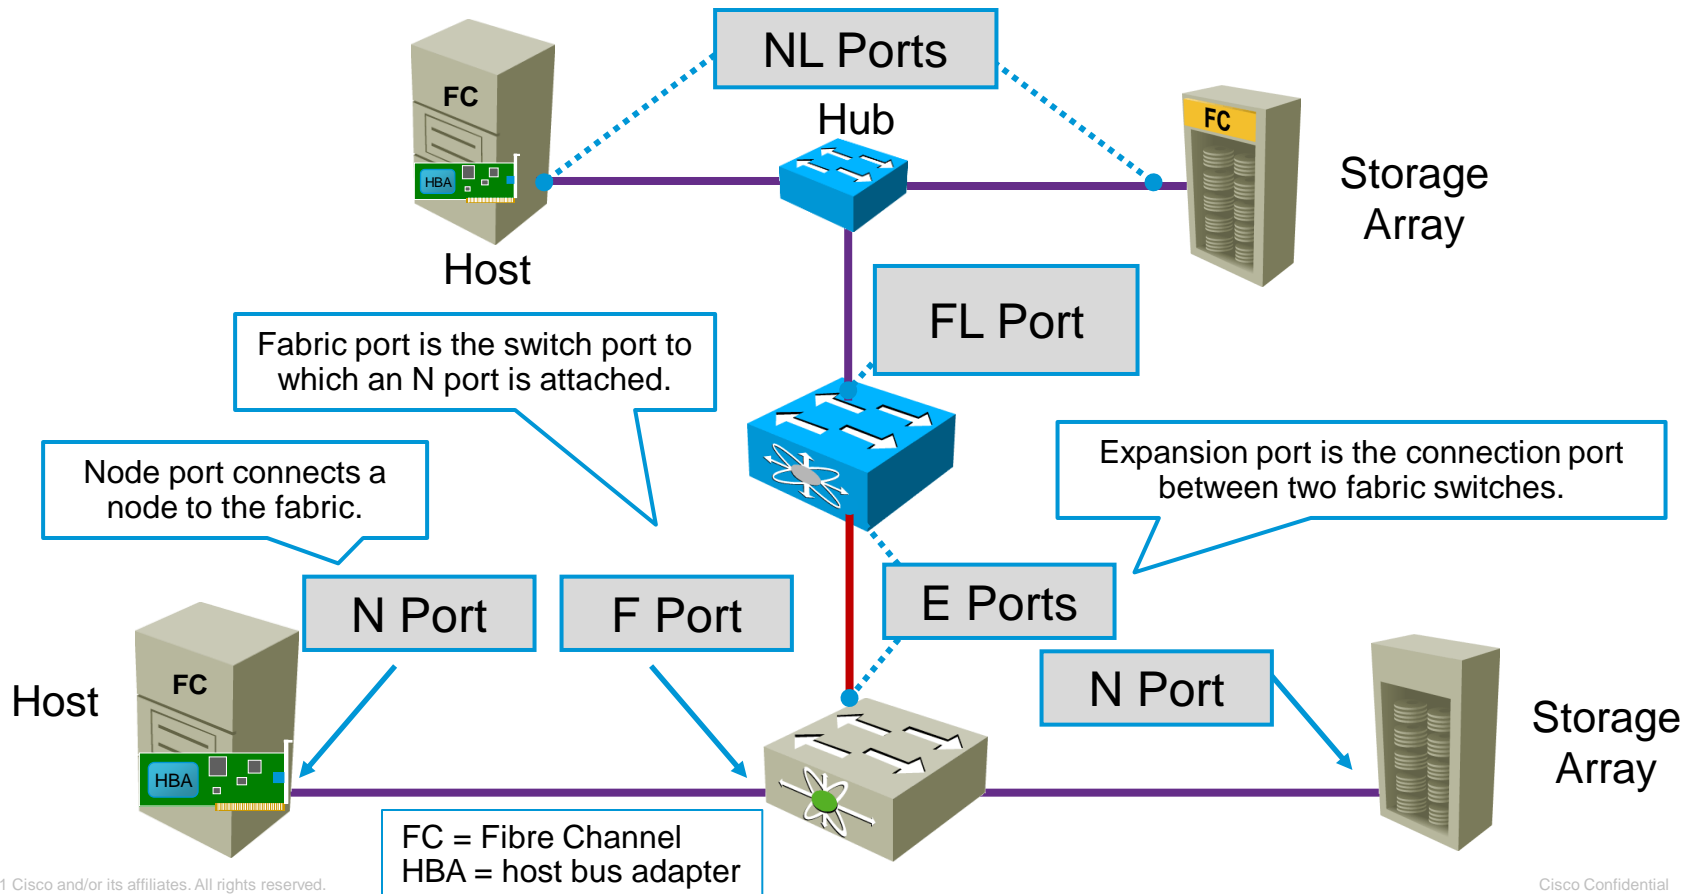
Fibre Channel

- Fibre Channel: industry standard used for storage networking
- Mapped to various protocols:
 - Fibre Channel Protocol (FCP)
 - SCSI
 - FICON
- Fibre Channel ports are intelligent interface points:
 - Embedded in an I/O adapter
 - Embedded in an array or tape controller
 - Embedded in a fabric switch



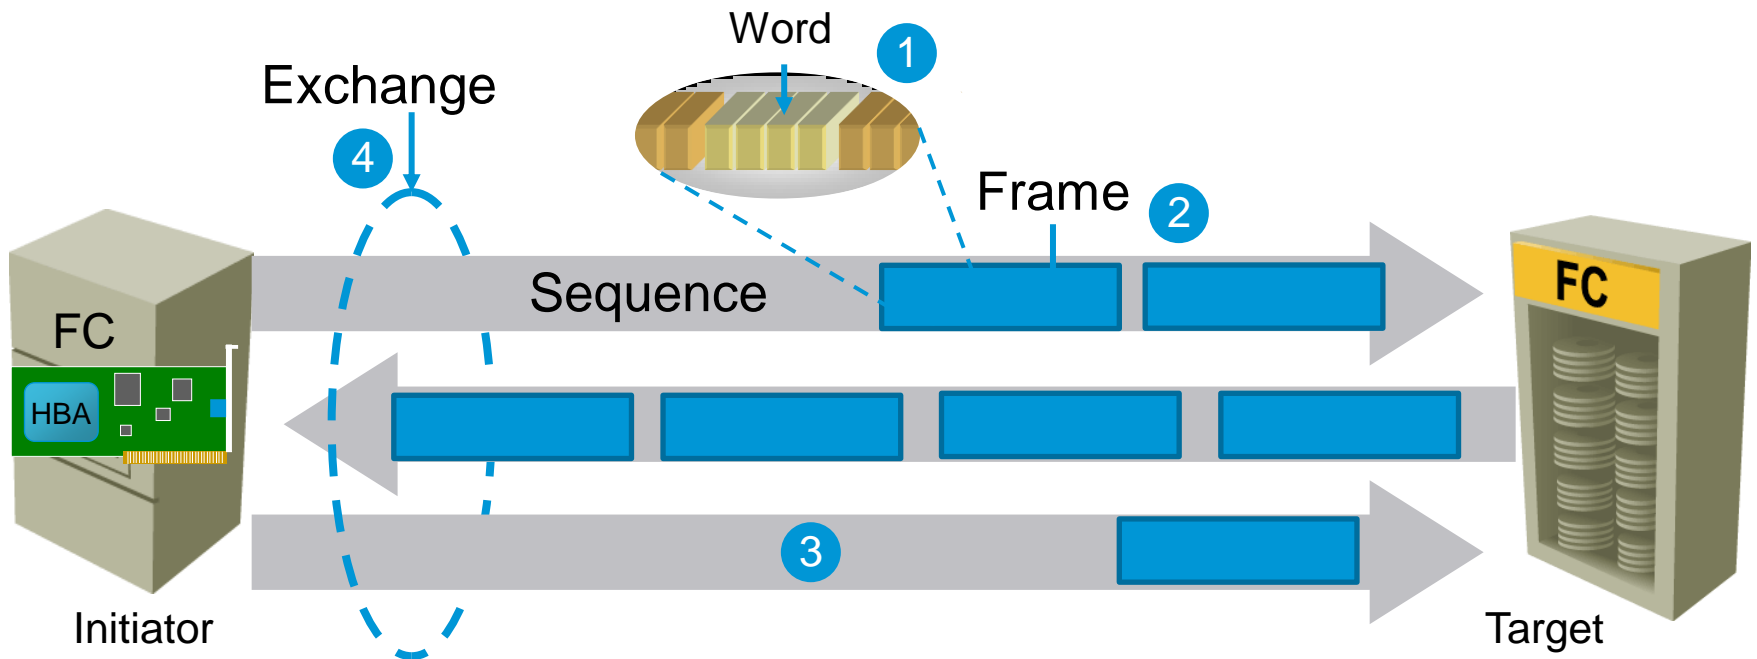
Fibre Channel Ports

- Common ports: node (N), fabric (F), expansion (E)
 - Other ports: node loop (NL) and fabric loop (FL)
- Node and fabric ports in one-way ring topology called arbitrated loop



Fibre Channel Framing

- Word: Smallest data unit of data, 32 bits encoded into a 40-bit form
- Frame: Package of multiple words, equivalent to an IP packet
- Sequence: Unidirectional series of frames
- Exchange: Series of sequences sent between two nodes




FC = Fibre Channel
HBA = host bus adapter

Fibre Channel Login and Communication

Registered State Change Notification

- Informs hosts about changes in the fabric
- Timely indication of one or more of events:
 - Disks joining or leaving the fabric
 - Name server registration change
 - New zone enforcement
- Resources must belong to the same VSAN and zone
- Hosts can receive this information by registering with the fabric controller
- Switch RSCN (SW-RSCN) sent to all reachable switches



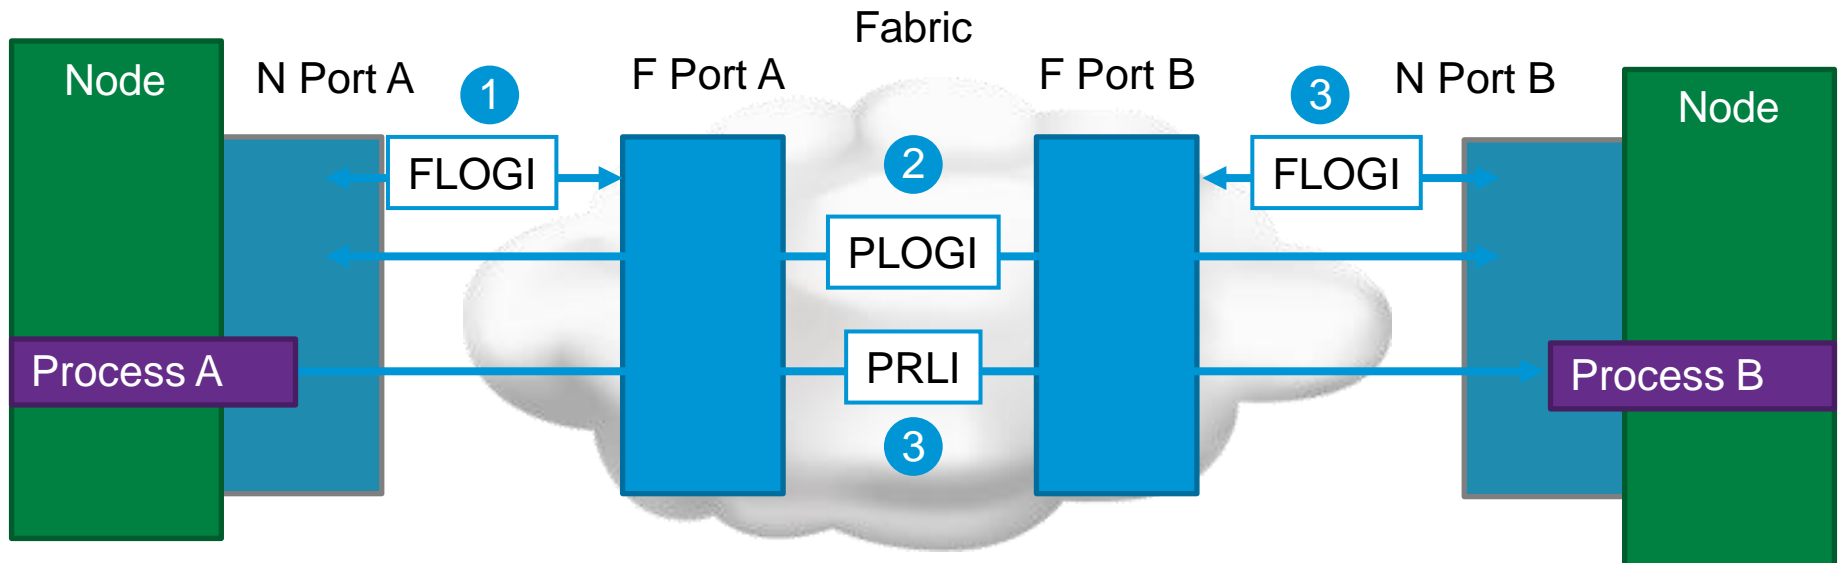
Fabric Login (FLOGI)

Three processes prerequisite to node port communication:

1. N Port must log into its attached F Port—fabric login (FLOGI)
2. N Port must log in to its target N Port—port login (PLOGI)
3. N Port must exchange ULP (SCSI-3, IP, FICON) support information with its target N Port:


Ensure that the initiator and target process can communicate.

