

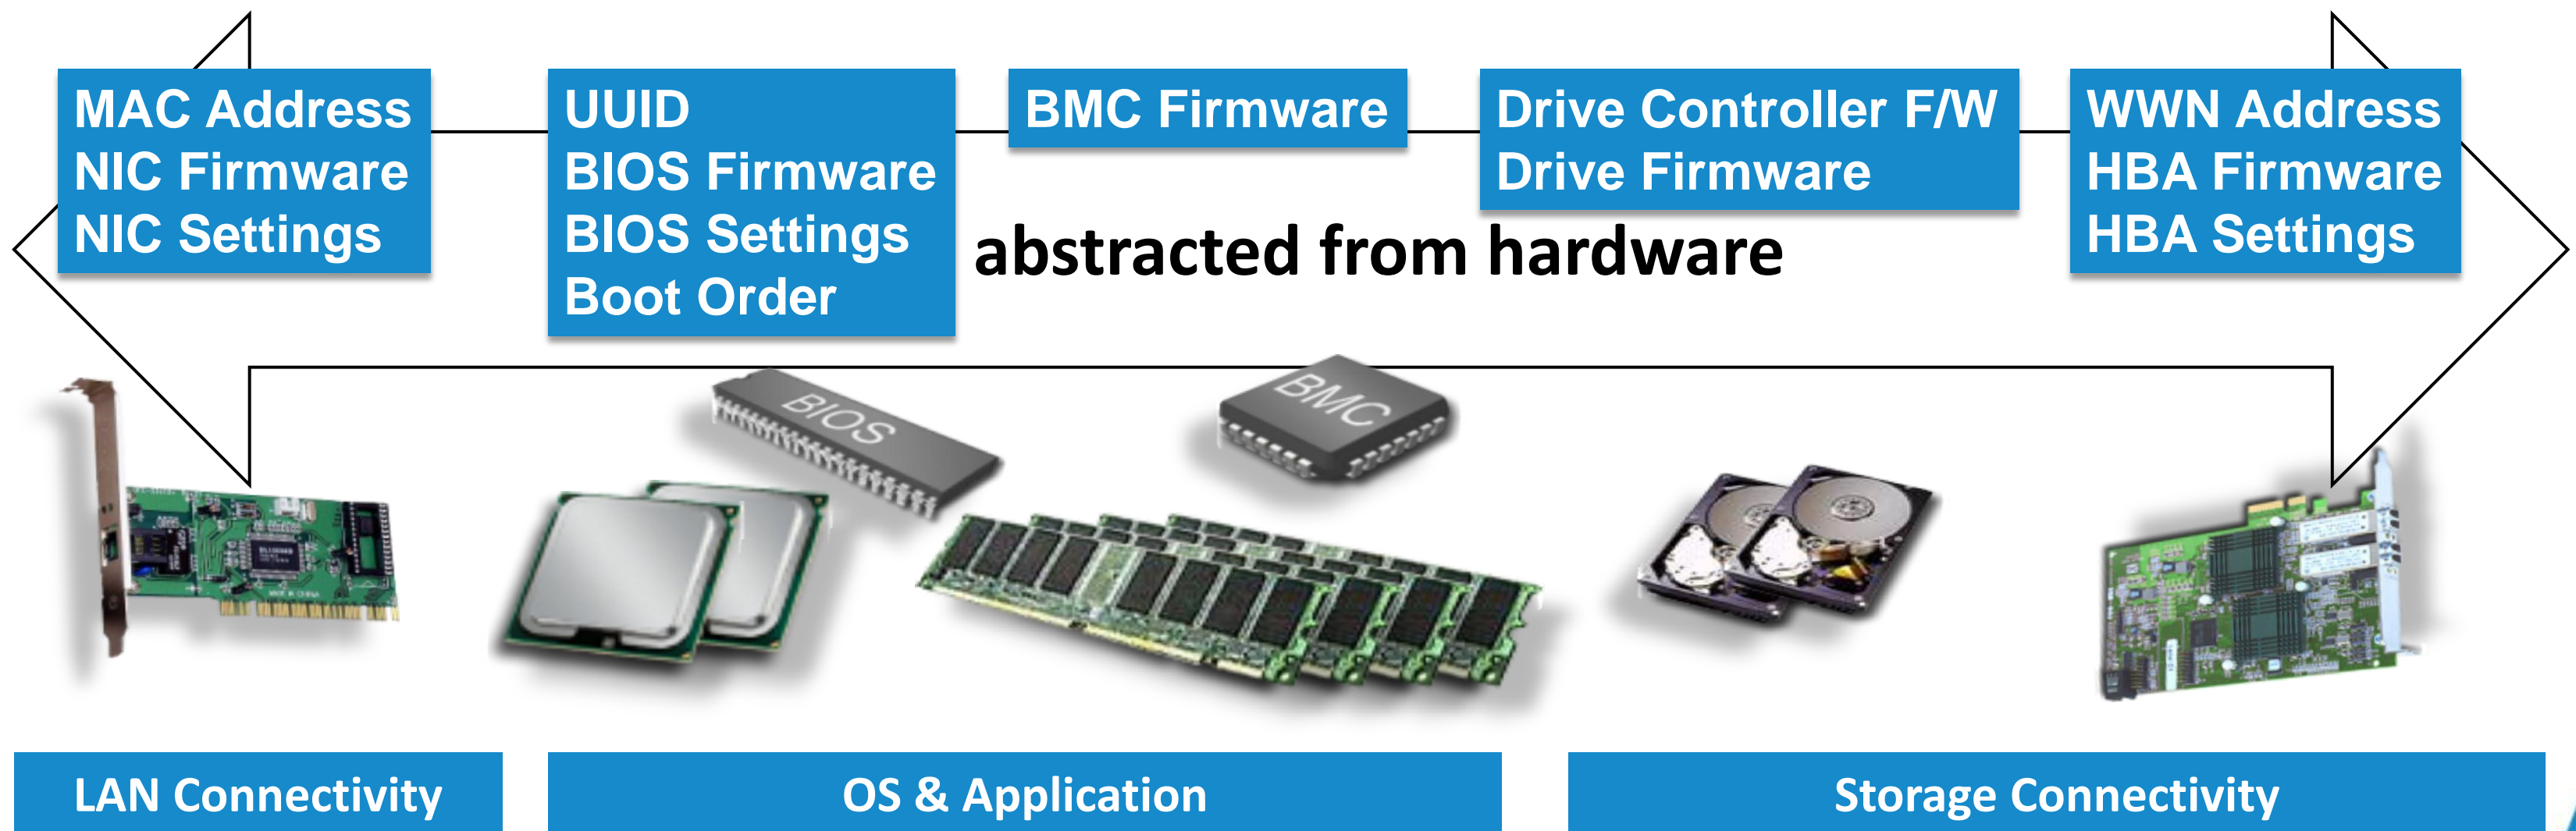
What You Make Possible



UCS Fabric Fundamentals

BRKCOM-1001

Hardware “State” Abstraction



Agenda

- Introduction
- UCS Key Features
- Hardware Overview
- Ethernet Connectivity
- Storage Connectivity

Introduction



Building Blocks of Cisco UCS

UCS Manager

Embedded – manages entire UCS



UCS Fabric Interconnect

10GE unified fabric switch



UCS Fabric Extender (I/O Module)

