
TOMORROW starts here.
cisco
Cisco live!

UCS Networking – Deep Dive

BRKCOM-2003

Sarva Chakravarthi
Cisco Services

Cisco live!

UCS LAN Deep Dive - Agenda

- High-level System Overview
- Fabric Forwarding Mode of Operations
- Uplink Pinning
- Chassis / Fabric Extender
- Server Connectivity Options
- Recommended Topologies
- C-Series Integration

High-level System Overview

System Components: High-level Overview

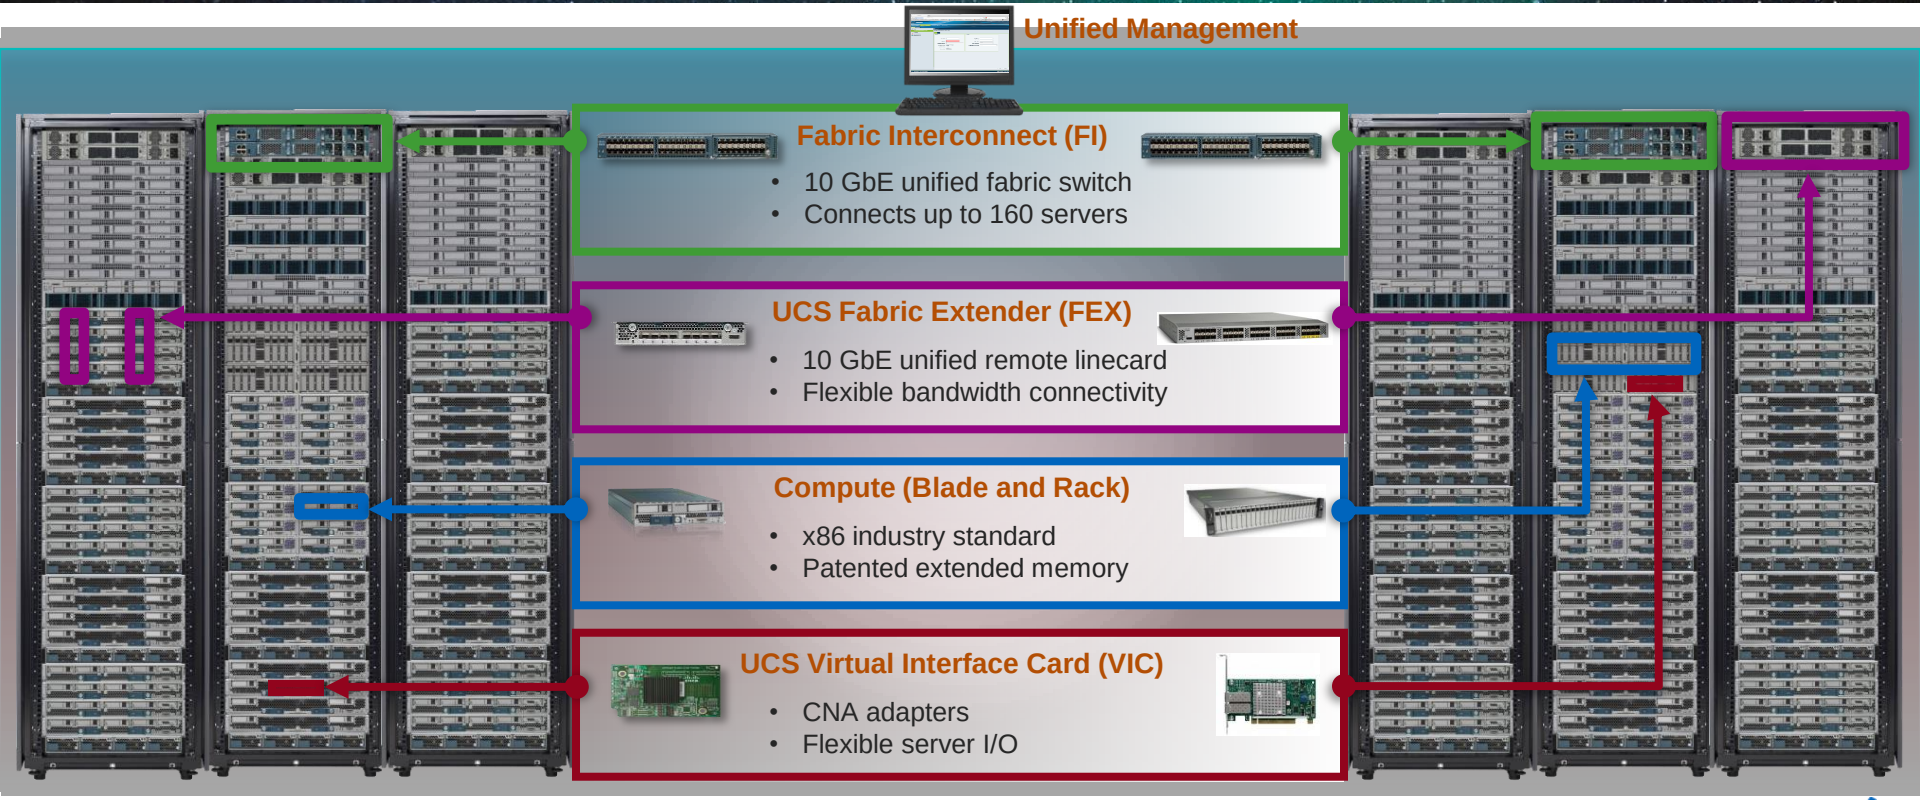

UCS Fabric Portfolio

UCS Fabric Interconnect

UCS Fabric Interconnect – UCS 6248

- 1RU
- 32 unified base ports and 1 expansion slot
- Line rate – 960Gbps

UCS Fabric Interconnect – UCS 6296

- 2RU
- 48 unified base ports and 3 expansion slots
- Line rate – 1920 Gbps

UCS Fabric Extender

UCS IO Module (IOM) – 2204 or 2208

- 4 or 8 10GbE fabric links (to Fabric Interconnect)
- 16 or 32 10GbE server links (to servers)

Nexus 2232PP or Nexus 2232TM

- 8 10GbE fabric links (to Fabric Interconnect)
- 32 10GbE server links (to servers)

UCS VIC Adapters

UCS VIC1240 plus Pass-through (PT) Expansion Card - Blades

- VIC1240: Up to 4 x 10 GbE
- PT: Expands VIC1240 up to 8 x 10GbE
- Up to 256 vPCIe

UCS VIC 1280 - Blades

- Up to 8 x 10GE ports
- Up to 256 vPCIe

UCS VIC 1225 - Racks

- Up to 2 x 10GE ports
- Up to 256 vPCIe

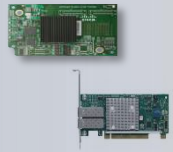

Cisco UCS 6200 Series Fabric Interconnects

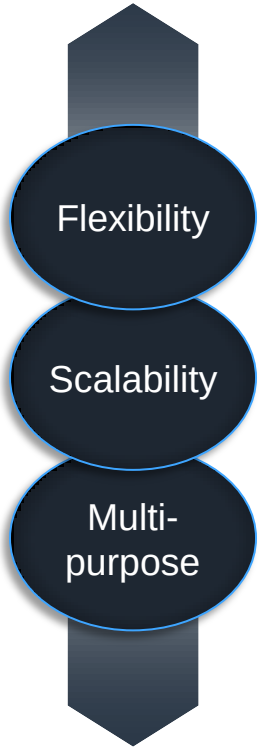

Product Features and Specs	UCS 6248UP
Switch Fabric Throughput	960 Gbps
Switch Footprint	1RU
1 Gigabit Ethernet Port Density	48
10 Gigabit Ethernet Port Density	48
8G Native FC Port Density	48
Port-to-Port Latency	2.0us
# of VLANs	4096*
Layer 3 Ready (future)	✔
40 Gigabit Ethernet Ready (future)	✔
Virtual Interface Support	63 per Downlink
Unified Ports (Ethernet or FC)	✔

*1024 with current 2.1 release

UCS 6248: Unified Ports

Dynamic Port Allocation: Lossless Ethernet or Fibre Channel

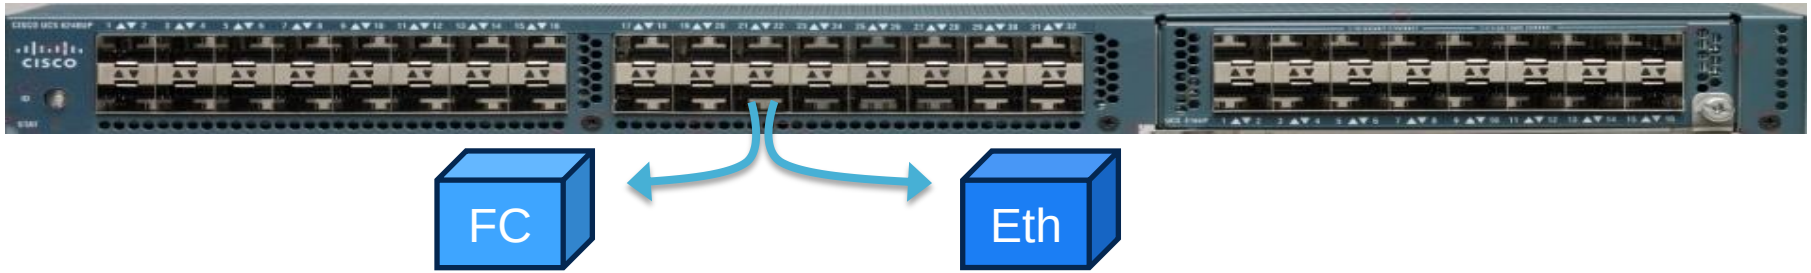

Native Fibre Channel

Lossless Ethernet:
1/10GbE, FCoE, iSCSI, NAS

Benefits

- Simplify switch purchase - remove ports ratio guess work
- Increase design flexibility
- Remove specific protocol bandwidth bottlenecks

Use-cases

- Flexible LAN & storage convergence based on business needs
- Service can be adjusted based on the demand for specific traffic

UCS 6248: Unified Ports

Dynamic Port Allocation: Lossless Ethernet or Fibre Channel

- Ports on the base card or the Unified Port GEM Module can be Ethernet or FC
- Only a continuous set of ports can be configured as Ethernet or FC
- Ethernet Ports have to be the 1st set of ports
- Port type changes take effect after next reboot of switch for Base board ports or power-off/on of the GEM for GEM unified ports.

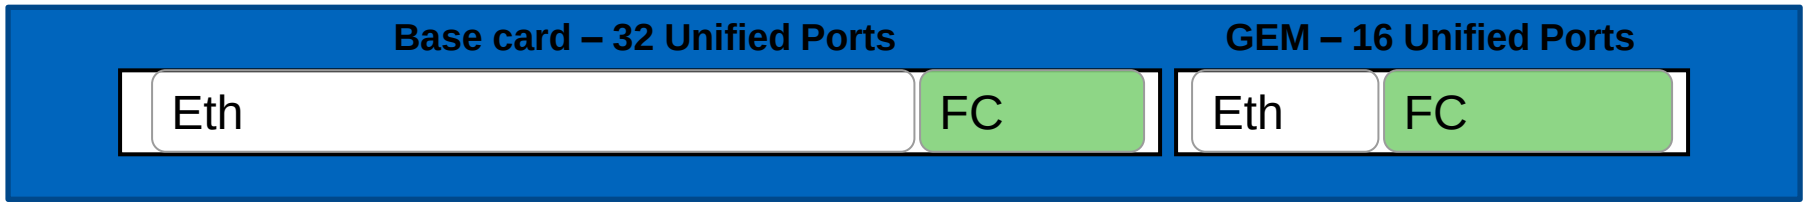

Generic Expansion Module (GEM)

Unified Port GEM for UCS 6200 Series

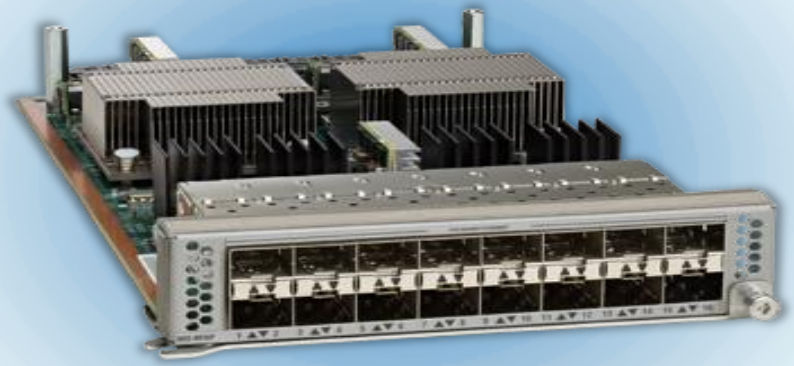

UCS-FI-E16UP

- 16 “Unified Ports”
- Ports can be configured as either Ethernet or Native FC Ports
- Ethernet operations at 1/10 Gigabit Ethernet
- Fibre Channel operations at 8/4/2/1G
- Uses existing Ethernet SFP+ and Cisco 8/4/2G and 4/2/1G FC Optics

Flexibility AND Simplicity

Unified Port Screen

- Configured on a per FI basis
- Slider based configuration
- Reboot is required for the new port personality to take into affect
- Recommendation is to configure GEM card, therefore GEM is only needed to be rebooted

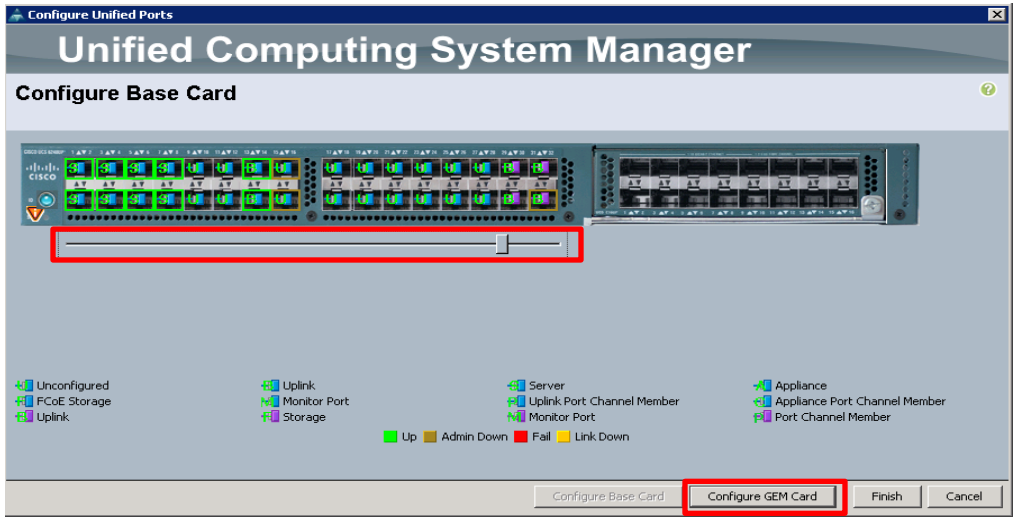

Logical Architecture

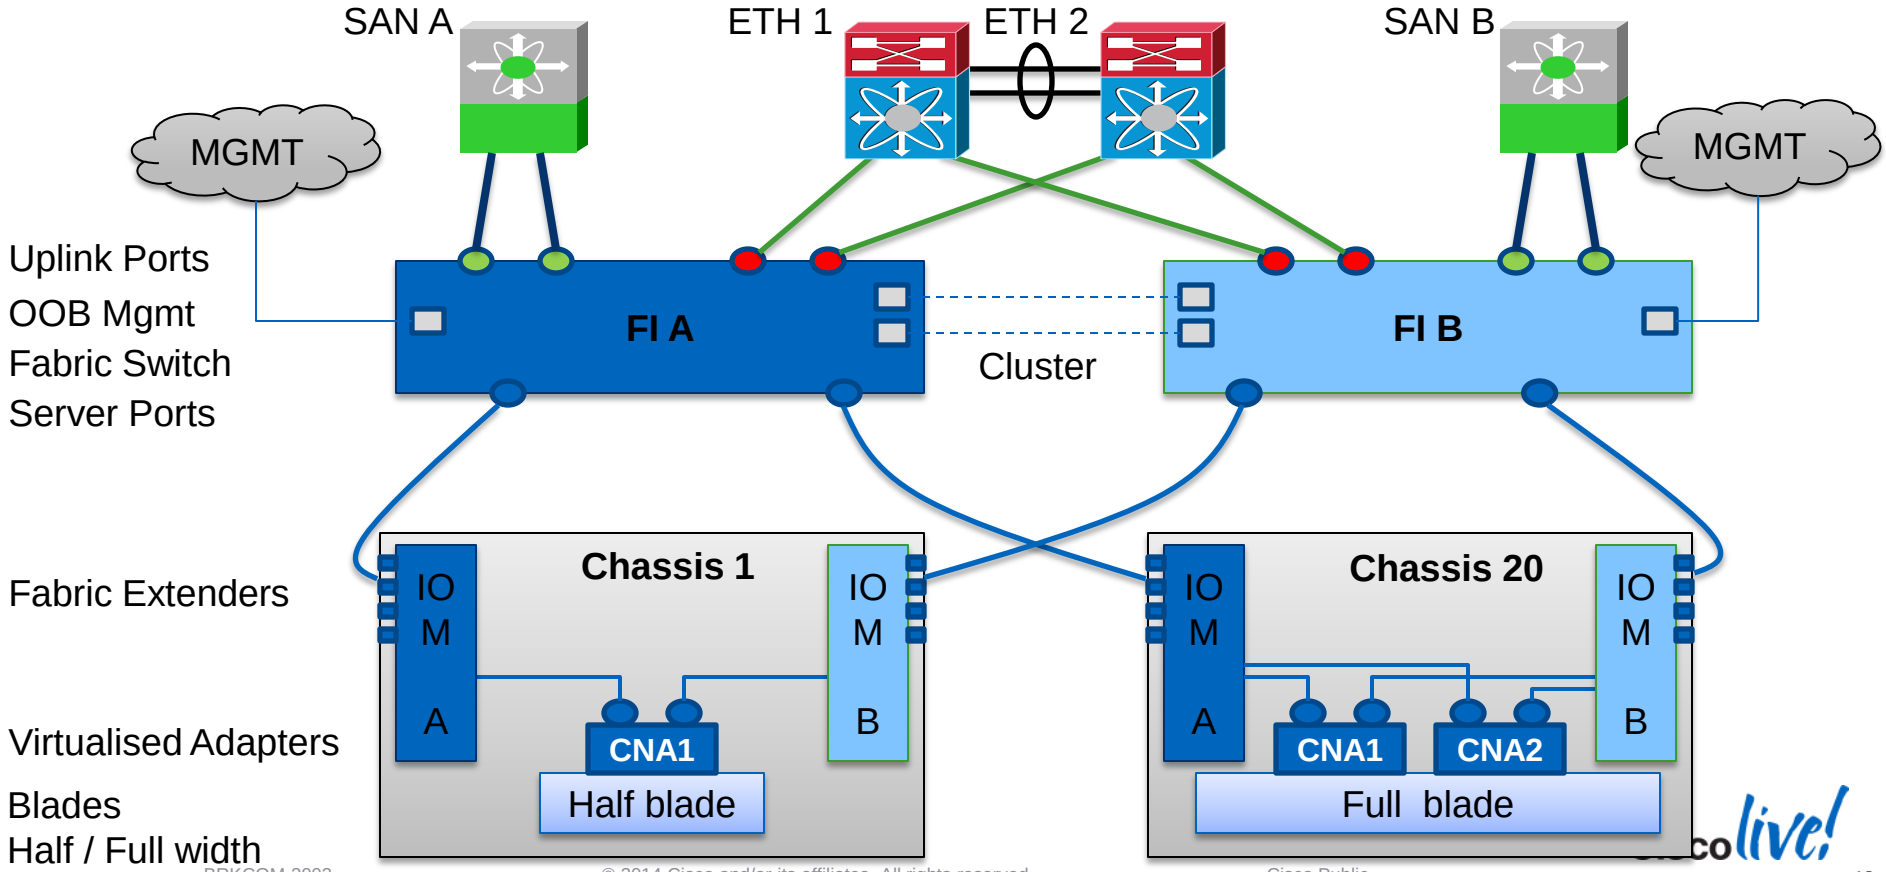

Abstracting the Logical Architecture

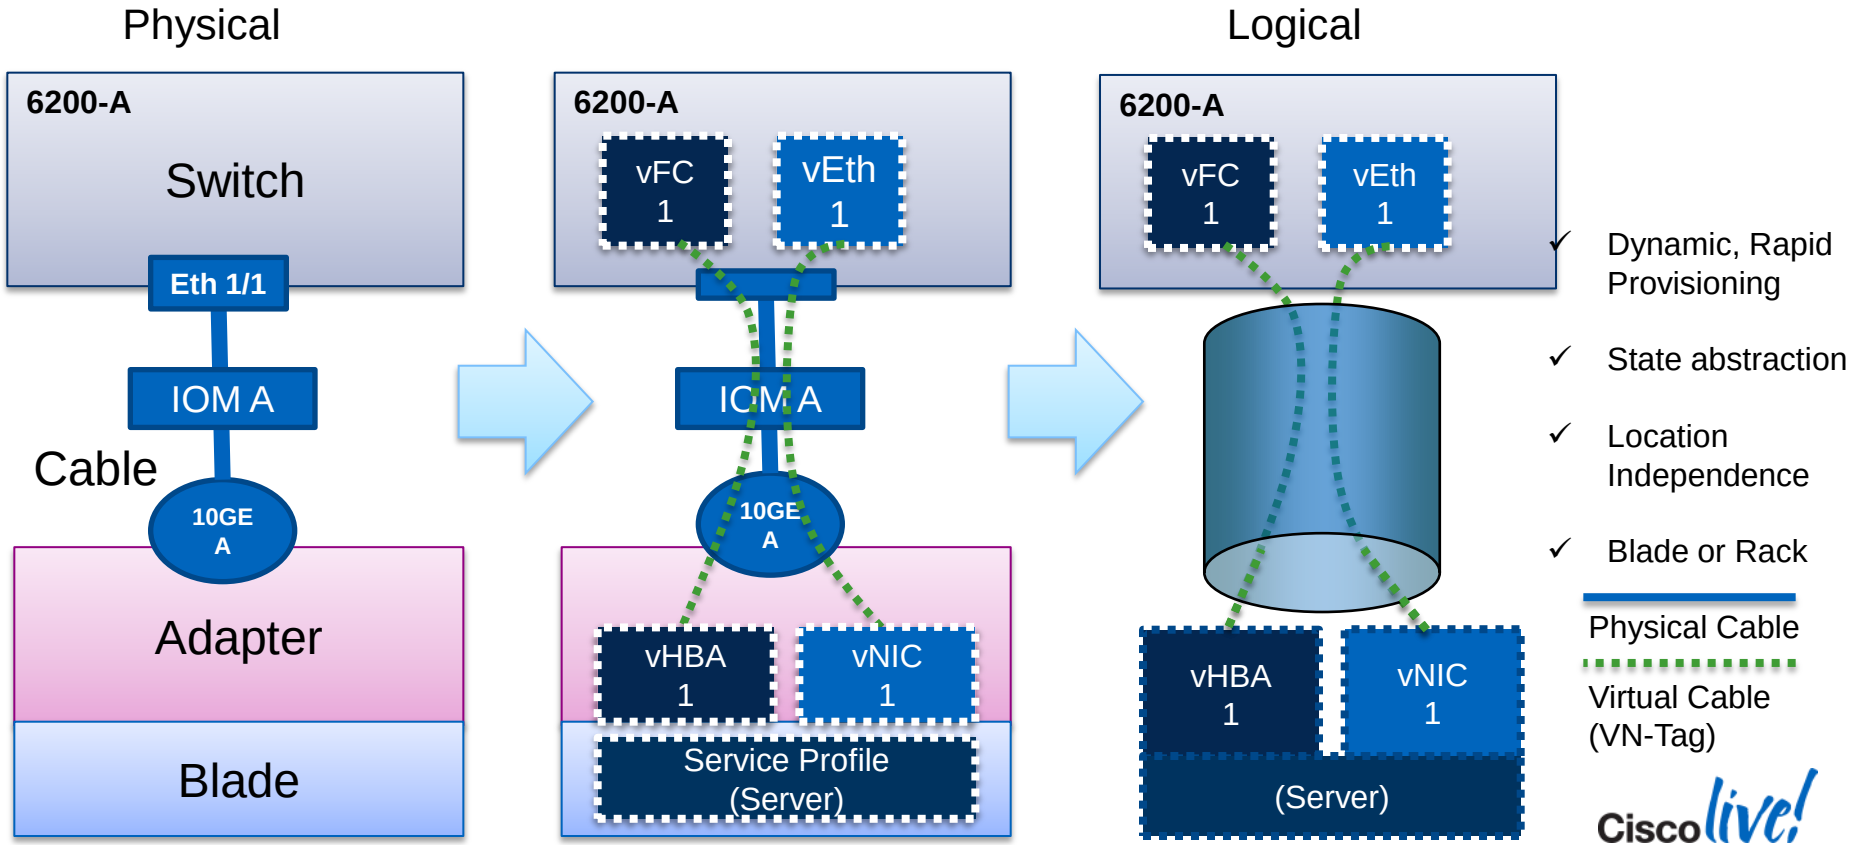

Fabric Forwarding Mode of Operations

Fabric Forwarding Mode of Operations

Modes of Operation

- End-host mode (EHM): Default mode
  - No spanning-tree protocol (STP); no blocked ports
  - Admin differentiates between server and network ports
  - Using dynamic (or static) server to uplink pinning
  - No MAC address learning except on the server ports; no unknown unicast flooding
  - Fabric failover (FF) for Ethernet vNICs (not available in switch mode)
- Switch mode: User configurable
  - Fabric Interconnects behave like regular ethernet switches
  - STP parameters are lock

End Host Mode

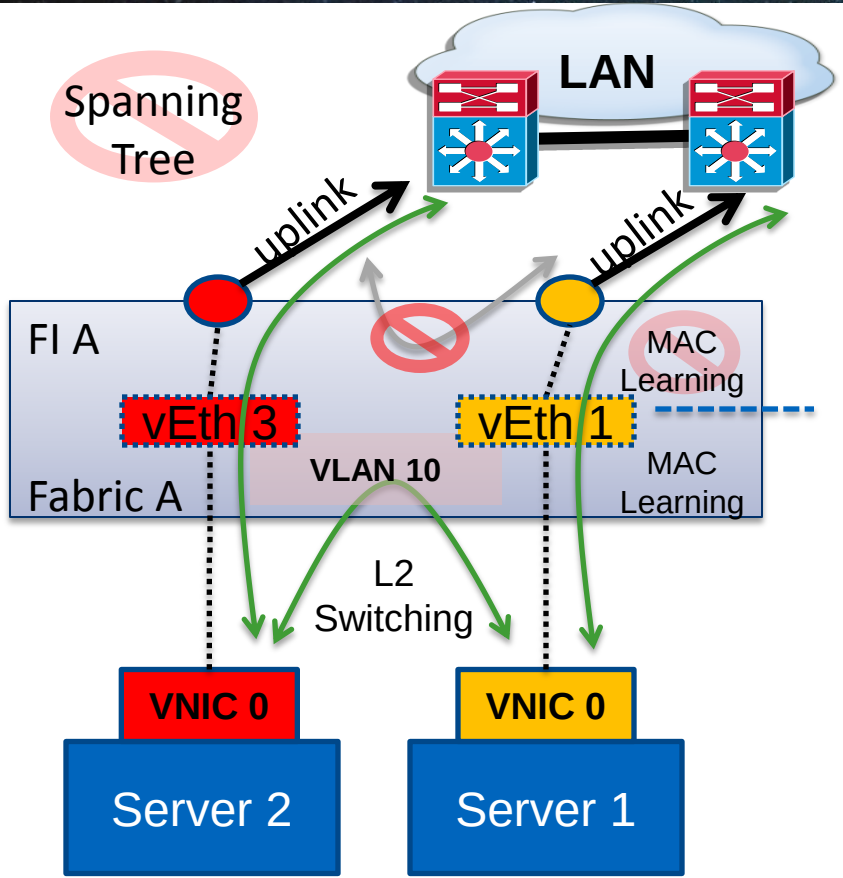

- Completely transparent to the network
  - Presents itself as a bunch of hosts to the network
- No STP – simplifies upstream connectivity
- All uplinks ports are forwarding – never blocked