Known as the process login (PRLI).



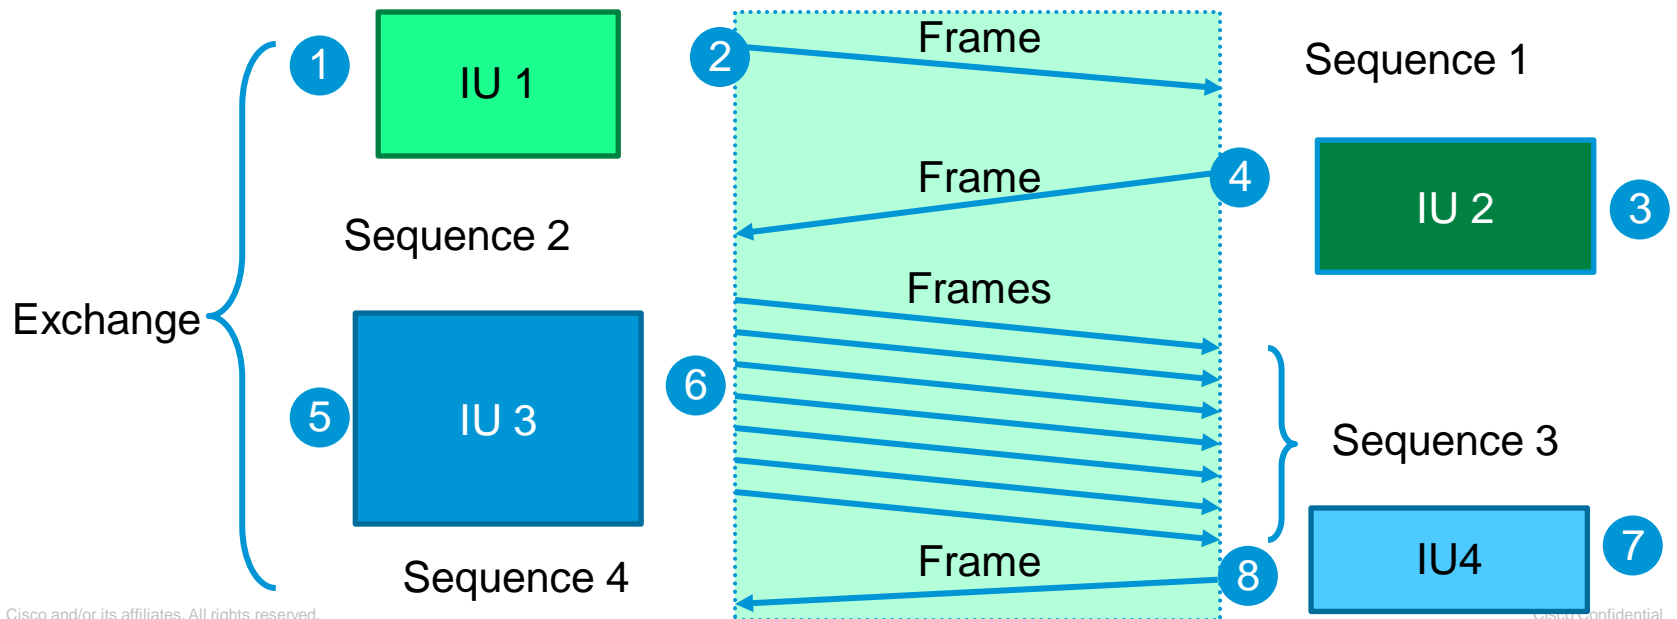
Fibre Channel Read Operation

1. Information Unit (IU) 1 carries SCSI read request (FCP_CMD).
2. Initiator FC-2 layer sends IU 1 as a single frame.
3. Target retrieves requested data (FCP_DATA) and packages as IU 2.
4. Target FC-2 layer sends IU 2 as frame sequence.
5. Target generates status command (FCP_RSP) and packages as IU 3.
6. Target FC-2 layer sends IU 3.



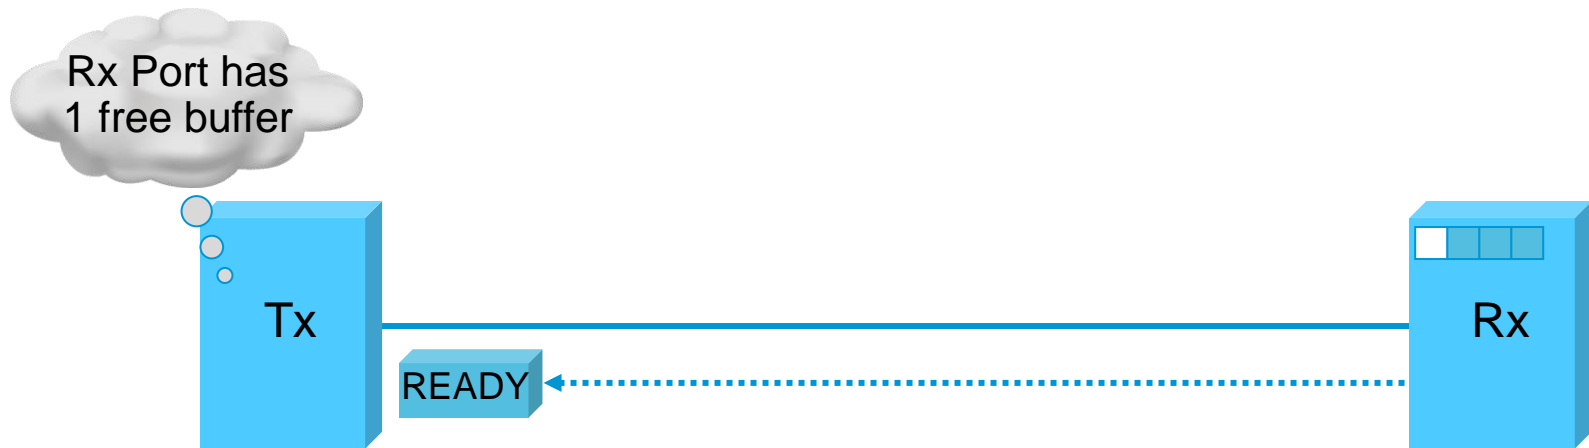
Fibre Channel Write Operation

1. Information IU 1 contains SCSI write request (FCP_CMD).
2. Initiator FC-2 layer sends IU 1 as a single frame.
3. Target responds with a SCSI write request response (FCP_XFR_RDY.)
4. Target FC-2 layer sends IU 2 as frame sequence.
5. Initiator retrieves data (FCP_DATA) from its upper-layer protocol (ULP) buffers.
6. Initiator FC-2 layer converts IU 3 to one or more data chunks and transmits.
7. Target generates status command (FCP_RSP) to confirm the end of the exchange.
8. Target FC-2 layer converts IU 4 to a single command chunk and as sequence 4.



Fibre Channel Flow Control

- Fibre Channel uses a credit-based strategy.
 - Transmitter does not send a frame until the receiver tells the transmitter that the receiver can accept another frame.
 - The receiver is always in control.
- Benefits
 - Prevents loss of frames due to buffer overruns.
 - Maximizes performance under high loads.



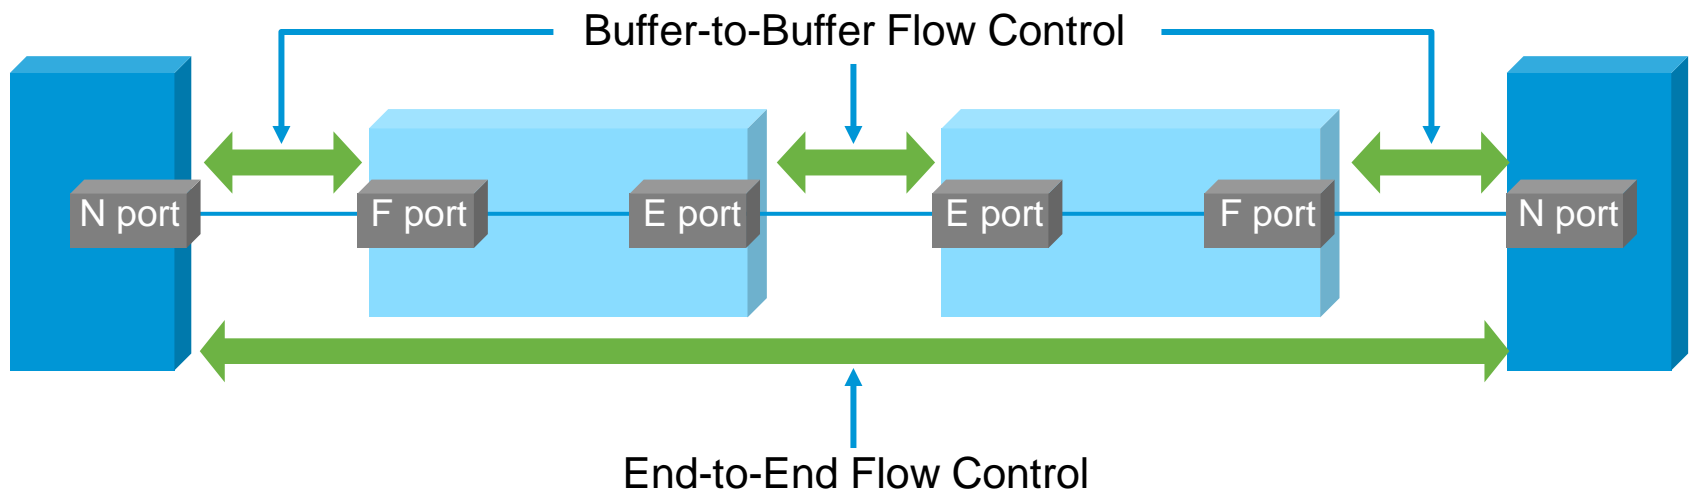
Flow Control in Fibre Channel

Fibre Channel Flow Control Types

Fibre Channel defines two types of flow control:

Buffer-to-buffer (port-to-port)


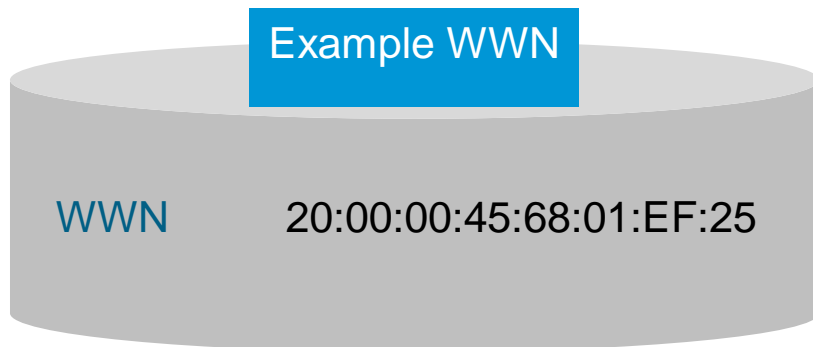
End-to-end (source-to-destination)



The Fibre Channel Addressing Schemes

Fibre Channel WWNs

- Every Fibre Channel port and node has a hard-coded address called a world wide name (WWN).
 - Allocated to manufacturer by IEEE
 - Coded into each device when manufactured
 - 64 or 128 bits (128 bits most common today)
- Switch name server maps WWNs to Fibre Channel addresses.
 - Node World Wide Names (nWWNs) uniquely identify devices
 - Port World Wide Names (pWWNs) uniquely identify each port in a device