Remote line card



UCS Blade Server Chassis

Flexible bay configurations



UCS Blade and Rack Servers

x86 industry standard

Patented extended memory

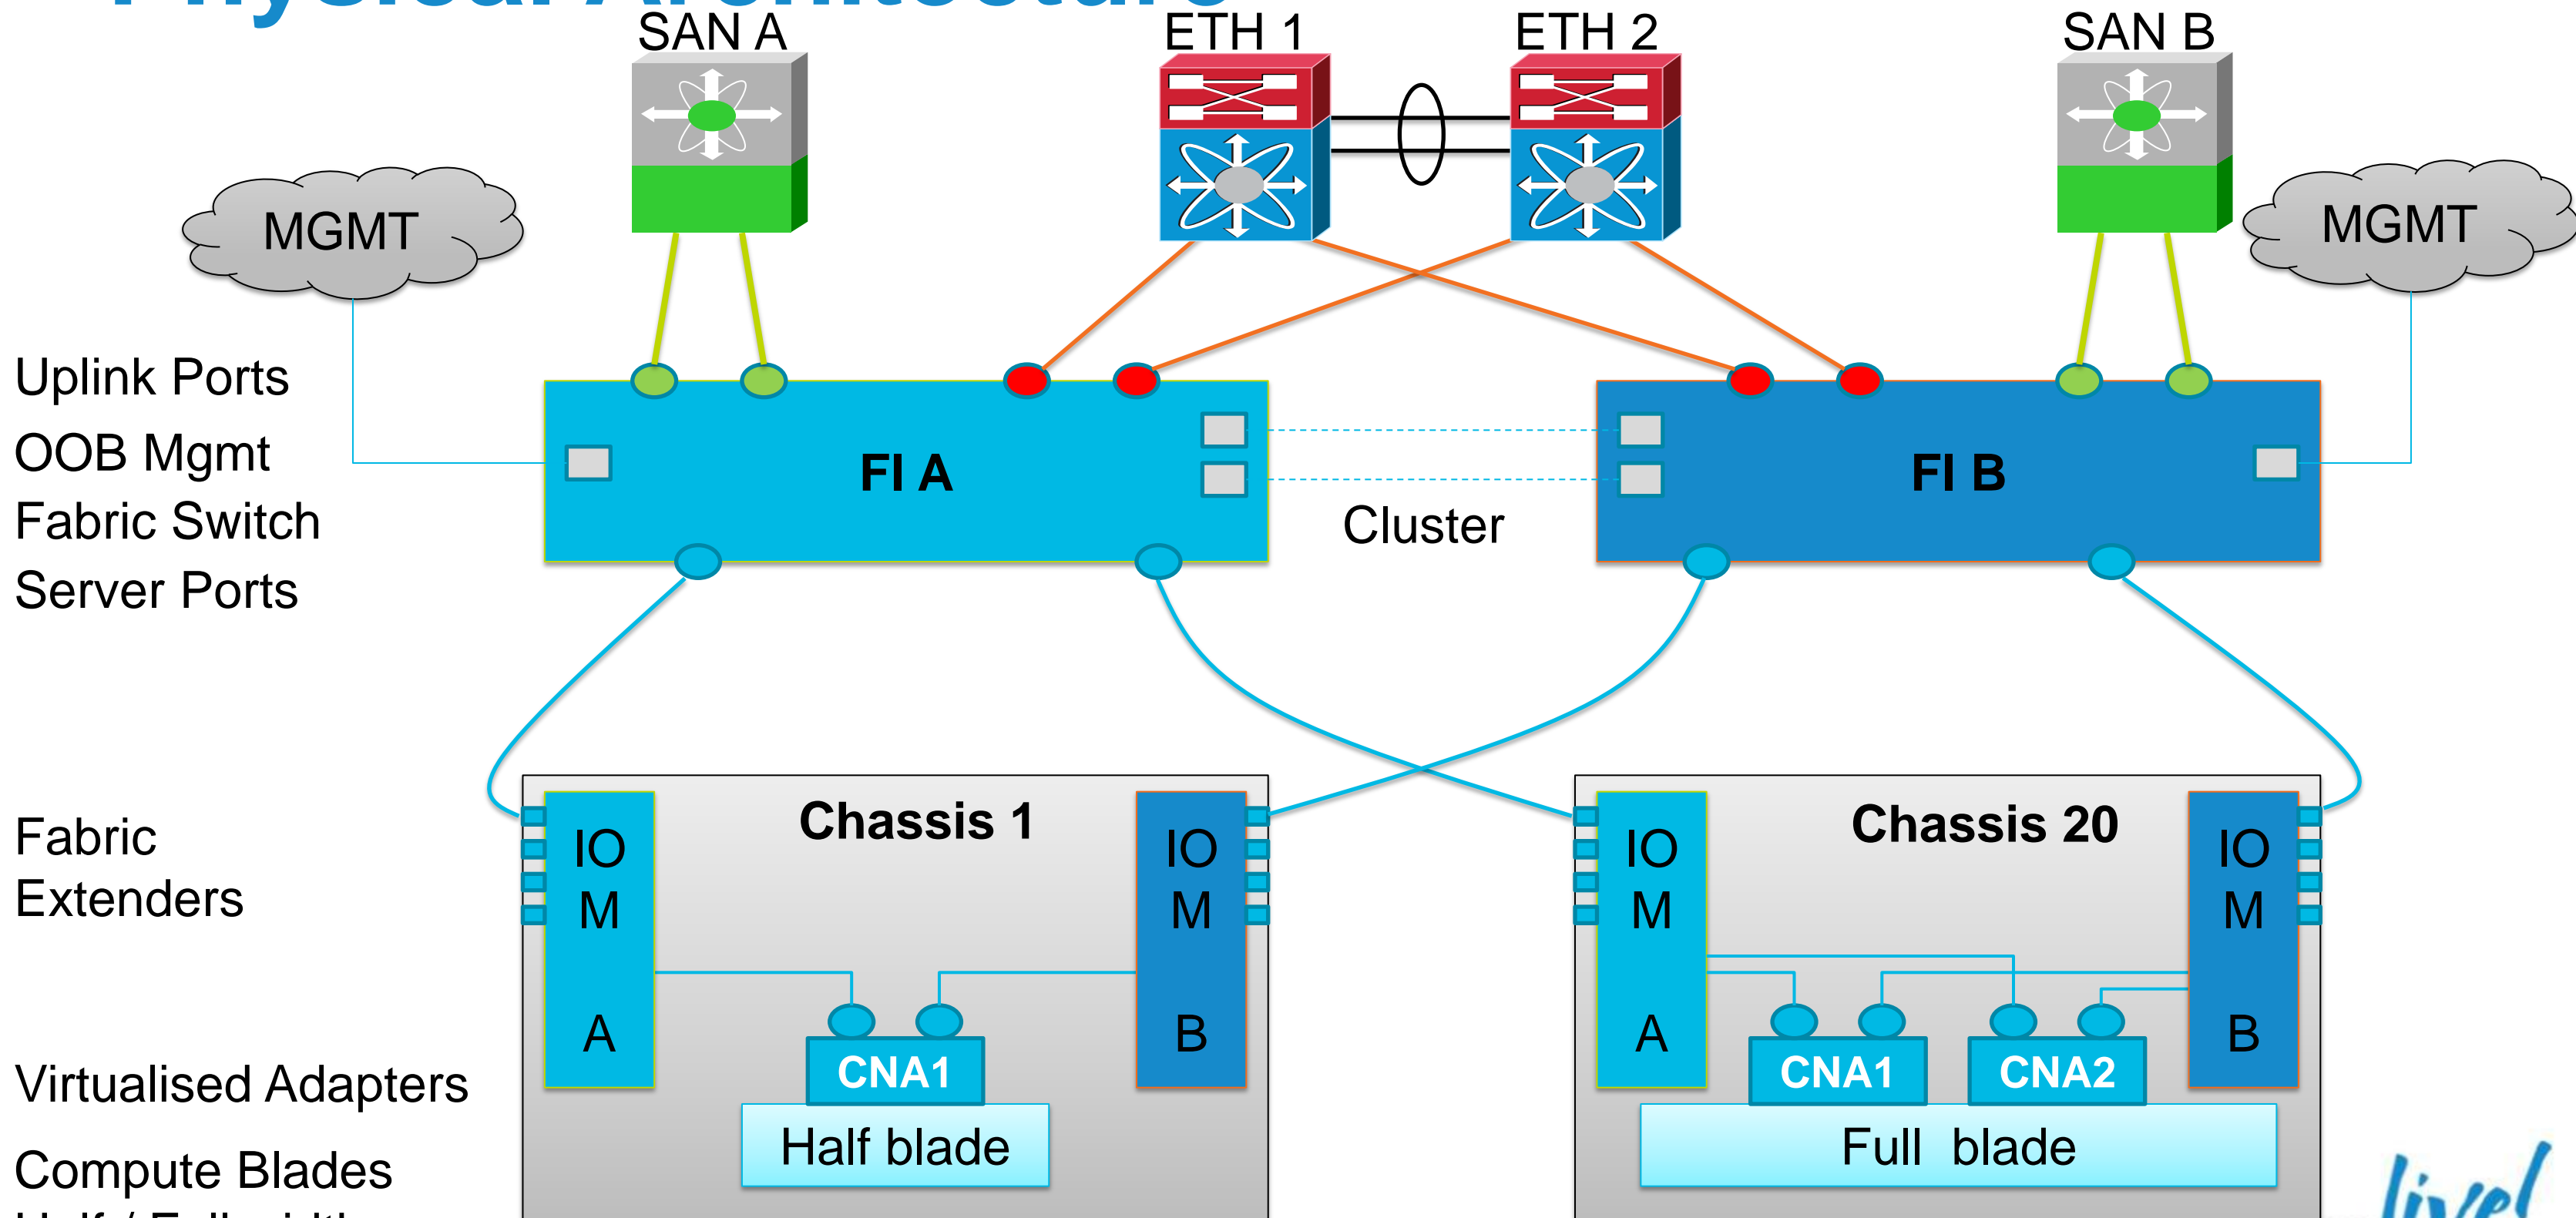


UCS I/O Adapters

Choice of multiple adapters



Physical Architecture



UCS Key Features



Cisco Unified Computing System



“Stateless” Computing

Hardware Abstraction

Service Profiles

Just-in-time Provisioning

Unified Fabric

Consolidated I/O

UCS Manager

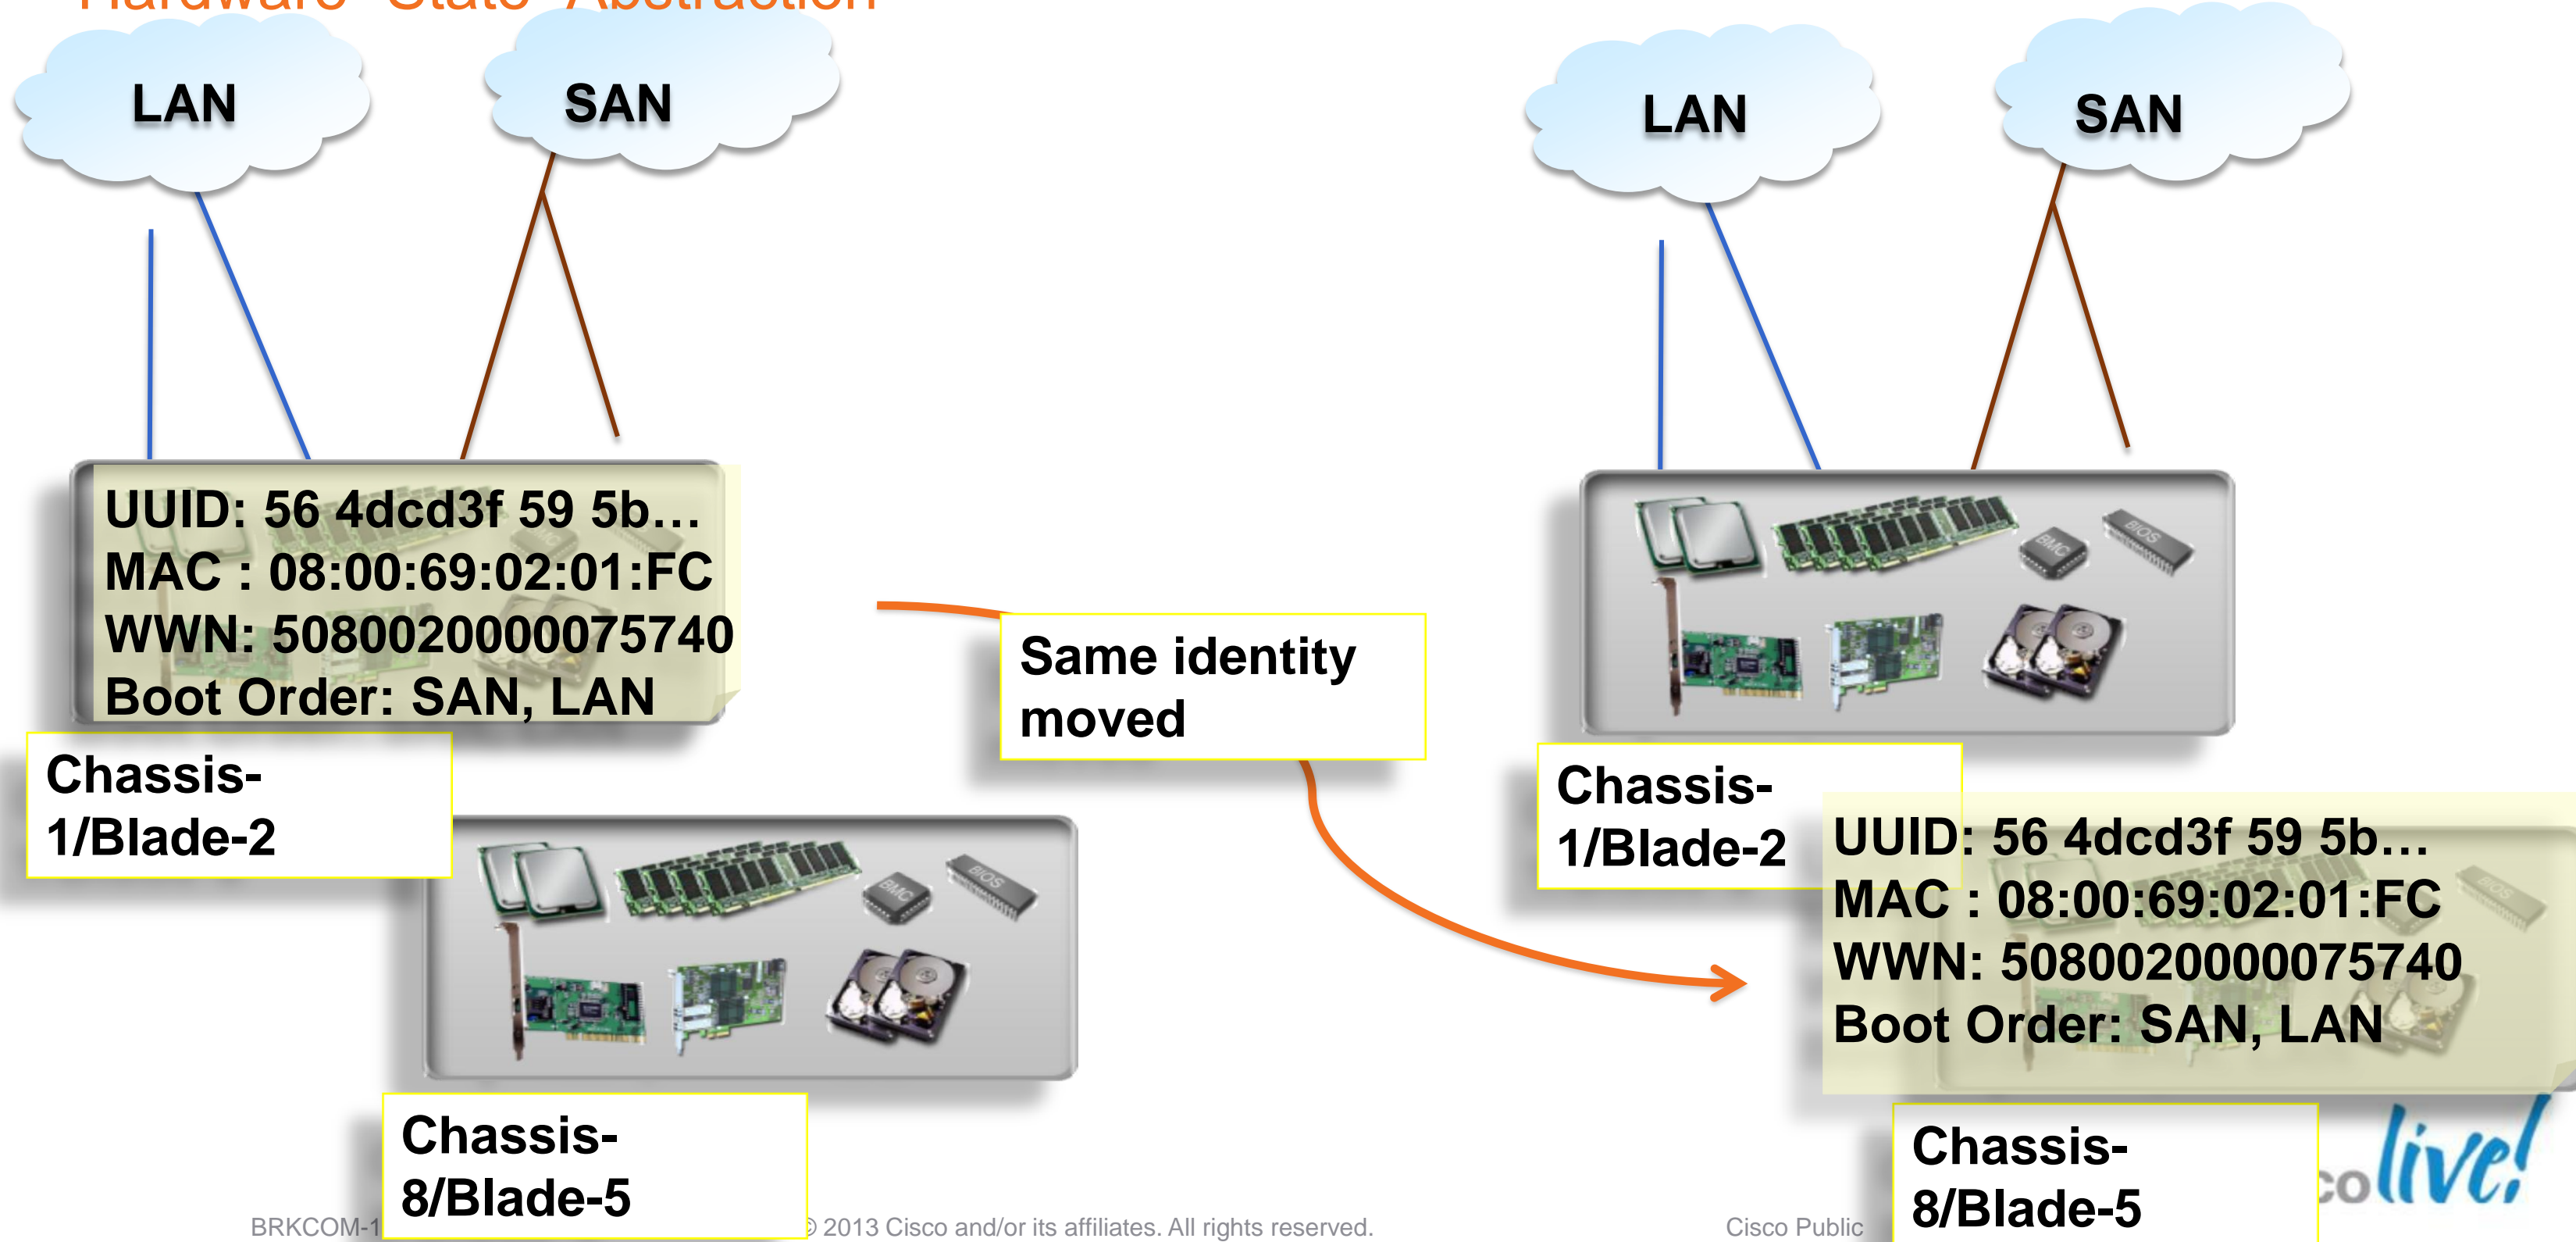
Single Management Domain

Virtual Adapters


Virtualised I/O

“Stateless” Computing

Hardware “State” Abstraction



Service Profiles



UUID; MAC,WWN
Boot info
LAN, SAN Config
Firmware...