End Host Mode
Unicast Forwarding

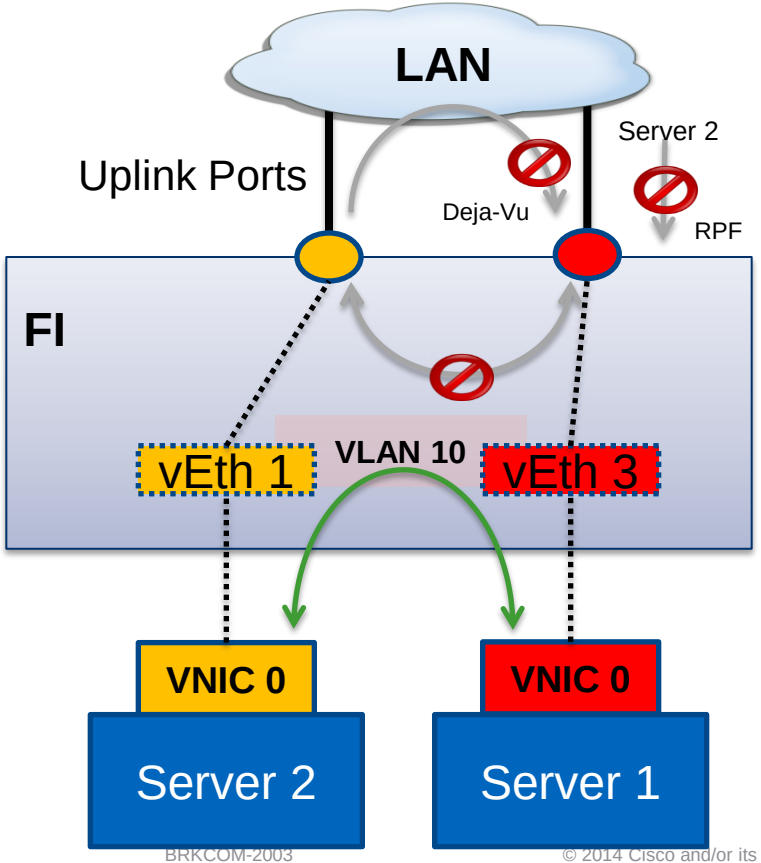

- MAC/VLAN plus policy based forwarding
  - Server pinned to uplink ports
- Policies to prevent packet looping
  - déjà vu check
  - RPF
  - No uplink to uplink forwarding
- No unknown unicast or multicast
  - igmp-snooping can be disable on per-VLAN basis

End Host Mode
Multicast Forwarding

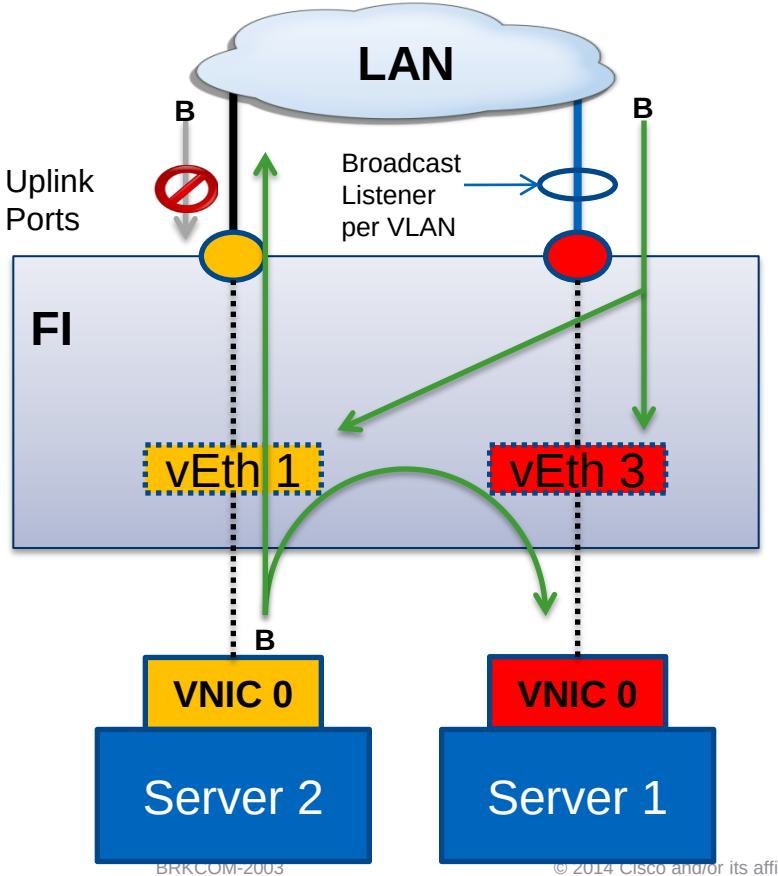

- Broadcast traffic for a VLAN is pinned on exactly one uplink port (or port-channel) i.e., it is dropped when received on other uplinks
- Server to server multicast traffic is locally switched
- RPF and déjà vu check also applies for multicast traffic

Switch Mode

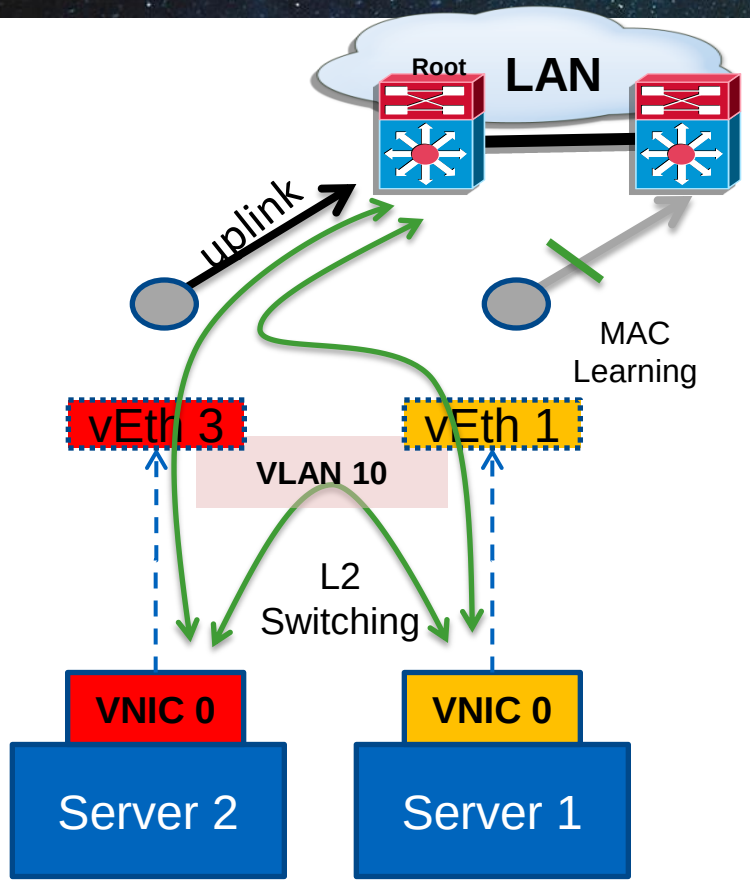

- Fabric Interconnect behaves like a normal L2 switch
- Rapid-STP+ to prevent loops
  - STP parameters are not configurable
- Server vNIC traffic follows STP forwarding states
  - Use VPC to get around blocked ports
- VTP is not supported
- MAC address learning on both uplinks and server links

Uplink Pinning

End Host Mode - Dynamic Pinning

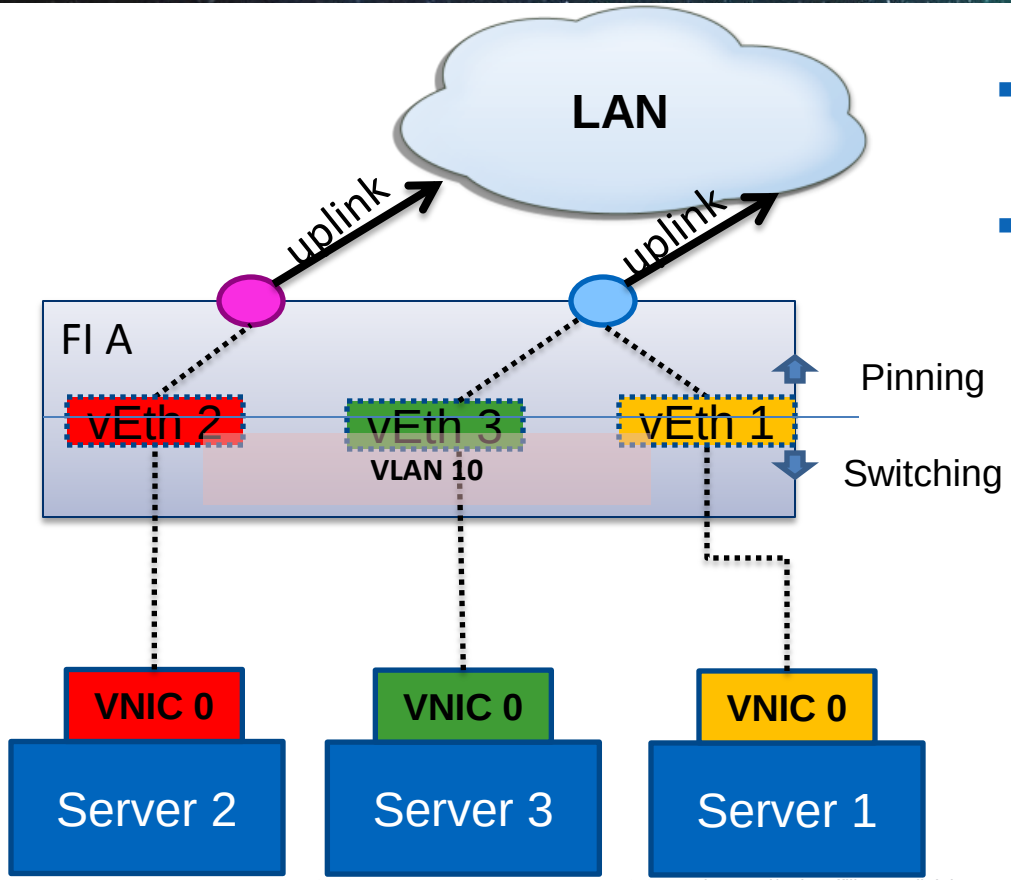

- UCSM manages the vEth pinning to the uplink
- UCSM will periodically vEth distribution and redistribute the vEths across the uplinks

End Host Mode – Individual Uplinks

- ✓All uplinks forwarding for all VLANs
- ✓GARP aided upstream convergence
- ✓No STP
- ✓Sub-second re-pinning
- ✓No server NIC disruption

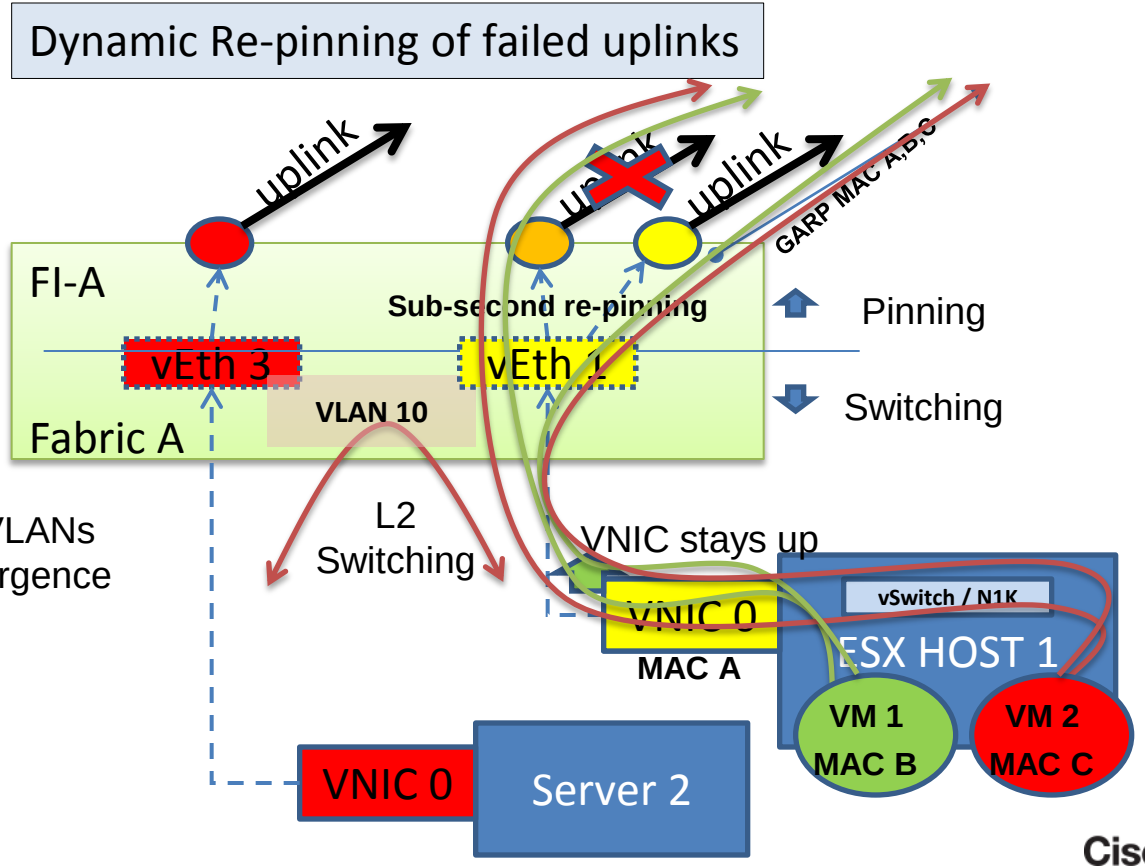

End Host Mode – Port Channel Uplinks

Recommended: Port Channel Uplinks

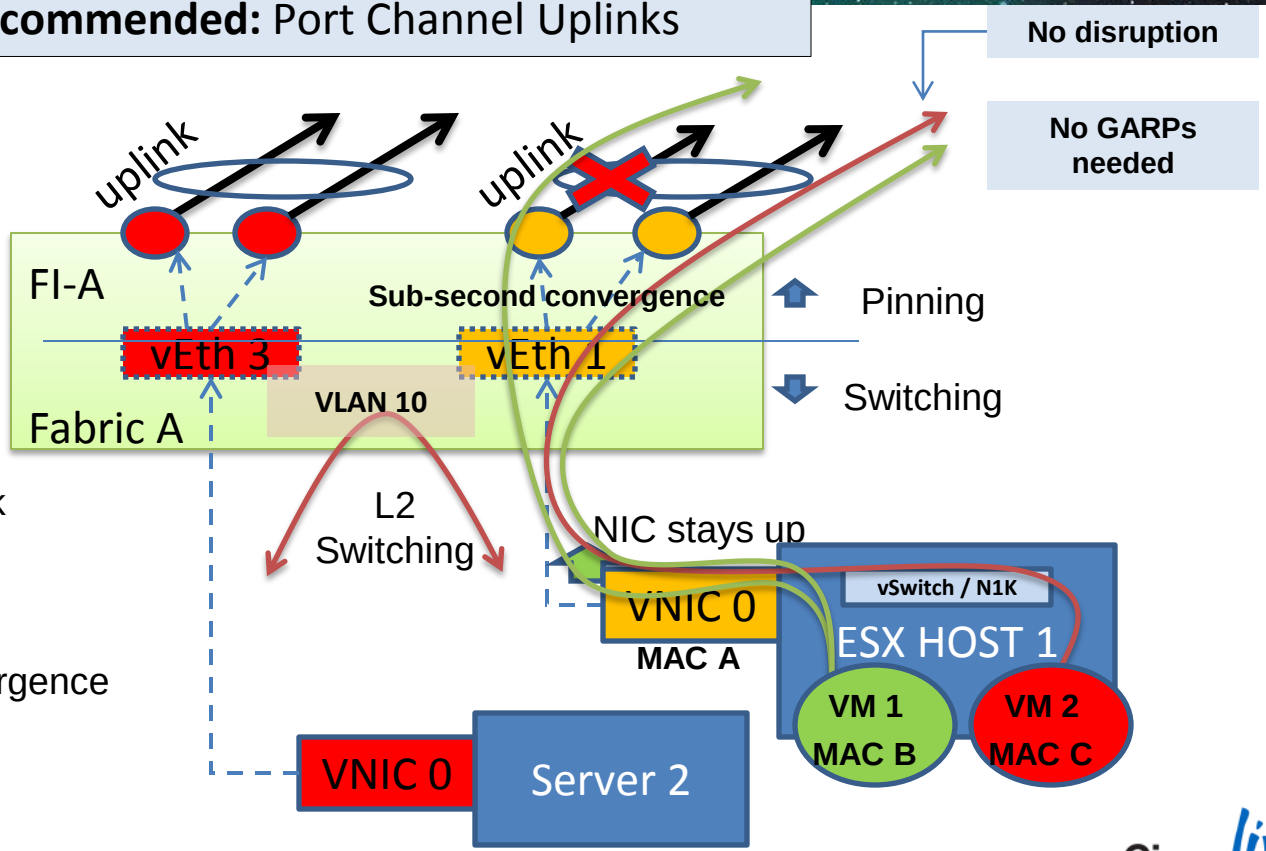

- ✓More Bandwidth per Uplink
- ✓Per flow uplink diversity
- ✓No Server NIC disruption
- ✓Fewer GARPs needed
- ✓Faster bi-directional convergence
- ✓**Fewer moving parts**

RECOMMENDED

End Host Mode – Static Pinning

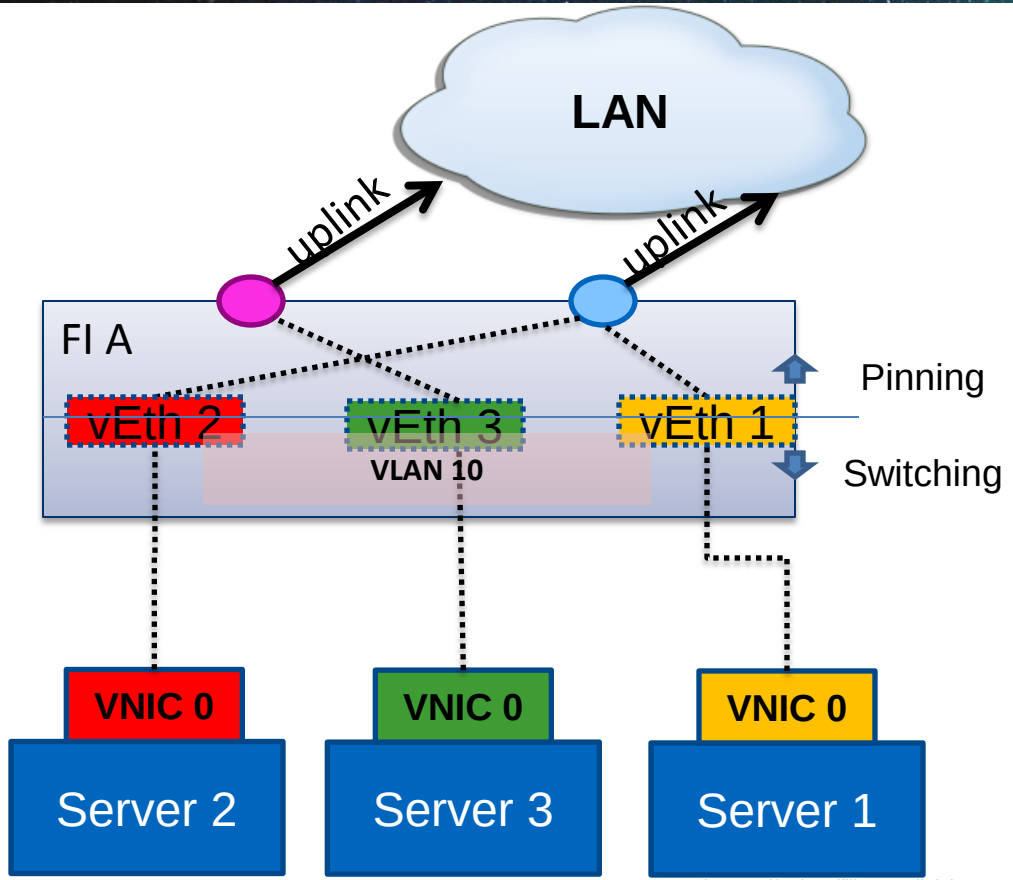

Administrator Pinning Definition

vEth Interfaces	Uplink
vEth 1	Blue
vEth 2	Blue
vEth 3	Purple

- Administer controls the vEth pinning
- Deterministic traffic flow
- Pinning configuration is done under the LAN tab -> LAN Pin groups and assigned under the vNIC
- No re-pinning with in the same FI
- Static and dynamic pinning can co-exist




Chassis / Fabric Extender

UCS Fabric Topologies
Chassis Bandwidth Options

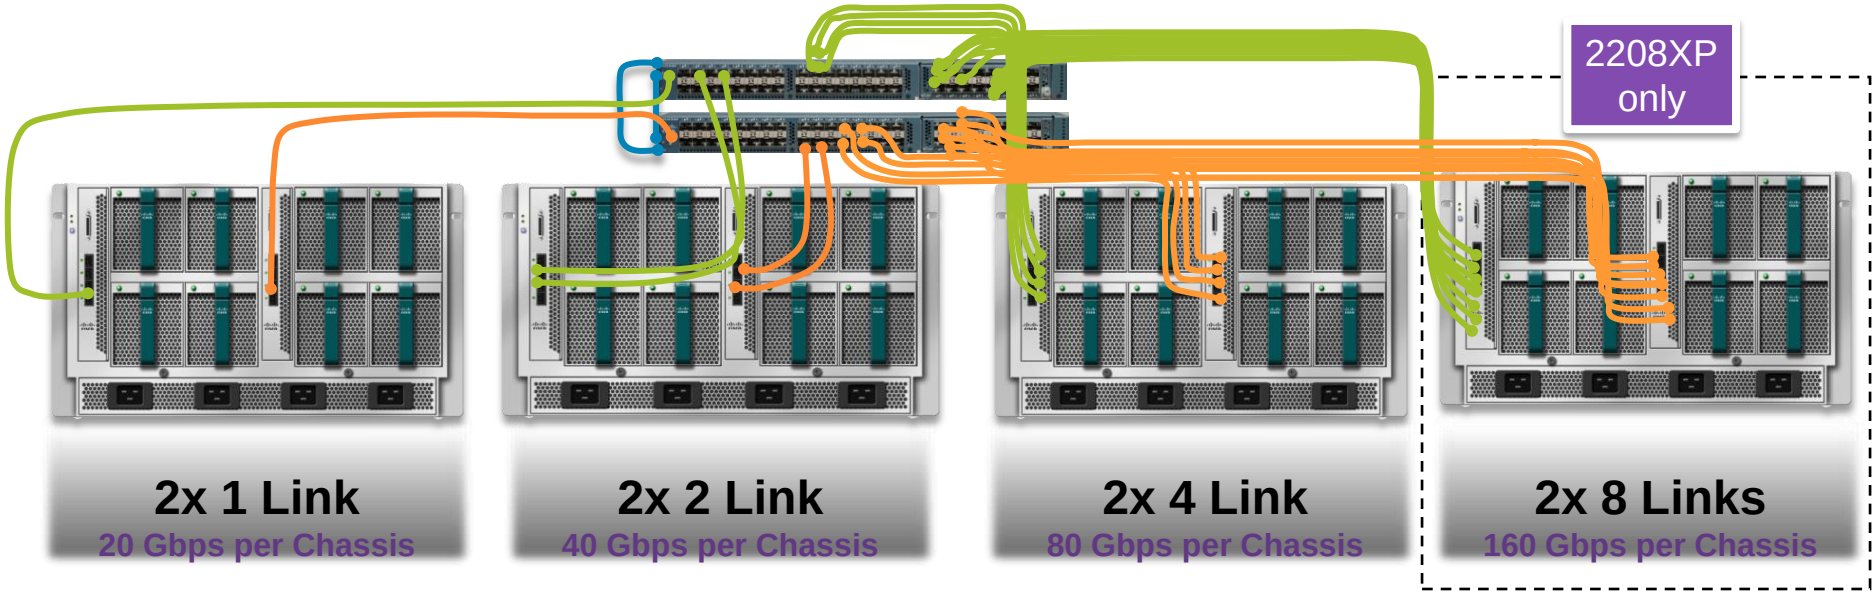

What are Those IOMs?