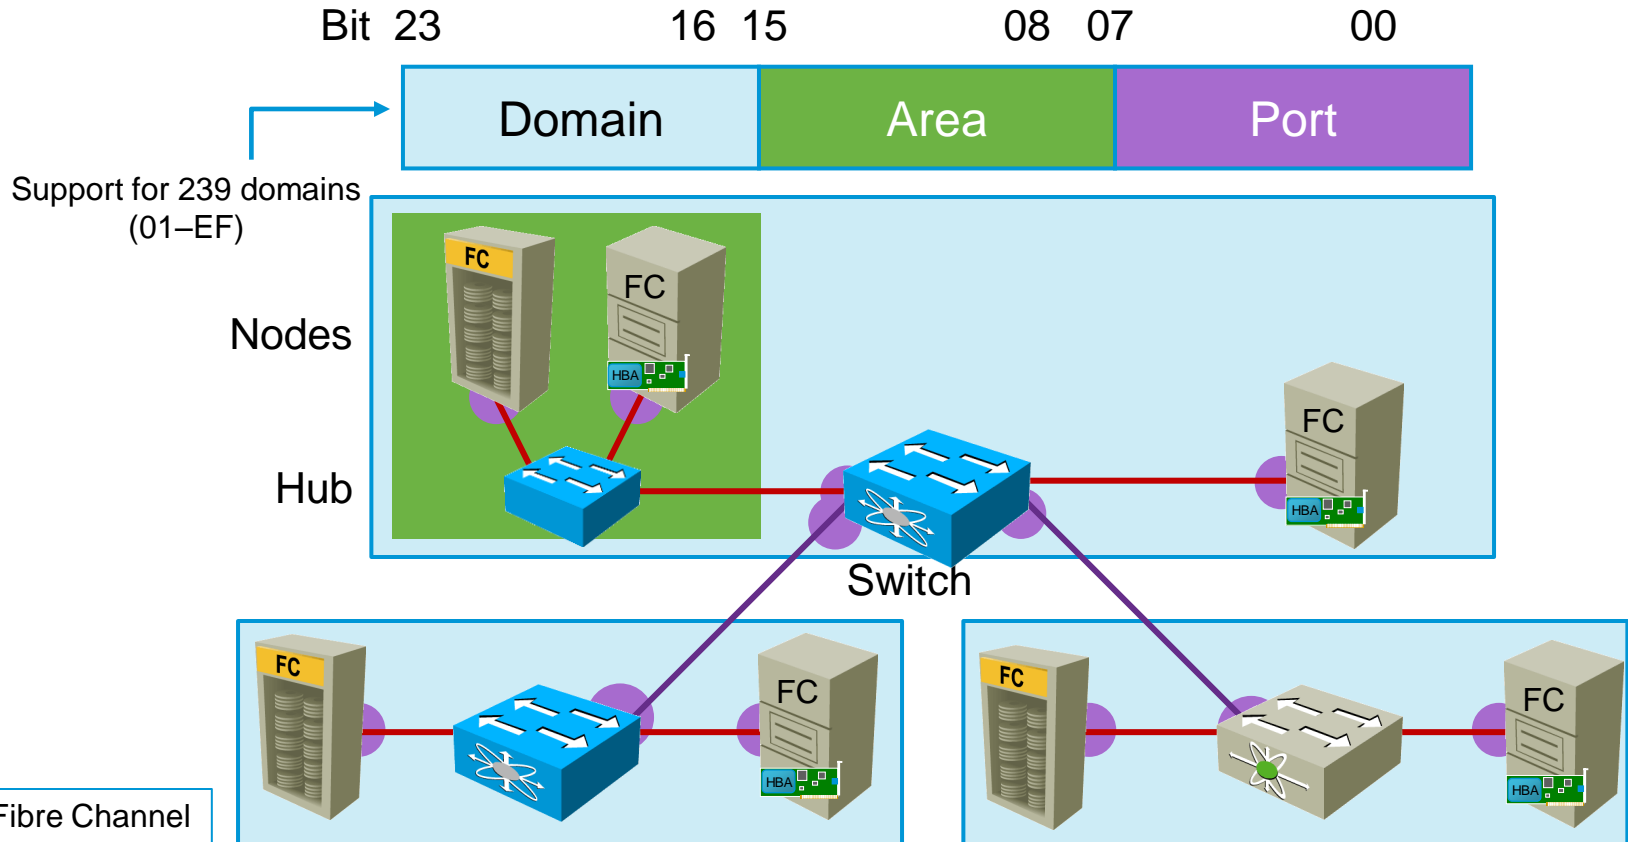


Fibre Channel Address Format

Domain ID defines a switch. Each switch receives a unique domain ID.

Area ID identifies groups of ports within a domain.

Port ID identifies individual device(s) on that port.




N-Port ID Virtualization

N-Port ID virtualization (NPIV) provides a means to assign multiple FCIDs to a single N Port.

NPIV allows multiple applications to share the same HBA.

The use of different pWWNs allows access control, zoning, and port security to be implemented at the application level.

Usage applies to virtual server applications such as VMware, Microsoft Hyper-V, and Citrix XenServer.

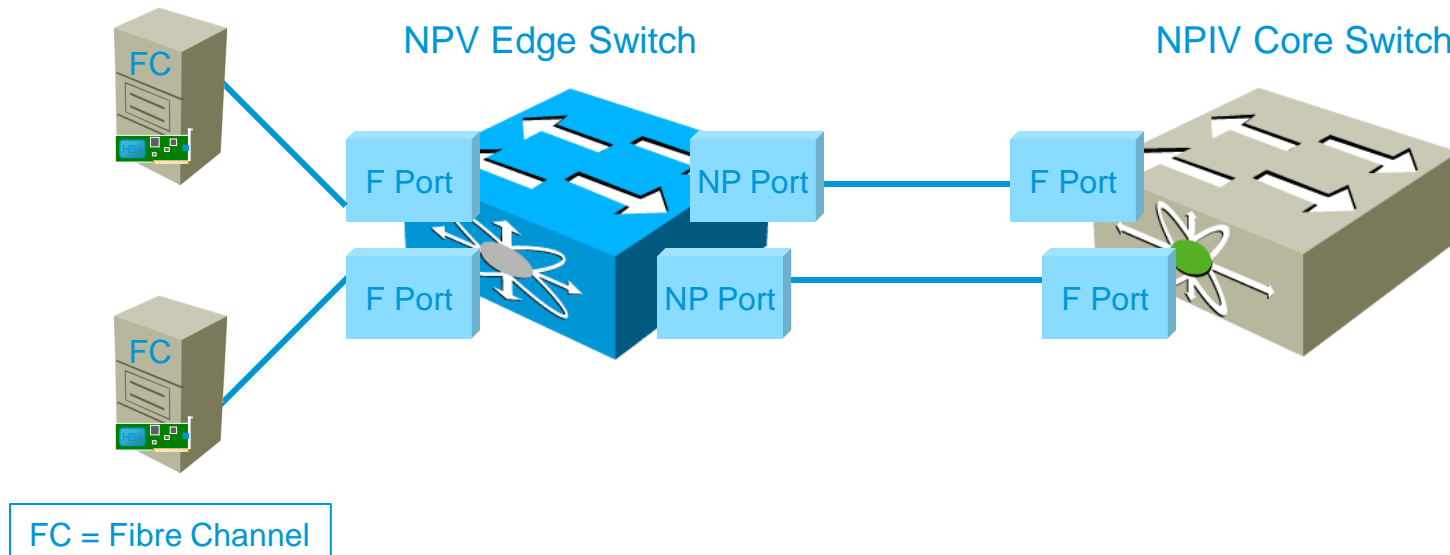


N-Port Virtualization

Extension to NPIV

Allows the blade switch or fabric device to behave as an NPIV-based HBA to the core Fibre Channel switch

NPV device aggregates the locally connected host ports (N Ports) into one or more uplinks (pseudo-interswitch links) to the core switches



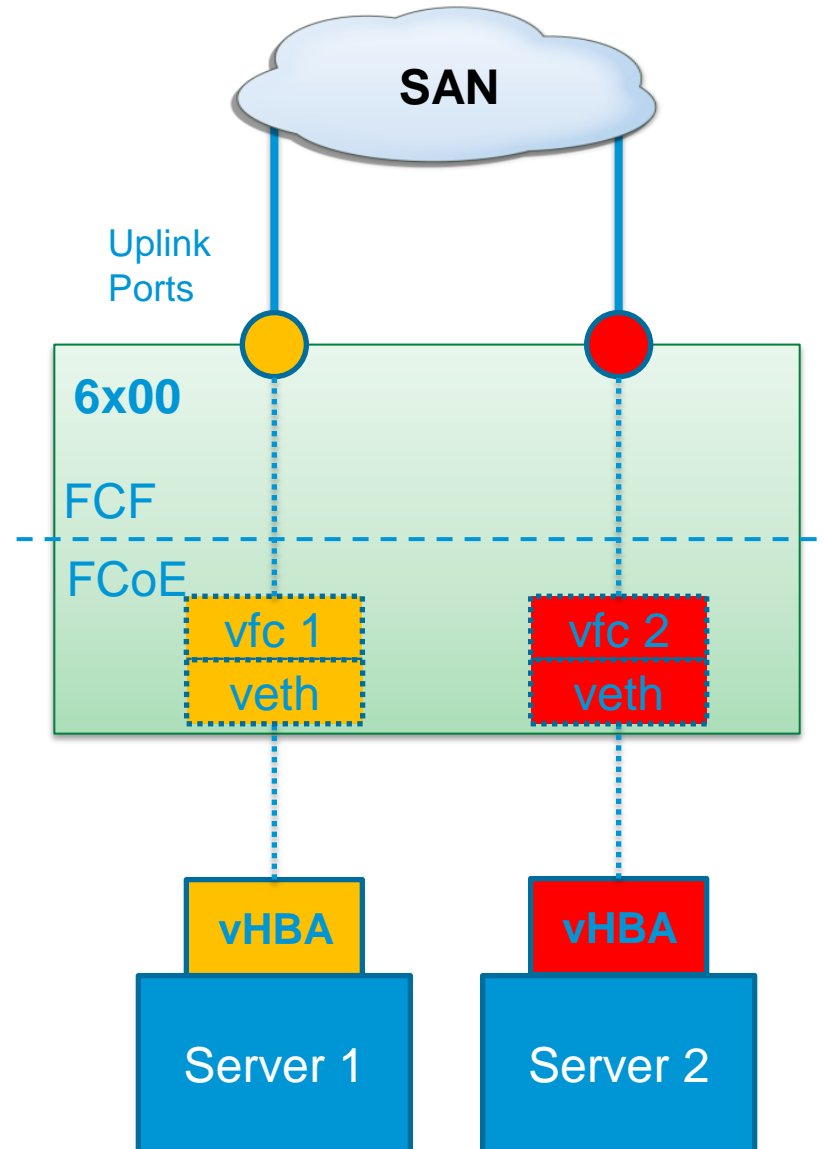
NPV vs. Regular Switching

Feature	Regular switching	NPV
Fibre Channel services	<ul style="list-style-type: none"> All Fibre Channel services are provided (FLOGI, name server, zoning, domain server, FSPF, and management). FSPF, zoning, and name server databases are distributed among connected switches. 	Most Fibre Channel services are switched off.
Switching operation	Local switching enabled	<ul style="list-style-type: none"> Device acts as proxy Subordinate to a Fibre Channel switch
Domain ID	Each switch consumes a domain ID.	<ul style="list-style-type: none"> Does not use a domain ID No domain ID limitation
Scalability and manageability	<ul style="list-style-type: none"> ISL between switches becomes a path within the FSPF routing table. Up to 16 ISLs may belong to a port channel. Up to 239 switches can be in a fabric (theoretical limit), 	<ul style="list-style-type: none"> Eliminates the need for server administrators to manage the SAN Third-party integration VSAN scalability
QoS	Available	Not available

UCS EHM SAN Data Path & Troubleshooting

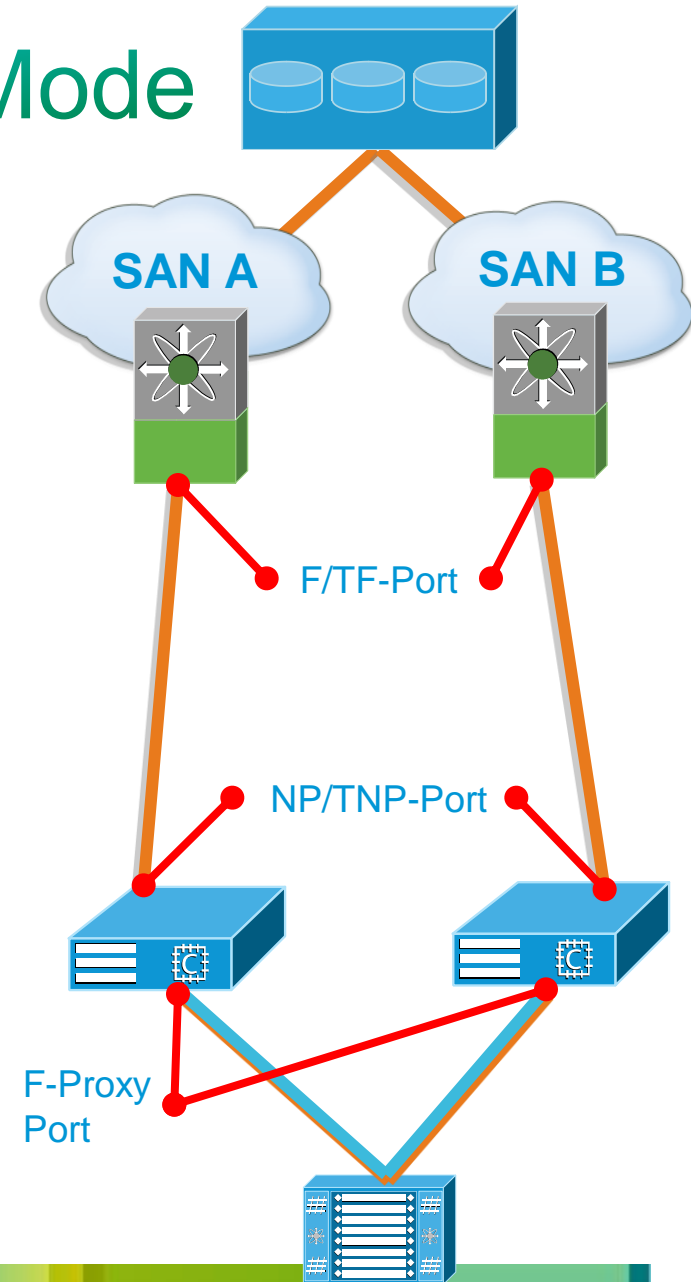
UCS EHM SAN Data Path

- EHM is NPV mode
- FI is working as a FCF (FC forwarder)
- vHBA of the server binds to a vfc on FI
- vfc binds to a vethernet which carry the fcoe traffic
- vfc pins to FC uplink by round-robin