Run-time association

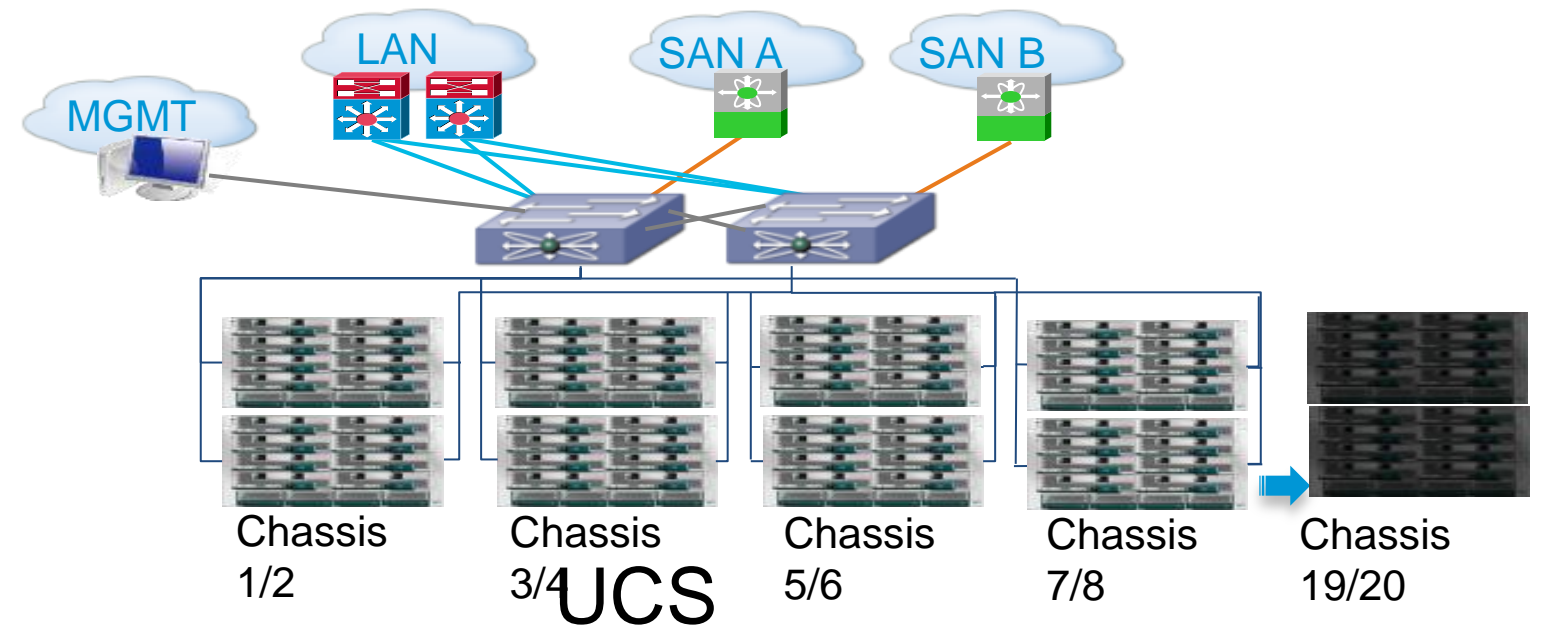
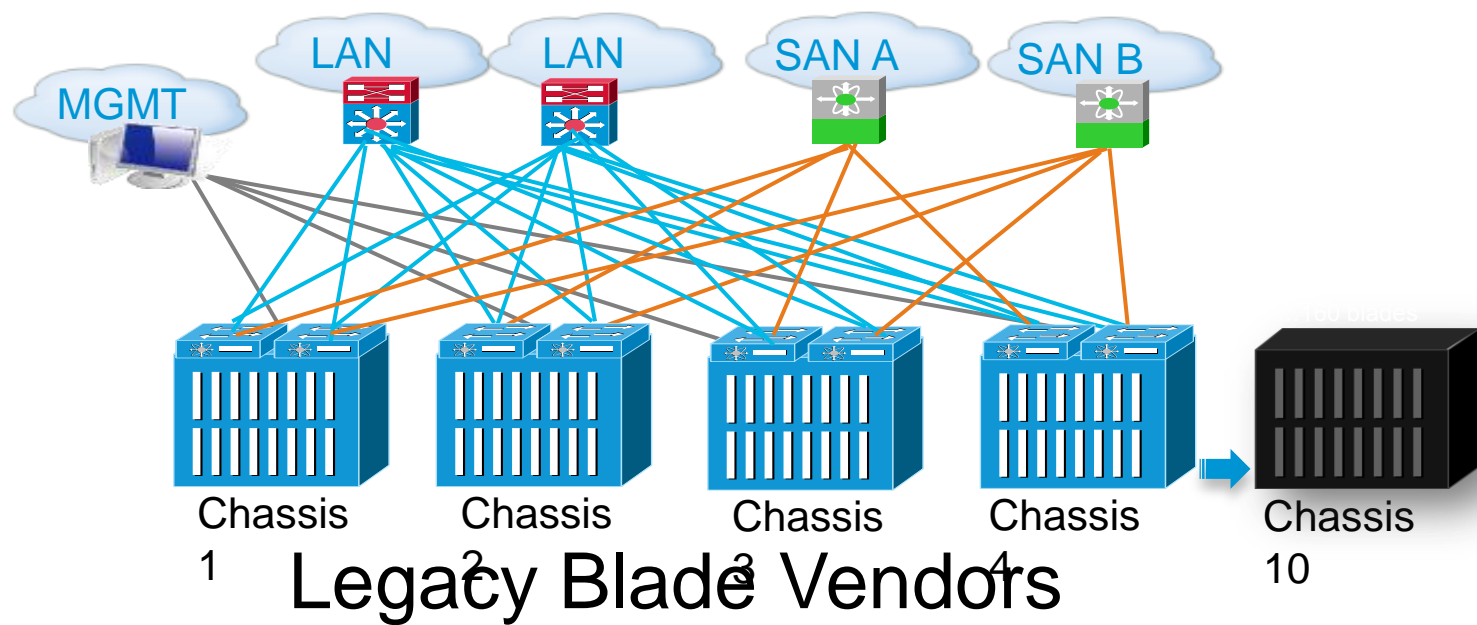
- Logical container of server state information
- User-defined
 - Each profile can be individually created
 - Profiles can be generated from a template
- Applied to physical blades at run time
 - Without profiles, blades are just anonymous hardware components

Benefits :

- Consistent and simplified server deployment
- Simplified server upgrades
- Enhanced server availability

Unified Fabric

Consolidated I/O



Chassis	4	10
Network Devices	8	20
Chassis Management Devices	8	20

Convergence inside each chassis

Chassis	4	10
Network Devices	2	2
Chassis Management Devices	2	2


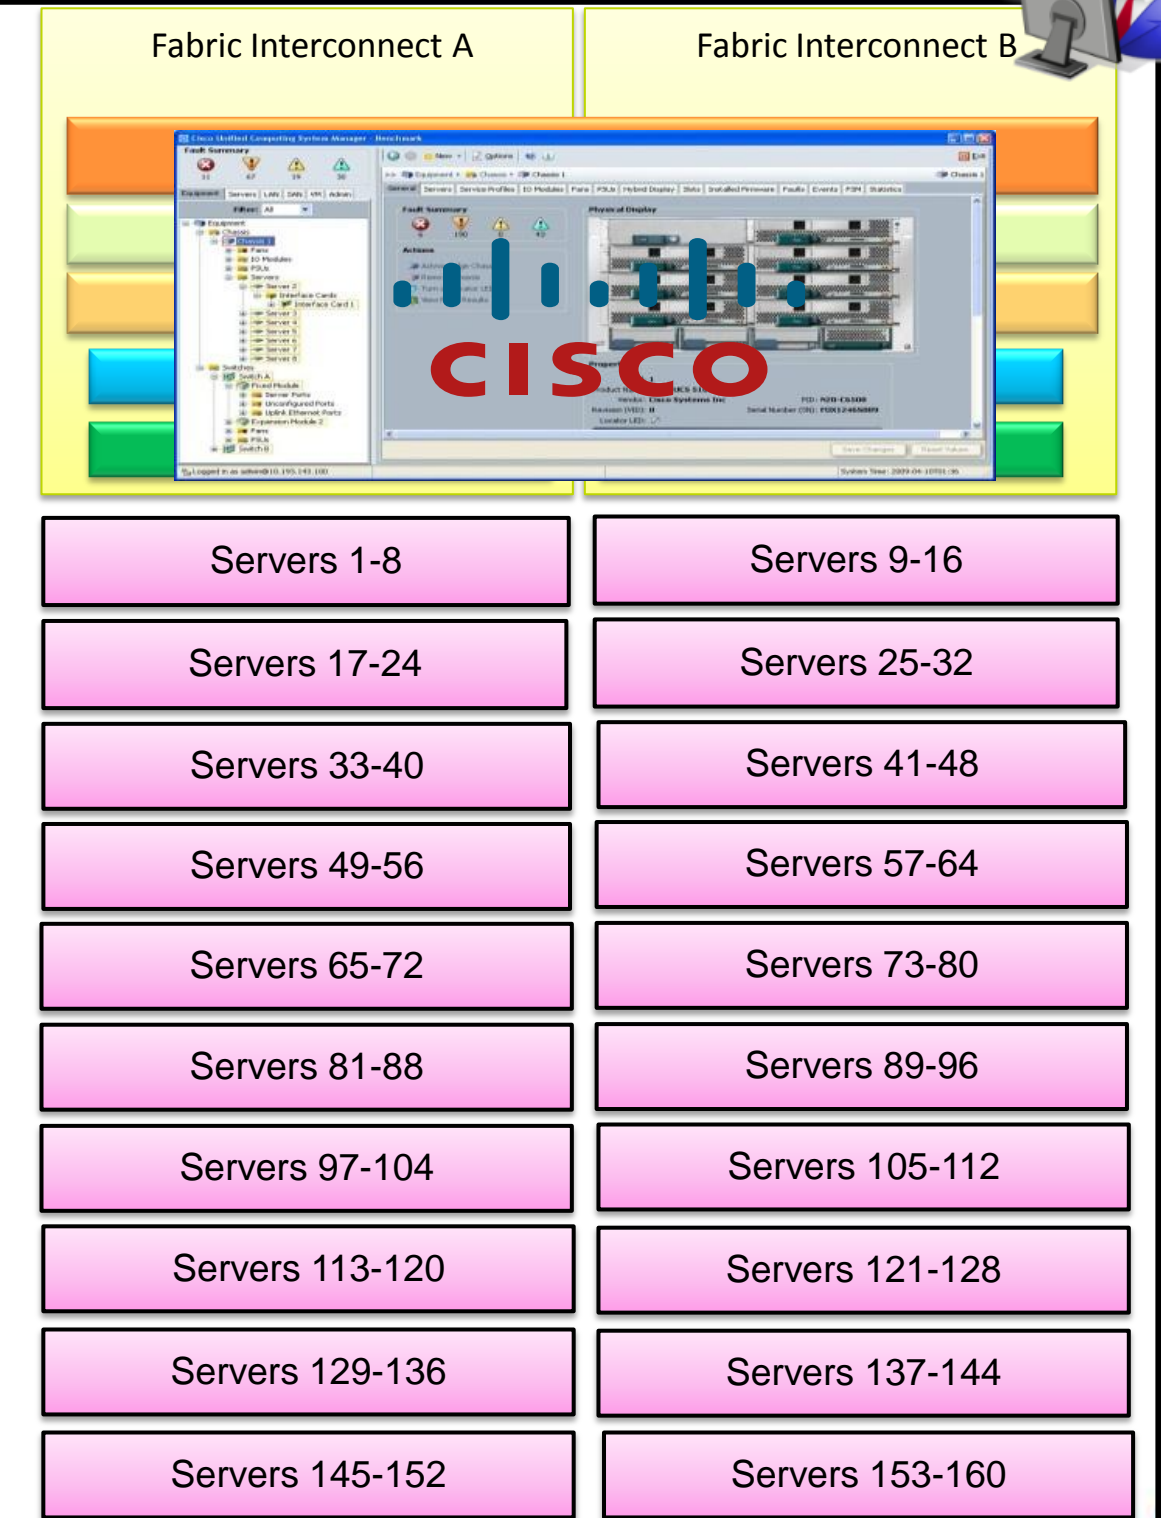
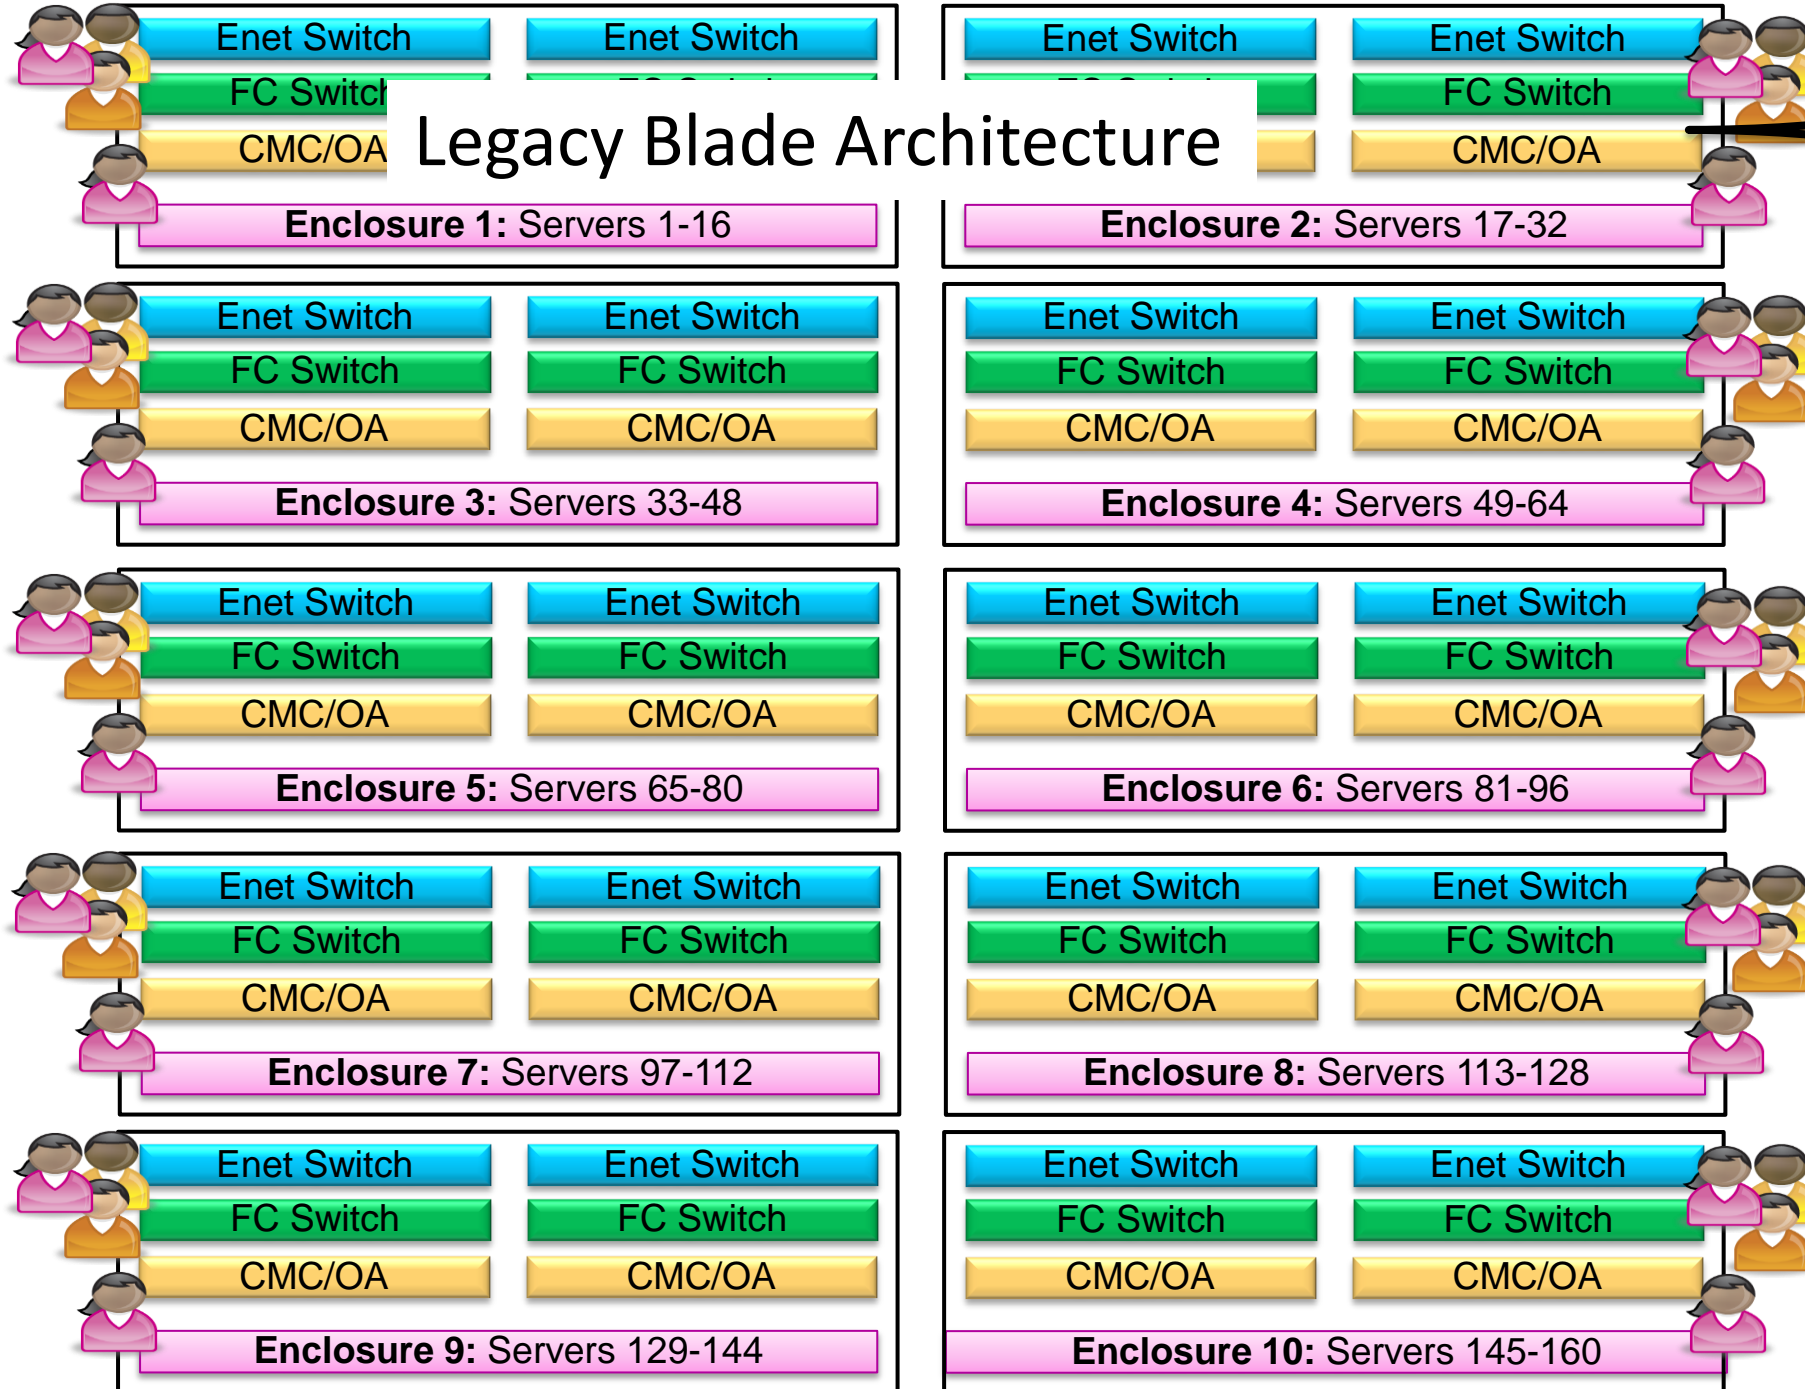
True FCoE of LAN/SAN & MGMT