- A IOM (sometimes called ‘Fabric Extender’) provides
  - A 1GE switch used for internal management (1GE per slot)
  - A number of 10G-KR sever facing links (HIF)
  - A number of Fabric links (NIF)
- NIC cards on the servers use those HIF ports for external connectivity
- Each IOM provides a separate dedicated IO channel for internal management connectivity
- There is no local switching on IOMs – traffic is always switched by the FIs

Let's go back in time a bit…

2104XP
Gen 1 IOM

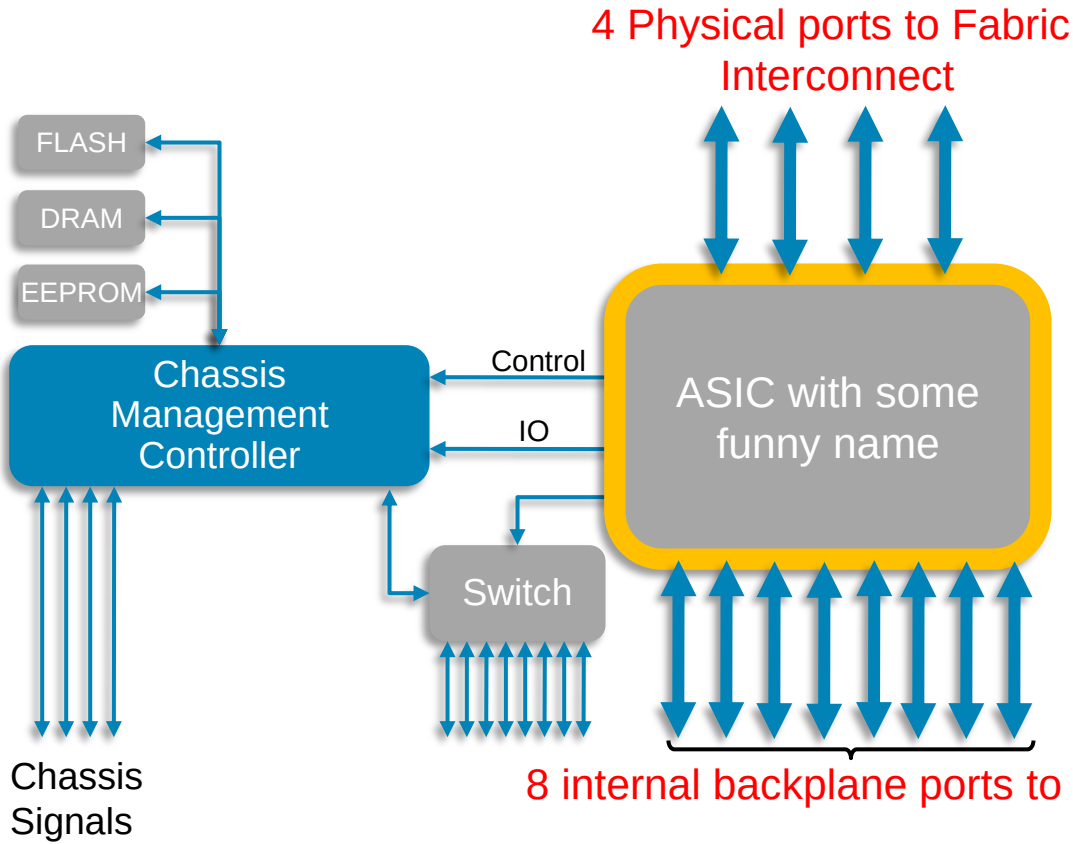

Up to 80Gbps per chassis

Up to 20Gbps per slot

So this is what we had with the 2104XP

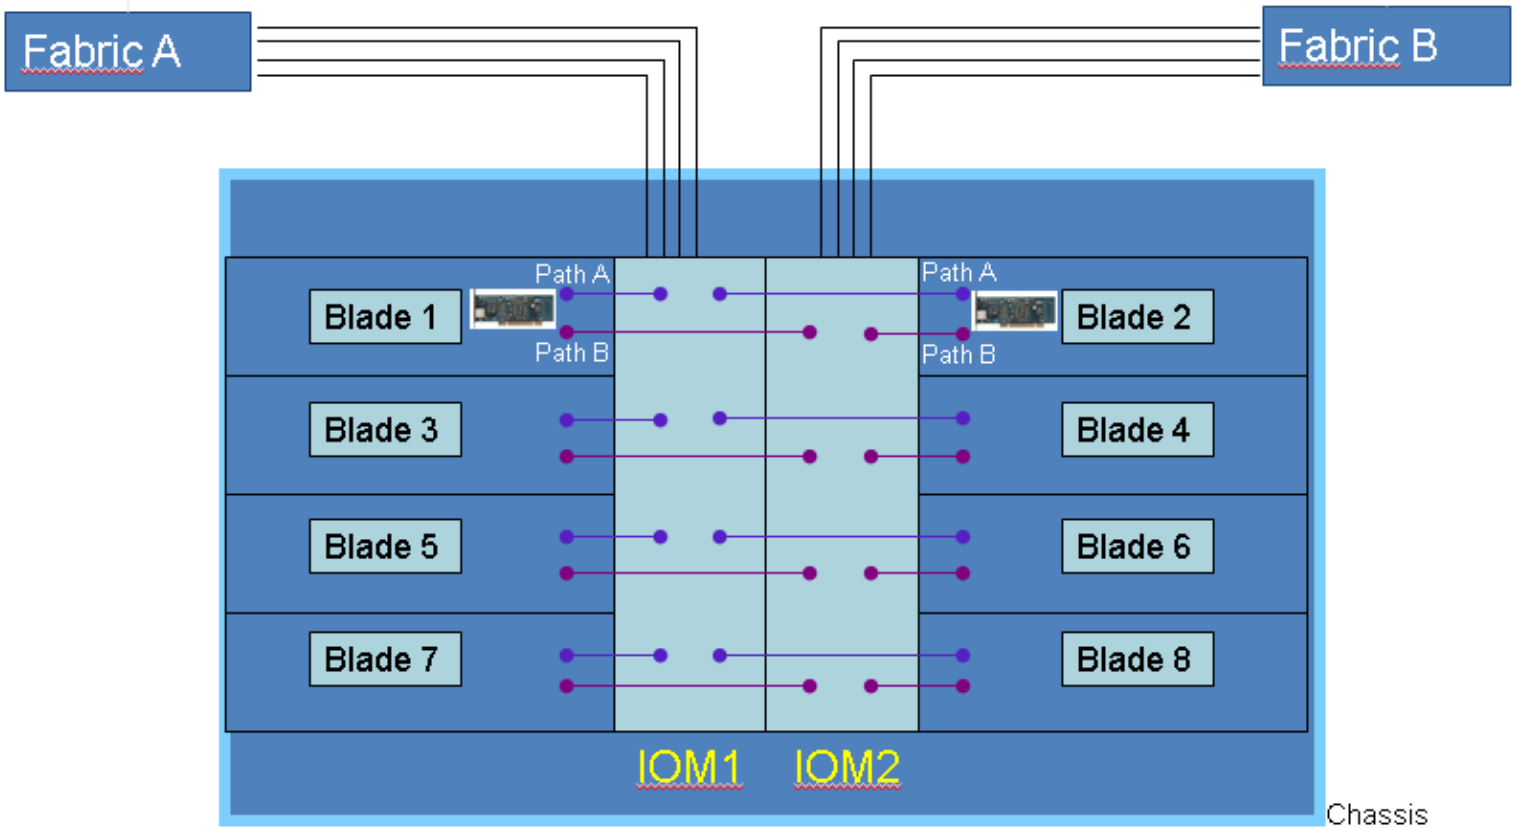

Today: UCS 2204 IO Module

Dual 20 Gbps to Each Blade Server

UCS-IOM-2204XP

- **Bandwidth increase**
  - 40G to the Network
  - 160G to the hosts (redundant)
    - (2x10G per half-width slot; 4x10G per Full-width slot)
- **Latency** Lowered to 0.5us within IOM
- **Investment Protection** with Backward Compatibility

UCS 2208 IO Module

Enable Dual 40 Gbps to Each Blade Server

UCS-IOM-2208XP

- **Bandwidth increase**
  - 80G to the Network
  - 320G to the hosts (redundant)
    - (4x10G per half-width slot; 8x10G per full-width slot)
- **Latency** Lowered to 0.5us within IOM
- **Investment Protection** with Backward Compatibility

220x-XP Architecture

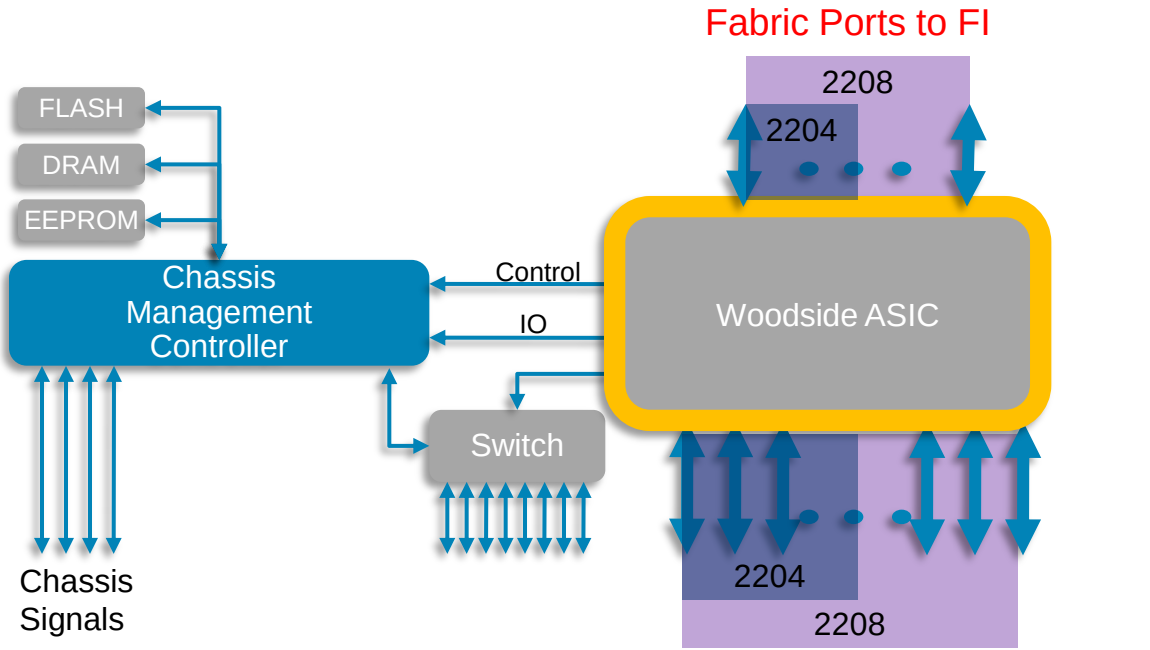

Feature	2204-XP	2208-XP
ASIC	Woodside	Woodside
Fabric Ports (NIF)	4	8
Host Ports (HIF)	16	32
CoS	8	8
Latency	~ 500ns	~ 500ns

No Local Switching – ever!
Traffic goes up to FI

Blade Northbound Ports

- These interfaces (show int brief – NXOS shell) are backplane traces
- Eth x/y/z nomenclature where
  - x = chassis number
  - y = is always 1
  - z = host interface port number

```
--------------------------------------------------------------------------------
Ethernet      VLAN    Type  Mode   Status  Reason                   Speed     Port
Interface                                                                     Ch #
--------------------------------------------------------------------------------
Eth1/1/1      1       eth   access down    Administratively down       10G(D) --
Eth1/1/2      1       eth   access down    Administratively down       10G(D) --
Eth1/1/3      1       eth   access down    Administratively down       10G(D) --
Eth1/1/4      1       eth   access down    Administratively down       10G(D) --
Eth1/1/5      1       eth   access down    Administratively down       10G(D) --
Eth1/1/6      1       eth   access down    Administratively down       10G(D) --
Eth1/1/7      1       eth   access up      none                        10G(D) --
Eth1/1/8      1       eth   access up      none                        10G(D) --
Eth1/1/9      1       eth   trunk  up      none                        10G(D) --
cali-spring-B#
```

UCS Internal Block Diagram

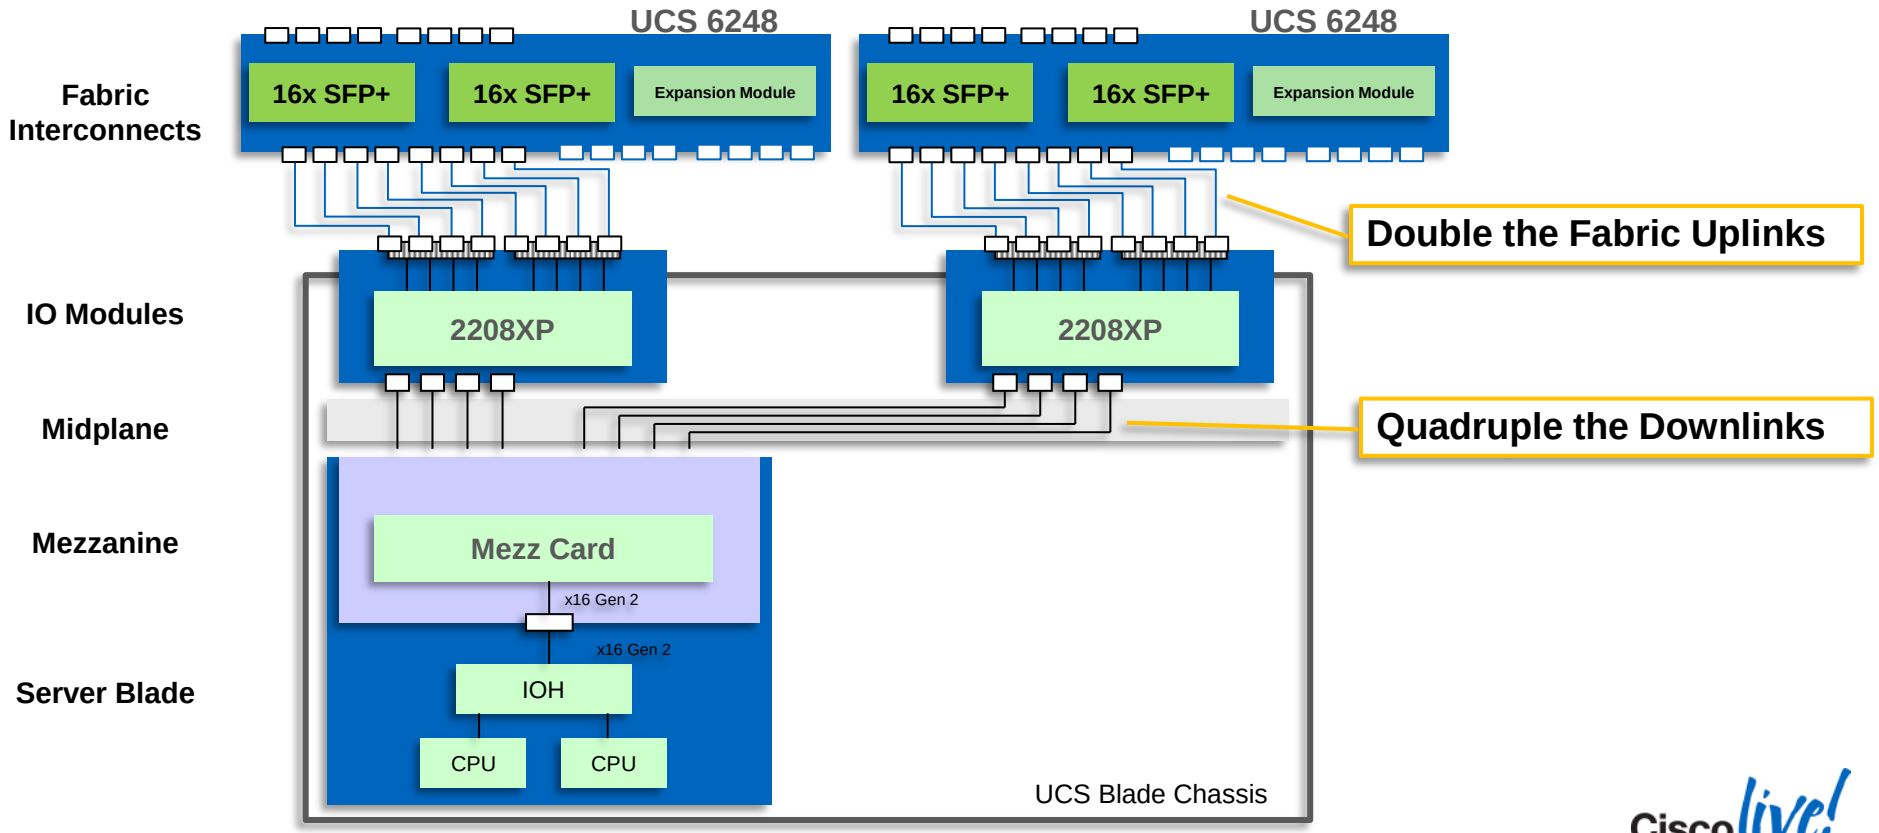

IO Module HIF to NIF Pinning

2208XP – 1 Link

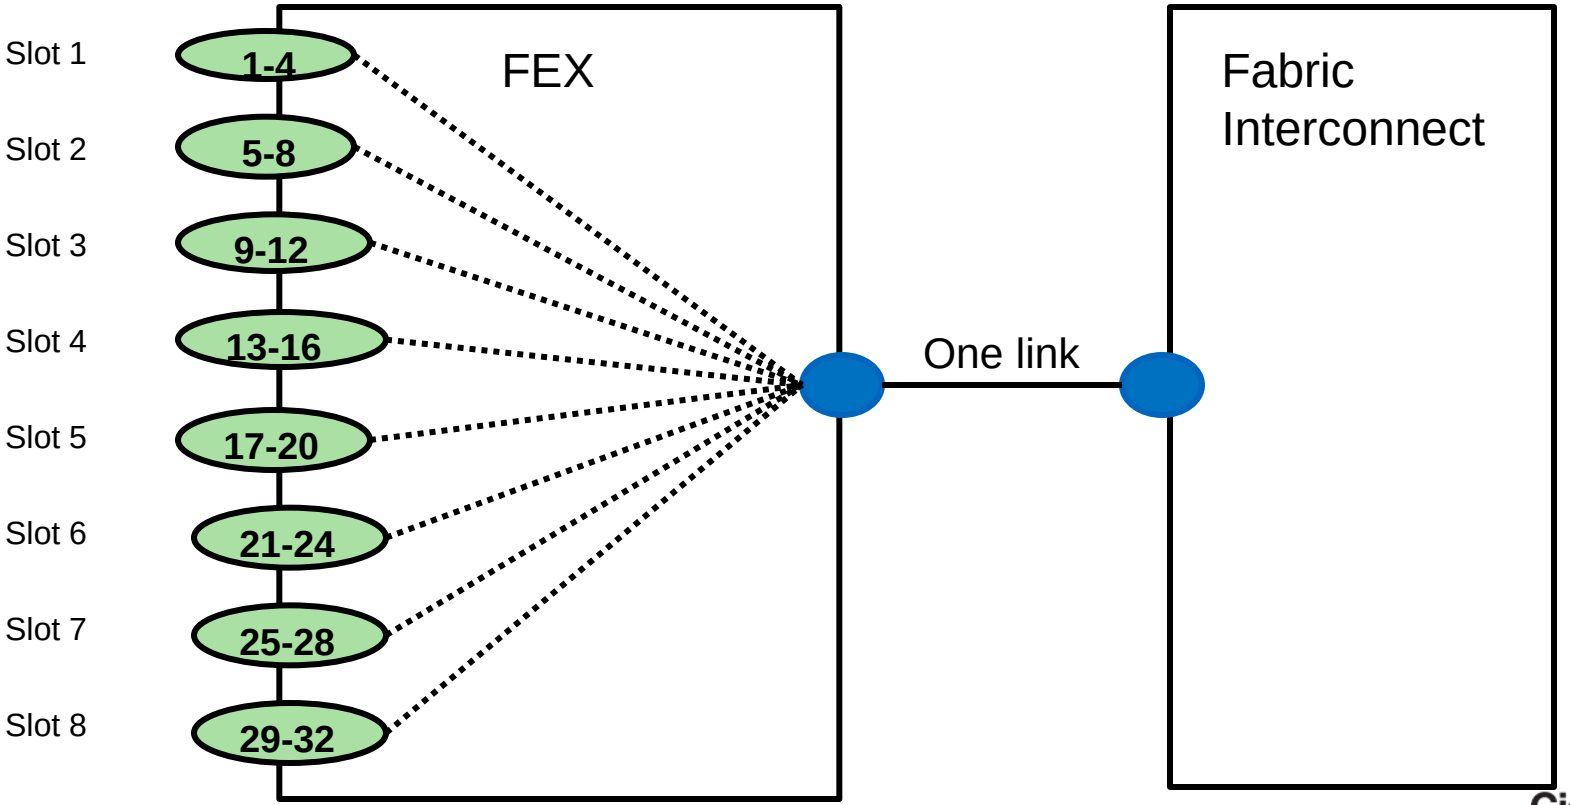

IO Module HIF to NIF Pinning

2208XP – 2 Link

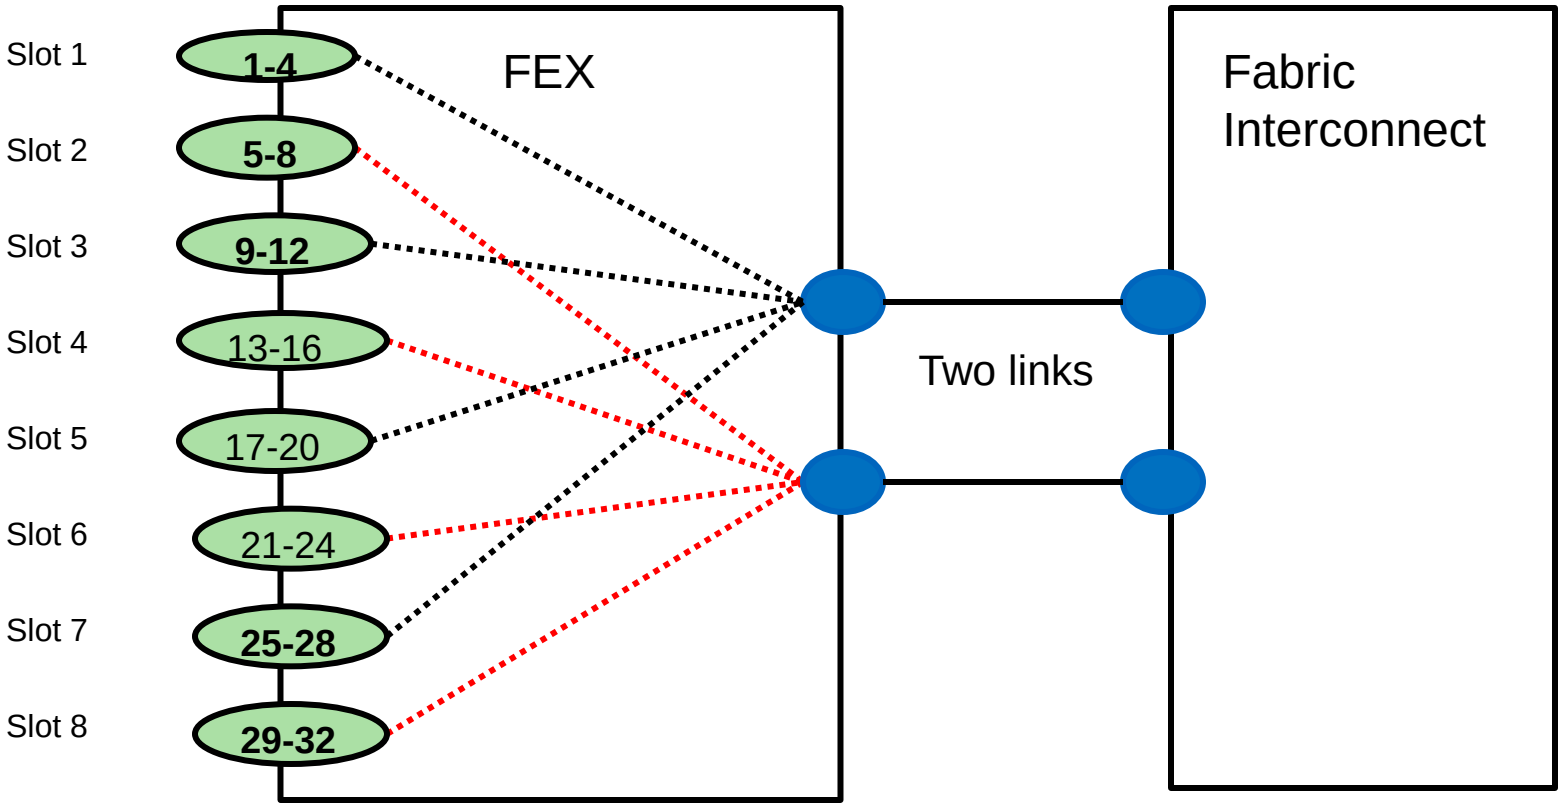

IO Module HIF to NIF Pinning

2208XP – 4 Link

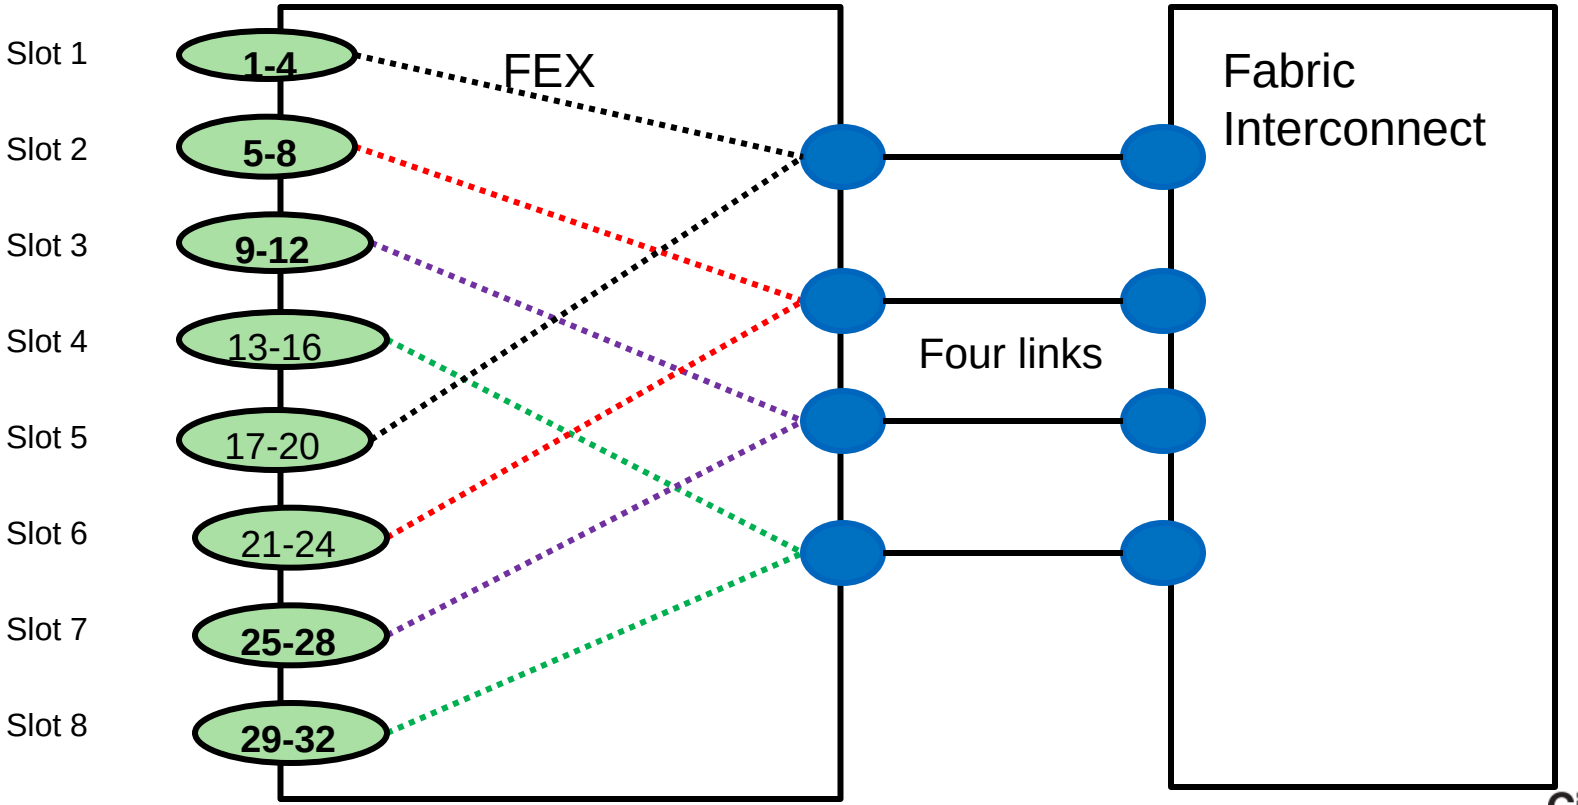

IO Module HIF to NIF Pinning

2208XP – 8 Link

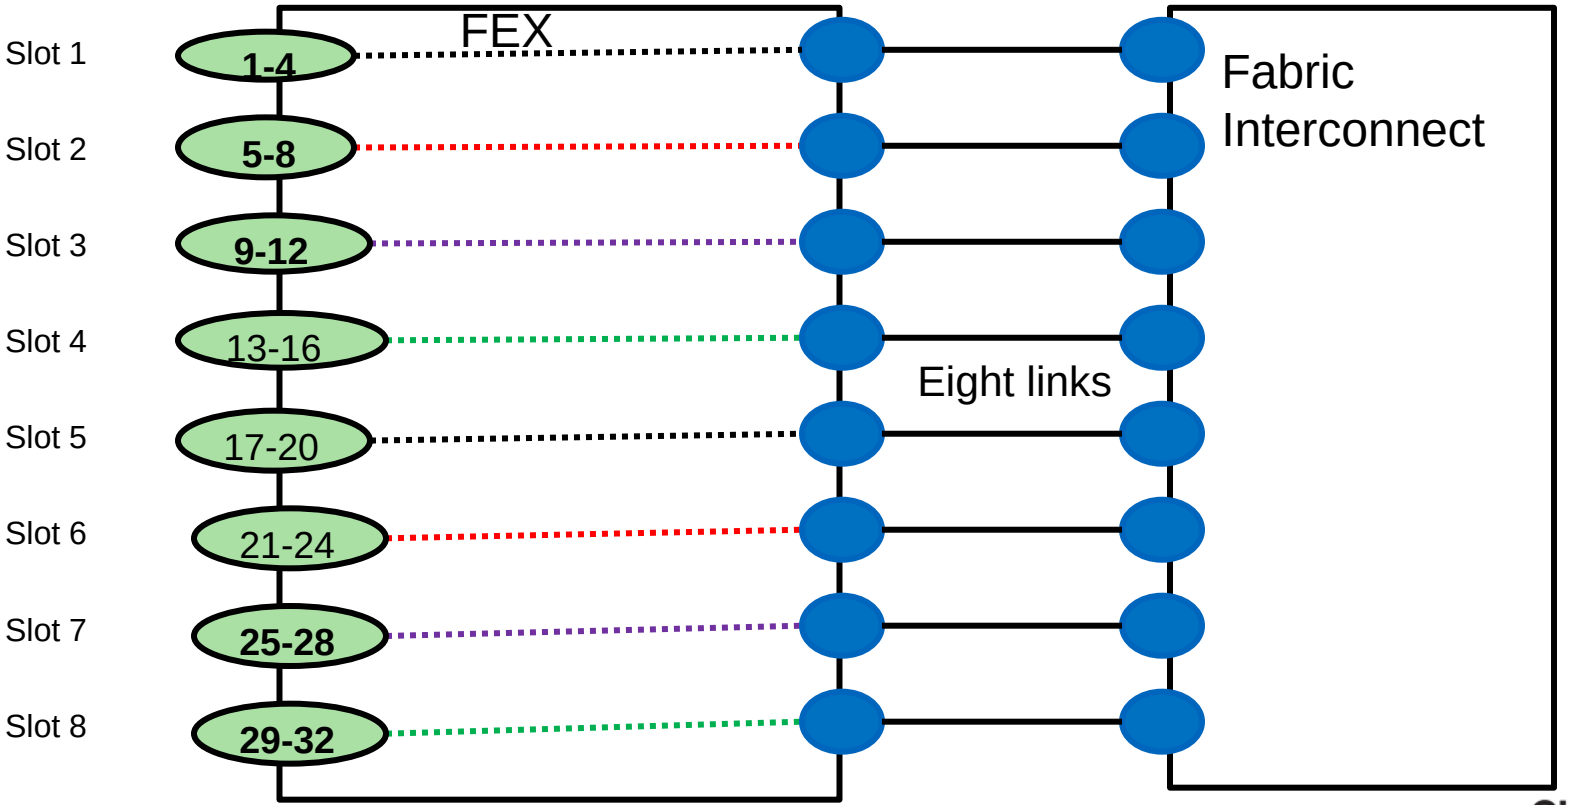

IOM and Failover

- What happens in a 4-link topology when you loose 1 link?
  - Server interfaces pinned to that link go down*
  - The remaining 3 links still pass traffic for the other blade servers
  - To recover the failed servers' vNICs, re-acknowledged of the chassis is required
  - After a re-ack UCS falls back to 2 links with regards to blade to fabric port mapping
    - That's because the link count must be a power of 2!