Fibre Channel End-Host-Mode

- N-Port Virtualization (NPV) mode
 - Core SAN NPIV are F (TF) ports
 - UCS NPV fabric are NP (TNP) ports
- Version 1.4 supports
 - f port-trunk, (VSAN Trunk, multiple VSANs)
 - f san-port-channel (link aggregation)
- Zoning done upstream on the NPIV switch
- No Fibre Channel fabric services
 - e.g. No Domain mgr, zone server, fabric login server, name server



FCID's - fabric channel login ID's

- FCID's are made of hex in the form of:
0XDDAAPP
- DD = Domain ID, unique to a physical switch and VSAN
- AA = Area ID is normally used in FICON, but also with NPV, all f port logins from the same UCS will have the same DomainID and AreaID (0XDDAA - -)
- PP = PortID
- We can use this to help understand what is connected to what.
Similar to show cdp neighbors
- Each device (switch, HBA, or vfc) that flogi's in will have a FCID value

VSAN NPV and NPIV Status

- What DomainIDs are in my fabric, 0XDD---- values

```
MDS-CiscoLive-1# show fcdomain domain-list
```

```
Number of domains: 1
```

```
Domain ID          WWN
-----
0x64(100)         20:64:00:0d:ec:24:5b:c1 [Local] [Principal]
VSAN 50
```

```
Number of domains: 1
```

```
Domain ID          WWN
-----
0x1c(28)          20:32:00:0d:ec:24:5b:c1 [Local] [Principal]
```

```
MDS-CiscoLive-1# show flogi database
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/8	100	0x640102	20:41:00:0d:ec:d3:5d:c0	20:64:00:0d:ec:d3:5d:c1
fc1/8	100	0x64012d	20:00:00:25:b5:30:20:01	20:00:00:25:b5:00:00:1d
fc1/8	100	0x640131	20:00:00:23:33:a9:20:6e	20:00:00:25:b5:10:10:0f
fc1/9	100	0x640203	20:42:00:0d:ec:d3:5d:c0	20:64:00:0d:ec:d3:5d:c1
fc1/9	100	0x64022b	20:00:00:25:b5:30:10:01	20:00:00:25:b5:00:aa:4f
fc1/13	100	0x640300	50:0a:09:84:86:78:3b:98	50:0a:09:80:86:78:3b:98
			[dev-alias-netapp]	
fc2/13	50	0x1c0000	50:0a:09:83:86:78:3b:98	50:0a:09:80:86:78:3b:98

VSAN NPV and NPIV Status

- Where is UCS connecting to the MDS?

```
MDS-CiscoLive-1# show flogi database
```

INTERFACE	VSAN	FCID	PORT NAME
fc1/8	100	0x640102	20:41:00:0d:ec:d3:5d:c0
fc1/8	100	0x64012d	20:00:00:25:b5:30:20:01
fc1/8	100	0x640131	20:00:00:23:33:a9:20:6e
fc1/9	100	0x640203	20:42:00:0d:ec:d3:5d:c0
fc1/9	100	0x64022b	20:00:00:25:b5:30:10:01
fc1/13	100	0x640300	50:0a:09:84:86:78:3b:98
			[dev-alias-netapp]
fc2/13	50	0x1c0000	50:0a:09:83:86:78:3b:98

```
Total number of flogi = 7.
```

```
UCS-CiscoLive-A(nxos)# show npv status
```

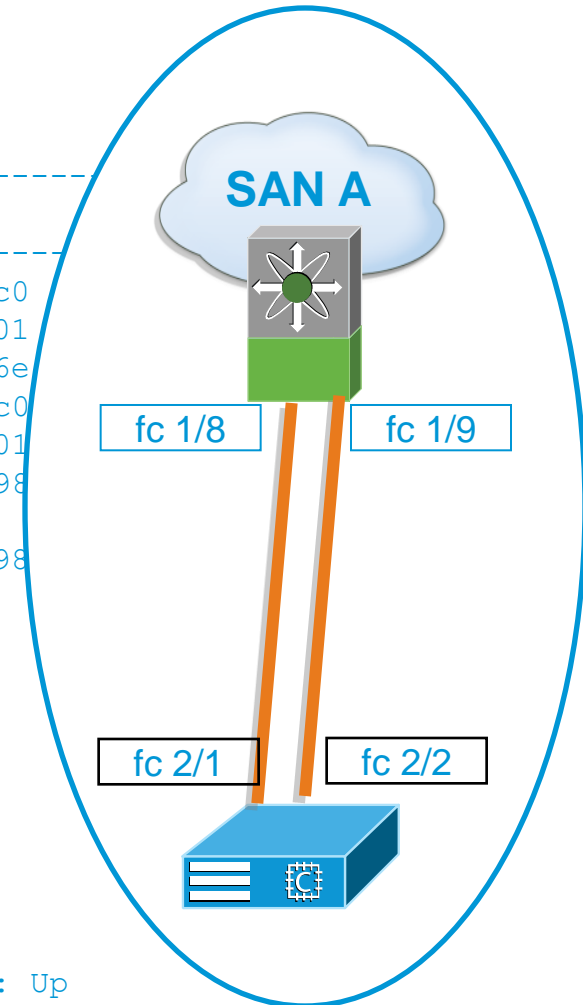
```
npiv is enabled
```

```
External Interfaces:
```

```
=====
```

```
Interface: fc2/1, VSAN: 100, FCID: 0x640102, State: Up
```

```
Interface: fc2/2, VSAN: 100, FCID: 0x640203, State: Up
```



NPV commands on Core MDS

- Another way to determine the connection between the core MDS and UCS in NPV mode

```
MDS-CiscoLive-1# show fcns database npv
```

```
VSAN 1:
```

```
-----  
NPV NODE-NAME          NPV IP_ADDR    NPV IF  CORE SWITCH WWN          CORE IF  
-----  
20:01:00:0d:ec:d3:5d:c1 172.25.183.133 Po22    20:00:00:0d:ec:24:5b:c0 Po44
```

```
MDS-CiscoLive-1# show interface port-channel 44
```

```
port-channel 44 is trunking
```

```
<truncated>
```

```
Member[1] : fc1/8
```

```
Member[2] : fc1/9
```

```
UCS-CiscoLive-A(nxos)# show interface san-port-channel 22
```

```
san-port-channel 22 is trunking
```

```
<truncated>
```

```
Member[1] : fc2/1
```

```
Member[2] : fc2/2
```

Fibre Channel Path Trace

```
UCS-CiscoLive-A(nxos) # sho interface vethernet 767  
vethernet767 is up  
Bound Interface is Ethernet1/1/2
```

```
UCS-CiscoLive-A(nxos) # sho interface vfc 768  
vfc768 is up  
Bound interface is vethernet8960  
Hardware is Virtual Fibre Channel  
Port WWN is 22:ff:00:0d:ec:d3:5d:ff  
Admin port mode is F, trunk mode is off  
Port mode is F, FCID is 0x640106  
Port vsan is 100
```

```
UCS-CiscoLive-A(nxos) # show interface vethernet 8960  
vethernet8960 is up  
Bound Interface is Ethernet1/1/2  
Hardware: Vethernet
```

Go full circle.

Query the virtual interfaces to display the bound physical interfaces and their status

Fibre Channel Path Trace

```
UCS-CiscoLive-A# show service-profile circuit server x/y
Service Profile: 530DC/Cehuiyuan38
Server: 2/2
```

Fabric ID: A

VIF	vNIC	Link State	Oper State	Prot State	Prot Role	Admin Pin	Oper Pin	Transport
1104	eth1	Up	Active	Passive	Backup	1/17	1/17	Ether
1105	eth0	Up	Active	Active	Primary	1/17	1/17	Ether
1108	fc0	Up	Active	No Protection	Unprotected	0/0	2/1	Fc
9300		Up	Active	No Protection	Unprotected	0/0	0/0	Ether

Fabric ID: B

VIF	vNIC	Link State	Oper State	Prot State	Prot Role	Admin Pin	Oper Pin	Transport
1103	eth1	Up	Active	Active	Primary	1/17	1/17	Ether
1106	eth0	Up	Active	Passive	Backup	1/17	1/17	Ether
1107	fc1	Up	Active	No Protection	Unprotected	0/0	2/1	Fc
9299		Up	Active	No Protection	Unprotected	0/0	0/0	Ether



Fibre Channel Path Trace (cont.)

```
UCS-CiscoLive-A(nxos)# show fex detail
FEX: 1 Description: FEX0001 state: Online
<TRUNCATED>
```

```
pinning-mode: static Max-links: 1
Fabric port for control traffic: Eth1/11
Fabric interface state:
  Eth1/11 - Interface Up. State: Active
  Eth1/12 - Interface Up. State: Active
  Eth1/13 - Interface Up. State: Active
  Eth1/14 - Interface Up. State: Active
```

IOM to FI port status

Fex Port	State	Fabric Port	Primary
Fabric			
Eth1/1/1	Up	Eth1/11	Eth1/11
Eth1/1/2	Up	Eth1/12	Eth1/11
Eth1/1/5	Down	Eth1/11	Eth1/11
Eth1/1/6	Down	Eth1/12	Eth1/11
Eth1/1/7	Up	Eth1/13	Eth1/11
Eth1/1/8	Down	Eth1/14	Eth1/11
Eth1/1/9	Up	Eth1/11	Eth1/11

Mezzanine to FI server port
Chassis1/Nic0/Slot#
Eth1/1/7 is
Chassis 1, Nic 0, Slot 7

Fibre Channel Path Trace (cont.)

```
UCS-CiscoLive-A(nxos)# show interface ethernet 1/12 fex-  
  intf
```

Fabric Interface	FEX Interfaces
Eth1/12	Eth1/1/2 Eth1/1/6

Alternate query of FI interface and attached mezz/IOM interfaces

```
UCS-CiscoLive-A(nxos)# show interface ethernet 1/12  
Ethernet1/12 is up
```

```
  Hardware: 1000/10000 Ethernet, address: 000d.ecd3.5dd3  
    (bia 000d.ecd3.5dd3)  
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,  
    reliability 255/255, txload 1/255, rxload 1/255  
  Encapsulation ARPA  
  Port mode is fex-fabric
```

Status of FI interface

Fibre Channel Path Trace (cont.)

```
UCS-CiscoLive-A(nxos) # show interface brief | grep -B 4 veth
```

Ethernet Interface	VLAN	Type	Mode	Status	Reason
veth767	1	eth	trunk	up	none
veth776	1	eth	trunk	up	none
veth8960	100	eth	access	up	none
veth8969	100	eth	access	up	none

Interface	Vsan	Admin Mode	Admin Trunk Mode	Status	SFP	Oper Mode
vfc768	100	F	off	up	--	F
vfc777	100	F	off	up	--	F
vfc780	100	F	off	up	--	F

Check FI interface status against the veth's defined in the Service Profile

Check FI interface status against the vfc's defined in the Service Profile

Fibre Channel Path Trace (cont.)