UCS Manager: Single Management Domain

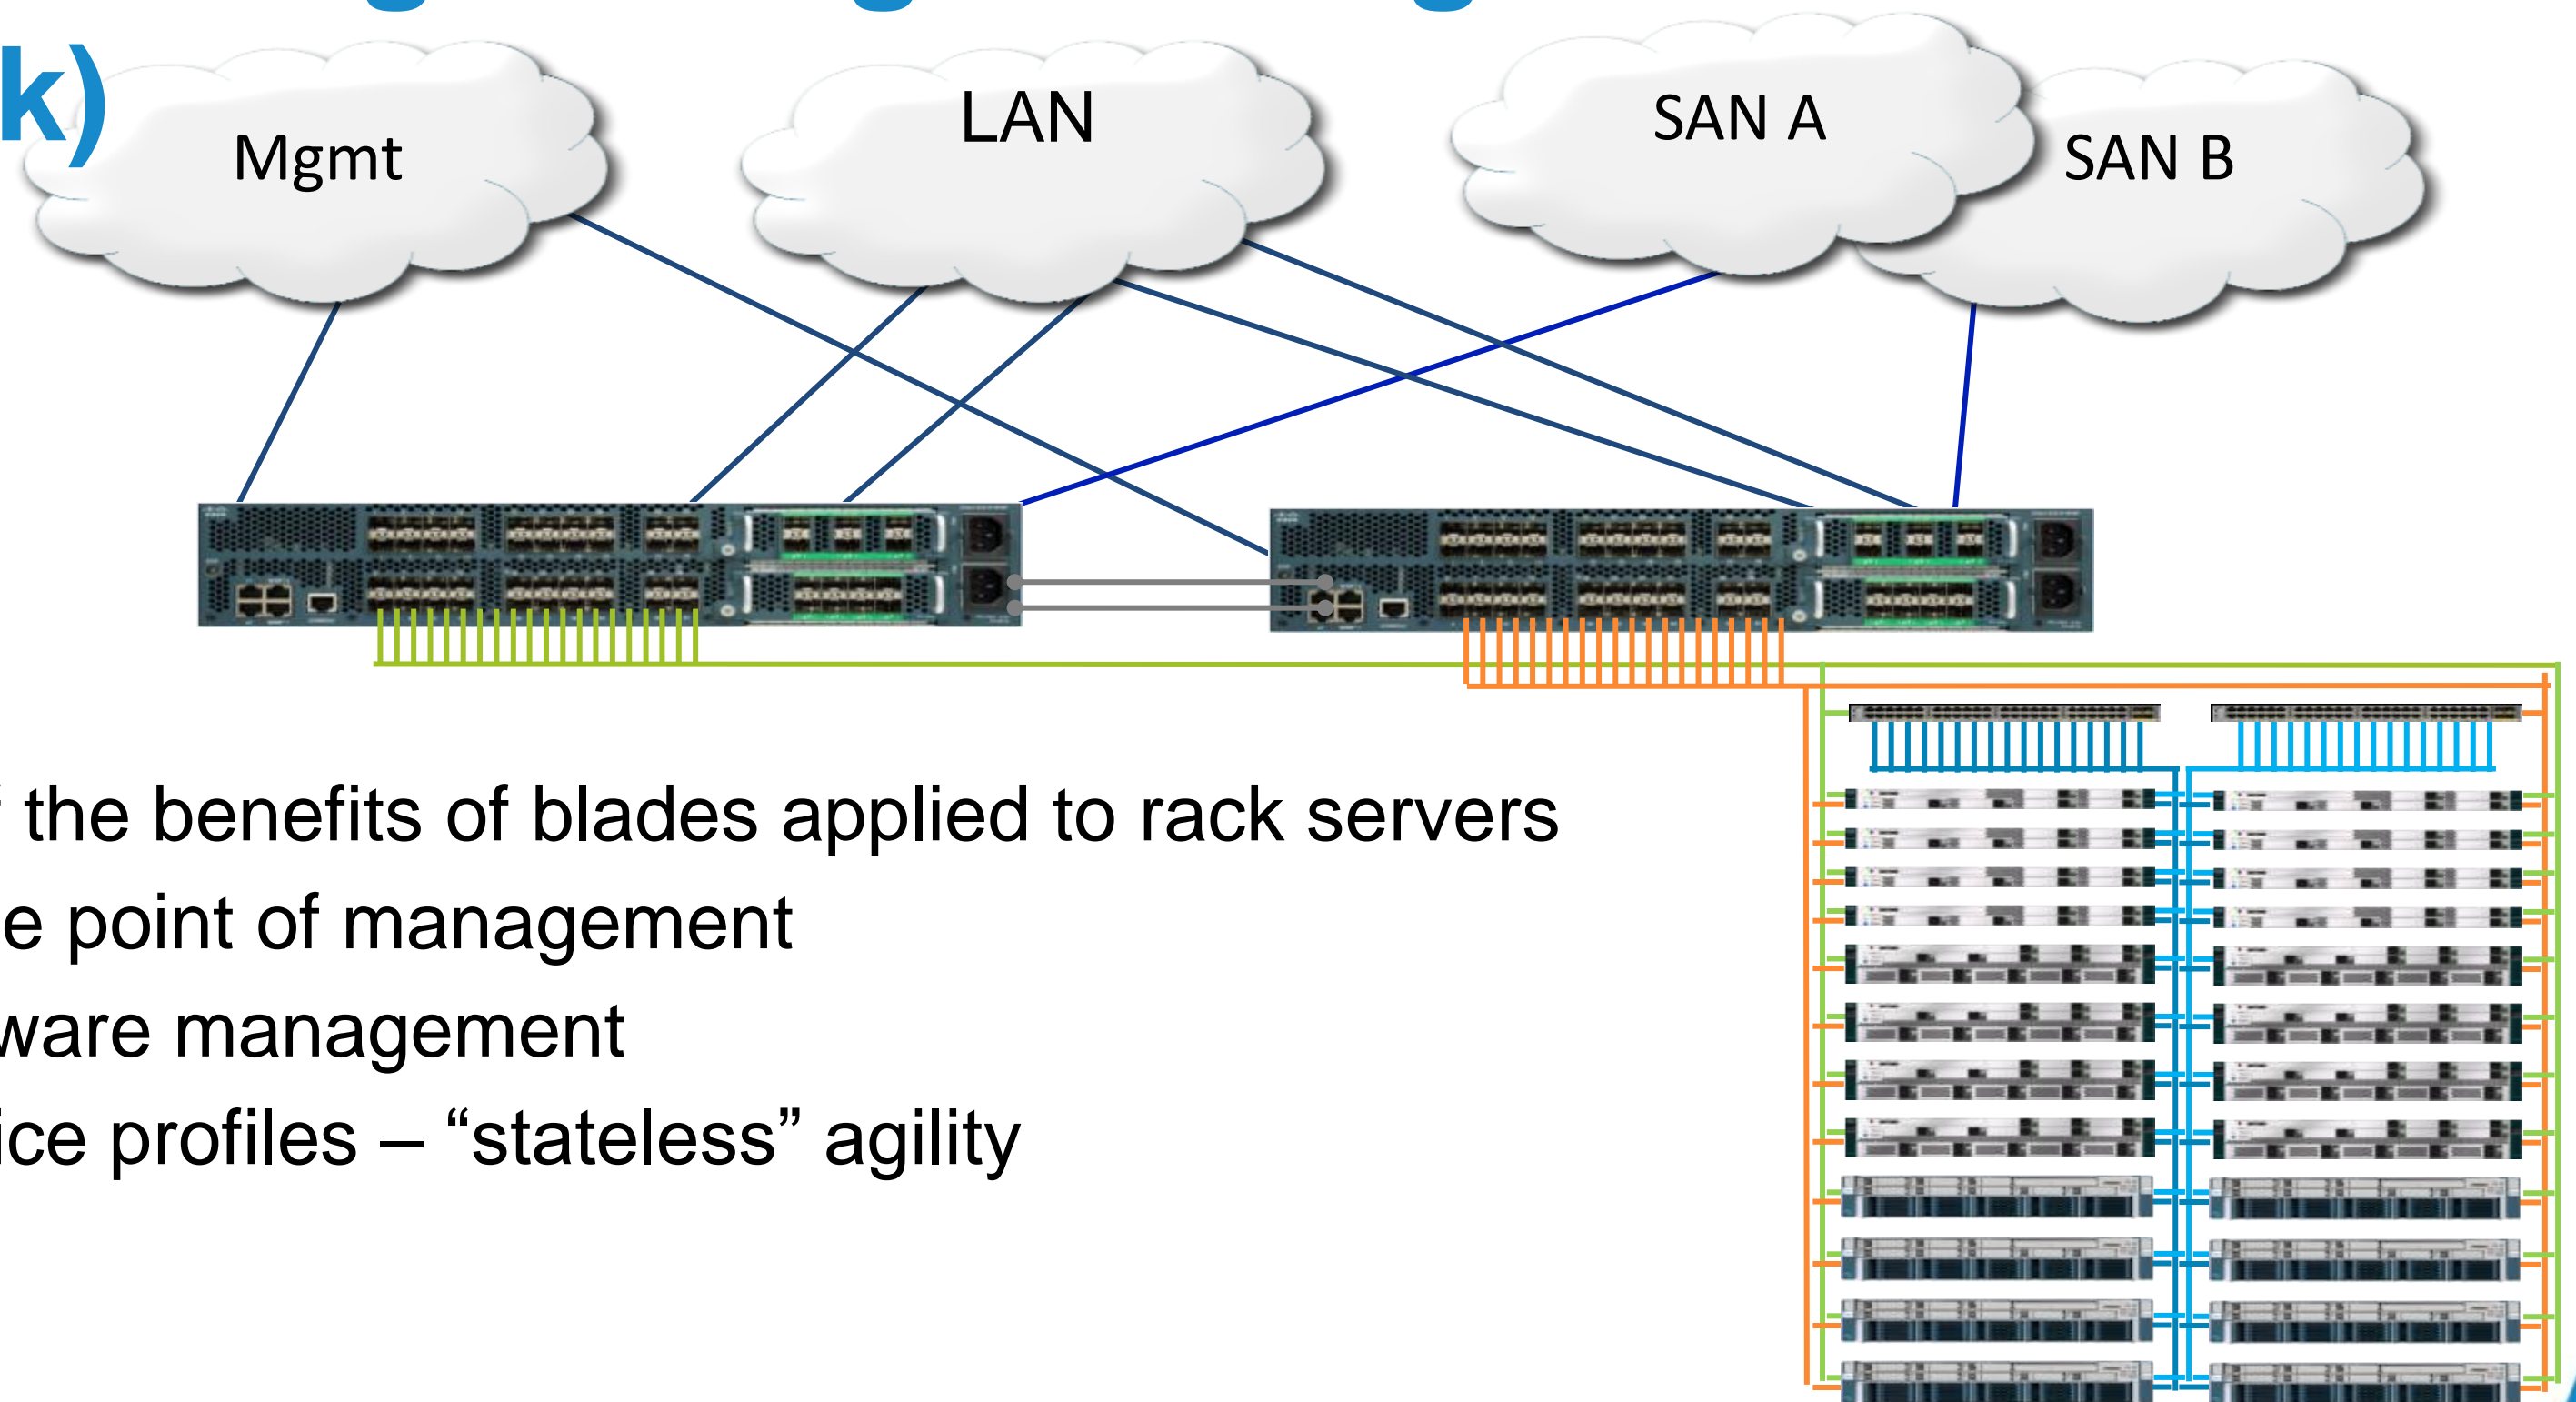


Multi-Chassis Server Identity Manager (VCEM)

Server Health Monitoring (SIM)



UCS Manager: Single Management Domain (Rack)



- All of the benefits of blades applied to rack servers
- Single point of management
- Firmware management
- Service profiles – “stateless” agility

UCS Manager: Multi-Domain Management

Do It Yourself

- SDK & Emulator
- GoUCS
- PowerShell & .NET
- Most flexible
- Most powerful

Open Source

- UCS Dashboard
-
- Flexible
 - Leverage power of community

Cisco Product

- Multi UCS Manager
 - Cisco CIAC
 - Cisco NSM
-
- Cisco support and roadmap

ISV Partners

- System Centre plug-in
 - BMC Integration
 - CA integration etc
-
- Multi-vendor support
 - Incumbency
 - Broadest use cases