* unless you enabled Fabric Failover

IOM and Failover

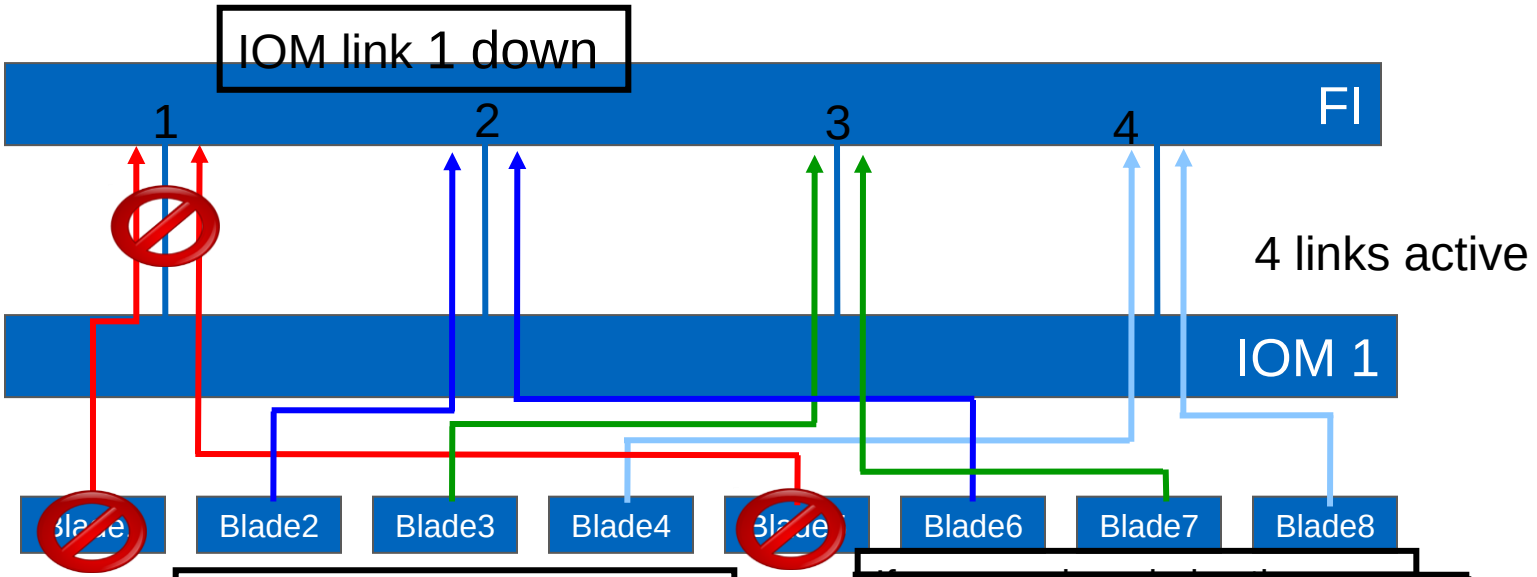

Lose connectivity on mezzanine port mapped to IOM 1 for blades 1 and 5

Increased Bandwidth Access to Blades

160 Gb (Discrete Mode)

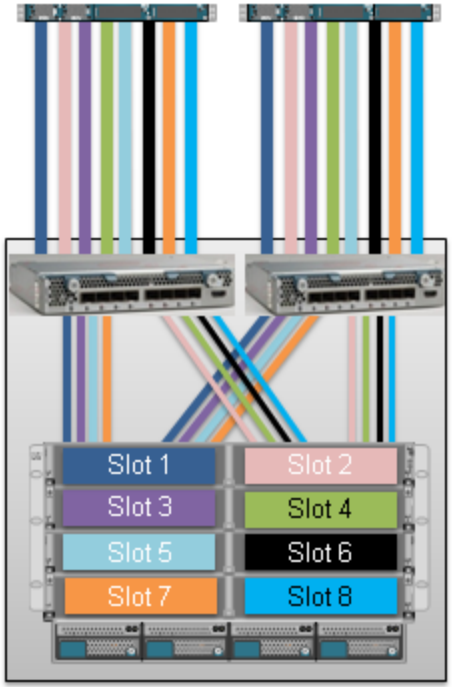

- 6100 to 2208
- 6200 to 2208

160 Gb (Port Channel Mode)

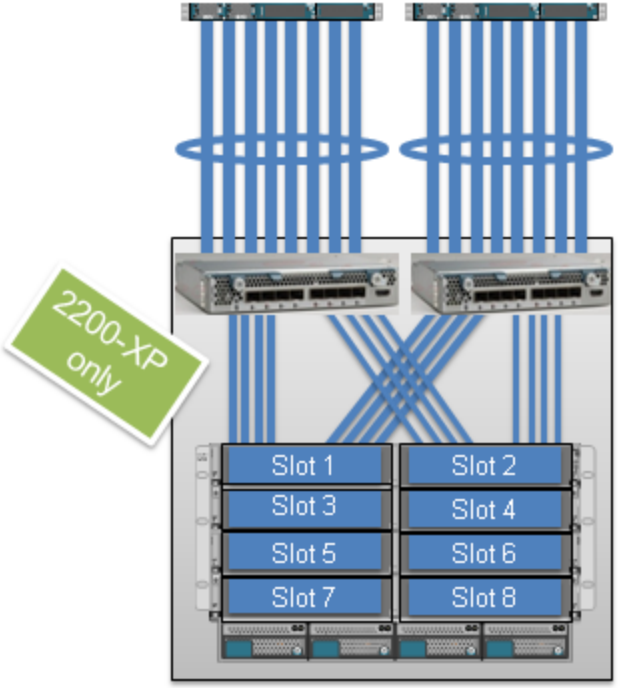

- 6200 to 2208

Port-channel Pinning

- No slot based pinning
- No invalid link count for NIF ports (no “power of 2” rule)

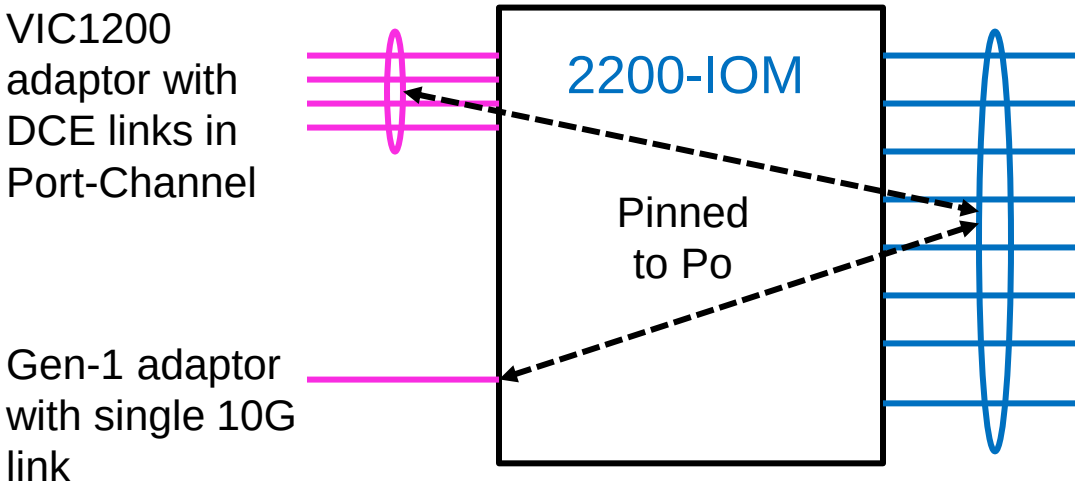

Port-channel Pinning

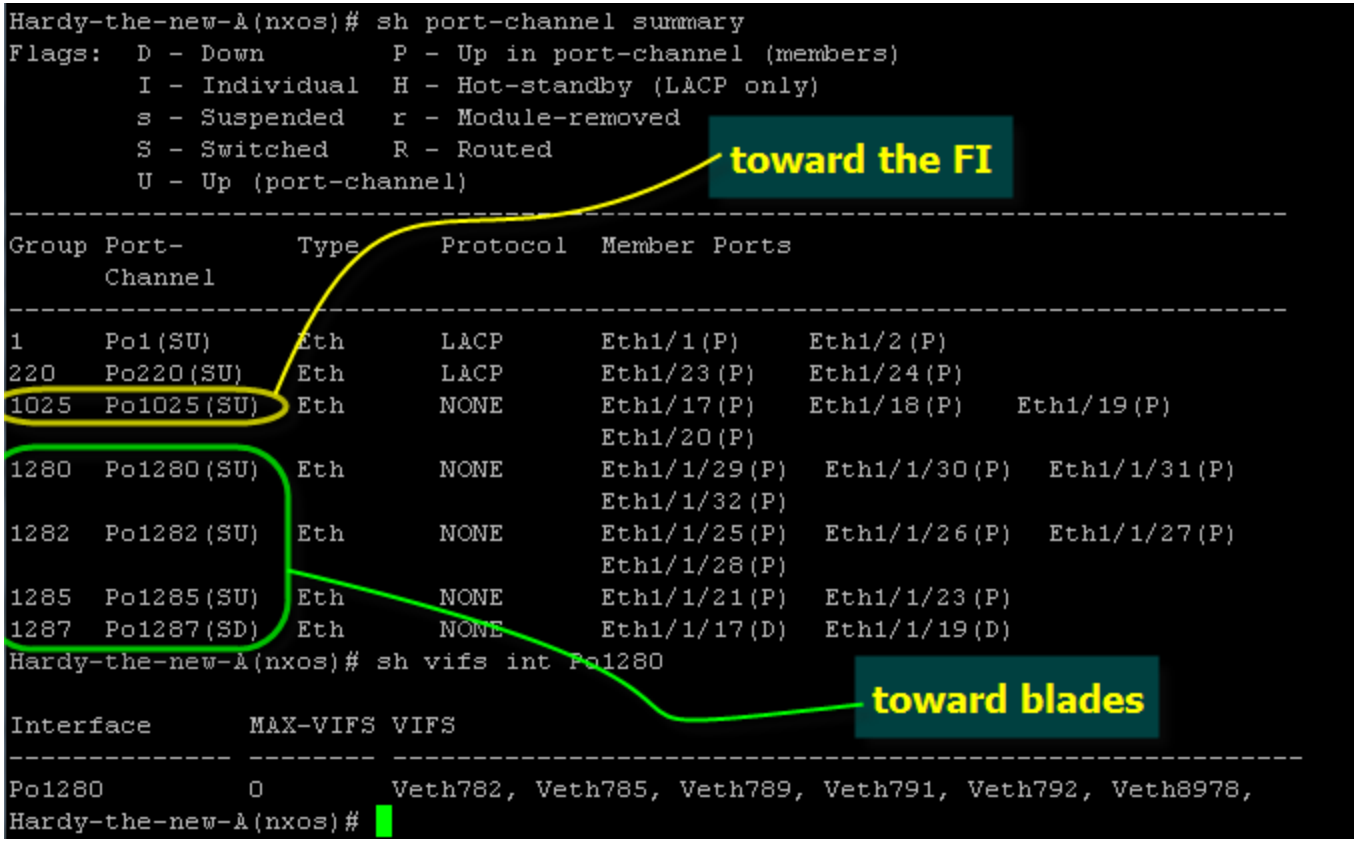




Server Connectivity

Let’s go back in time once more

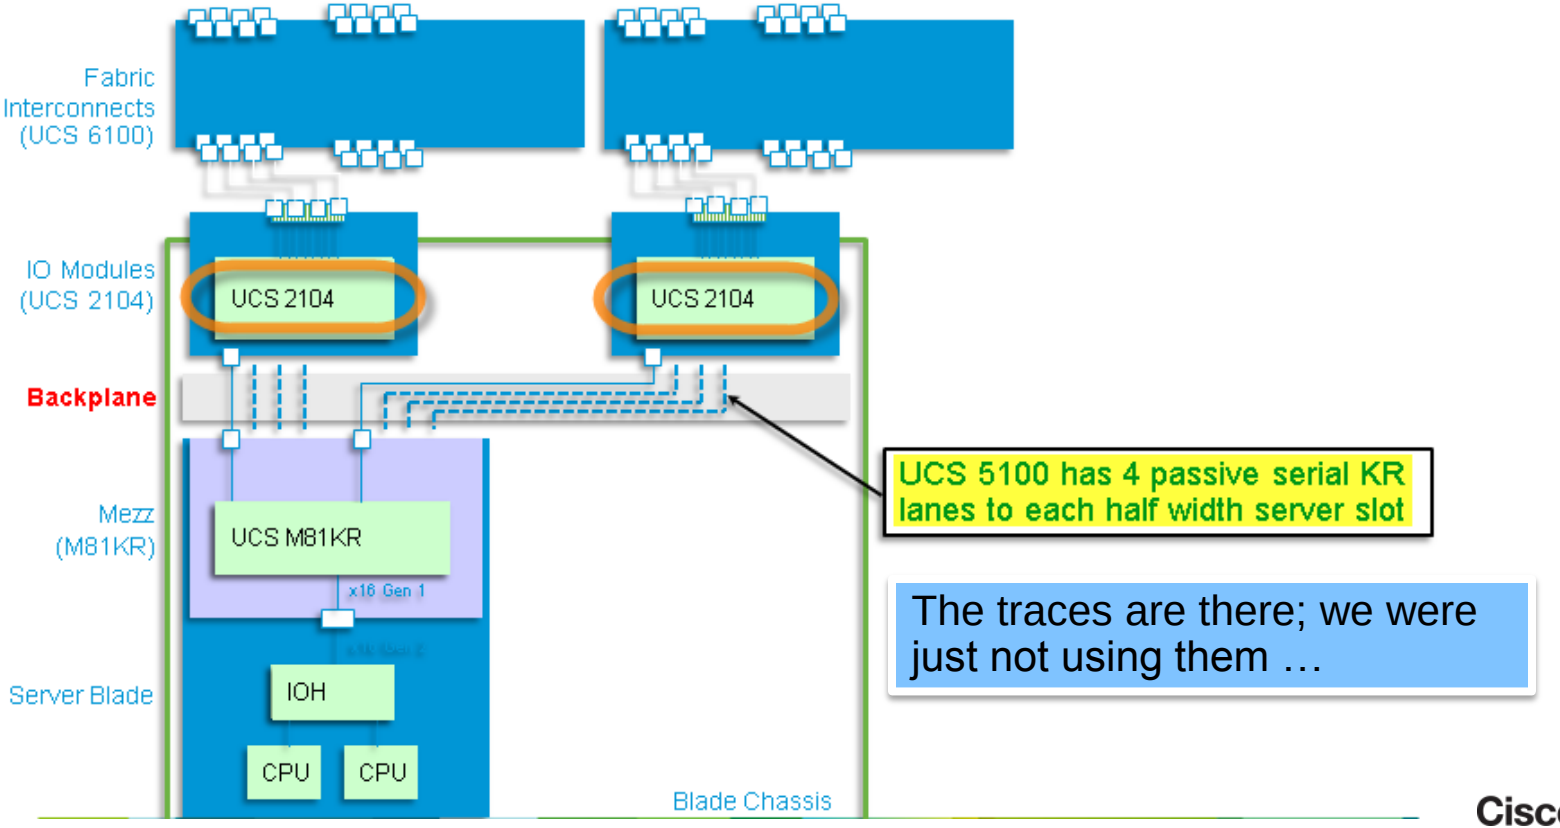

The traces are there; we were just not using them …

UCS Cisco 1280 VIC Adapter

Customer benefits

- **Dual 4x 10 GE (80 Gb per host)**
- **VM-FEX scale, up to 112 VM interfaces /w ESX 5.0**

Feature details

- Dual 4x 10 GE port-channels to a single server slot
- Host connectivity PCIe Gen2 x16
- PCIe Gen 2 x16 bandwidth limit is 32 Gbps
- HW Capable of 256 PCIe devices
  - OS restriction apply
- PCIe virtualisation OS independent (same as M81KR)
- Single OS driver image for both M81KR and 1280 VIC
- FabricFailover supported
- Eth hash inputs : Source MAC Address, Destination MAC Address, Source Pprt, Destination Port, Source IP address, Destination, P address and VLAN
- FC Hash inputs: Source MAC Address Destination MAC Address, FC SID and FC DID

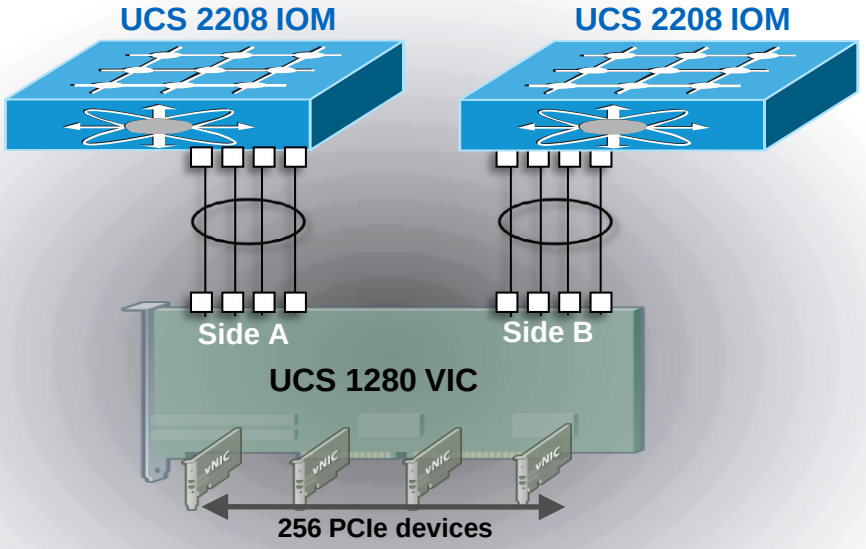

All M3 blades
& B230M2
& B440M2

Connectivity IOM to Adapter

Up to 32 Gbps (PCIe speed) throughput per vNIC using flow based port-channel hash

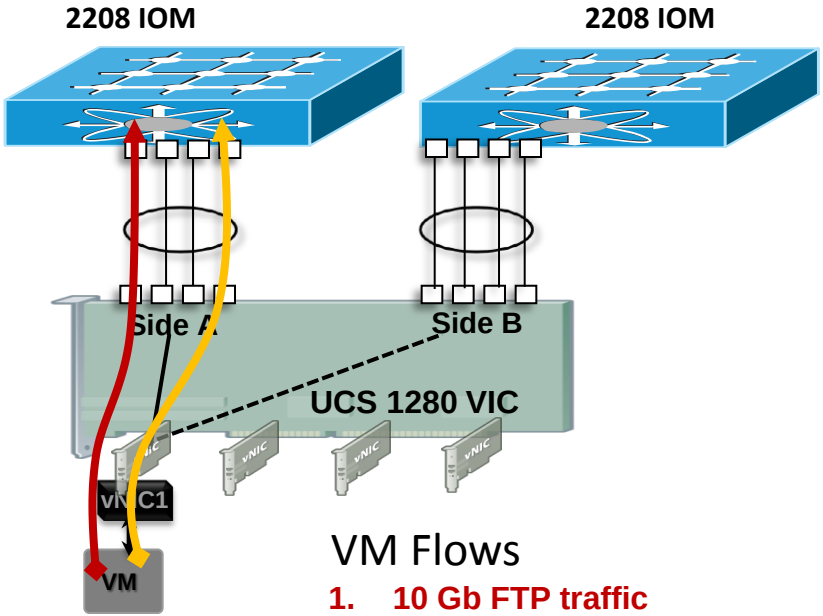

- Implicit Port-channel between UCS 1280 VIC adapter and UCS 2208 IOM
- 7-Tuple Flow based hash
- A vNIC is active on side A or B.
- A vNIC has access to up to 32 Gbps throughput .

What does the OS See?

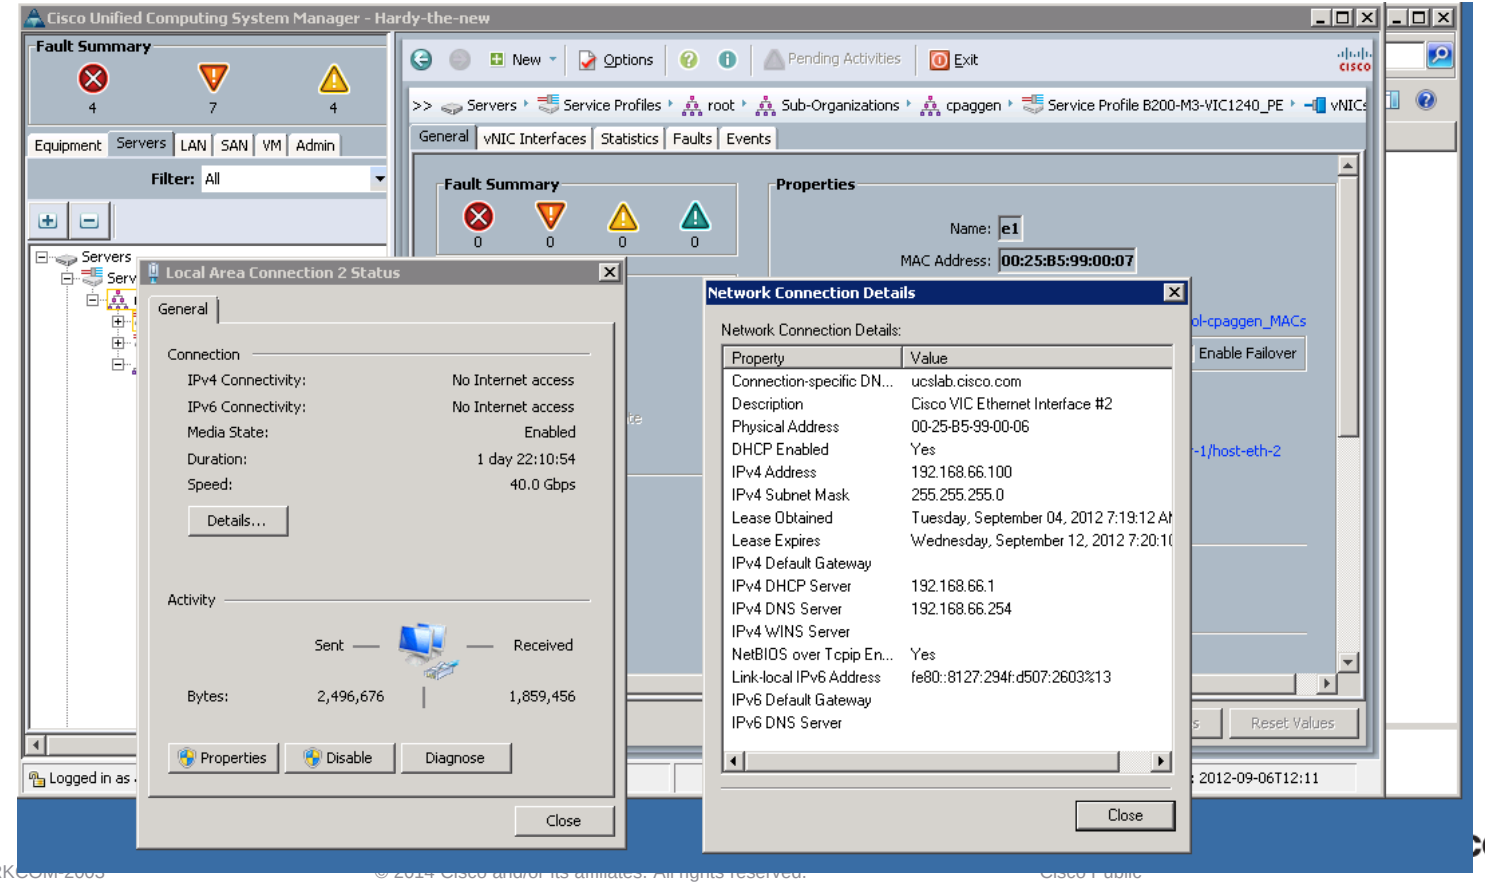

Fabric Failover
End Host Mode (only)

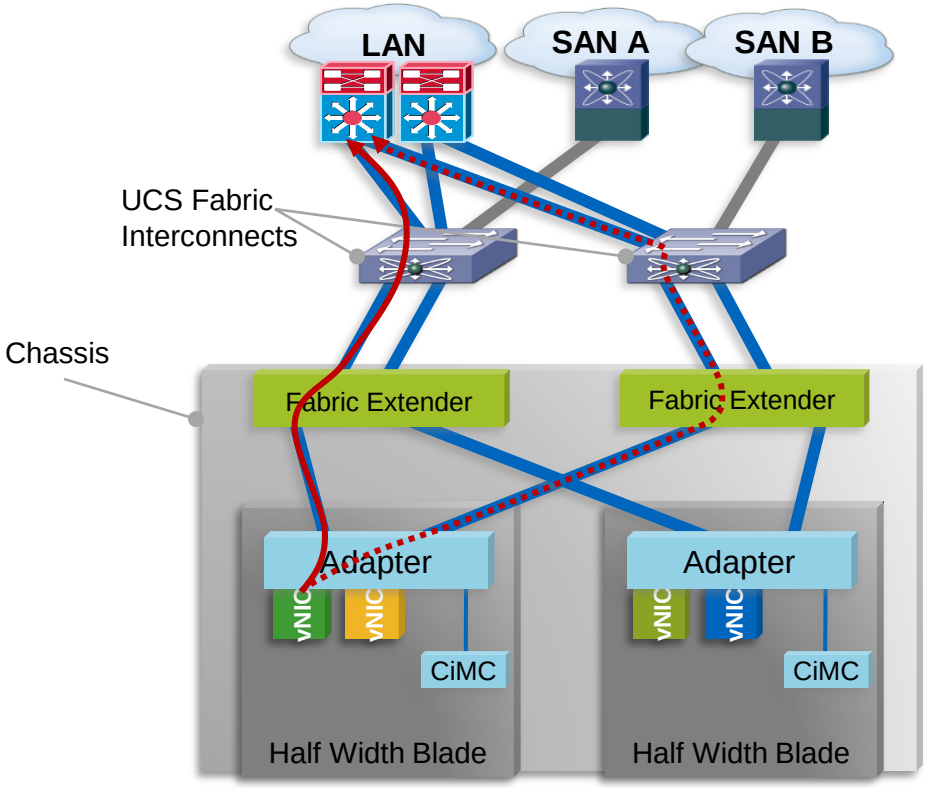

- Fabric provides NIC failover capabilities chosen when defining a service profile
- Traditionally done using NIC bonding driver in the OS
- Provides failover for both unicast and multicast traffic
- Works for any OS on bare metal and hypervisors