```
UCS-CiscoLive-A(nxos) # show interface brief
```

Ethernet Interface	VLAN	Type	Mode	Status	Reason
Eth1/1/1	1	eth	access	up	none
Eth1/1/2	1	eth	access	up	none
Eth1/1/3	1	eth	access	down	Administratively down
Eth1/1/4	1	eth	access	down	Administratively down
Eth1/1/5	1	eth	trunk	down	Link not connected
Eth1/1/6	1	eth	trunk	down	Link not connected
Eth1/1/7	1	eth	trunk	up	none
Eth1/1/8	1	eth	access	down	Link not connected
Eth1/1/9	4044	eth	trunk	up	none

Check FI interface status against the mezzanine ports

Mgmt vlan 4044

Adapter Debug attach-mcp

C-series

```
ucs-c220-m3# connect debug-shell <1|2> ← Only certain C-series will support 2 VICs
```

B-series

```
ucs-b-series-m3# scope adapter 1/5/1 ← e.g. Blade 5
```

```
adapter 2/2/1 # connect
```

```
adapter (top):1# <TAB|?>
```

```
attach-fls      attach-mcp      estat           exit            help
history         phy-read        show-acltab     show-fru        show-fwdtab
show-log        show-macstats   show-pcisw
```

```
adapter (top):1# attach-mcp
```

```
adapter (mcp):1# <TAB>
```

```
adv_uifetscfg      amp-dump          amp-env
amp-stats          bmc_chan          bmc_env
bmc_macs           cfgblk            dcbx-env
dcbx-port          dcem-env          dcem-macstats
dcem-port          dcem-showlinkval def_uifetscfg
exit              fipd              fwdtab-show
fwdtab-tcamreg    help              history
iscsi_get_config  iscsi_ping        lif
lifstats          lifstats_lifbase  lifstats_lifdelta
lifstats_logical_uplink lifstats_uifbase  lldp-env
```

Adapter Debug attach-mcp

lldp-port	ncsi-dump	ncsi-stats
ptifs	uif	uifcfg
uifenv	uifetscfg	uiflldpcnt
uiflldppktinfo	uifportprofile	uifportstatus
uifqoscfg	vic-ccstats	vic-dump
vic-env	vic-mstats	vic-stats
vic_enum	vic_get_negotiation	vic_stats_get
vicapp	vicappstats	vicapptimer
vif	vif_long	viflist
viflist_stats	vnic	vnicfind
vnicl	vnicpci	vnicpcibr

- Attach-mcp - Command syntax continued

Adapter Debug attach-mcp

```
adapter (mcp):1# vnic
vnic id      : internal id of vnic, use for other vnic cmds
vnic name/mac : ucsd provisioned name (-n) or mac address (-m)
vnic type    : enet=ethernet, enet_pt=dynamic ethernet, fc=fcoe
vnic bb:dd.f : host pci bus/device/function id
vnic state   : state of vnic
lif          : internal logical if id, use for other lif/vif cmds
lif state    : state of lif
vif uif      : bound uplink 0 or 1, =:primary, -:secondary, >:current
vif ucsd     : ucsd id for this vif
vif idx      : switch id for this vif
vif vlan     : default vlan for traffic
vif state    : state of vif
```

```
-----
id  name      v n i c      l i f      v i f
   name      type    bb:dd.f state lif state uif  ucsd  idx vlan state
-----
 6 eth0      enet    0e:00.0 UP    2 UP  =>0    0    ce 183 UP
 7 eth1      enet    0f:00.0 UP    3 UP  =>1    0    ce 183 UP
 8 fc0       fc      10:00.0 UP    4 UP  =>0    0    ce 100 UP
 9 fc1       fc      11:00.0 UP    5 UP  =>1    0    ce 200 UP
10 mgmt      mgmt    0b:00.0 INIT
```

Adapter Debug attach-fls

```
adapter (mcp):5# exit
adapter (top):2# attach-fls
adapter (fls):1#
d          exit          fwactive  fwcqs      fwexch     fwlif      fwvnic      help
history   lif            login      lunlist    lunmap     vnic
adapter (fls):3# vnic
-----
vnic ecpu type state   lif
-----
8     1     fc   active  4
9     2     fc   active  5
adapter (fls):1# lunmap
lunmapid: 0  port_cnt: 1
  lif_id: 4
    PORTNAME          NODENAME          LUN          PLOGI
    50:0a:09:86:86:78:39:66  00:00:00:00:00:00:00  0019000000000000  Y
lunmapid: 0  port_cnt: 1
  lif_id: 5
    PORTNAME          NODENAME          LUN          PLOGI
    50:0a:09:85:86:78:39:66  00:00:00:00:00:00:00  0019000000000000  Y
```

PLOGI – fc port login
LunID = 0x19 (25 dec)

Adapter Debug attach-fls Working example

```
adapter (fls):2# lunlist
```

```
vnic : 8 lifid: 4
```

```
- FLOGI State : flogi est (fc_id 0x500000)
```

```
- PLOGI Sessions
```

```
- WWNN 50:0a:09:86:86:78:39:66 WWPN 50:0a:09:86:86:78:39:66 fc_id 0x640004
```

```
- LUN's configured (SCSI Type, Version, Vendor, Serial No.)
```

```
  LUN ID : 0x0019000000000000 (0x0, 0x4, NETAPP , C5CKe4iqMM-A)
```

```
- REPORT LUNs Query Response
```

```
  LUN ID : 0x0019000000000000
```

```
- Nameserver Query Response
```

```
- WWPN : 50:0a:09:86:86:78:39:66
```

Target and initiator have an FCID value



Configured and actual LunID match



NS target wwpn response matches configured



Configured LunID does not match Array

- Incorrect LunID
- Configured Boot LunID = dec 19 (0x25)
- Array has masked LunID = dec 5 (0x05)

CIMC configured LunID

```
adapter (fls):2# lunlist
vnic : 8 lifid: 4
- FLOGI State : flogi est (fc_id 0x500000)
- PLOGI Sessions
  - WWNN 50:0a:09:86:86:78:39:66 WWPN 50:0a:09:86:86:78:39:66 fc_id 0x640004
    - LUN's configured (SCSI Type, Version, Vendor, Serial No.)
      LUN ID : 0x0025000000000000 access failure
    - REPORT LUNs Query Response
      LUN ID : 0x0005000000000000
- Nameserver Query Response
  - WWPN : 50:0a:09:86:86:78:39:66
```

Can't access configured lun

LunID that is configured on the array

HBA initiator not configured in Array

```
zone name Z_c220m3_6_v100 vsan 100
* fcid 0x640004 [pwwn 50:0a:09:86:86:78:39:66]
* fcid 0x500000 [pwwn 20:00:e8:b7:48:4d:ca:61]
```

```
adapter (fls):1# lunmap
lunmapid: 0 port_cnt: 1
lif_id: 4
```

PORTNAME	NODENAME	LUN	PLOGI
50:0a:09:86:86:78:39:66	00:00:00:00:00:00:00:00	0019000000000000	Y

```
adapter (fls):2# lunlist
```

```
vnic : 8 lifid: 4
```

```
- FLOGI State : flogi est (fc_id 0x500000)
```

```
- PLOGI Sessions
```

```
- WWNN 50:0a:09:86:86:78:39:66 WWPN 50:0a:09:86:86:78:39:66 fc_id 0x640004
```

```
- LUN's configured (SCSI Type, Version, Vendor, Serial No.)
```

```
  LUN ID : 0x0019000000000000 access failure
```

```
- REPORT LUNs Query Response
```

```
- Nameserver Query Response
```

```
- WWPN : 50:0a:09:86:86:78:39:66
```

Can't
access lun

No LunID in
response

SAN mis-configuration

```
show zoneset active
```

```
<snip>
```

```
zone name Z_c220m3_6_v100 vsan 100
  pwnn 50:0a:09:86:86:78:39:66
  * fcid 0x500000 [pwnn 20:00:e8:b7:48:4d:ca:61]
```

```
adapter (fls):1# lunmap
lunmapid: 0 port_cnt: 1
lif_id: 4
```

PORTNAME	NODENAME	LUN	PLOGI
50:0a:09:86:86:78:39:66	00:00:00:00:00:00:00:00	0019000000000000	N

```
adapter (fls):2# lunlist
```

```
vnic : 8 lifid: 4
```

```
- FLOGI State : flogi est (fc_id 0x500000)
```

```
- PLOGI Sessions
```

```
- WWNN 50:0a:09:86:86:78:39:66 WWPN 50:0a:09:86:86:78:39:66 fc_id 0x000000
```

```
- LUN's configured (SCSI Type, Version, Vendor, Serial No.)
```

```
  LUN ID : 0x0019000000000000 access failure
```

```
- REPORT LUNs Query Response
```

```
- Nameserver Query Response
```

Target not logged in

Target plogi failed

Initiator logged in

No target fcid

Can't access lun

No LunID in response

Nameserver Query empty, target is not in NS

SAN error, wrong target wwpn in CIMC

```
zone name Z_ca73_c1b2_v100 vsan 100
* fcid 0x640004 [pwwn 50:0a:09:86:86:78:39:66]
* fcid 0x500000 [pwwn 20:00:e8:b7:48:4d:ca:61]
```

86:86 is configured in the zoneset

```
adapter (fls):1# lunmap
lunmapid: 0 port_cnt: 1
lif_id: 4
```

77:86 is configured in CIMC for the P81e

plugi failed

PORTNAME	NODENAME	LUN	PLOGI
50:0a:09:77:86:78:39:66	00:00:00:00:00:00:00:00	0019000000000000	N

```
adapter (fls):2# lunlist
vnic : 8 lifid: 4
```

P81e flogi'd

No target wwpn

```
- FLOGI State : flogi est (fc_id 0x500000)
- PLOGI Sessions
```

```
- WWNN 50:0a:09:77:86:78:39:66 WWPN 50:0a:09:77:86:78:39:66 fc_id 0x000000
- LUN's configured (SCSI Type, Version, Vendor, Serial No.)
```