Multi-UCS Management Solutions

UCS XML API

UCS XML API

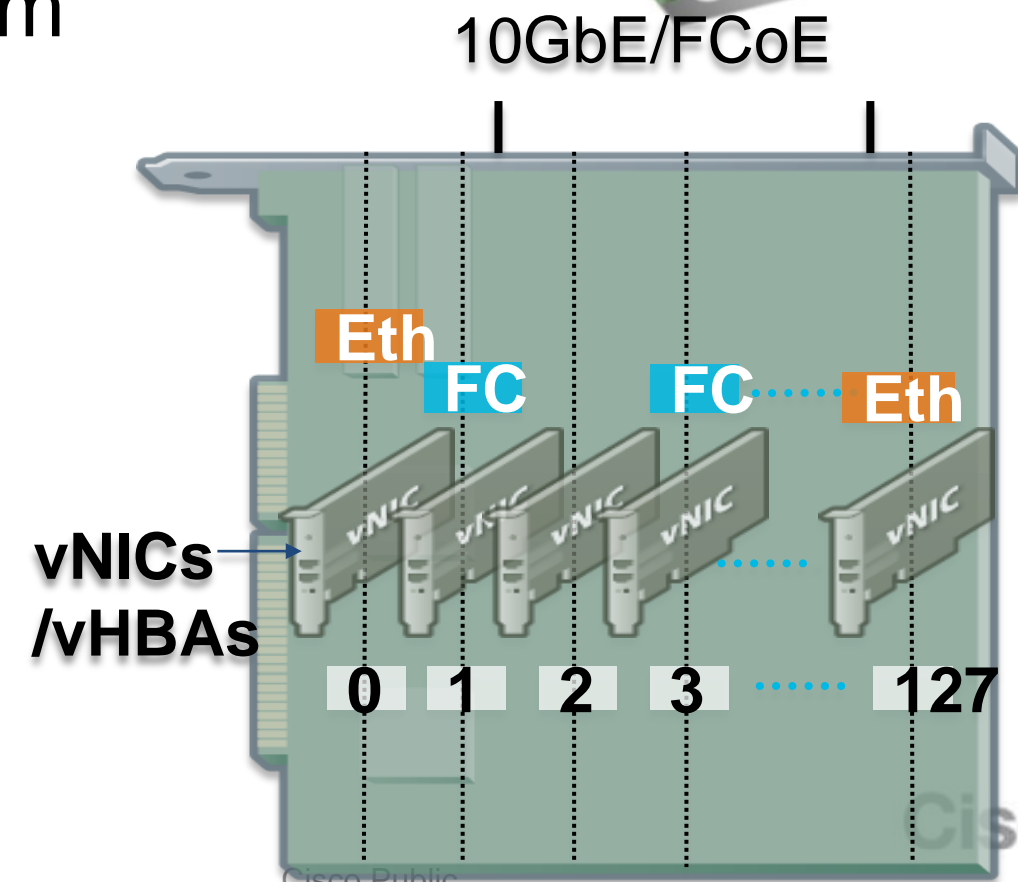
UCS XML API

UCS XML API




Cisco UCS Virtual Interface Card

- Creates multiple x16 PCIe Ethernet and FC interface adapters
- Provides redundant path to fabric interconnect using dual 10 Gigabit Ethernet ports
- Centralised management of virtual interfaces from UCS Manager
- Offers Virtualisation aware networking with best performance (VM-FEX)




VM-FEX with VMware



Emulated Mode

- Each VM gets a dedicated PCIe device
- 12%-15% CPU performance improvement
- Appears as distributed virtual switch to hypervisor
- vMotion supported



PCIe Pass-Thru or VMDirectPath

- Co-exists with Standard mode
- Bypasses Hypervisor layer
- 30% improvement in I/O performance
- Appears as distributed virtual switch to hypervisor
- Currently supported with ESX 5.0 only
- vMotion supported

Hardware Overview



Unified Ports


Customer benefits

Simplify fabric purchase by removing port ratio estimates

Increase design flexibility

Feature details

- Flexibility to configure any port at Ethernet (1/10 Gigabit with SFP+) or FCoE or Native FC Ports (8/4/2/1G with FC Optics)
- All ports on UCS 6200 Series
- 16-port expansion module (E16UP)
- Uses existing Ethernet / Fibre Channel SFP and SFP+ transceiver
- Compared to 6100 Series, Fibre Channel is standard. No requirement to buy FC module



UCS 6248 Fabric Interconnect

Customer benefits

Higher density and performance

Unified ports

Feature details

- Double the port density at 1RU compared to UCS 6100 Series
- 32 fixed unified ports and 1 expansion module
- All ports are usable as uplinks/downlinks
- Reduced port latency to 2 μ s
- Dual power supplies
- DC option introduced
- Layer 3 ready via daughter card



UCS 6248 Fabric Interconnect

Rear Panel



32 x Fixed ports: 1/10 GE or 1/2/4/8 FC

Expansion Module (GEM)

Fabric Interconnect Cluster Connectivity

Out of Band Mgmt
10/100/1000

Front Panel



Console

Fan Module

Fan Module

Power Entry

Power Entry

N + N Redundant Fans

N + N Power Supplies

UCS 6296 Fabric Interconnect

Customer benefits

Higher density and performance

Unified ports

Feature details

- High density, 96 ports in 2RU
- 48 fixed unified ports and 3 expansion modules
- Increased 2Tbps Switching Performance
- All ports are usable as uplinks/downlinks
- Reduced port latency to 2 μ s
- Dual power supplies, AC/DC option
- Layer 3 ready via daughter card



UCS FI 6296 Front Panel

Fabric Interconnect

Out of Band Mgmt
10/100/1000

Console

USB Flash



Power Entry

Power Entry

Fan Module

Fan Module

Fan Module

Fan Module

N + N Power Supplies

N + 1 Redundant Fans

UCS FI 6296 Rear Panel

Expansion Module

Expansion Module

Expansion Module



48 x Fixed ports 1/10 GE or 1/2/4/8 FC

UCS 6200 Series Expansion Module

Customer benefits

Simplify fabric purchase by removing port ratio estimate

More ports

Feature details

- 16 unified ports can be configured as either Ethernet/FCoE 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
- Not backwards compatible with 6100 Series
- Slider based FC port configuration



UCS 5108 Blade Chassis

- Up to 8 half slot blades
- Up to 4 full slot blades
- 4x power supplies, N+N grid redundant
- 8x fans included
- 2x UCS Fabric Extender
- All items hot-pluggable
- AC / DC



UCS 5108 Blade Chassis Parts

Rear

6U, 19" Rack



2 Fabric Extenders
(I/O modules)

8 Fan Modules

4 Power
Connectors

Front

6U, 19" Rack




4 to 8 Blades

4 Power
Supplies

UCS 5108 Blade Chassis Backplane

Blade Connectors



PSU Connectors

I/O Modules

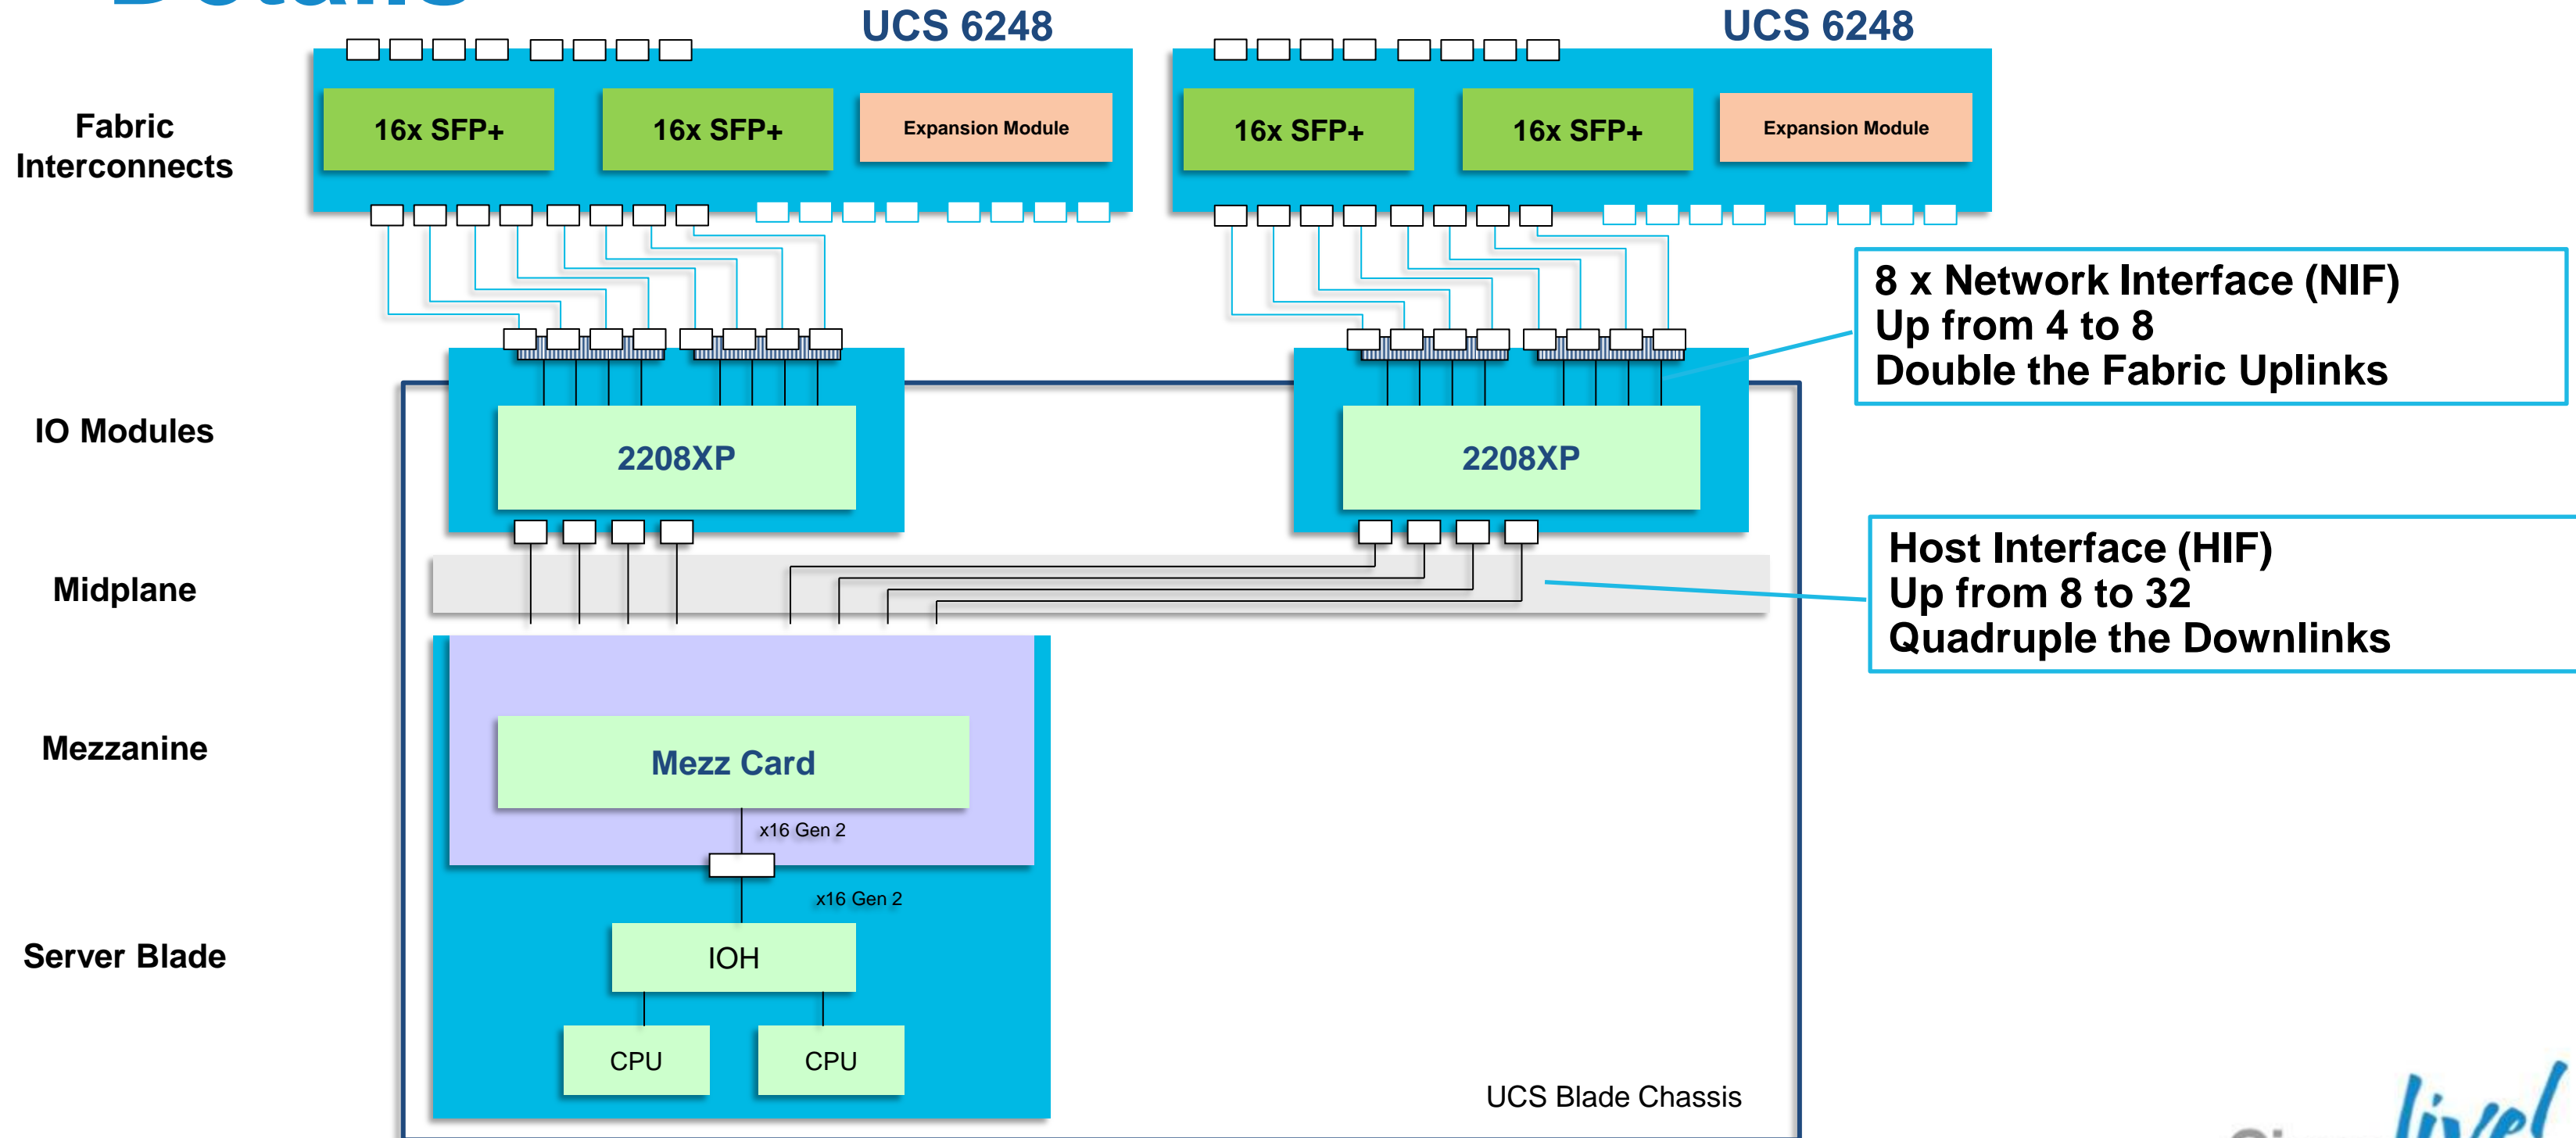
Redundant data and management paths

UCS 2104 Fabric Extender (I/O Module)