Fabric Failover
Bare Metal OS

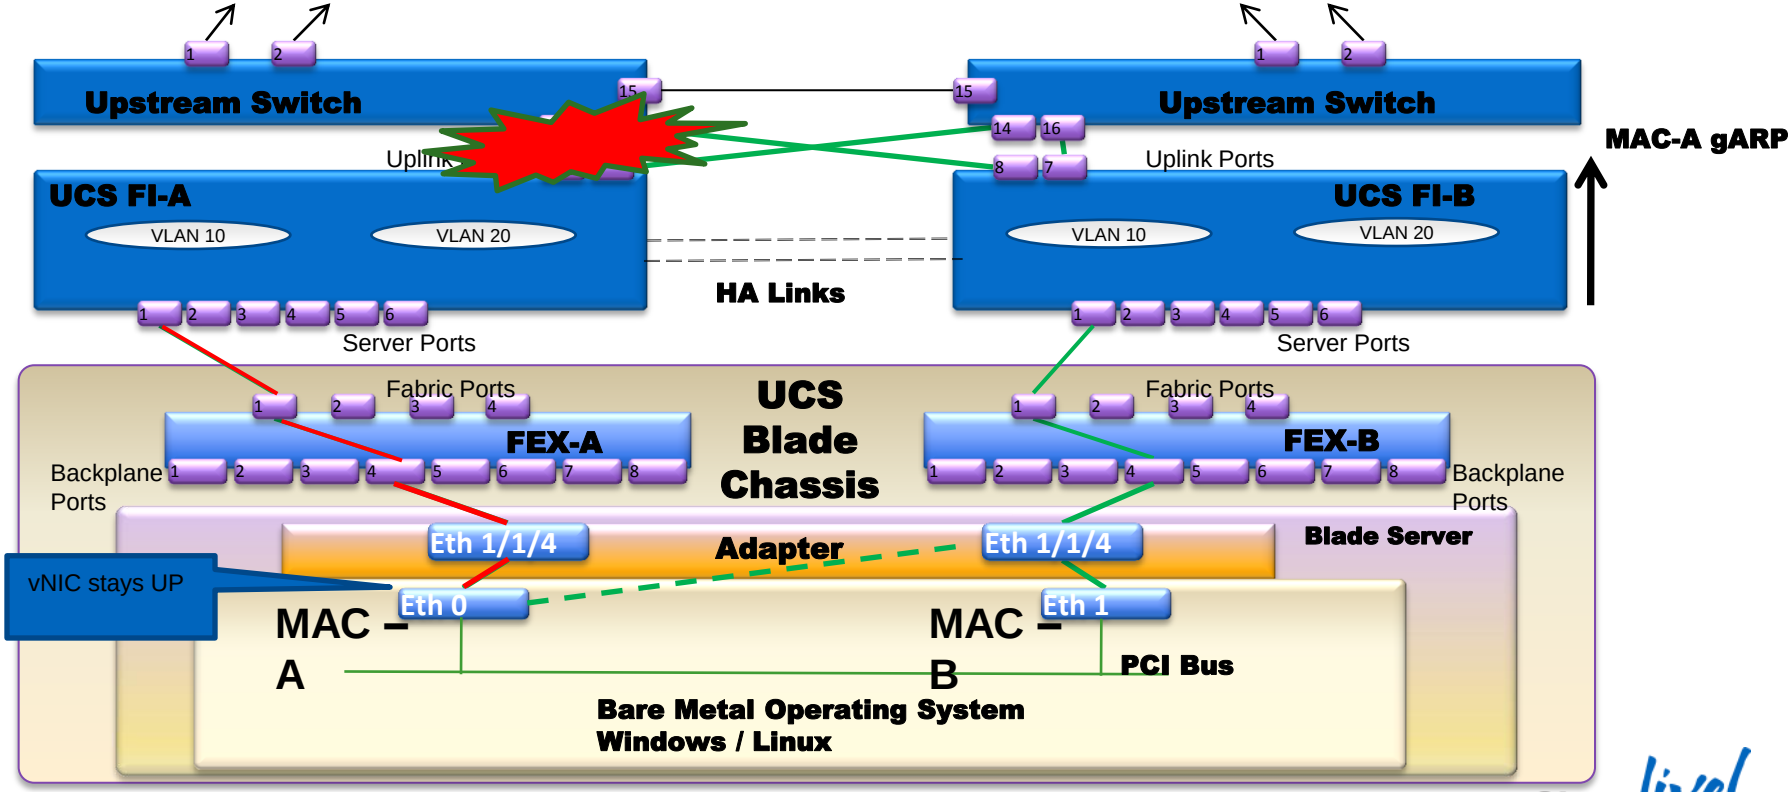

IOM 2104 with M81KR in M1/M2 Blades

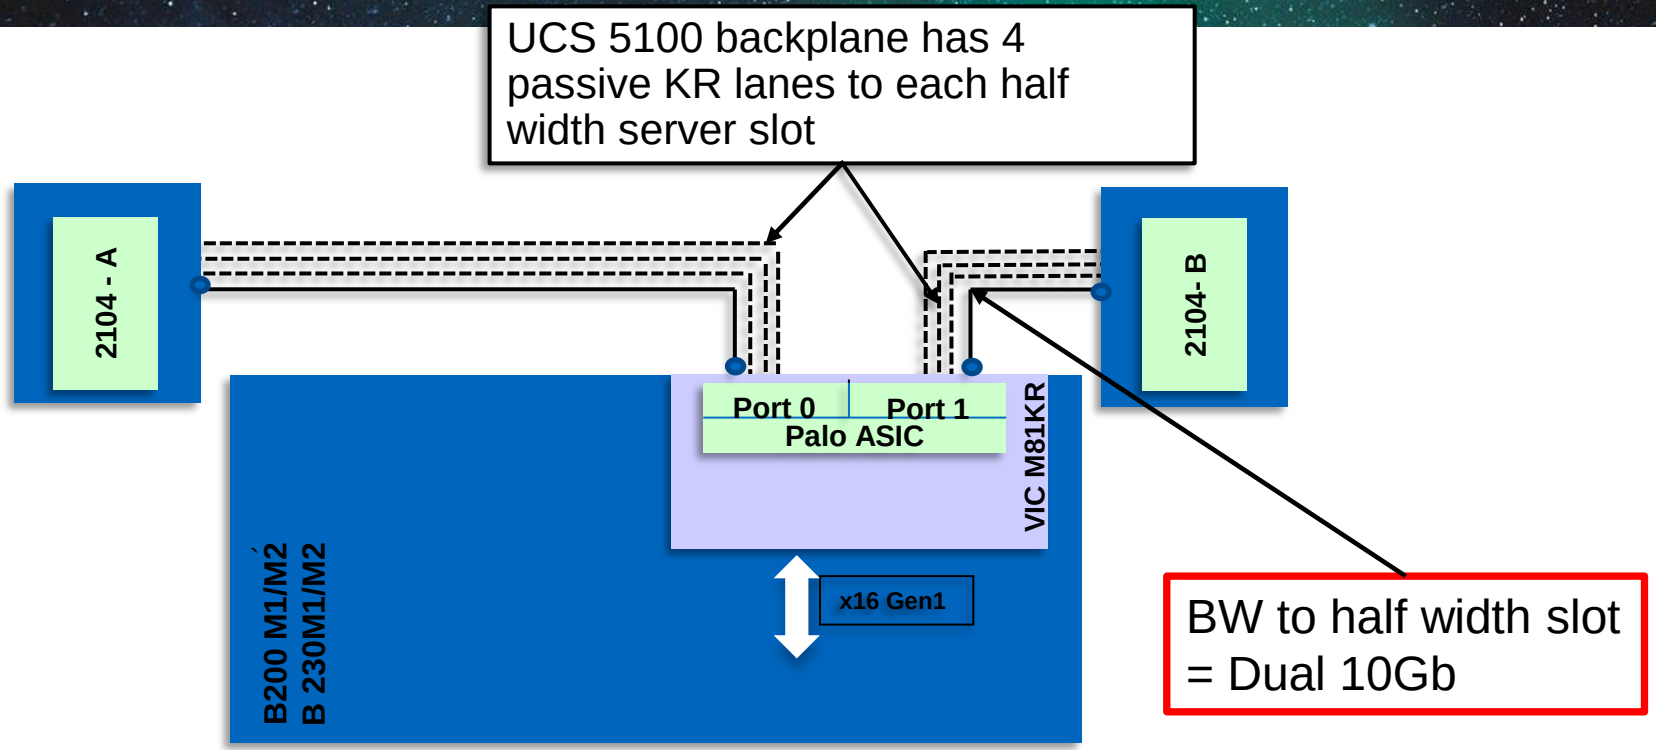

IOM 2104 with VIC 1280 in M2 Blades

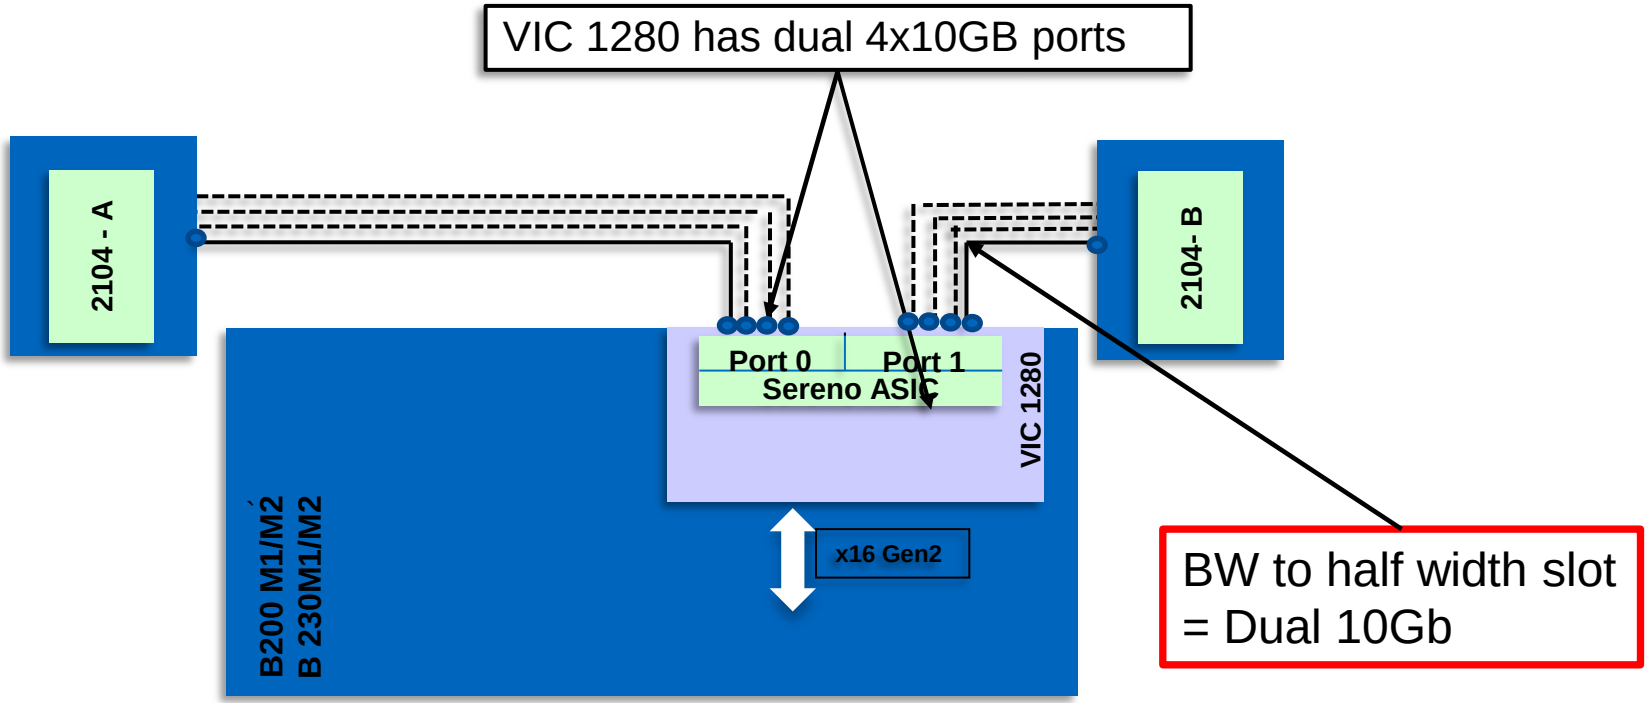

IOM 2208 with M81KR in M1/M2 Blades

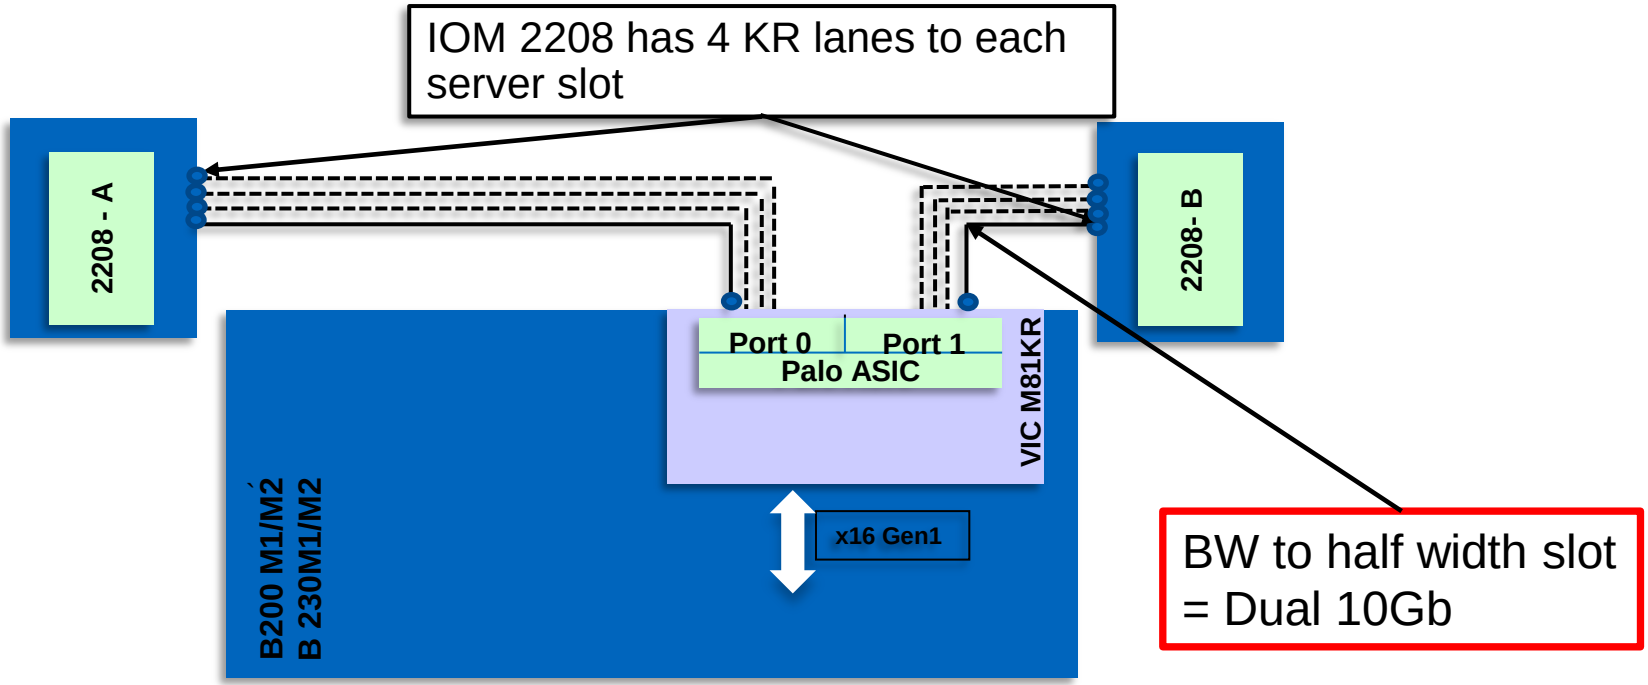

IOM 2208 with VIC 1280 in M2 Blades

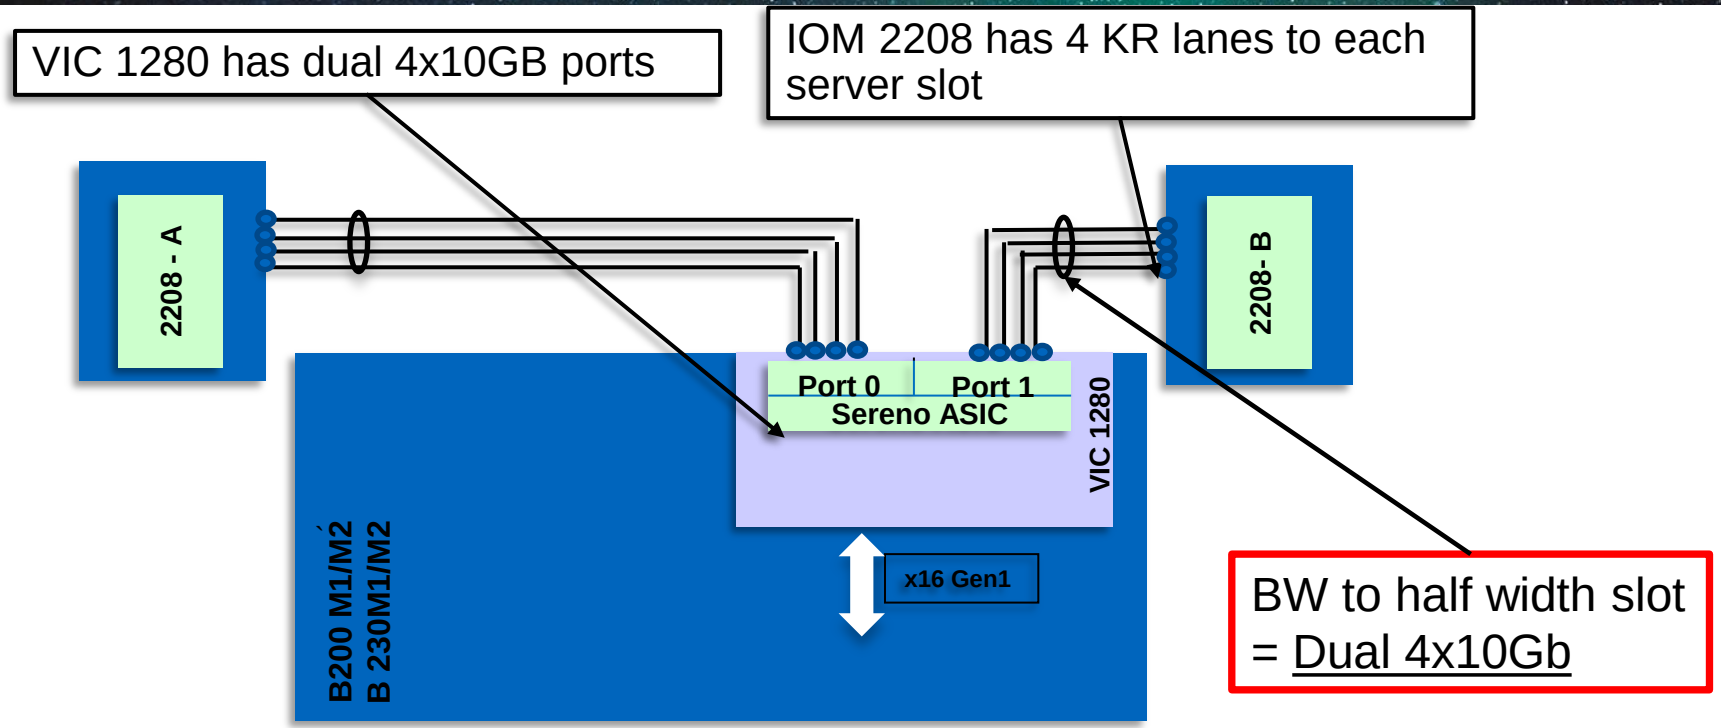

IOM 2204 with M81KR in M1/M2 Blades

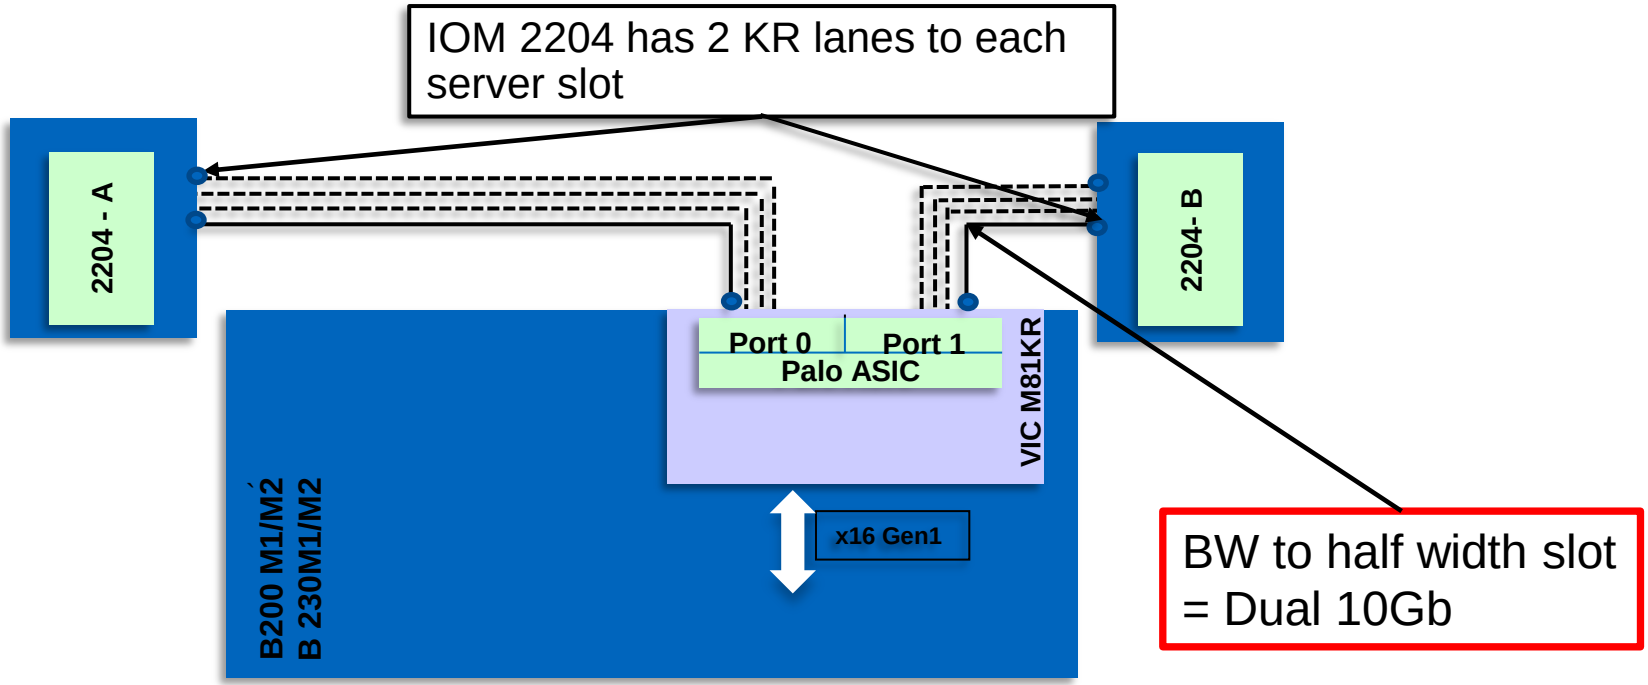

IOM 2204 with VIC 1280 in M2 Blades

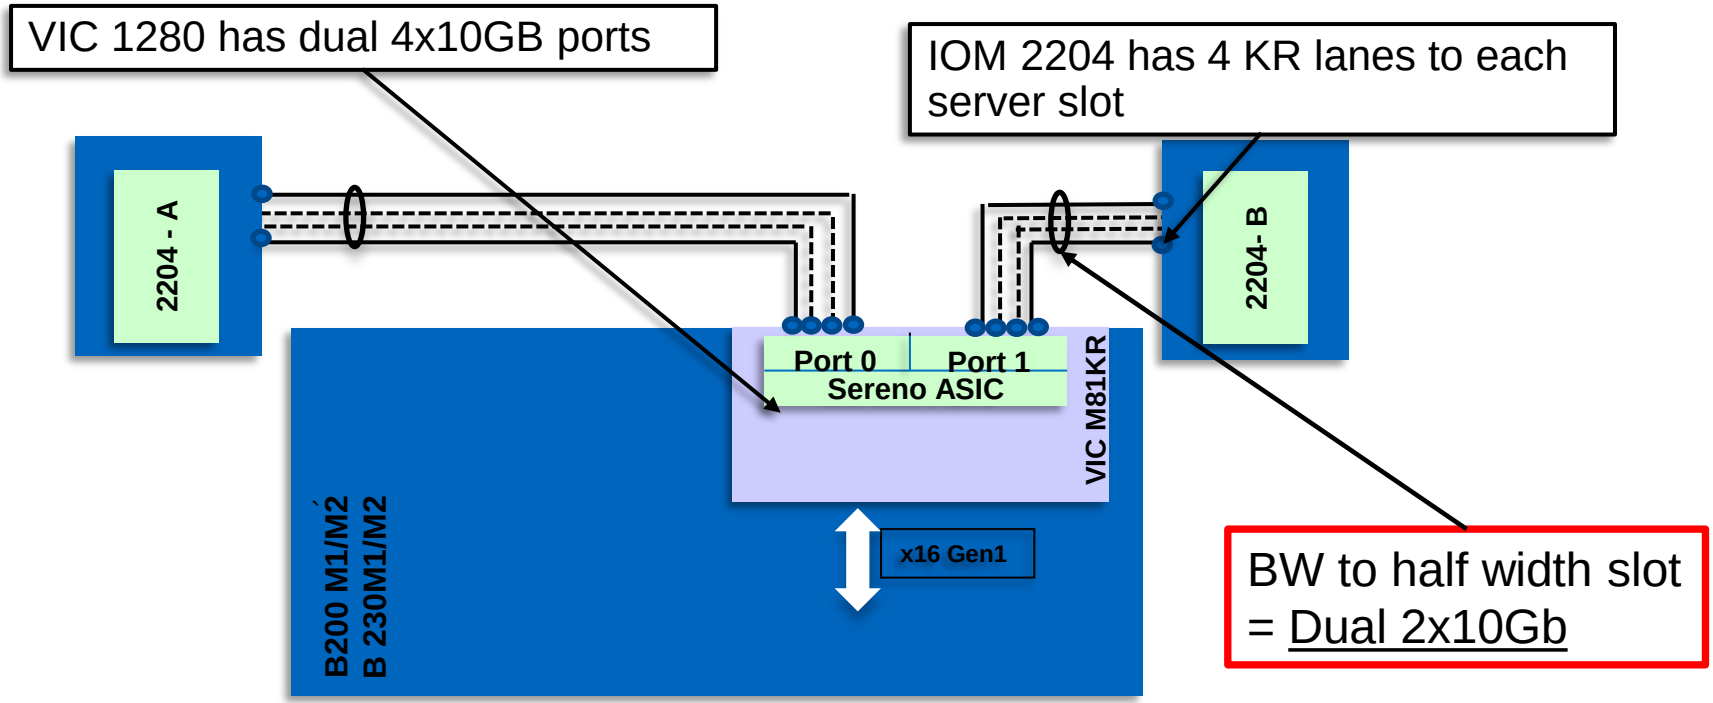

UCS B200 M3

Half width blade form factor

Modular LOM Slot

Mezzanine Slot

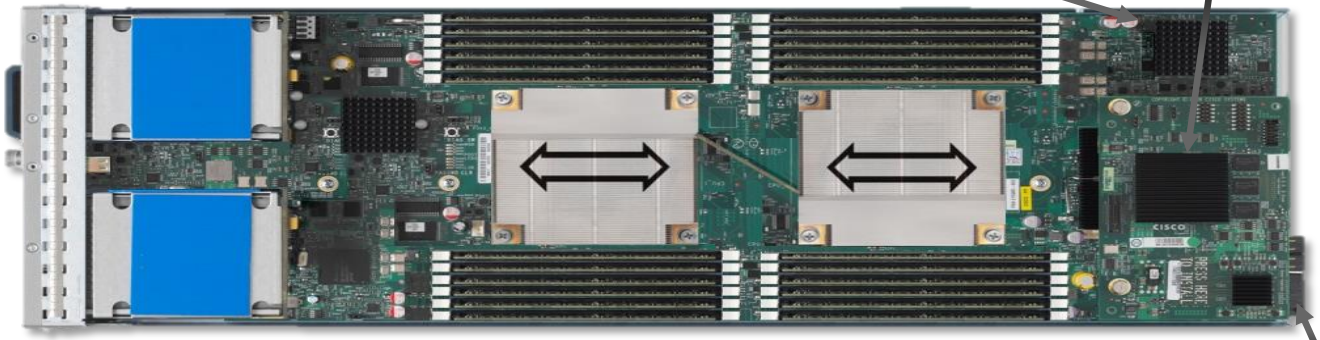

Connectors to the backplane

1. mLOM Slot: Supports only the VIC 1240
2. Mezz Slot: Supports VIC 1280 & 3rd party Gen 3 Mezz cards
3. M81KR *NOT* supported on the M3 (Romley blades)