```
LUN ID : 0x0019000000000000 access failure
```

Can't access lun

```
- REPORT LUNs Query Response
```


```
- Nameserver Query Response
- WWPN : 50:0a:09:86:86:78:39:66
```

No LunID in response

UCS EHM Ethernet Data Path

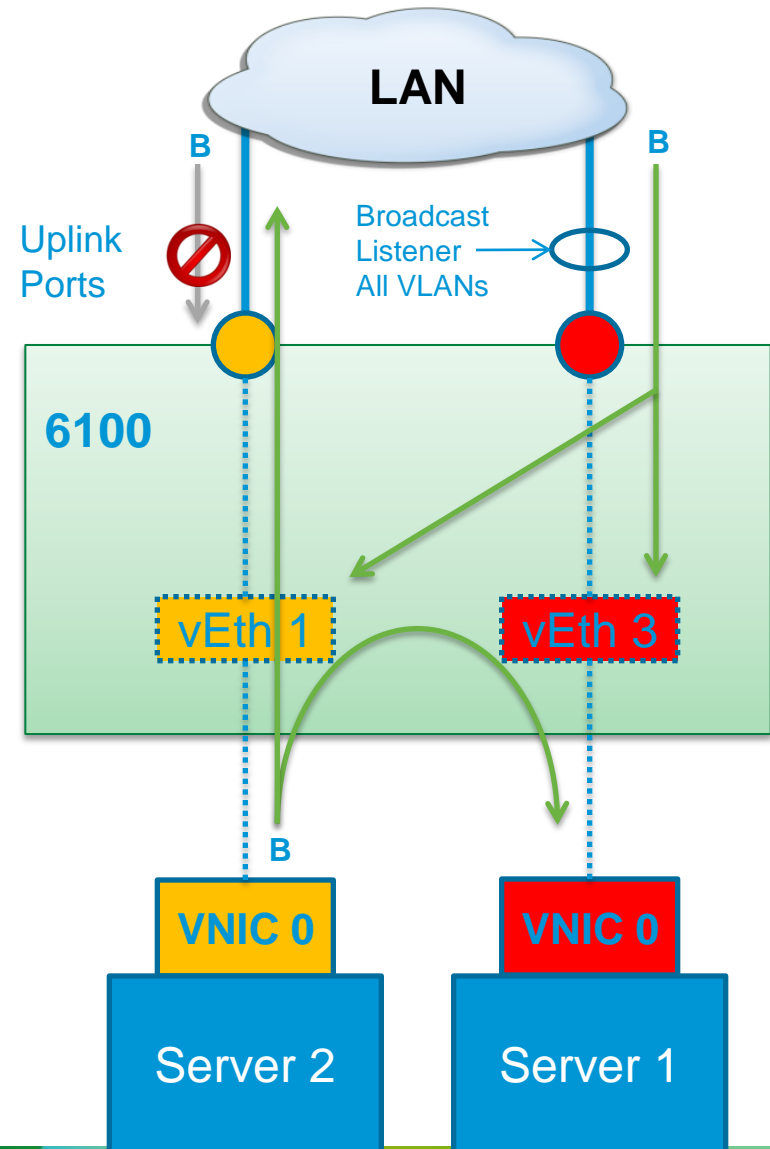
End Host Mode Unicast Forwarding

- Server to server traffic on the same VLAN is locally switched
- Uplink port to Uplink port traffic not switched
- Each server link is pinned to an uplink port / port-channel
- Network to server unicast traffic is forwarded to server only if it arrives on pinned uplink port. This is termed as the Reverse Path Forwarding—(RPF) check
- Packet with source MAC belonging to a server received on an uplink port is dropped (Deja-Vu Check)




End Host Mode Multicast Forwarding

- **Broadcast traffic is pinned on exactly one uplink port** (or port-channel) i.e., it is dropped when received on other uplinks
- All multicast groups are pinned to same uplink port (port-channel)
- Server to server multicast traffic is locally switched
- RPF and deja-vu check also applies for multicast traffic



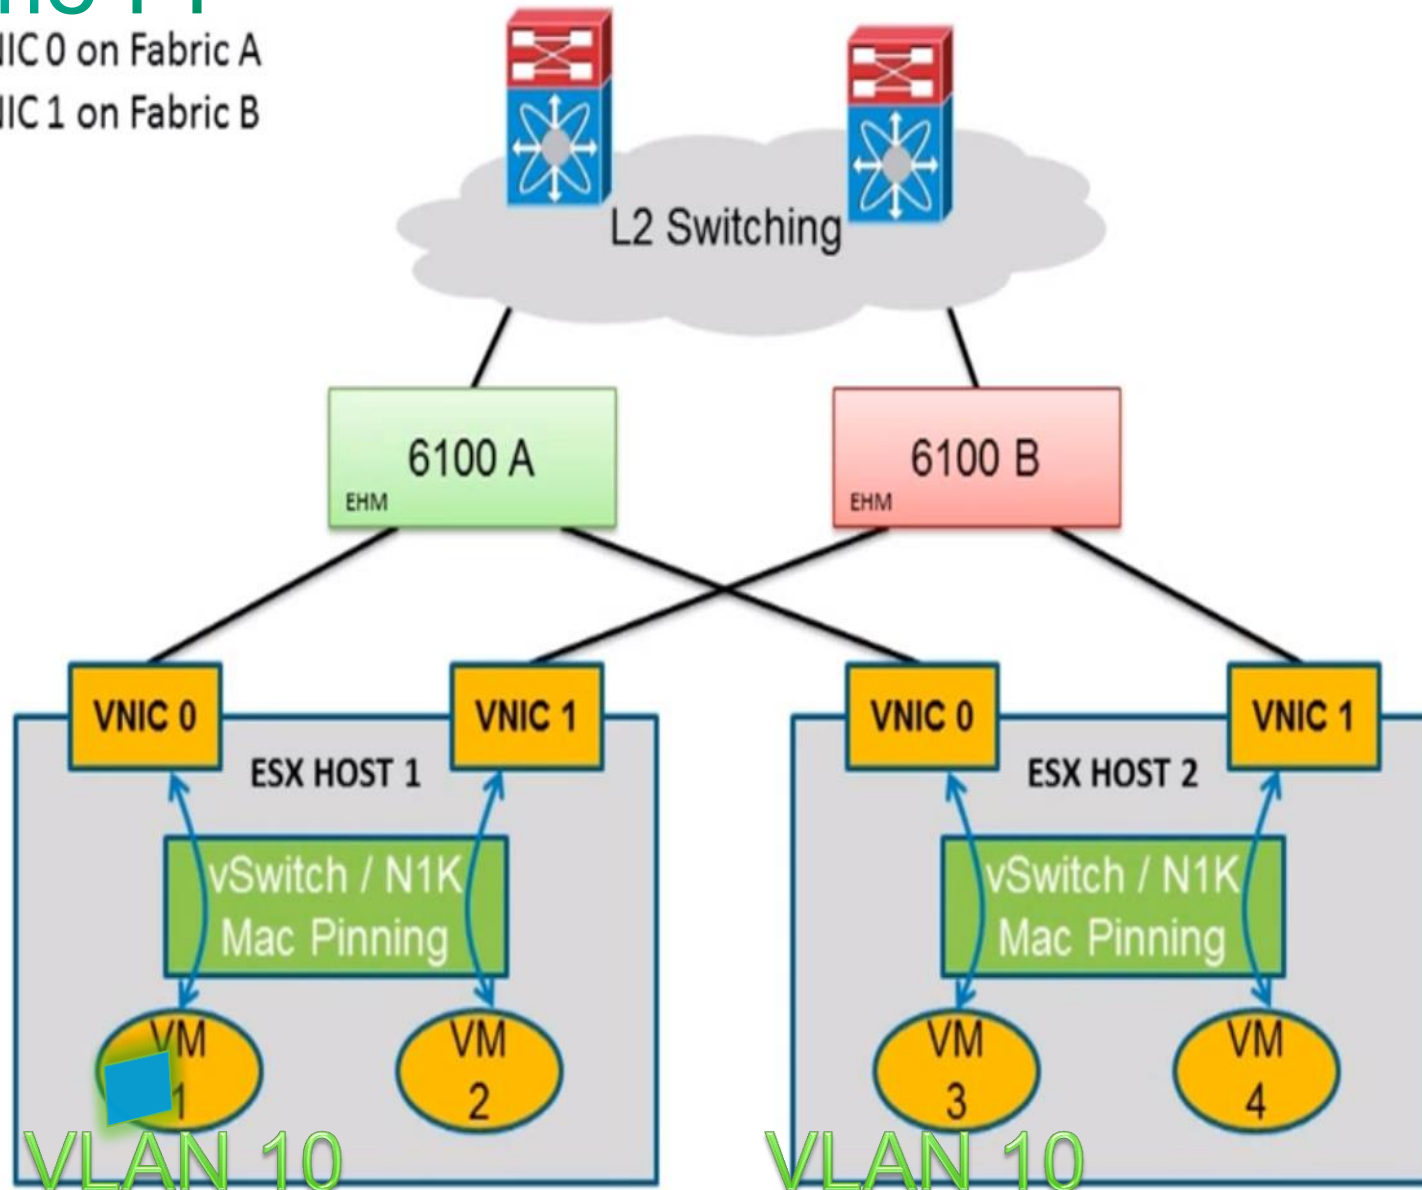
2 VMs on same vSwitch/VEM same VLAN

- VNIC 0 on Fabric A
- VNIC 1 on Fabric B




2 VMs on different host, same VLAN, same FI

- VNIC 0 on Fabric A
- VNIC 1 on Fabric B




2 VMs same VLAN, different host, different FI

- VNIC 0 on Fabric A
- VNIC 1 on Fabric B




2 VMs different VLAN, same host, different FI

- VNIC 0 on Fabric A
- VNIC 1 on Fabric B



2 VMs different VLAN, different host, different FI

- VNIC 0 on Fabric A
- VNIC 1 on Fabric B



Pinning

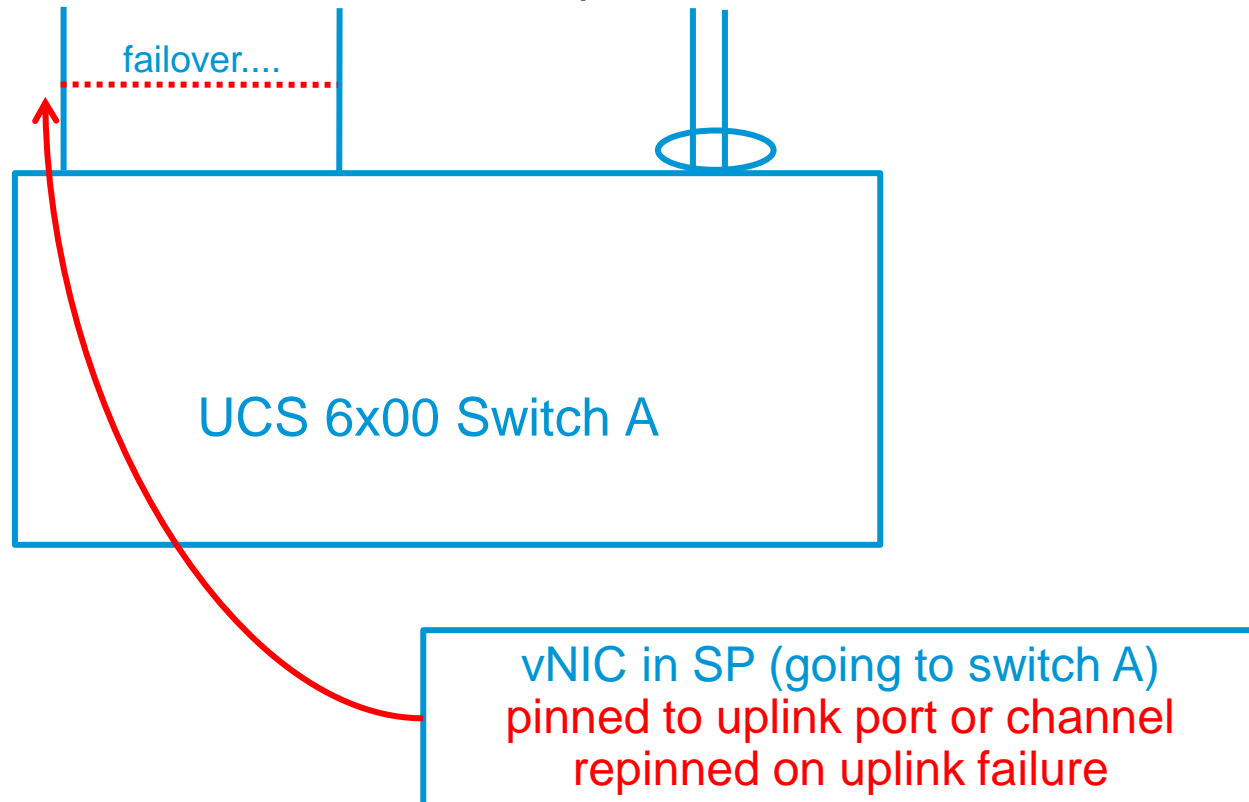
Automatic Pinning

- UCS chooses uplink port or port channel for vNIC

Choice *not ever made on basis of VLAN*


Remember all uplinks support same VLANs

Assigns different vNICs to different uplinks in round-robin fashion



Static (Manual Pinning)

Pin Group: ZippythePinGroup



UCS 6x00 Switch A

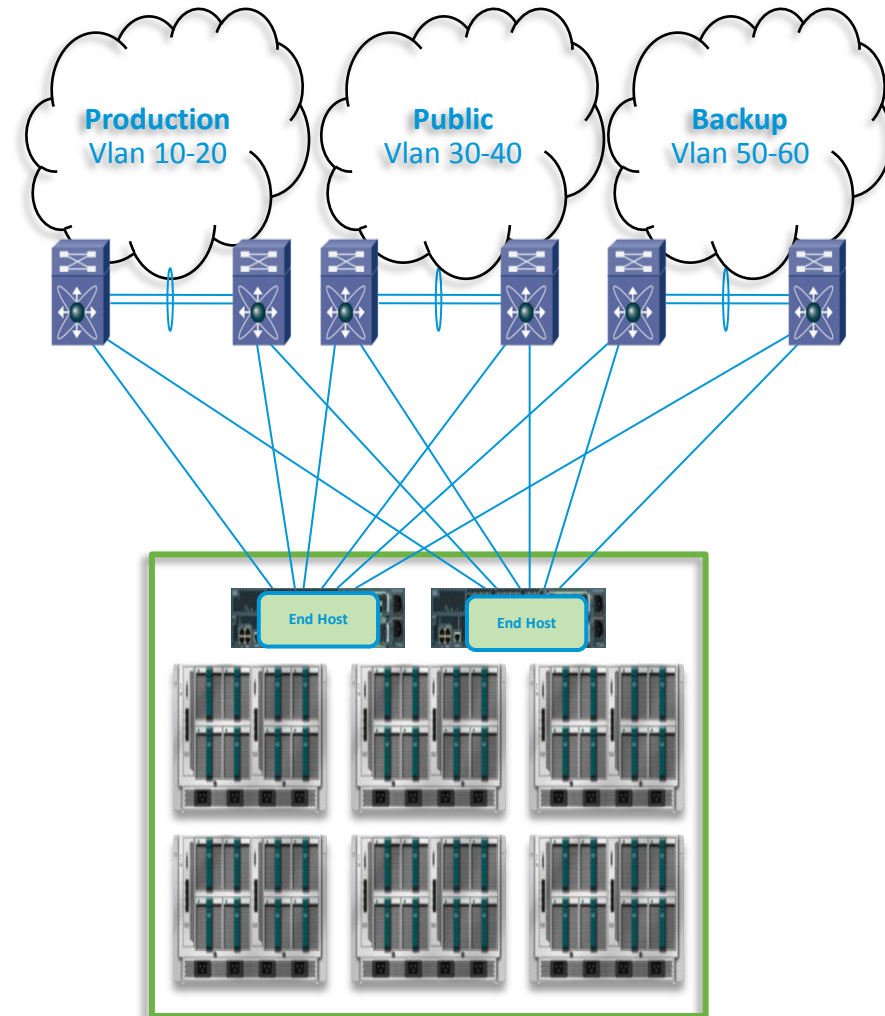
UCS 6x00 Switch B

vNIC in SP (going to switch A)
pinned to port/port channel on A
can use no other uplink on A
can failover to B

DLU (Disjoint L2 Upstream)

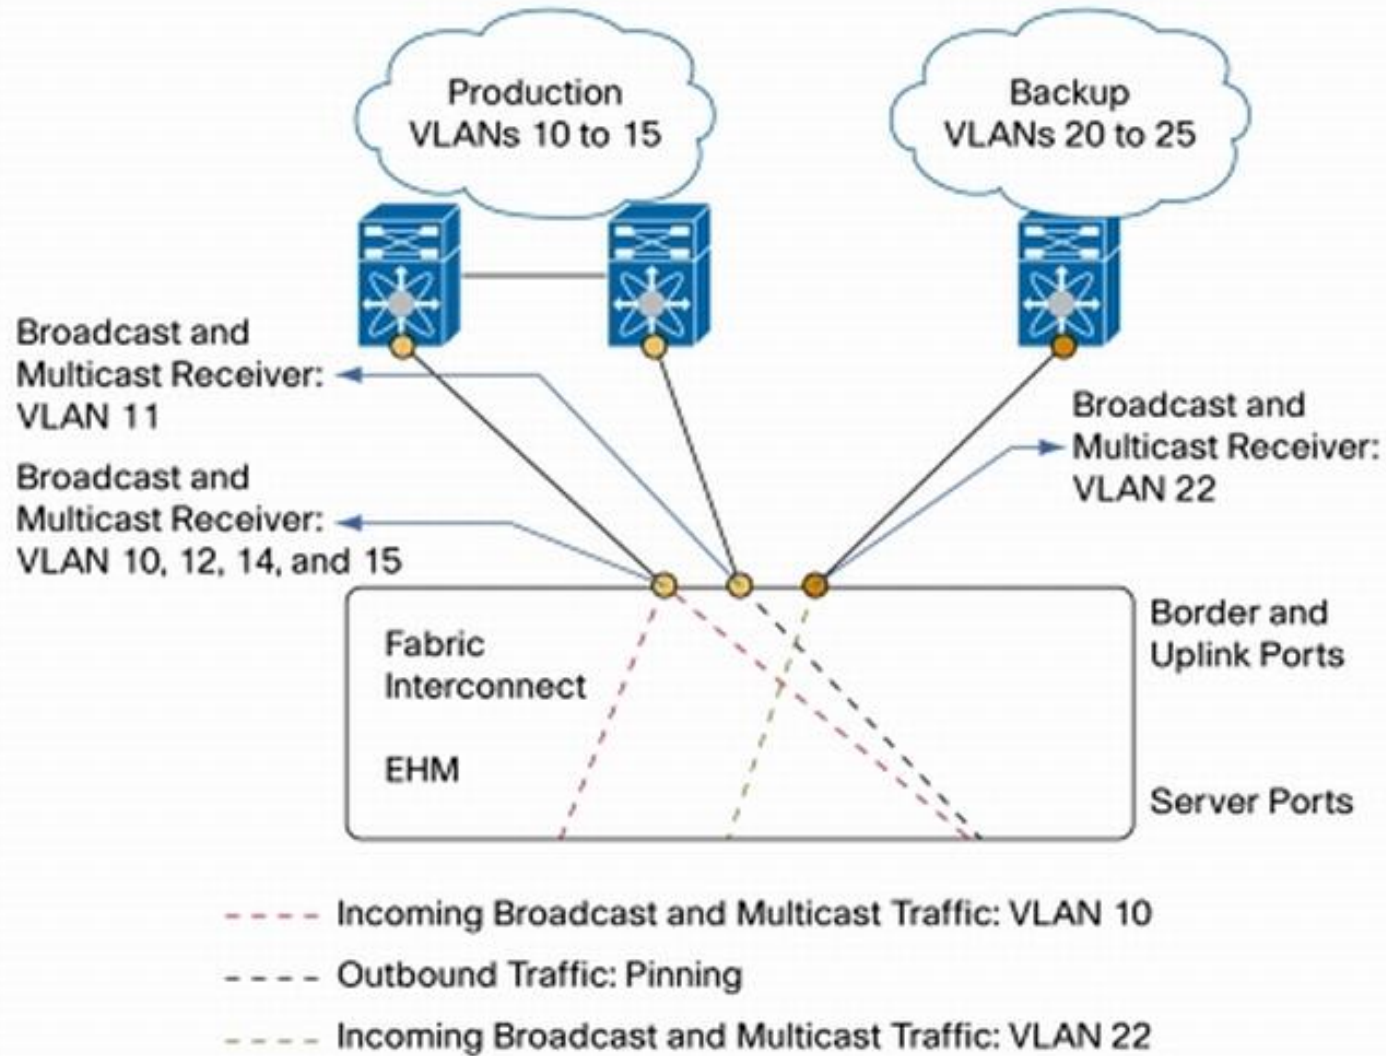
Disjoint L2 Upstream

- EHM built on the premise that the L2 upstream is NOT disjoint.
- Incoming broadcast/multicast received only on 1 uplink for ALL VLANs (pre 2.0)
- DR selected per-VLAN (>2.0 version)



Recommendation:
< 2.0 release Switch Mode
> 2.0 release End Host Mode

Disjoint L2



Pre 2.0

- Spanning Tree Protocol is not run on both the uplink ports
- MAC address learning occurs only on the server ports and appliance ports
- MAC address aging is not supported; MAC address changes are fully supported
- All uplink ports connect to the same Layer 2 cloud
- A single Ethernet uplink port (or PortChannel) on each fabric interconnect is chosen to be the broadcast and multicast traffic receiver for all VLANs, and incoming and broadcast traffic is dropped on the other uplinks. This port is called the G-pinned.