- Connects UCS blade chassis to the Fabric Interconnect
- Four 10 Gigabit Ethernet, FCoE capable, SFP+ ports
- Up to 2 Fabric Extenders per chassis for redundancy and up to 80 Gbps of bandwidth per chassis
- Built-in chassis management functionality
- Fully managed by UCS Manager through Fabric Interconnect
- No local switching



Block Diagram: Next Gen UCS Fabric Details



2208XP Fabric Extender (I/O Module)

Customer benefits

Double the uplink bandwidth to the FI

Quadruple the downlink bandwidth to the server slots

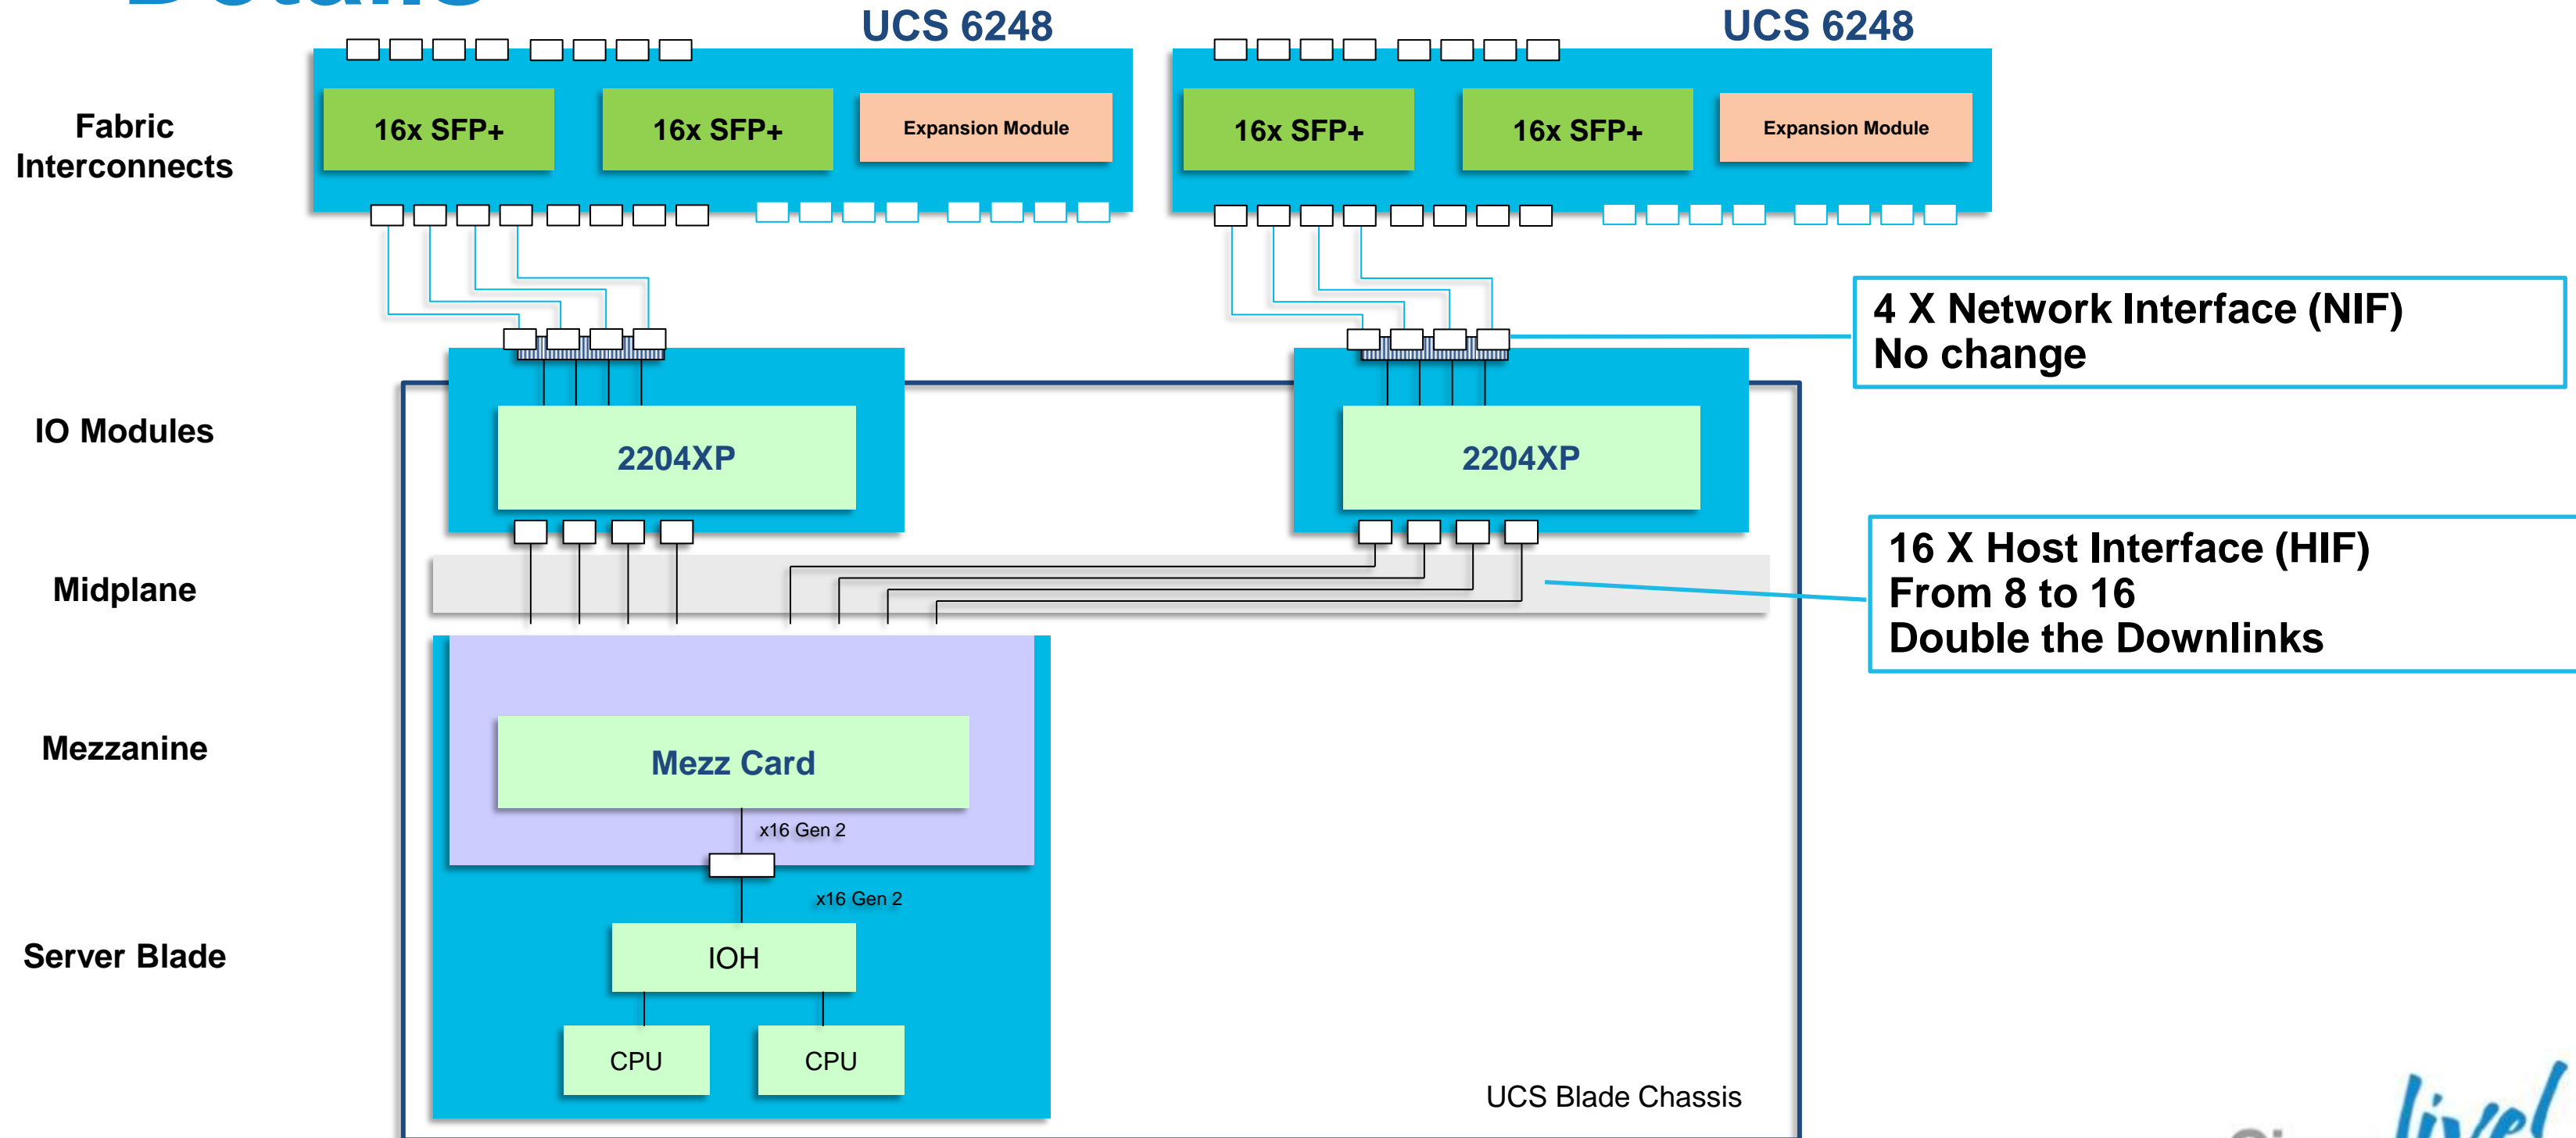
Lower latency and better QoS

Feature details

- Double the uplinks
 - 8x 10GE uplinks to FI
 - Total 160 Gbps per chassis
- Quadruple the downlinks
 - 32x 10GE to blades
- Lower latency (0.5 μ s within IOM)
- Backward compatible