B200 M3 I/O Block Diagram

Modular LOM and Mezzanine slot

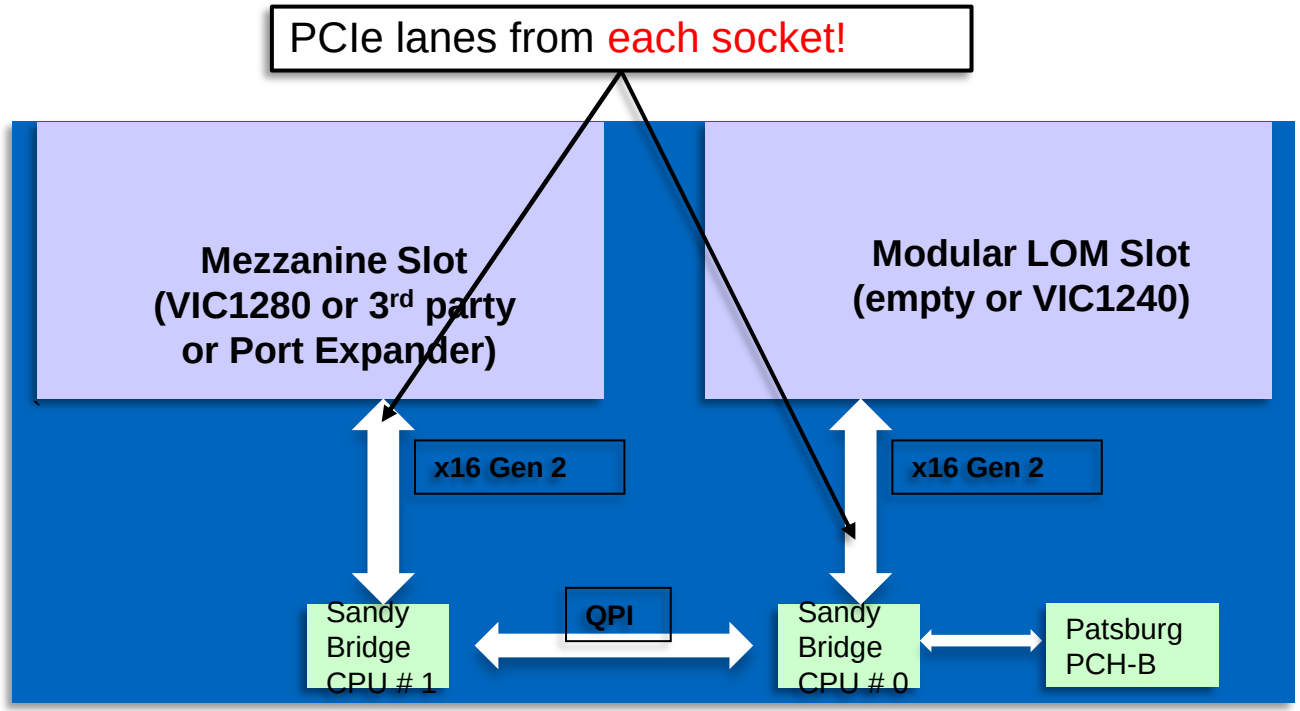

Introducing UCS VIC 1240 Modular LOM

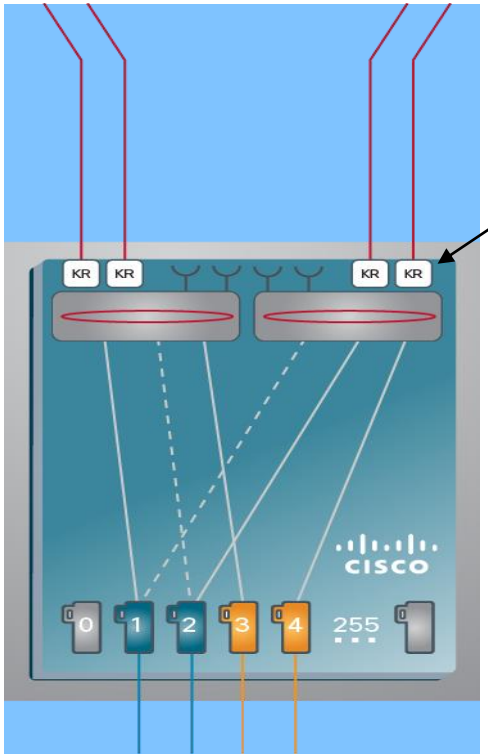

- Based on Sereno ASIC (VIC 1280 uses the same ASIC)
- PCIe Devices 256 (vNICs or vHBA)
- Support VM-FEX
- Base option supports dual 2x10Gb (2 lanes to each FI)

Port Expander Card for VIC 1240

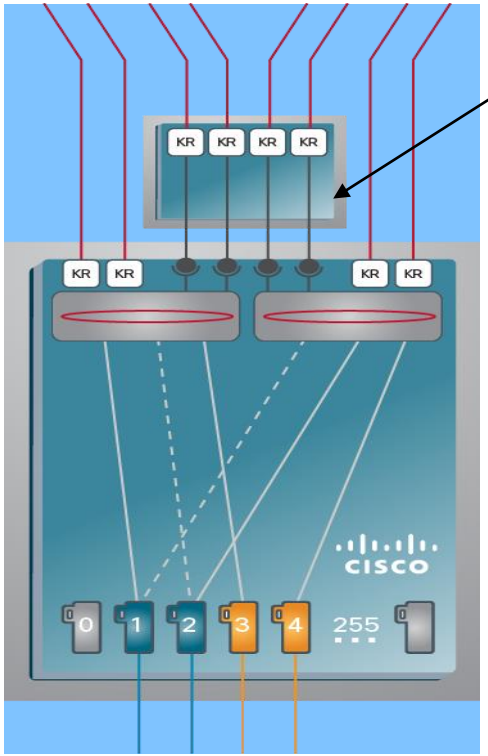

- Option to enable all port of 2nd Gen VIC ASIC (Sereno)
- 4 ports to each FI
- Fits in the Mezzanine slot of B200M3
- Port Expander has no PCIe presence
- It is a “passive connector” device

B200 M3 I/O Block Diagram

VIC 1240 and Port Expander Card for VIC 1240

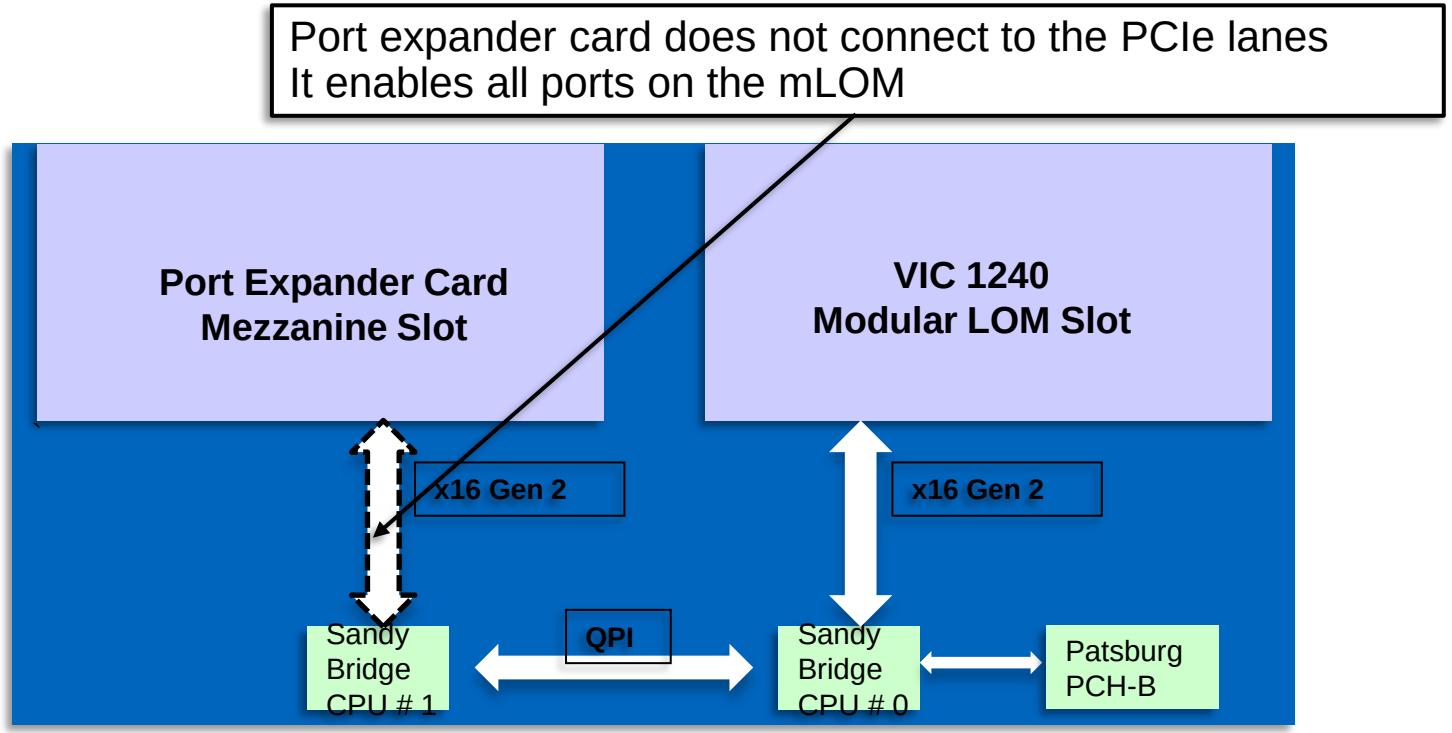

Backplane Lanes for B200 M3

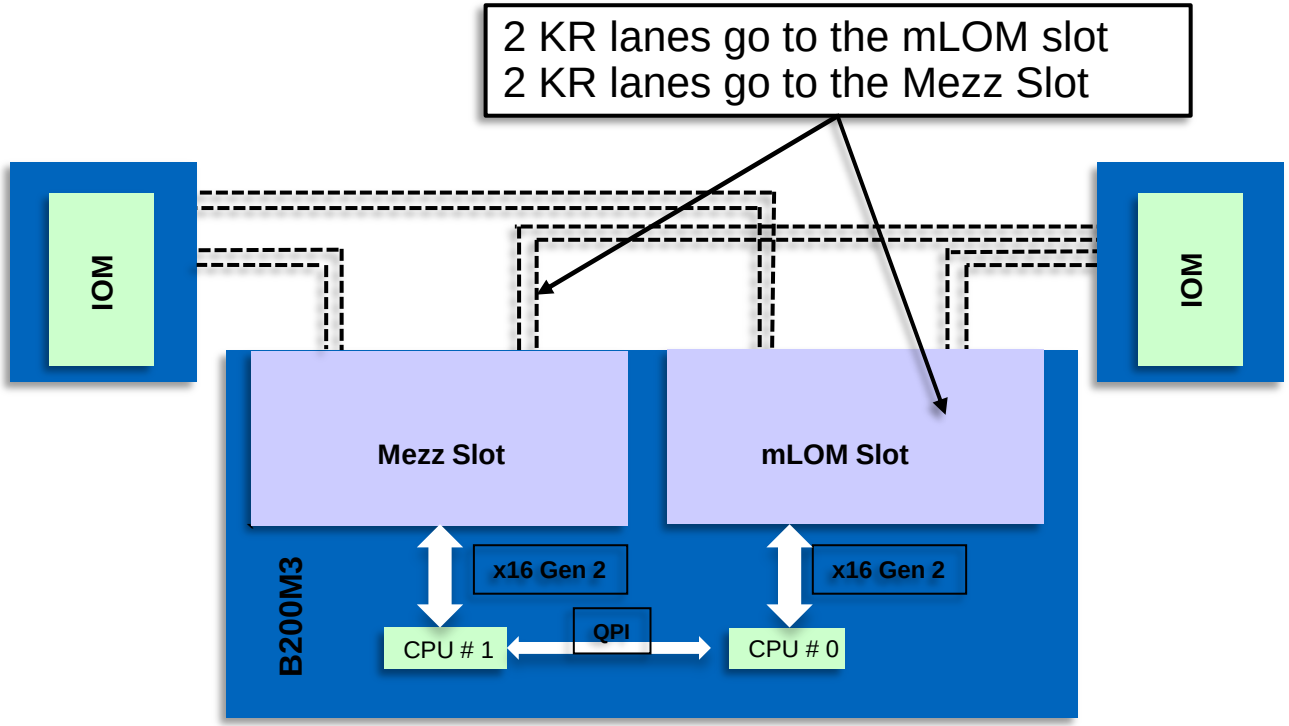

IOM 2208 with VIC1240 in B200M3

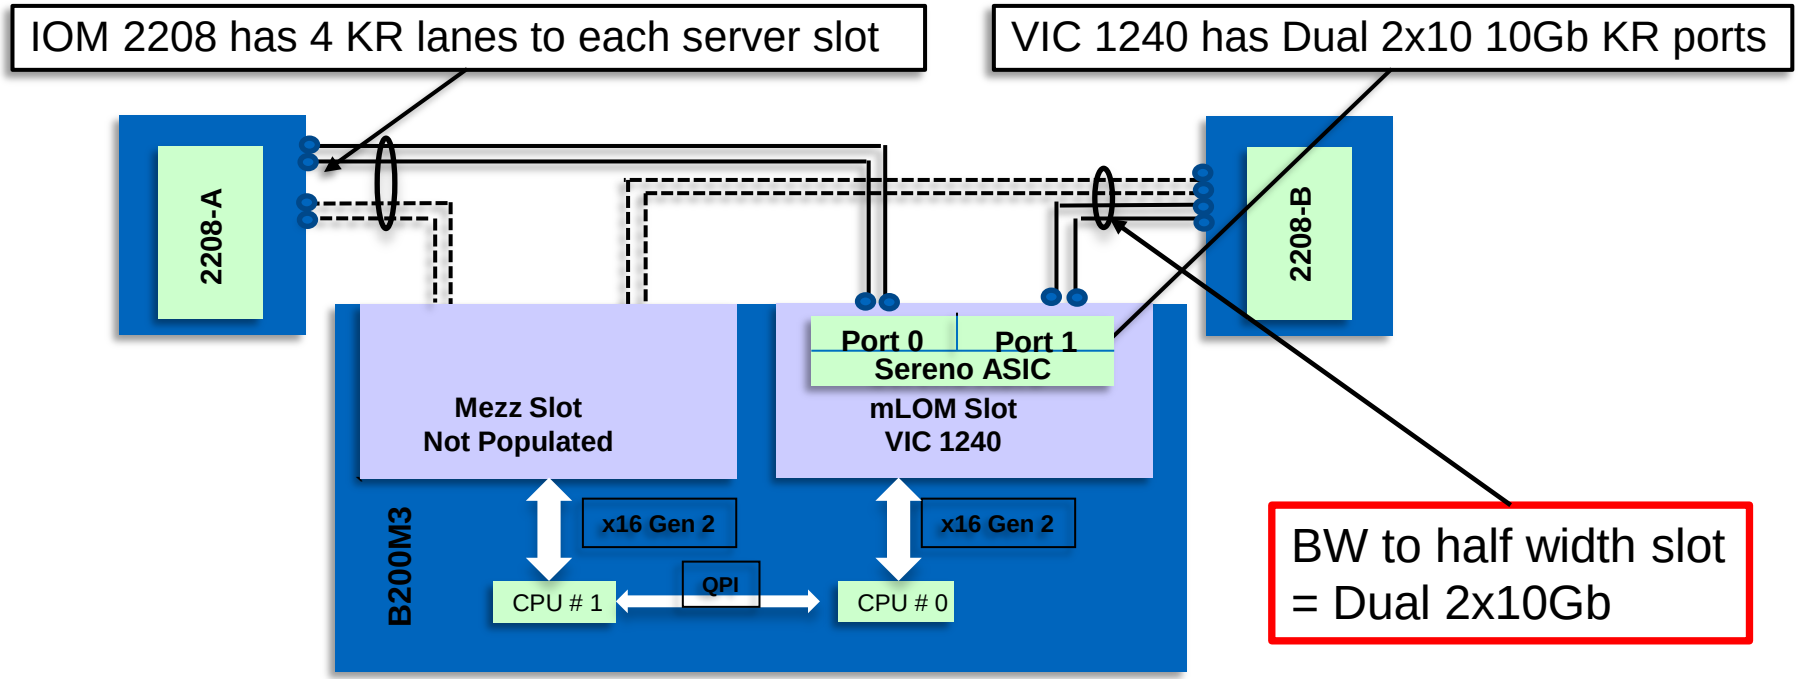

IOM 2208 with VIC1240 & Port Expander in B200M3

Port Expander for VIC 1240 enables 2 additional ports of Sereno ASIC to each fabric

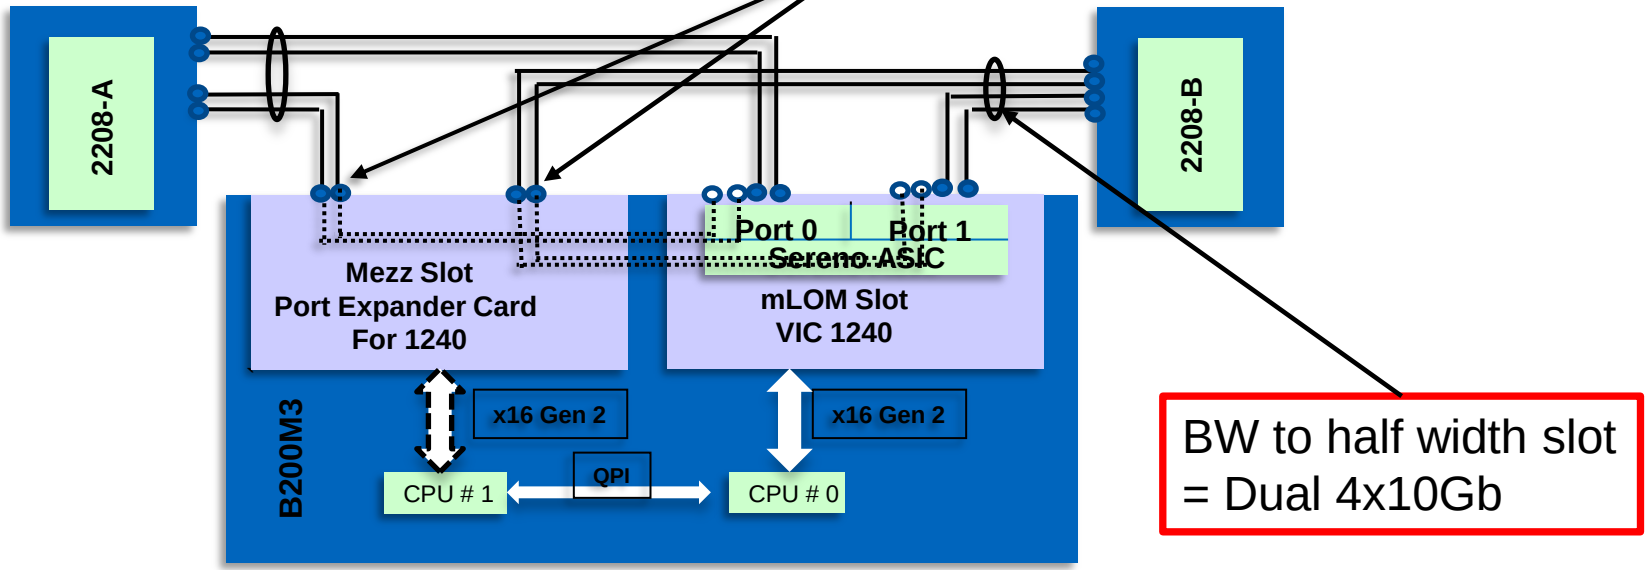

BW to half width slot = Dual 4x10Gb

Full BW of 2nd Gen VIC ASIC exposed

IOM 2208 with VIC1240 & VIC 1280 in B200M3

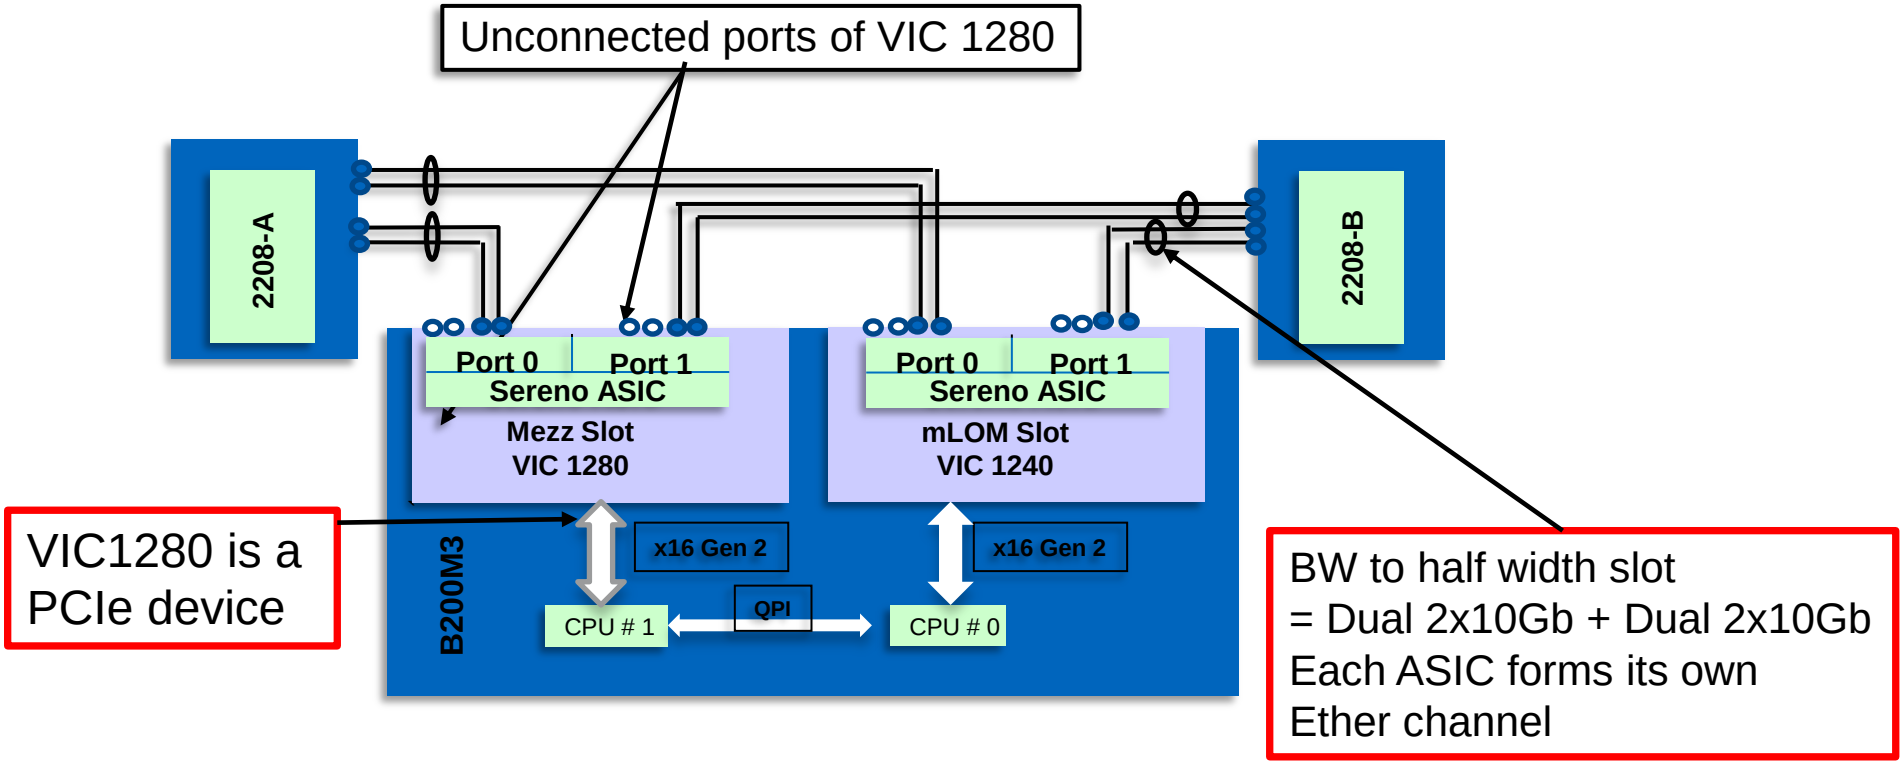

IOM 2204 with VIC1240 in B200M3

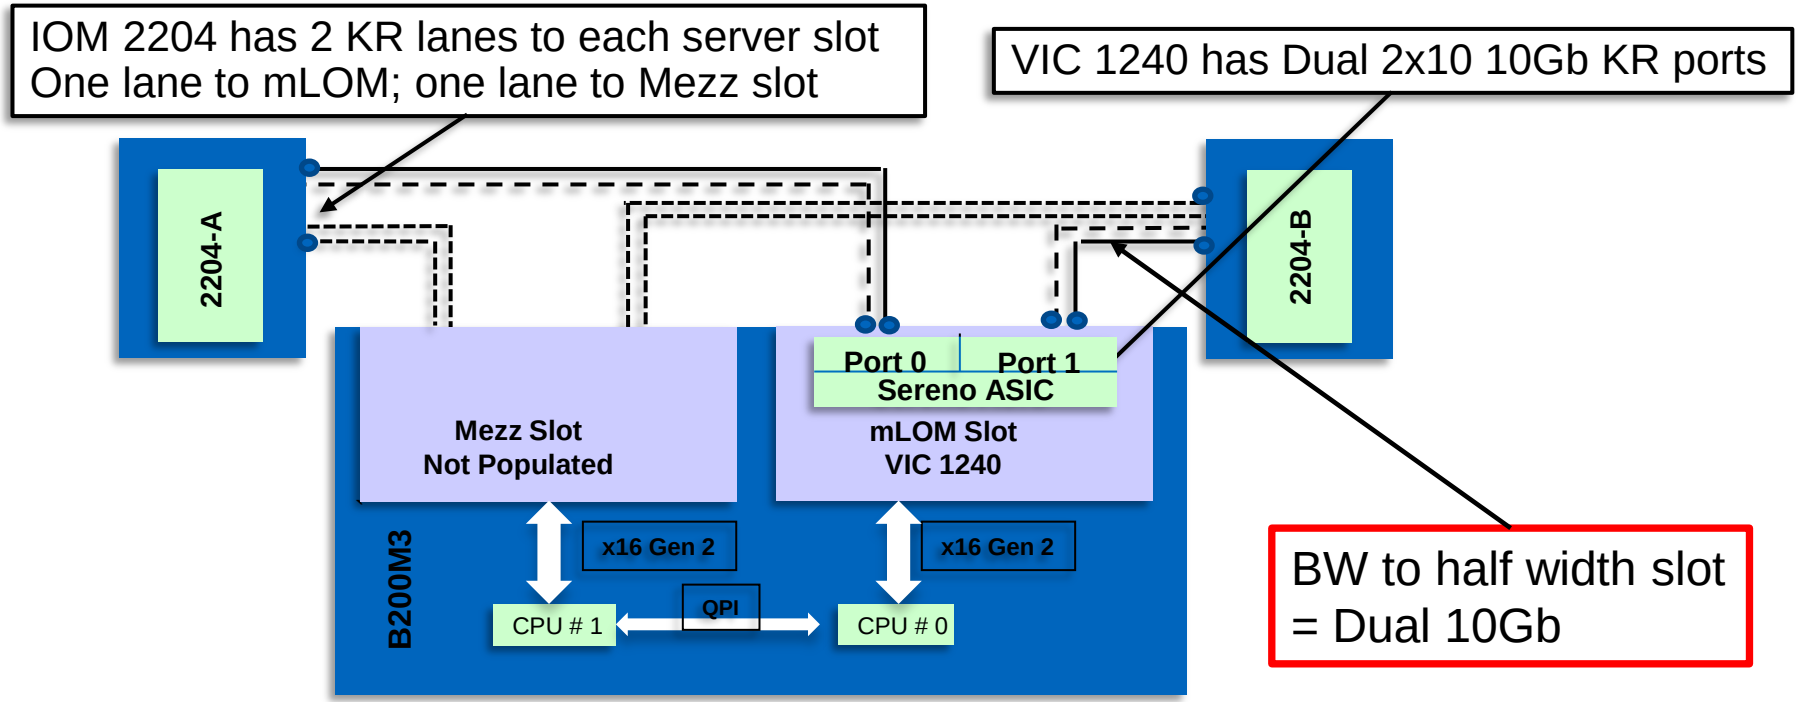

Topology Designs For Maximum Bandwidth

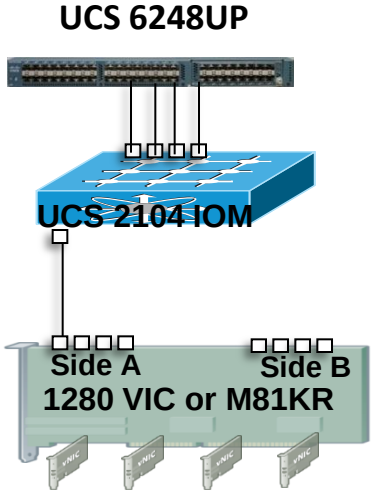

- Shared IOM uplink bandwidth of 10Gbps
- vNIC Burst up to 10Gbps
- Shared IOM Uplink with 1 server
- Host port pinned to a discrete IOM uplink