```
UCS-A(nxos)# show platform software enm internal info global | grep -A 6 `Global Params`
```

Other Global Params:

```
broadcast-if 0x88c1f04(Ethernet1/16)
```

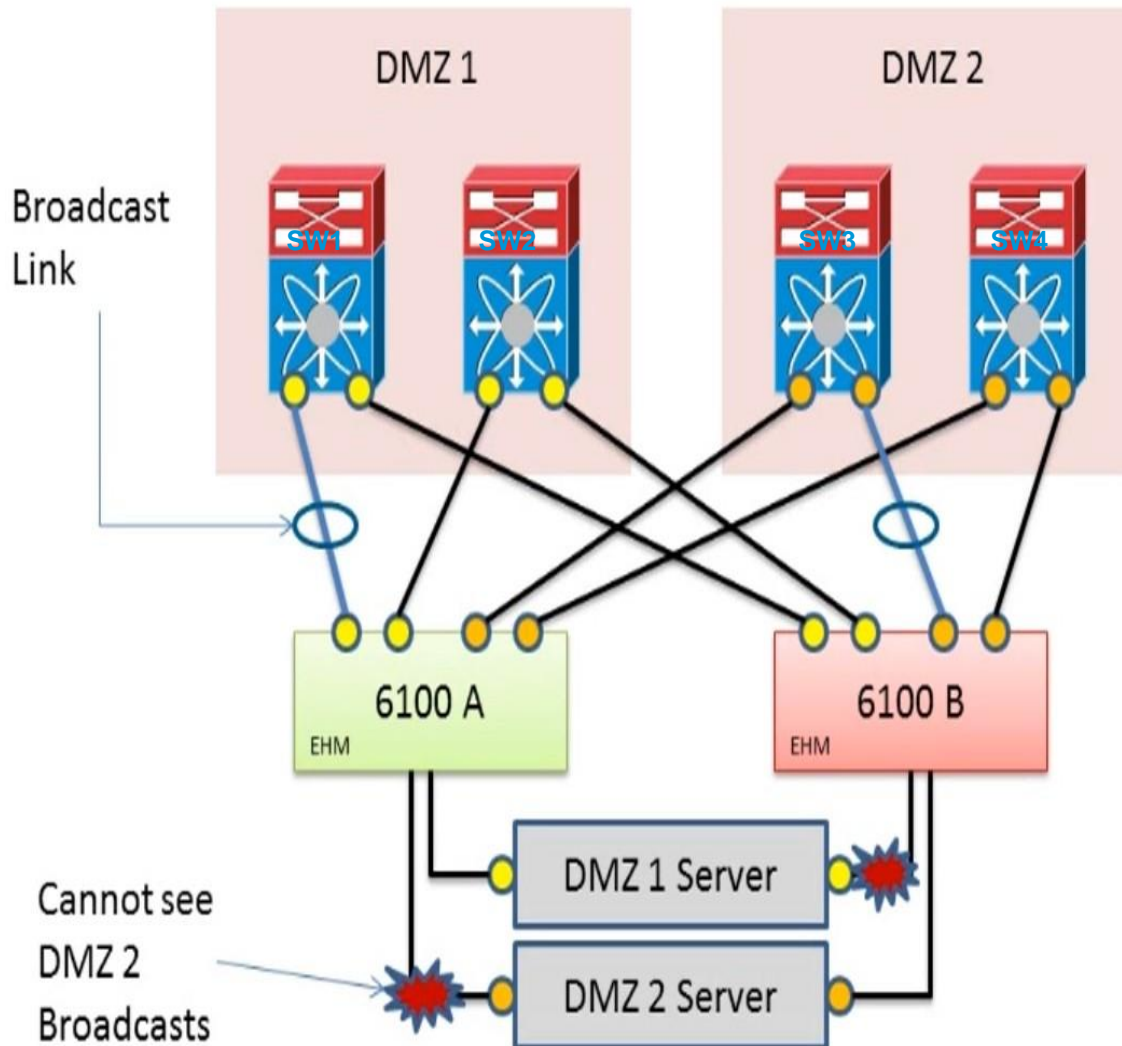
```
multicast-if 0x88c1f04(Ethernet1/16)
```

```
ip_multicast-if 0x88c1f04(Ethernet1/16)
```

```
end-host-mode: Enabled
```


EHM – Pre 2.0 Disjoint L2 Issue

Each 6100 picks ONE uplink for Broadcast/Mcst Processing



Post 1.4

- The ability to selectively assign VLANs to uplinks
- vNIC pinning decision based on - VLAN membership on uplink ports
- Allocating a designated broadcast/multicast receiver on a per VLAN rather than global basis
- Max 31 disjoint Layer 2 domains supported (Theoretical)

```
UCS (nxos)# show platform software enm internal info vlandb id 11  
vlan_id 11
```

```
-----
```

```
Designated receiver: Eth1/13
```

```
Membership:
```

```
Eth1/13
```

Overlapping VLANs are not supported.

Disjoint L2 Configuration Overview

- Launch LAN Uplink Manager

The screenshot displays the Cisco Unified Computing System Manager - CA-MAM interface. The top left shows a 'Fault Summary' with 2 critical, 19 warning, 7 error, and 35 info faults. The 'LAN' tab is selected in the top navigation bar. The left sidebar shows a tree view of the LAN configuration, with 'LAN Cloud' highlighted. The main content area shows the 'LAN Cloud' configuration page, with the 'Ports and Port Channels' tab selected. The 'Uplink 1' entry is highlighted in the table. Below the table, there are 'Enable' and 'Disable' buttons, and a red box containing the text: 'To configure the LAN, launch the LAN Uplinks Manager.' The bottom status bar shows 'Logged in as admin@10.29.177.15' and 'System Time: 2011-06-27T18:21'.

1 LAN

2 LAN Cloud

3 Enable Disable

To configure the LAN, launch the **LAN Uplinks Manager**.

Disjoint L2 Configuration Overview

- Associate VLANs with Uplinks

LAN Uplinks | VLANs | Server Links | QoS | Global Policies | Faults | Events | FSM

All | Dual Mode | Fabric A | Fabric B | VLAN Manager

Fabric A | Fabric B

Ports and Port Channels

Name	Administrative State
Port Channels	
Fabric A	
Port-Channel 7 (Fabric A)	Enabled
Port-Channel 8 (Fabric A)	Enabled
Uplink Interfaces	
Fabric A	
Eth Interface 1/3	Enabled

VLANs

Name	ID	VLAN Sharing	Native VLAN
VLANs			
VLAN 11 (11)	11	None	
VLAN 12 (12)	12	None	
VLAN 13 (13)	13	None	
VLAN 14 (14)	14	None	
VLAN 15 (15)	15	None	
VLAN 16 (16)	16	None	
VLAN 17 (17)	17	None	
VLAN 18 (18)	18	None	
VLAN 183 (183)	183	None	
VLAN 177 (177)	177	None	
Port-Channel 7 (Fabric A)			
VLAN 178 (178)	178	None	

Add Members to VLAN

The traffic on the VLAN : VLAN177 will now flow only on the selected uplink interface(s), and does not flow on any other uplink interfaces.

OK

Add to VLAN | Remove from VLAN

OK | Apply | Cancel | Help

1. Select the Port-Channel 7 and it will be highlighted

2. Select the VLAN 177 on the right side pane

3. Select Add to VLAN Button

4. Select OK

Disjoint L2 Configuration Overview

Verifying Configuration

- From UCSM CLI – NXOS Context:

`show interface trunk`

- Should expect to see your associated VLANs on the assigned interfaces only.
- Ex. Eth 1/16-17 are Uplinks

Port	Vlans Allowed on Trunk

Eth1/16	1,5-7,1000,3000 ← Only Eth1/16 allows VLAN 5
Eth1/17	1,6-7,1000,3000
Veth705	1000
Veth709	1,1000
<snip>	

Service Profile Circuit Paths

- A server's virtual interface (VIF) mapping can be displayed with the 'show service-profile circuit' command. The VIF number can be used in the NXOS level of the FI to trace a path to the upstream.