Block Diagram: Next Gen UCS Fabric Details



2204 Fabric Extender (I/O Module)

Customer benefits

Double the downlink bandwidth to the server slots

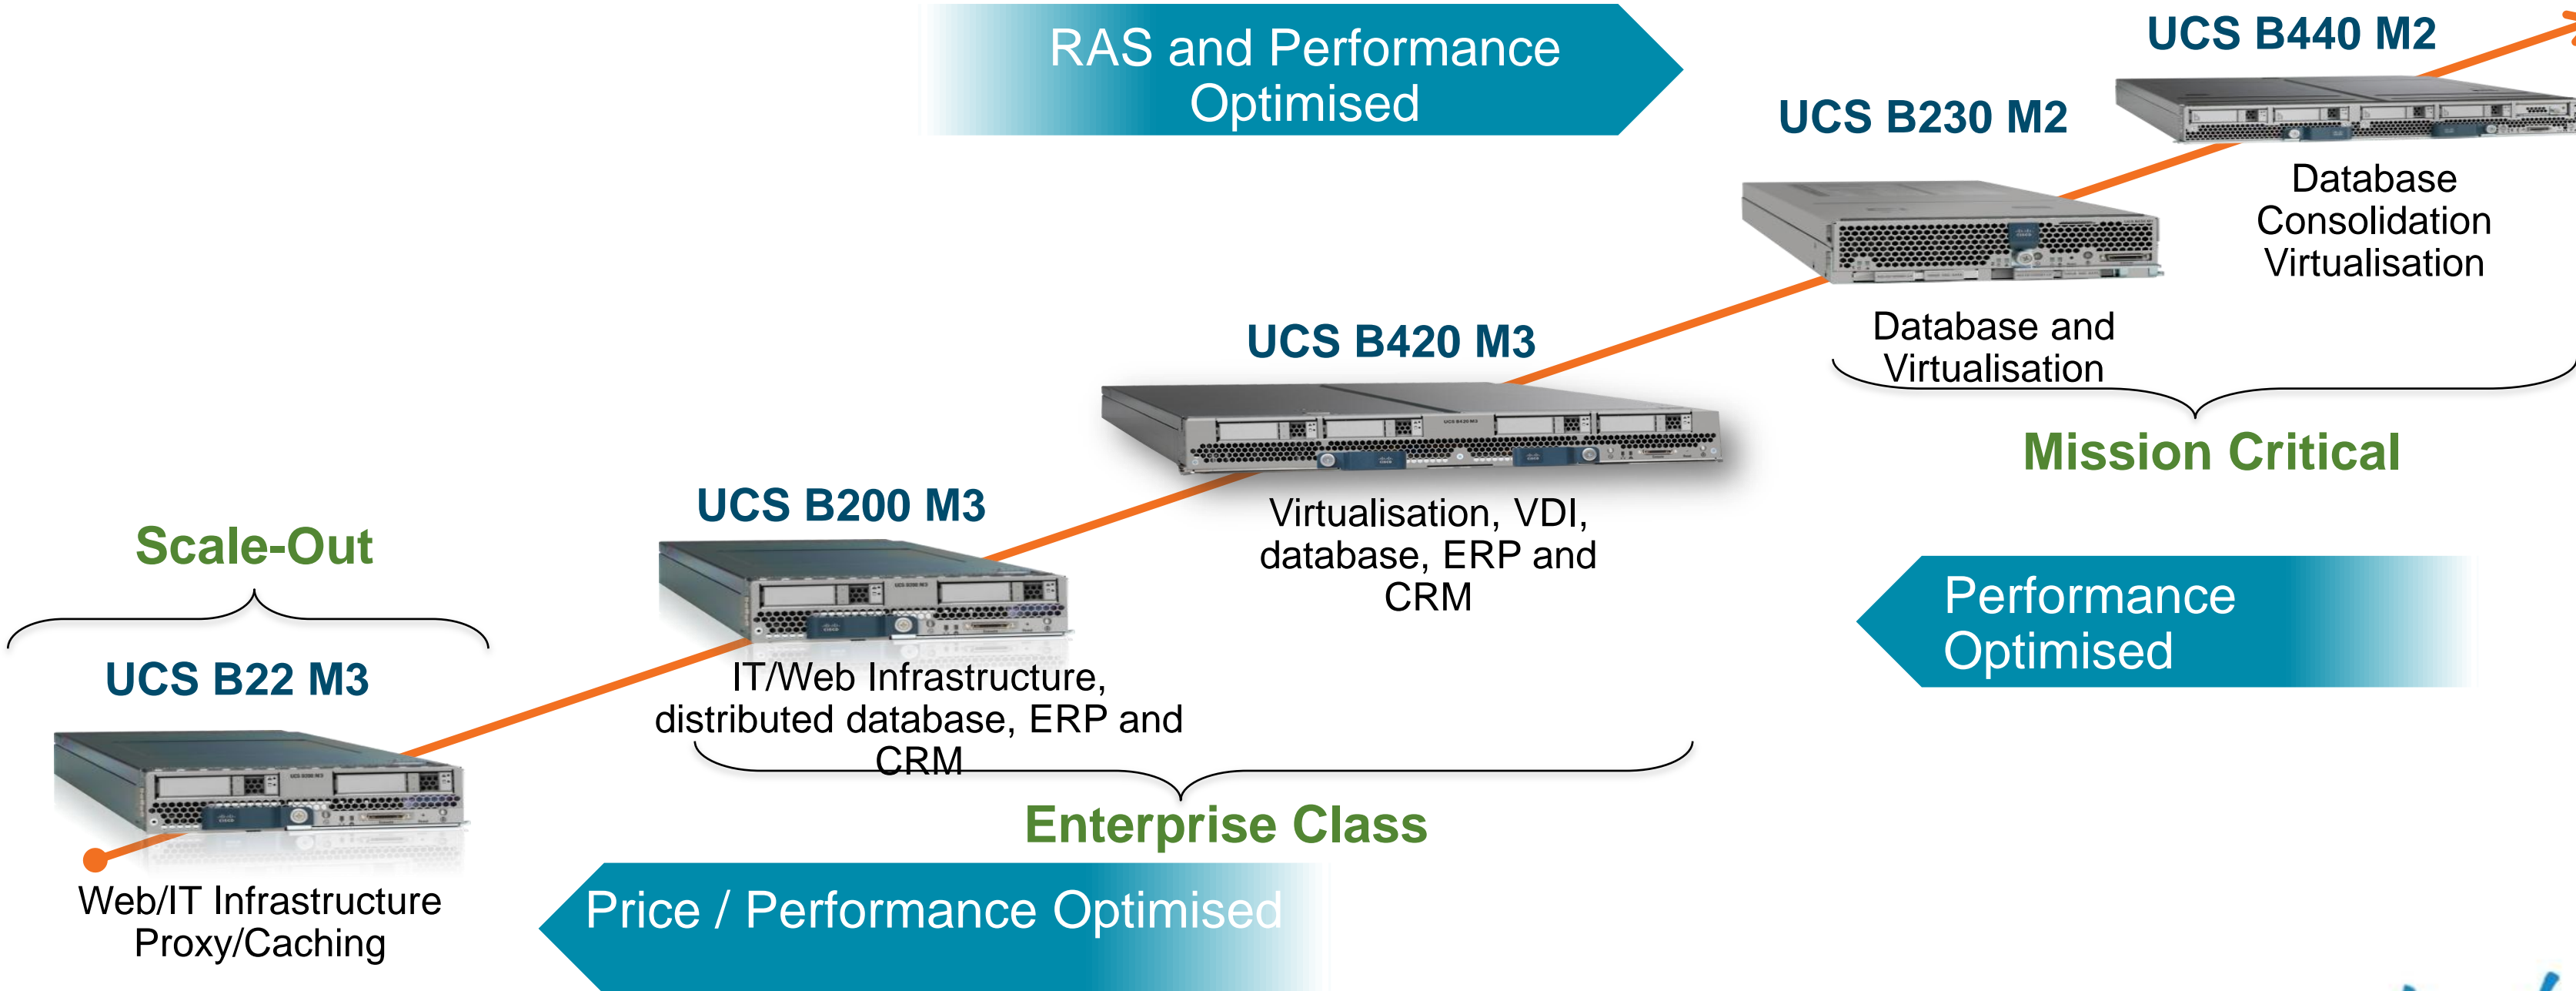
Lower latency and better QoS

Feature details

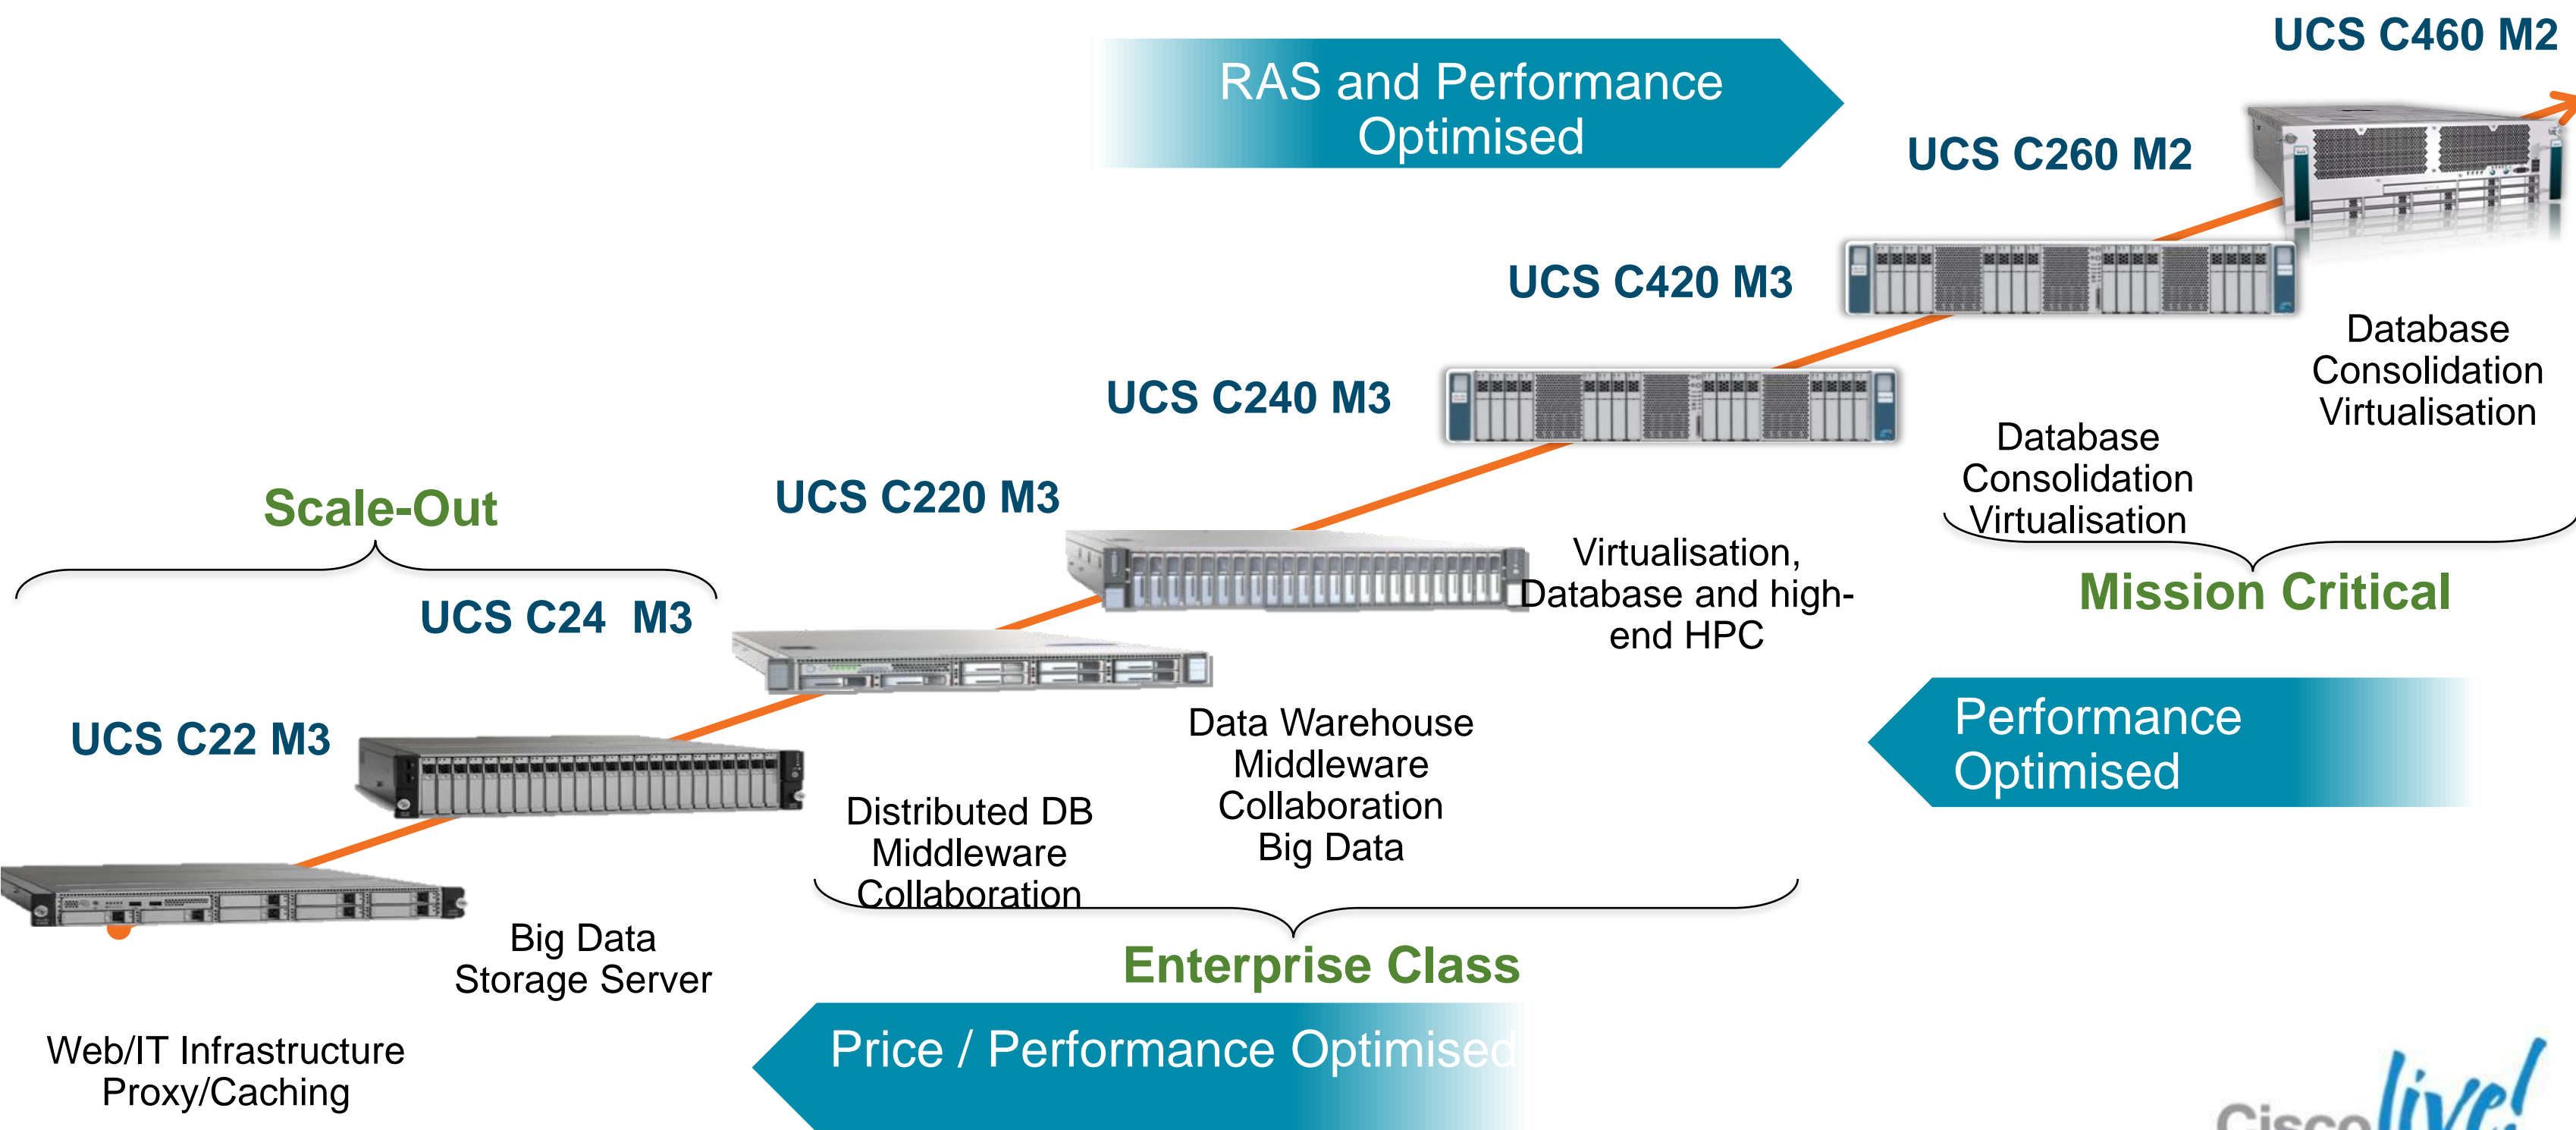
- 80 Gbps per chassis
- Double the downlinks
 - 16x 10GE to blades
- Increased support for 8 egress CoS queues
- Lower latency (0.5 μ s within IOM)
- Backwards compatible



Compute Servers : B series



Compute Servers : C-Series



Virtualisation Interface Card 1200


Customer benefits

Dual 4x10 GE (80 Gb per host)