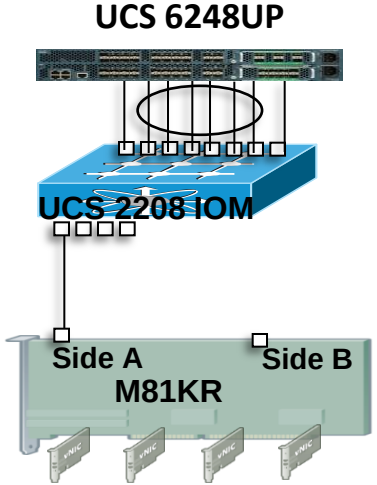

- Shared IOM uplink bandwidth of 80Gbps
- vNIC Burst up to 10Gb
- Shared IOM Port-Channel with 8 servers
- Host port pinned to a discrete IOM port-channel

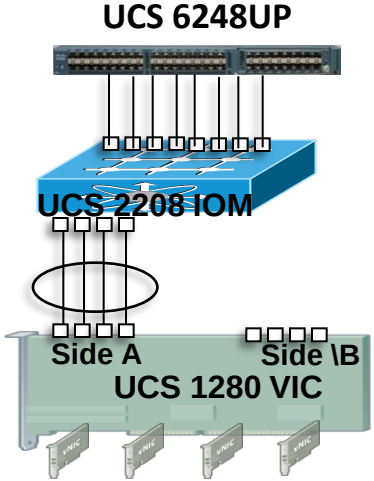

- Dedicated IOM uplink bandwidth of 10Gbps
- vNIC Burst up to 10Gbps *(IOM uplink limitation)
- Dedicated IOM Uplink
- Host port-channel pinned to discrete IOM uplink

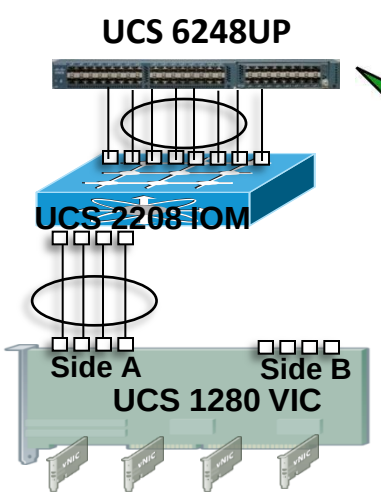

- Shared IOM uplink bandwidth of 80Gbps
- vNIC Burst up to 32Gbps *(PCIe Gen 2 limitation)
- Shared IOM Port-Channel with 8 servers
- Host port-channel pinned to the IOM port-channel

Virtual Interfaces (VN-TAG)

How Do Servers Communicate?

- We know servers with one mezz card present two (M81KR and non-Cisco adapters) or 2 x (4 x 10G) Base-KR ‘external’ or ‘northbound’ interfaces
- The OS knows nothing of this
- The OS sees PCI devices on the bus and loads device drivers for those devices
  - In UCS, the Service Profile controls the interfaces the OS sees
  - E.g.: a blade can be shown 6 x 10GE NICs and 2 x HBAs while another sees 8 x 10GE NICs and no HBAs
- This means the northbound physical interfaces between the adapter and the IOM can carry both Ethernet and FC traffic for several vNICs. We need a mechanism to identify the origin server
  - → Concept of Virtual Interface or VIF (see next slide)

Server Interface Virtualisation (Adapter FEX)

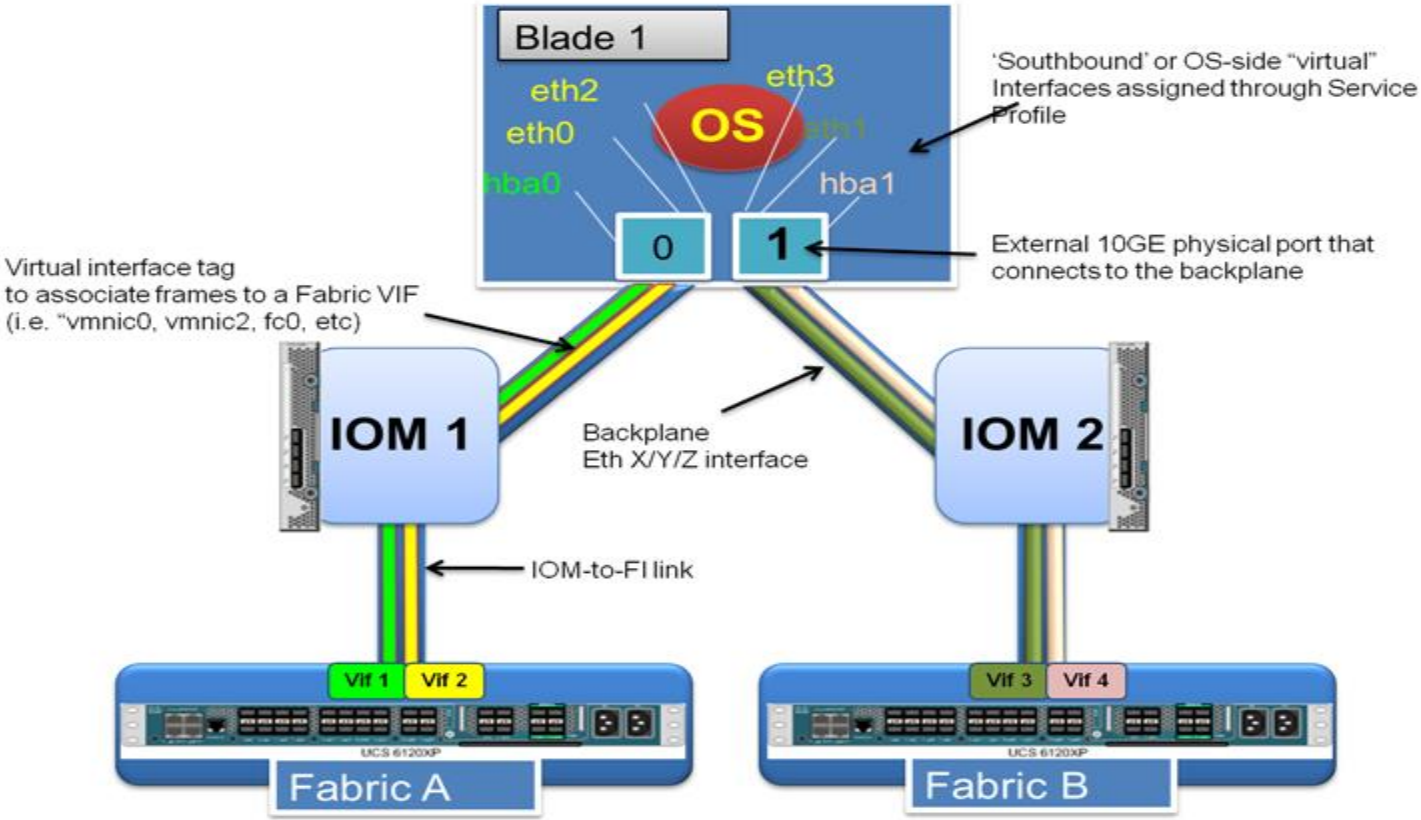

VN-Tag: Instantiation of Virtual Interfaces

- Virtual interfaces (VIFs) help distinguish between FC and Eth interfaces
- They also identify the origin server
- VIFs are instantiated on the FI and correspond to frame-level tags assigned to blade mezz cards
- A 6-byte tag (VN-Tag) is preprended by Palo and Menlo as traffic leaves the server to identify the interface
  - VN-Tag associates frames to a VIF
- VIFs are ‘spawned off’ the server’s EthX/Y/Z interfaces (examples follow)

VN-Tag at the Adapter (Mezz Card) Level

```
Hardy-the-new-A# connect adapter 1/1/1
adapter 1/1/1 # connect
adapter 1/1/1 (top):1# attach-mcp
adapter 1/1/1 (mcp):1# vnic
vnic id       : internal id of vnic, use for other vnic cmds
vnic name/mac : ucsm 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 ucsm      : ucsm id for this vif
vif idx       : switch id for this vif
vif vlan      : default vlan for traffic
vif state     : state of vif
-------------------------------------- --------- --------------------------
                v n i c                   l i f            v i f
id  name           type    bb:dd.f state lif state uif  ucsm   idx vlan state
--- -------------- ------- ------- ----- --- ----- --- ----- ----- ---- -----
  5 vnic_1         enet    08:00.0 UP      2 UP    - 0   970   860    1 UP
                                                   =>1   969   873    1 UP
  6 vnic_2         enet    09:00.0 UP      3 UP    - 0   972   861    1 UP
                                                   =>1   971   874    1 UP
  7 vnic_3         enet    0a:00.0 UP      4 UP    - 0   974   862    1 UP
                                                   =>1   973   875    1 UP

adapter 1/1/1 (mcp):2#
```

VIFs

- Ethernet and FC are muxed on the same physical links → concept of virtual interfaces (vifs) to split Eth and FC
- Two types of VIFs: veth and vfc
  - Veth for Ethernet and FCoE; vfc for FC traffic
- Each EthX/Y/Z or Po interface typically has multiple vifs attached to it to carry traffic to and from a server
- To find all *vifs* associated with a EthX/Y/Z or Po interface, do this:

```
UCS-TME-LAB-A(nxos)# sh vifs interface port-channel 1282

Interface       MAX-VIFS VIFS
--------------- -------- ----------------------------------------------------------
Po1282          0        Veth2046, Veth2047, Veth2048, Veth10244,
UCS-TME-LAB-A(nxos)#
```

FCoE: a Vfc Bound to a Veth

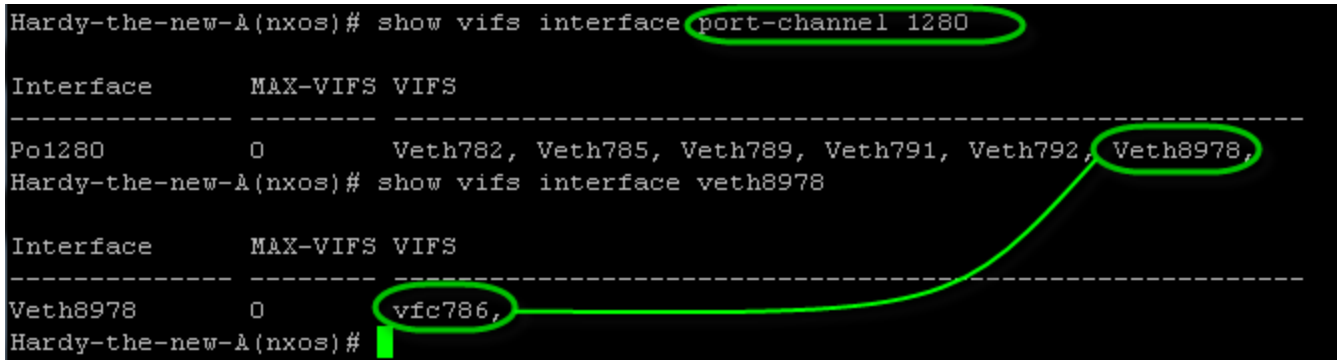

Another Way to Find VIFs:

```
UCS-BRU-STAGING-A# show service-profile circuit name STATICSETUP-WIN2K3
Service Profile: STATICSETUP-WIN2K3
Server: 1/1
    Fabric ID: A
        VIF         vNIC        Link State  Overall Status  Admin Pin  Oper Pin   Transport
        ----------  ----------  ----------  --------------  ---------  ---------  ---------
                41              Up          Active          0/0        1/20       Unknown
               697  eth0        Up          Active          0/0        0/0        Ether
               698  vhba0       Up          Active          0/0        0/0        Fc
              8890              Up          Active          0/0        0/0        Ether
UCS-BRU-STAGING-A#
```

```
UCS-BRU-STAGING-A(nxos)# sh run int vet697
version 4.0(1a)N2(1.1e)

interface vethernet697
  switchport trunk native vlan 27
  switchport trunk allowed vlan 27
  bind interface Ethernet1/1/1
  no pinning server sticky
  pinning server pinning-failure link-down

UCS-BRU-STAGING-A(nxos)# sh int ve697
vethernet697 is up
    Bound Interface is Ethernet1/1/1

  Description: server 1/1, VNIC
  Encapsulation ARPA
  Port mode is trunk
  Last clearing of "show interface" counters never
  5 minute input rate 0 bytes/sec, 0 packets/sec
  5 minute output rate 0 bytes/sec, 0 packets/sec
  Rx
    11459643 input packets     7313025615 bytes
  Tx
    24186526 output packets      11202354806 bytes

UCS-BRU-STAGING-A(nxos)#
```

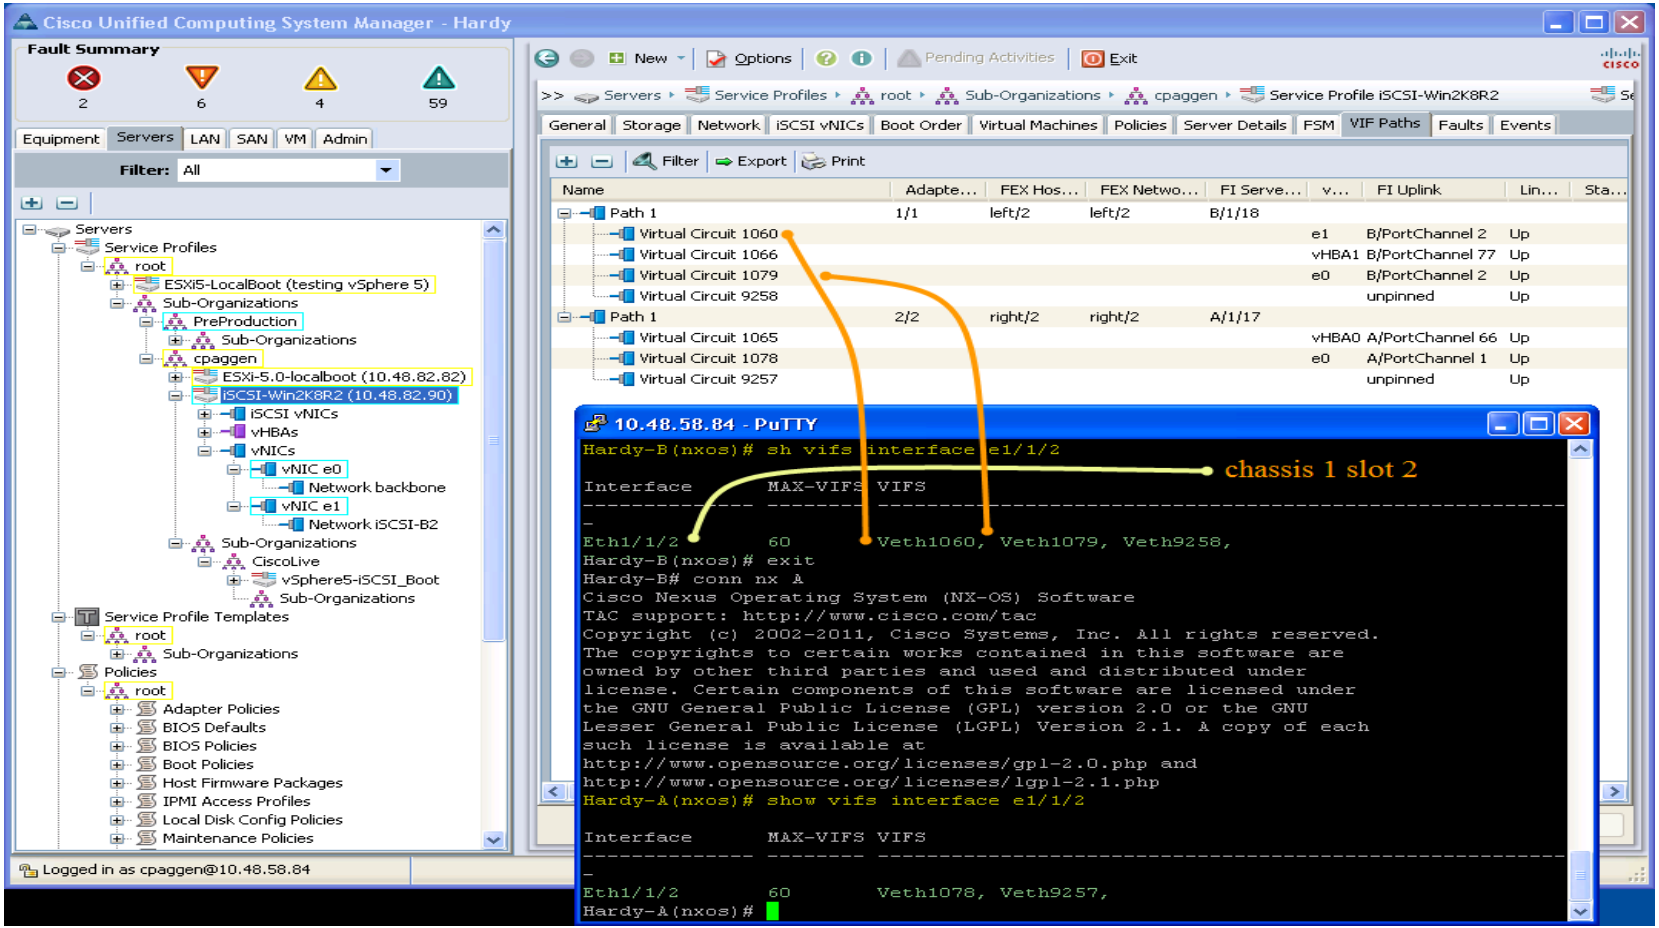

UCS 6248 Hardware Diagram

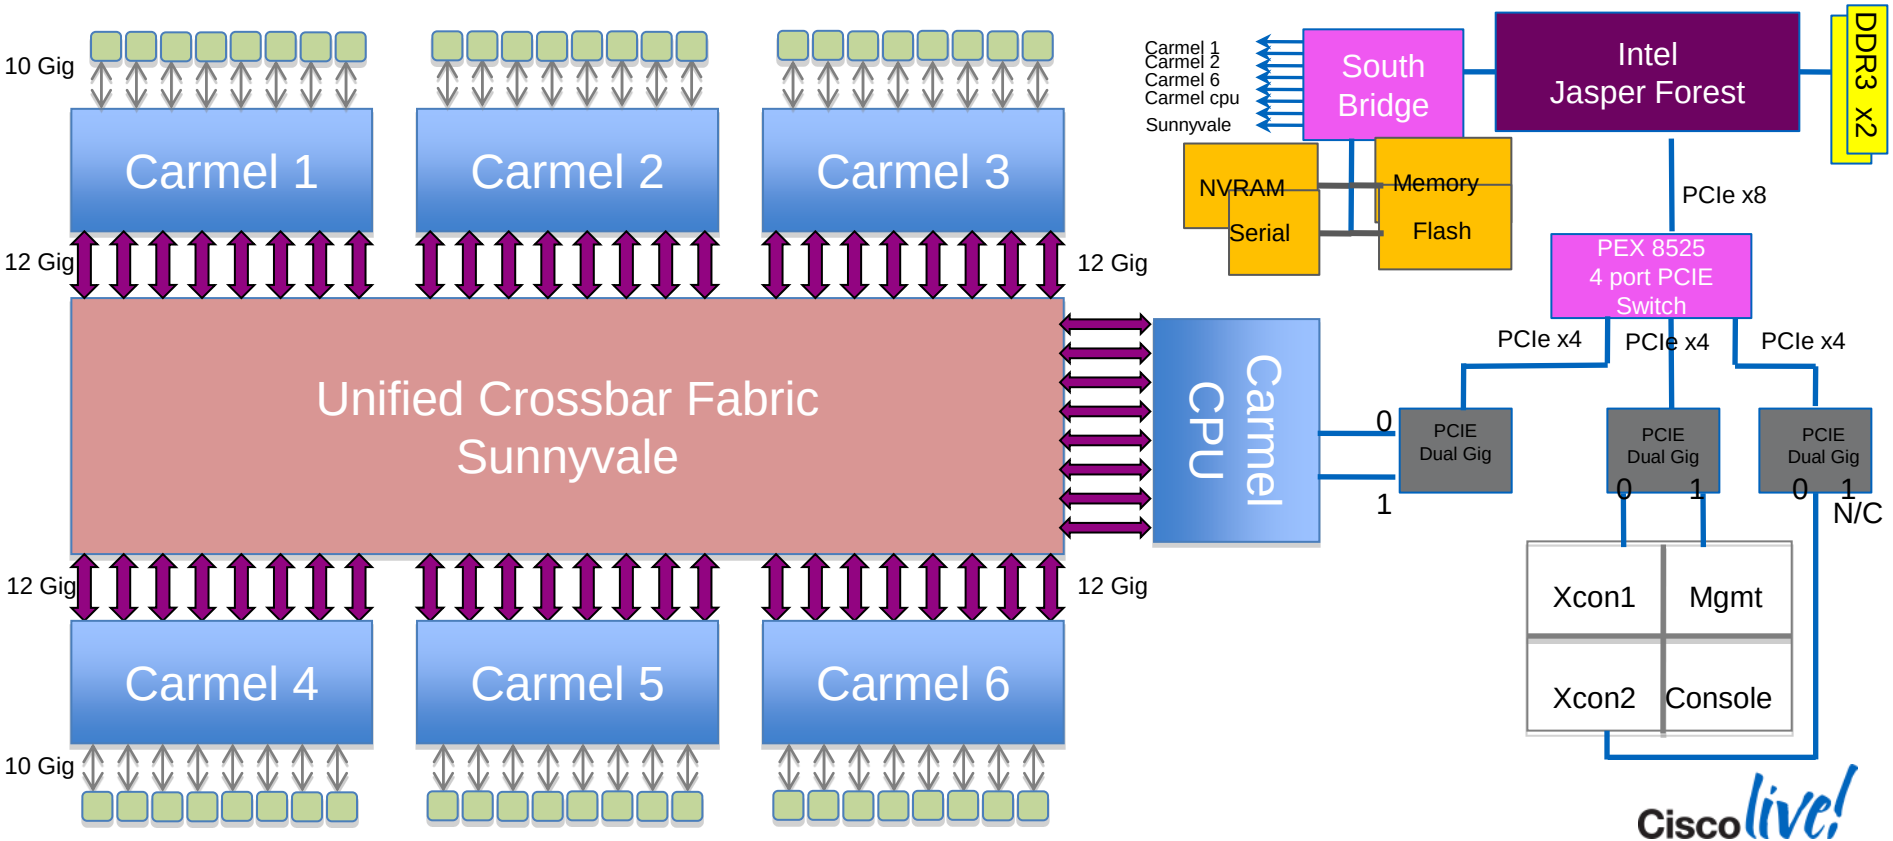

Maximising the VIF Count

Fabric Interconnect VIF calculation

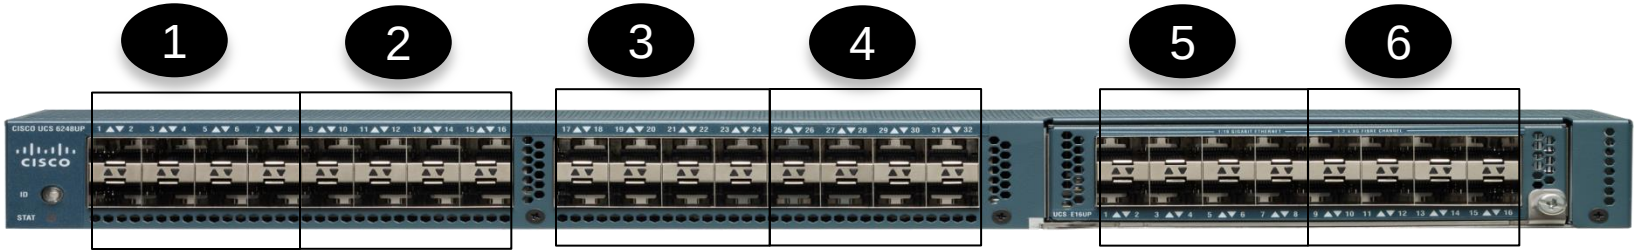

- Every 8 10GbE ports (on FI) are controlled by the same Unified Port ASIC
- Connect fabric links from IOM to the FI to the same UPC
- Virtual Interface (VIF) namespace varies depending on number and how the fabric links are connected to the FI ports.
  - Connecting to the same UPC (a set of eight ports), Cisco UCS Manager maximises the number of VIFs used in service profiles deployed on the servers.
  - If uplink connections are distributed across UPC, the VIF count is decreased. For example, if you connect seven (IOM) fabric links to (FI) ports 1-7, but the eighth fabric link to FI port 9, the number of available VIFs is based on 1 link – IOM port 8 to FI port 9.