```
UCS-B# show service-profile circuit server 1/1
```

```
Service Profile: root/ESXi_3
```

```
Server: 1/1
```

```
Fabric ID: A
```

VIF	vNIC	Link State	Overall Status	Admin Pin	Oper Pin	Transport
41		Unknown	Unknown	0/0	0/0	Unknown
9158		Up	Active	0/0	0/0	Ether
964	eth0	Up	Active	0/0	0/3	Ether
966	fc0	Up	Active	0/0	0/0	Fc

```
Fabric ID: B
```

VIF	vNIC	Link State	Overall Status	Admin Pin	Oper Pin	Transport
42		Unknown	Unknown	0/0	0/0	Unknown
9159		Up	Active	0/0	0/0	Ether
965	eth1	Up	Active	0/0	0/3	Ether
967	fc1	Up	Active	0/0	0/0	Fc

Service Profile Circuit Paths (cont)


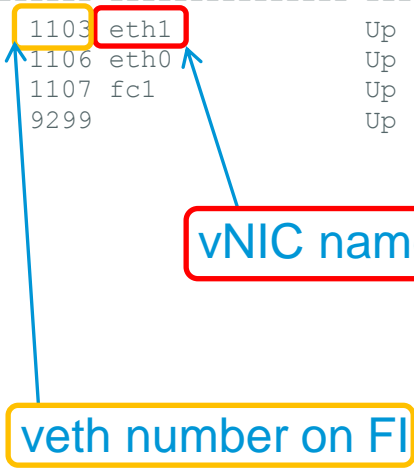
```
UCS-CiscoLive-A# show service-profile circuit server x/y
Service Profile: 530DC/Cehuiyuan38
Server: 2/2
```

Fabric ID: A

VIF	vNIC	Link State	Oper State	Prot State	Prot Role	Admin Pin	Oper Pin	Transport
1104	eth1	Up	Active	Passive	Backup	1/17	1/17	Ether
1105	eth0	Up	Active	Active	Primary	1/17	1/17	Ether
1108	fc0	Up	Active	No Protection	Unprotected	0/0	2/1	Fc
9300		Up	Active	No Protection	Unprotected	0/0	0/0	Ether

Fabric ID: B

VIF	vNIC	Link State	Oper State	Prot State	Prot Role	Admin Pin	Oper Pin	Transport
1103	eth1	Up	Active	Active	Primary	1/17	1/17	Ether
1106	eth0	Up	Active	Passive	Backup	1/17	1/17	Ether
1107	fc1	Up	Active	No Protection	Unprotected	0/0	2/1	Fc
9299		Up	Active	No Protection	Unprotected	0/0	0/0	Ether



Active Border Interfaces

- The interface a VIF will use to reach the upstream network can be determined by analyzing the active border-interfaces on the FI:

```
UCS-B(nxos)# show pinning border-interfaces active
```

Border Interface	Status	SIFs
Eth1/7	Active	Veth729, Veth964
Eth1/8	Active	Veth731
Eth1/9	Active	Veth732
Eth1/16	Active	Veth733
Eth1/17	Active	
Eth1/18	Active	
Eth1/19	Active	Eth2/1/29
Eth1/20	Active	Veth706
Eth1/21	Active	
Eth1/22	Active	

```
Total Interfaces : 10
```


Determining DR for Bcast & Mcast

```
UCS-B(nxos)# show platform software enm internal  
info vlandb [id x] | all
```

```
vlan_id 1
```

```
-----
```

```
Designated receiver: Eth1/8
```

```
Membership:
```

```
Eth1/8 Eth1/9 Eth1/16 Eth1/18 Eth1/19 Eth1/20 Eth1/21 Eth1/22  
Eth1/17
```

```
vlan_id 5
```

```
-----
```

```
Designated receiver: Eth1/16
```

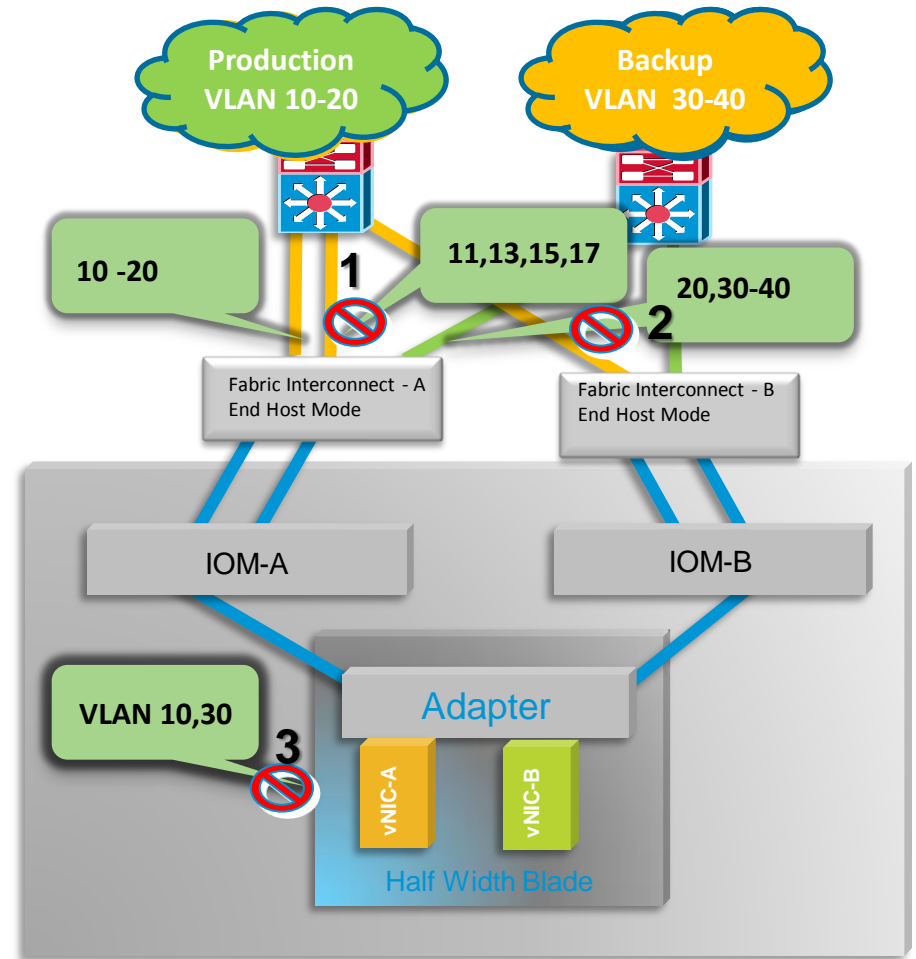
```
Membership:
```

```
Eth1/16
```

```
<snip>
```

Disjoint L2 – Do NOT

1. configure different range of VLANs across uplinks of a disjoint network
2. overlap VLANs across L2 networks
3. configure vNIC with VLANs across multiple L2 networks. vNIC does not get pinned if the uplink VLAN membership is partial, vNIC stays down



Quiz and Q&A

Q1

Q1. 怎样才可以使vfc1157变成正常的UP状态？下面是log的截取信息。多选。

```
`show npv status`  
npiv is enabled  
disruptive load balancing is disabled  
External Interfaces:  
=====  
Interface: fc2/1, VSAN: 2, FCID: 0xa20000, State: Up  
Number of External Interfaces: 1  
Server Interfaces:  
=====  
Interface: vfc957, VSAN: 2, State: Up  
.....  
Interface: vfc1157, VSAN: 1, State: Waiting for External Interface  
Interface: vfc1180, VSAN: 2, State: Up
```

- A. 将服务器的service profile配置中vfc1157的vsan从1变更为2
- B. 重置vfc1157
- C. 在操作系统上更新vfc HBA的驱动到最新版本
- D. 配置FC UPLINK fc2/1 为trunk模式，并trunk vsan 1和vsan 2

Q2

Q2. SAN Boot失败，下面为VIC的log信息截取，如果修复这个问题使SAN Boot可以正常启动？

```
adapter (fls):1# lunmap
lunmapid: 0 port_cnt: 1
lif_id: 4
  PORTNAME          NODENAME          LUN          PLOGI
  50:0a:09:77:86:78:39:66 00:00:00:00:00:00:00:00 0019000000000000 N
adapter (fls):2# lunlist
vnic : 8 lifid: 4
- FLOGI State : flogi est (fc_id 0x500000)
- PLOGI Sessions
- WWNN 50:0a:09:77:86:78:39:66 WWPN 50:0a:09:77:86:78:39:66 fc_id 0x000000
- LUN's configured (SCSI Type, Version, Vendor, Serial No.)
  LUN ID : 0x0019000000000000 access failure
- REPORT LUNs Query Response
- Nameserver Query Response
- WWPN : 50:0a:09:86:86:78:39:66
```

- A. 在San boot policy中配置Lun ID 参数为 0x0019000000000000
- B. 在上层SAN交换机上重新配置相应的zone并删除成员wwpn 50:0a:09:86:86:78:39:66
- C. 将san boot policy中的target wwpn参数配置变更为50:0a:09:86:86:78:39:66
- D. 更换从UCS到上层SAN交换机之间的线缆

Q3

UCS环境中配置了Eth1/1和Eth1/2为以太网UPLINK，并配置有 3 个 vlans VLAN 10,20,30。Eth1/1 被配置为Vlan10和20的广播组播DR端口。针对Eth1/2没有特殊的DR配置。一台服务器中配置2个以太网网卡分别为 vnic0 和 vnic1，而且针对2个vnic没有特殊的pin group的配置。vNIC0 配置为trunk vlan 10， 20和30. vNIC1 配置为trunk vlan 10 和 30.

Vlan 10， 20 和 30的 DR端口是怎样的？

vNIC0 和 vNIC1的链路状态分别是怎样的？是UP还是Down？并分别pin到那个UPLINK上面？

- A. vlan10和20的DR端口是Eth1/1， vlan30的DR端口是Eth1/1与Eth1/2. vNIC0 是UP并pin到Eth1/1. vNIC1 是UP并pin到Eth1/1
- B. Vlan10和20的DR端口是Eth1/1， Vlan30的DR端口是Eth1/2. vNIC0 是UP并pin到Eth1/1. vNIC1 是UP并pin到Eth1/1
- C. vlan10和20的DR端口是Eth1/1， vlan30的DR端口是Eth1/1与Eth1/2. vNIC0 是UP并pin到Eth1/1. vNIC1 是UP并pin到Eth1/2
- D. vlan10和20的DR端口是Eth1/1， vlan30的DR端口是Eth1/1与Eth1/2. vNIC0 是UP并pin到Eth1/1. vNIC1 是 DOWN.

Q4

uplink eth1/1 and eth1/2, 3 vlans VLAN 10,20,30. Eth1/1 is selected as DR port of Vlan 10 & 20. Eth1/2 is assigned to vln 30 as DR. Vnic0 is soft pin configuration with trunking vln 10,20,30. vnic1 is static pin to eth1/2 and trunking vln 10 and 30. What is the link status of vnic0 and vnic1? Up or Down?

UCS环境中配置了Eth1/1和Eth1/2为以太网UPLINK，并配置有 3 个 vlans VLAN 10,20,30。Eth1/1 被配置为Vlan10和20的广播组播DR端口。Eth1/2被配置为Vlan30的DR端口。一台服务器中配置2个以太网网卡分别为 vNIC0 和 vNIC1。vNIC0没有pin group的配置，并配置为trunk vln 10, 20 和 30. vNIC1 是配置了pin group 并静态pin到Eth1/2，同时trunk vln 10 和 30.

vNIC0 和 vNIC1的链路状态是怎样的？是UP还是Down？

- A. vNIC0是UP状态，但是只能承载vln 10和20的数据，vNIC1是Down的状态
- B. vNIC0是UP状态。vNIC1也是UP状态，但是 vln10在vNIC1上面不能正常通讯
- C. vNIC0 和 vNIC1 都是Down的状态
- D. vNIC0是DOWN状态。vNIC1是UP，但是vln10在vNIC1上面不能正常通讯

Quiz and Q&A 正确答案

- Q1—A D
- Q2—C
- Q3—A
- Q4—D