UCS FI and IOM Connectivity

Fabric Interconnect VIF calculation cont'd

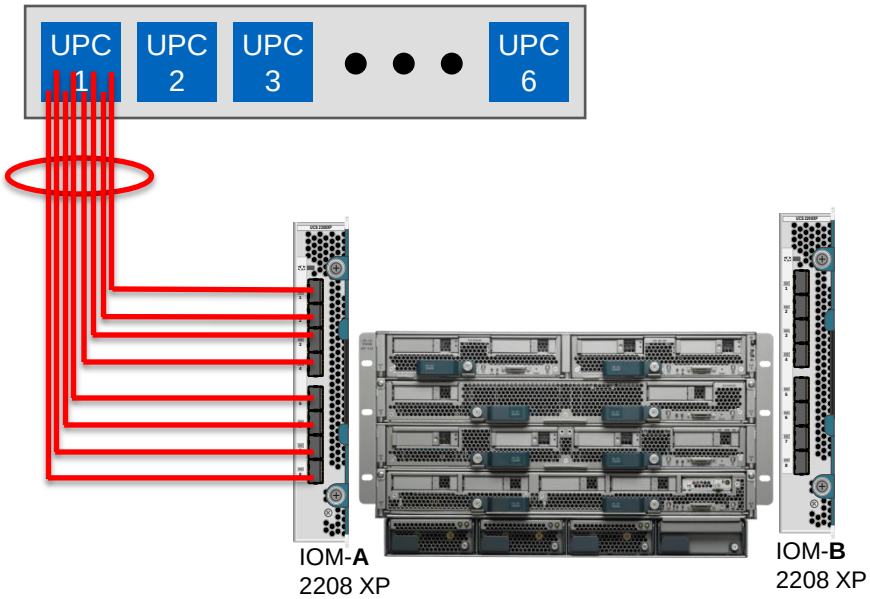

- Recommended
- Maximise number of available VIFs to the host

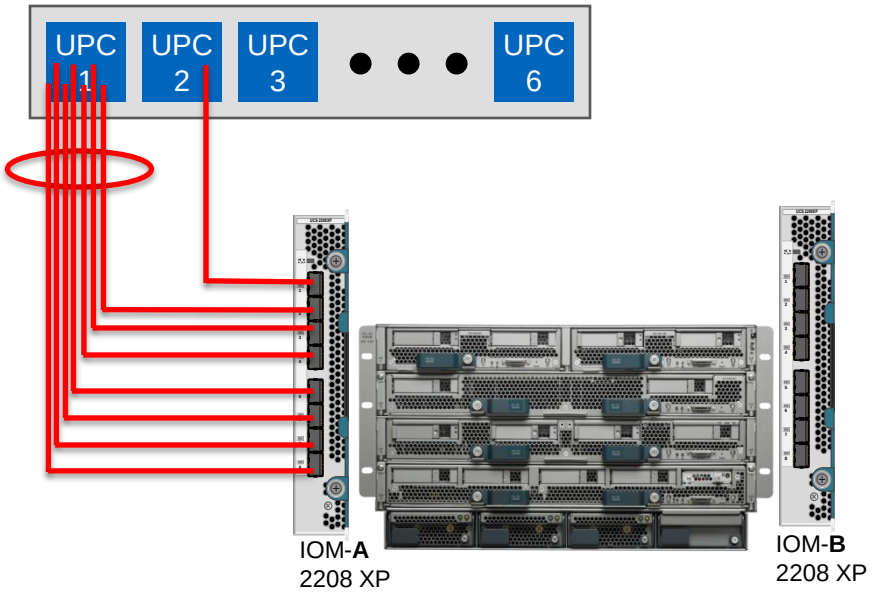

- Not recommended
- Minimal number of VIFs to the host




Topologies

Recommended Topology for Upstream Connectivity

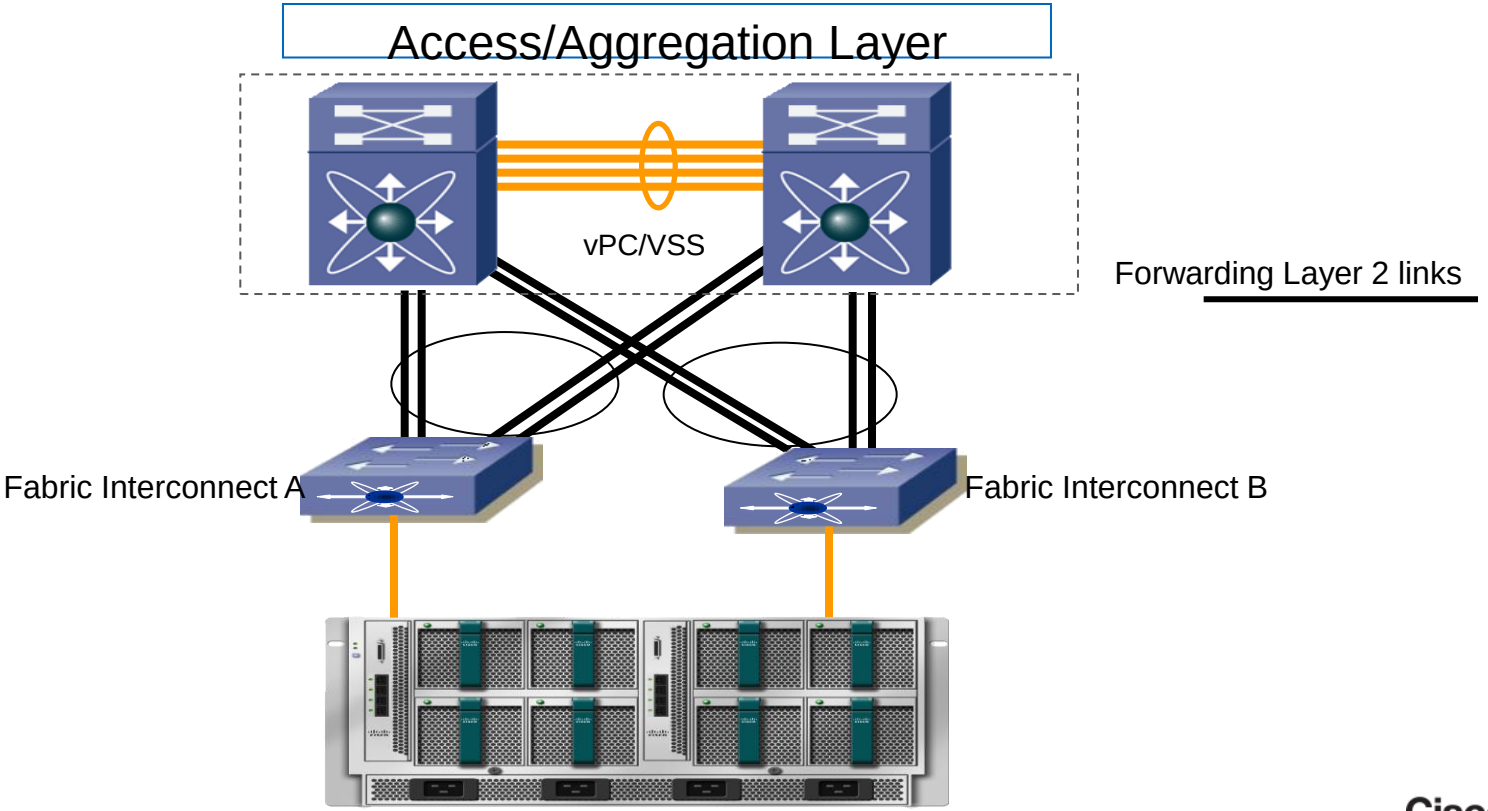

Layer 2 Disjoint Topology

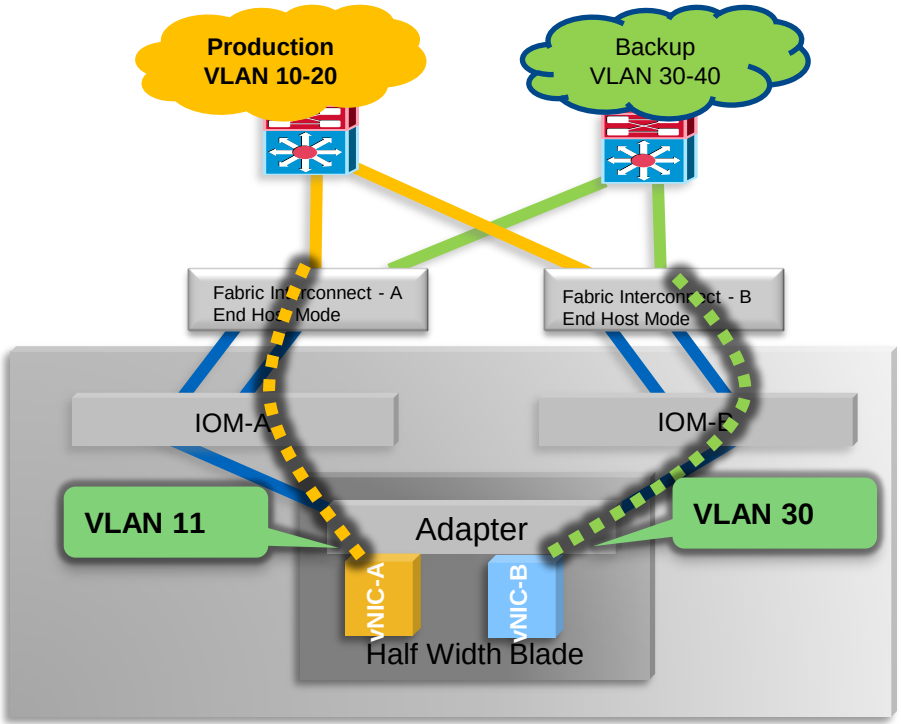

- A vNIC can only participate in one L2 network upstream
- Both dynamic and static pinning methods are supported

Inter-Fabric Traffic Example

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

- VM1 Pinned to VNIC0
- VM4 Pinned to VNIC1

- VM1 on VLAN 10
- VM4 on VLAN 10

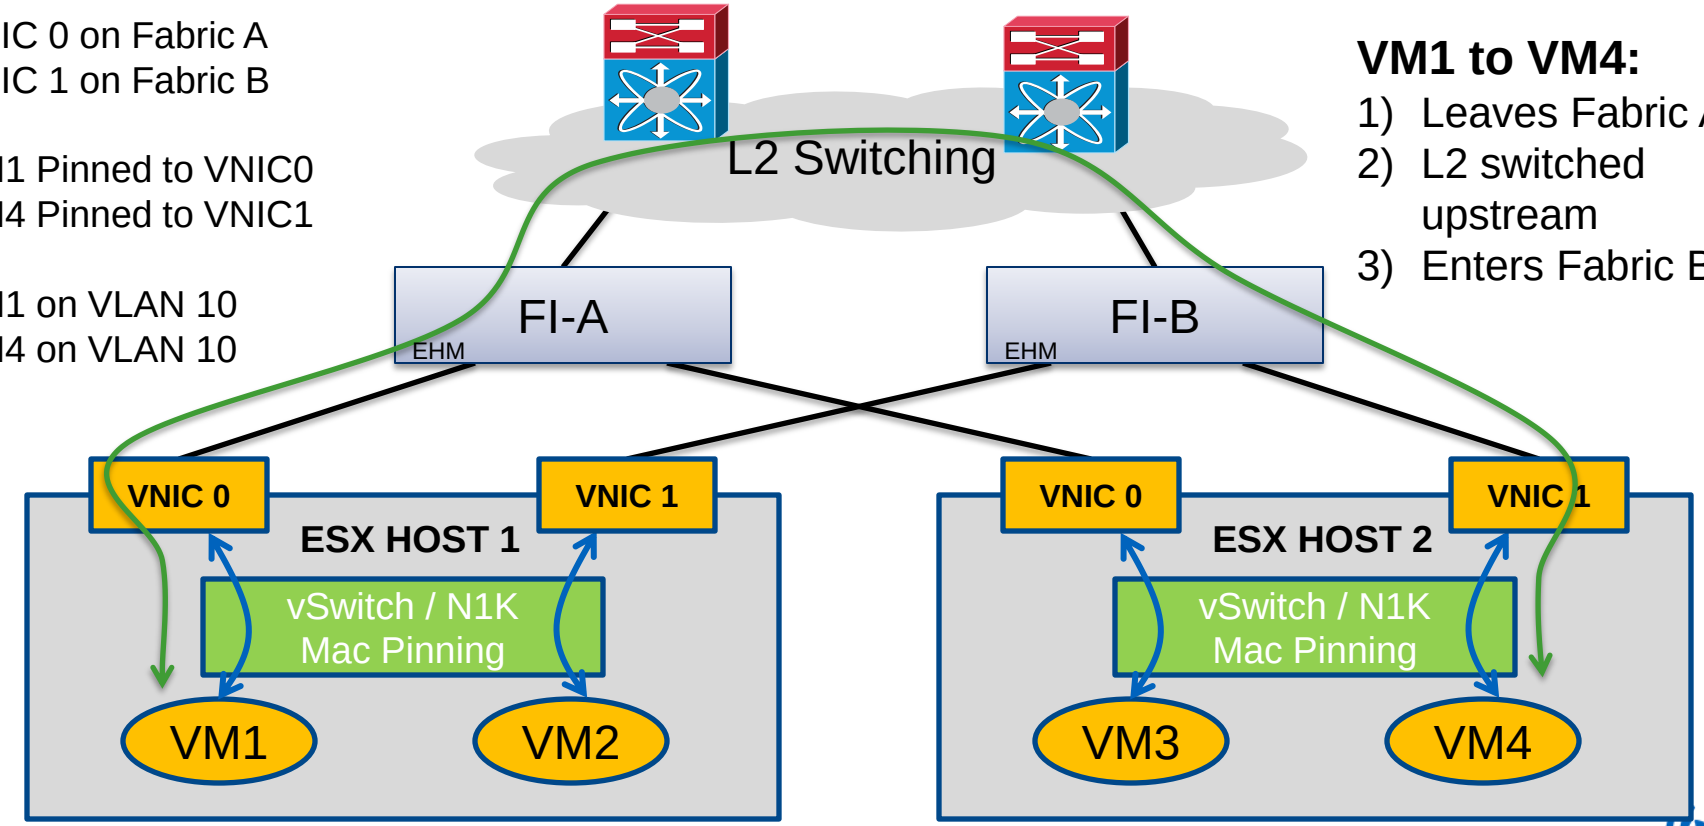

VM1 to VM4:

1) Leaves Fabric A
2) L2 switched upstream
3) Enters Fabric B

C-Series Integration into UCSM

C-Series UCSM Integration

Before UCSM 2.1 or without VIC1225

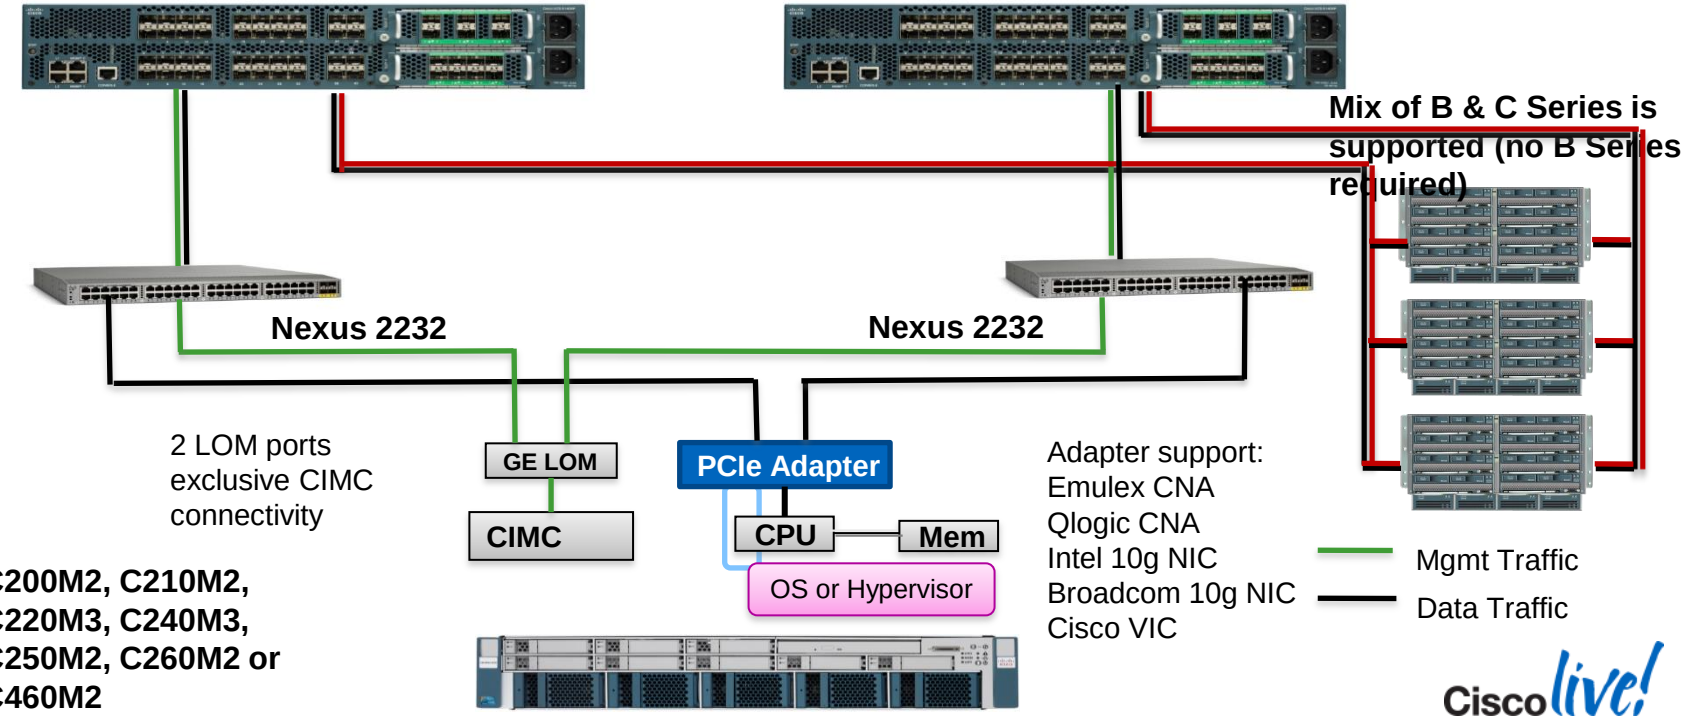

C-Series UCSM Integration

Single Wire Management with VIC1225 and UCSM 2.1

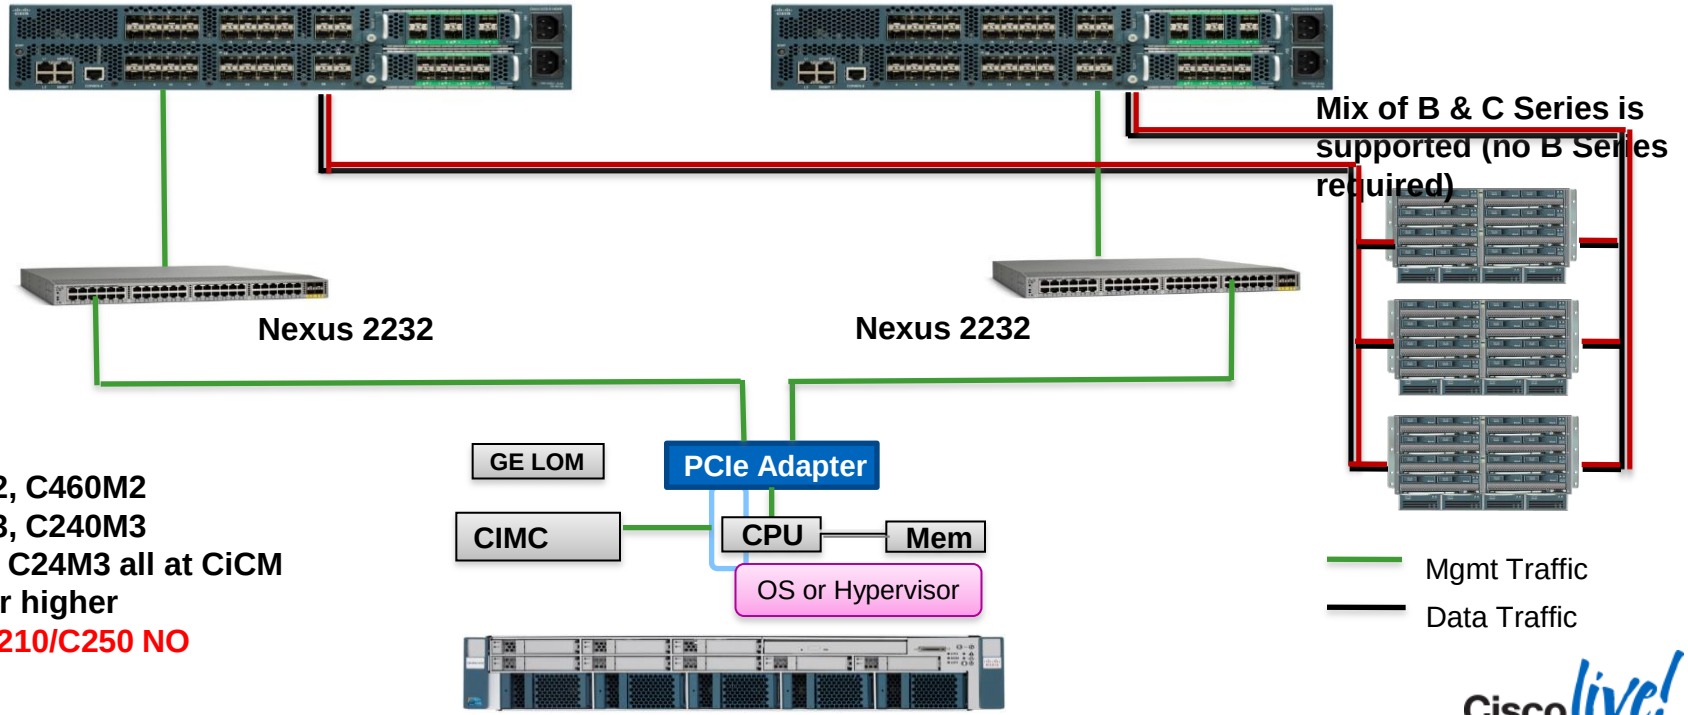

C260M2, C460M2
C220M3, C240M3
C22M3, C24M3 all at CiCM 1.4(6) or higher
C200/C210/C250 NO

UCSM 2.2: Direct Connect with VIC1225 (no FEX)

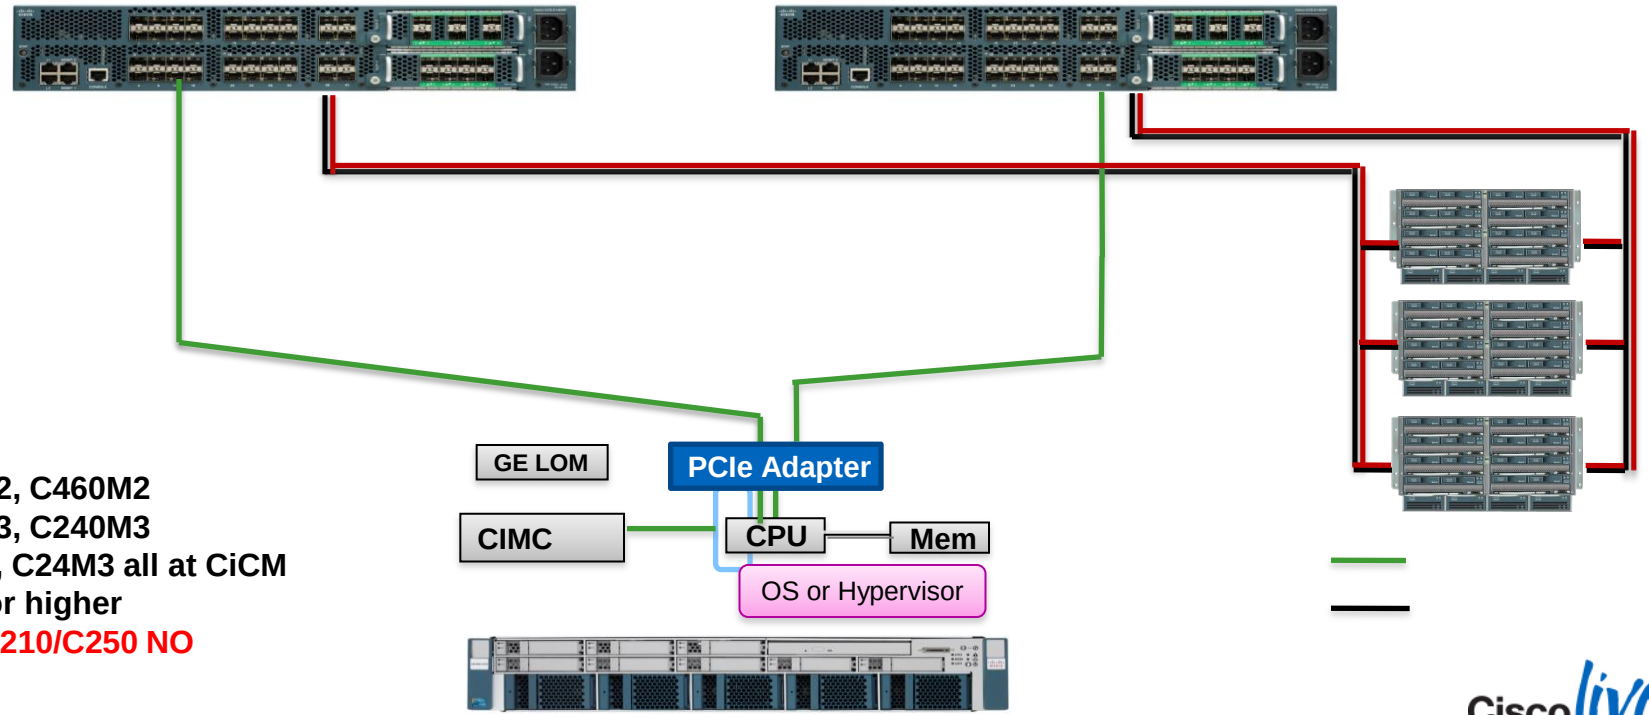

C260M2, C460M2
C220M3, C240M3
C22M3, C24M3 all at CiCM 1.4(6) or higher
C200/C210/C250 NO




Wrapping Up

Recap

- UCS is first and foremost a server, not a switch
- Gen 1 vs Gen 2 components
  - 2208 and VIC 1280 allow port-channels
- VIC 1280 with 2208XP for maximum bandwidth
- End-host mode forwarding rules: dynamic pinning
- Preferred mode of operation should always be end-host mode
  - Very much plug and play, scalability, L2 multipathing, fabric failover
  - Switch mode: spanning-tree, practically no user configuration possible
- Operational consistency: C-series integration




Q & A

Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2014 Polo Shirt!

Complete your Overall Event Survey and 5 Session Evaluations.

- Directly from your mobile device on the Cisco Live Mobile App
- By visiting the Cisco Live Mobile Site www.ciscoliveaustralia.com/mobile
- Visit any Cisco Live Internet Station located throughout the venue

Polo Shirts can be collected in the World of Solutions on Friday 21 March 12:00pm - 2:00pm

Learn online with Cisco Live!

Visit us online after the conference for full access to session videos and presentations.
www.CiscoLiveAPAC.com


CISCO
TM