VM-FEX scale, up to 112 VM interfaces /w ESX 5.0

Feature details

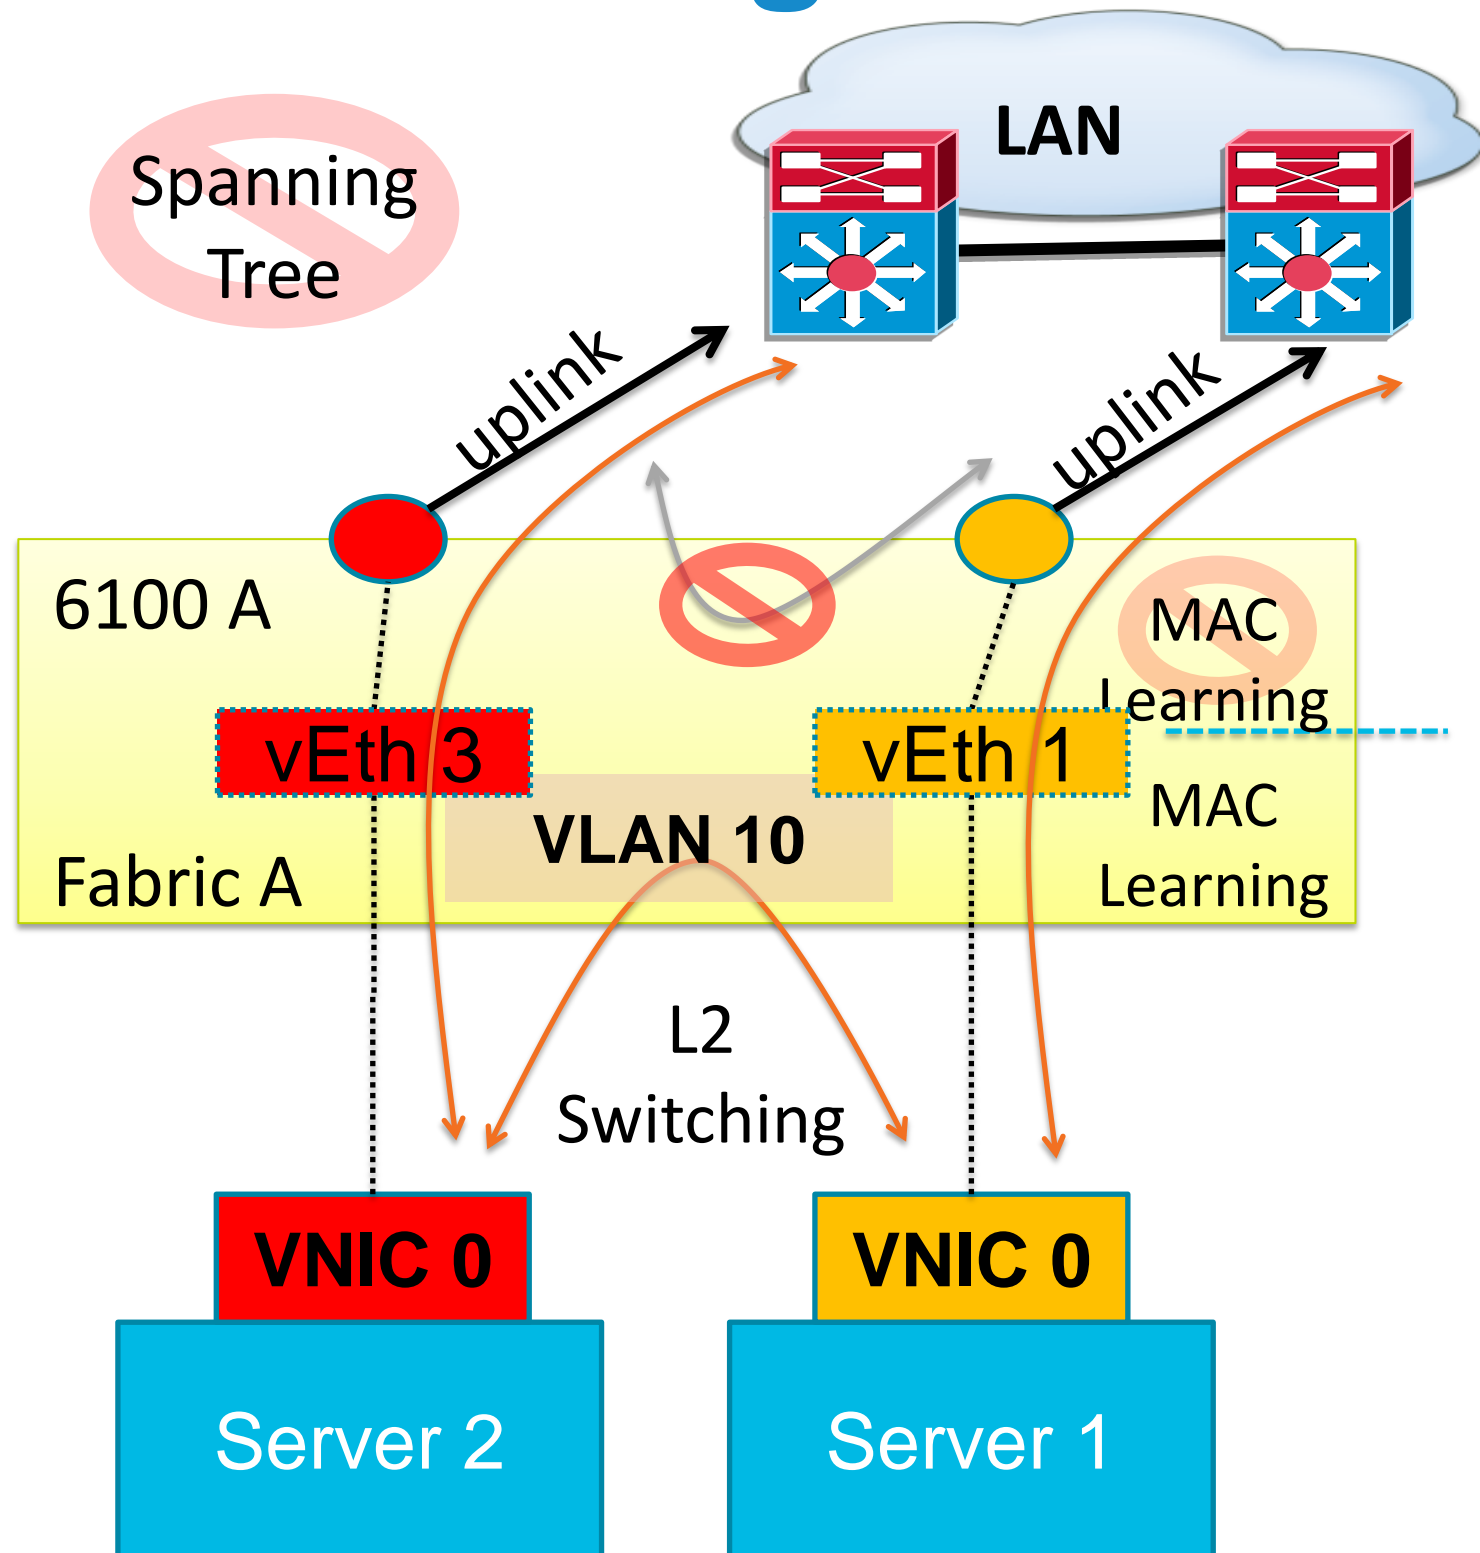
- Dual 4x10 GE port-channels to a single server slot
- Host connectivity PCIe Gen2 x16
- HW Capable of 256 PCIe devices
 - OS restriction apply
- PCIe virtualisation OS independent (same as M81KR)
- Single OS driver image for both M81KR and 1280 VIC
- Fabric Failover supported



Ethernet Connectivity



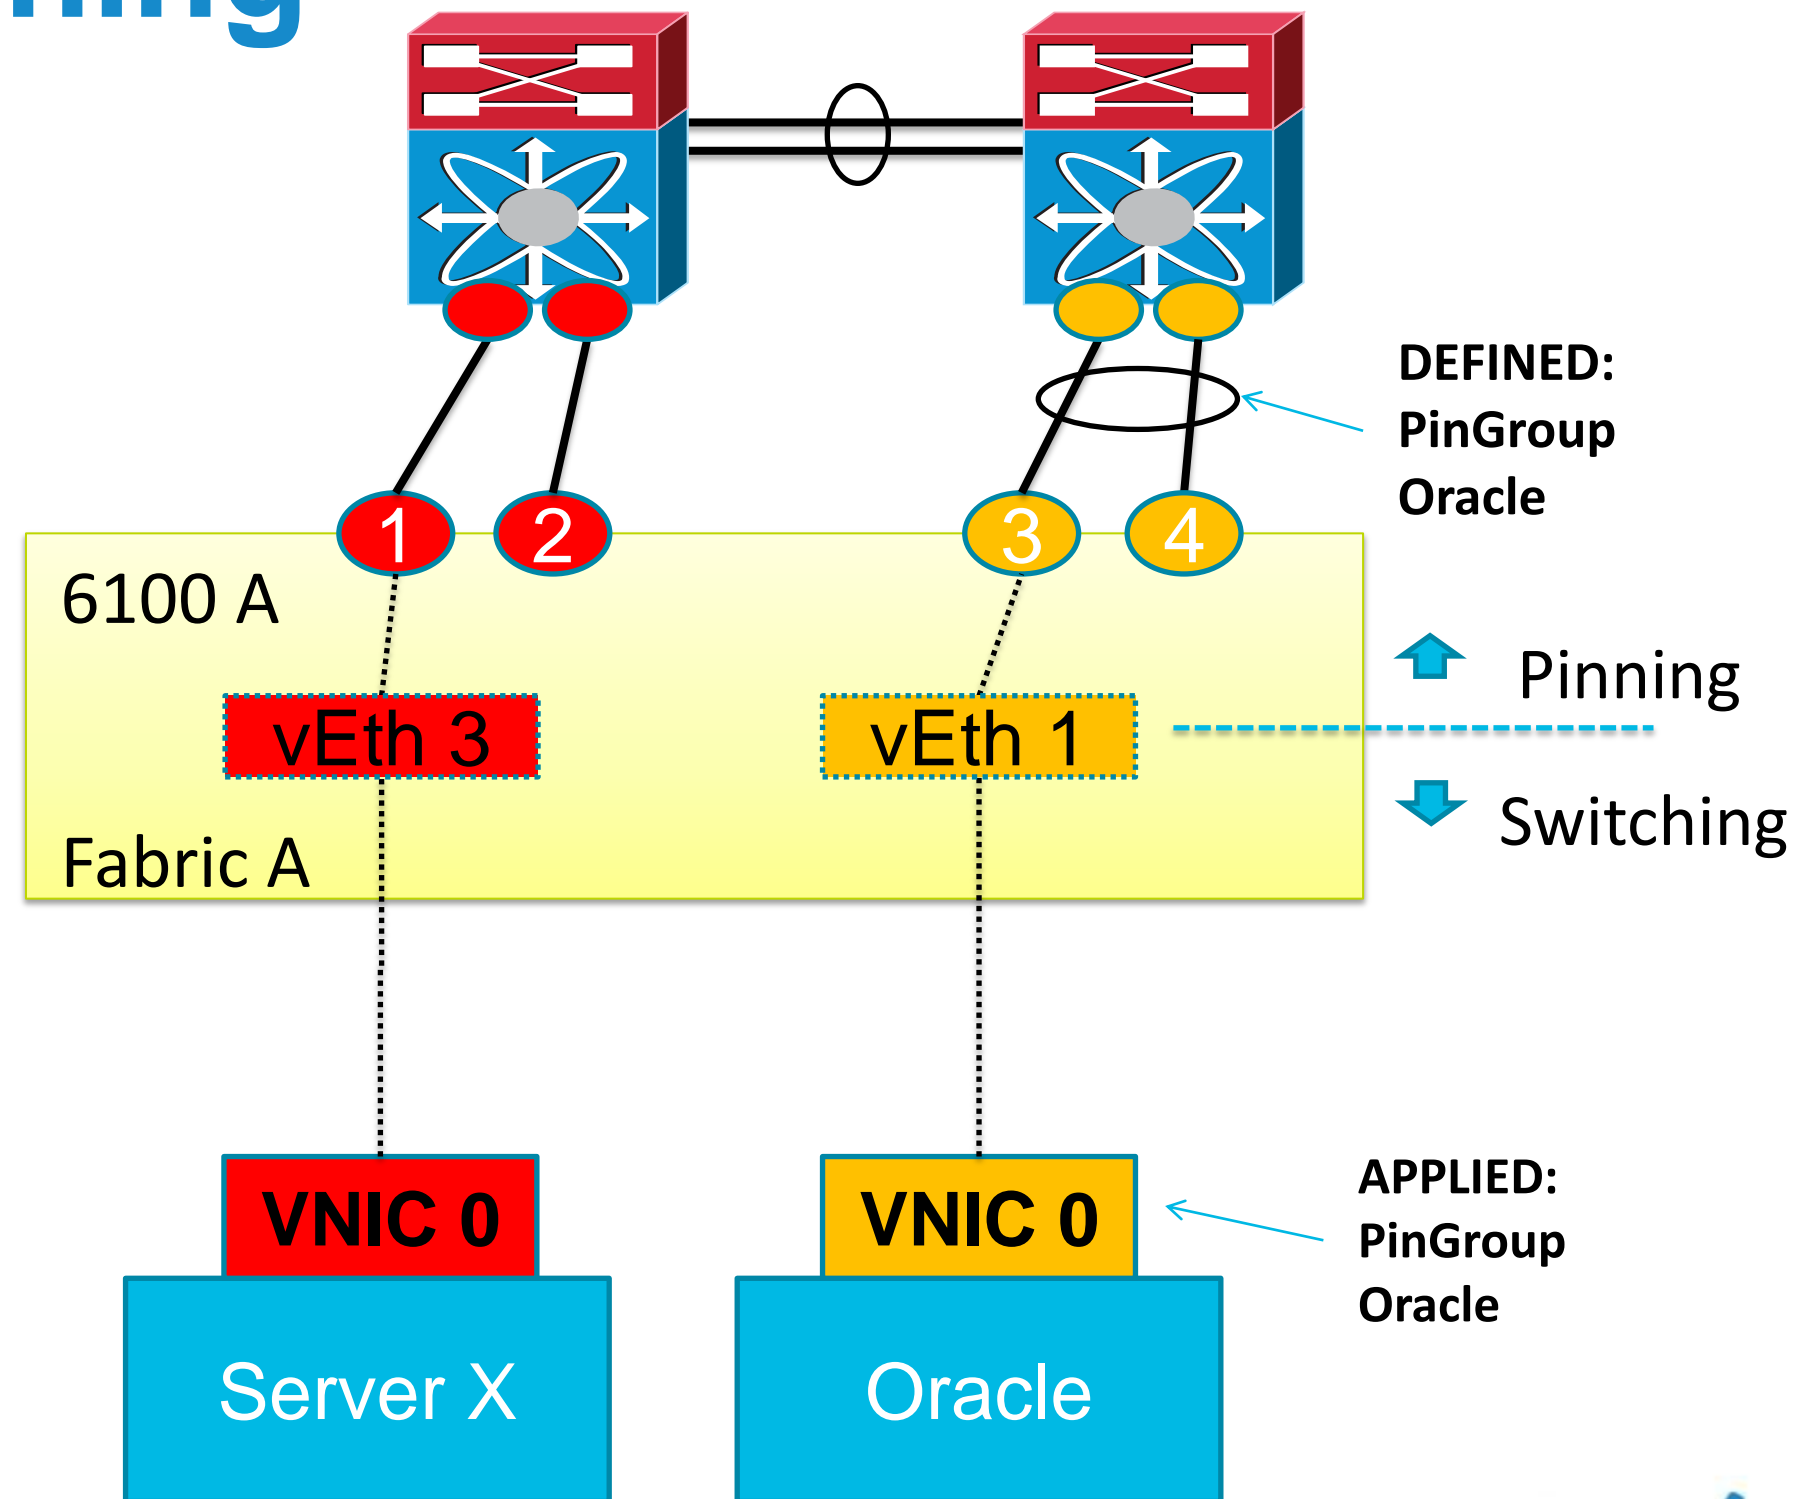
Switching Modes: End Host




- UCS connects to the LAN like a server, not like a switch
- Server vNIC pinned to an uplink port
 - No Spanning Tree Protocol
- Reduces control plane load on FI
- Simplified upstream connectivity
- Maintains MAC table for servers only
 - Eases MAC table sizing in the access layer
- Allows multiple active uplinks per VLAN
- Doubles effective bandwidth vs STP
- Prevents loops by preventing uplink-to-uplink switching
- Completely transparent to upstream LAN
- Traffic on same VLAN switched locally

End Host Mode : Pinning

- Dynamic pinning
 - Server ports pinned to an uplink port/port-channel automatically
- Static pinning
 - Specific ping groups created and associated with adapters
 - Static pinning allows traffic management, if required for certain applications/servers



End Host Mode: Upstream Connectivity




End Host Mode: Disjoint Layer 2 Networks

Customer benefits

Ability to support multiple layer 2 disjoint networks upstream to UCS in End Host Mode

Feature details

- Static or dynamic vNIC pinning based on VLAN membership of the uplink
- A VLAN can exist only in one L2 disjoint network, i.e. no overlap
- A vNIC is mutually exclusive to a L2 network upstream, i.e. a L2 network per vNIC
- More than two L2 disjoint networks supported per host with virtual interface card

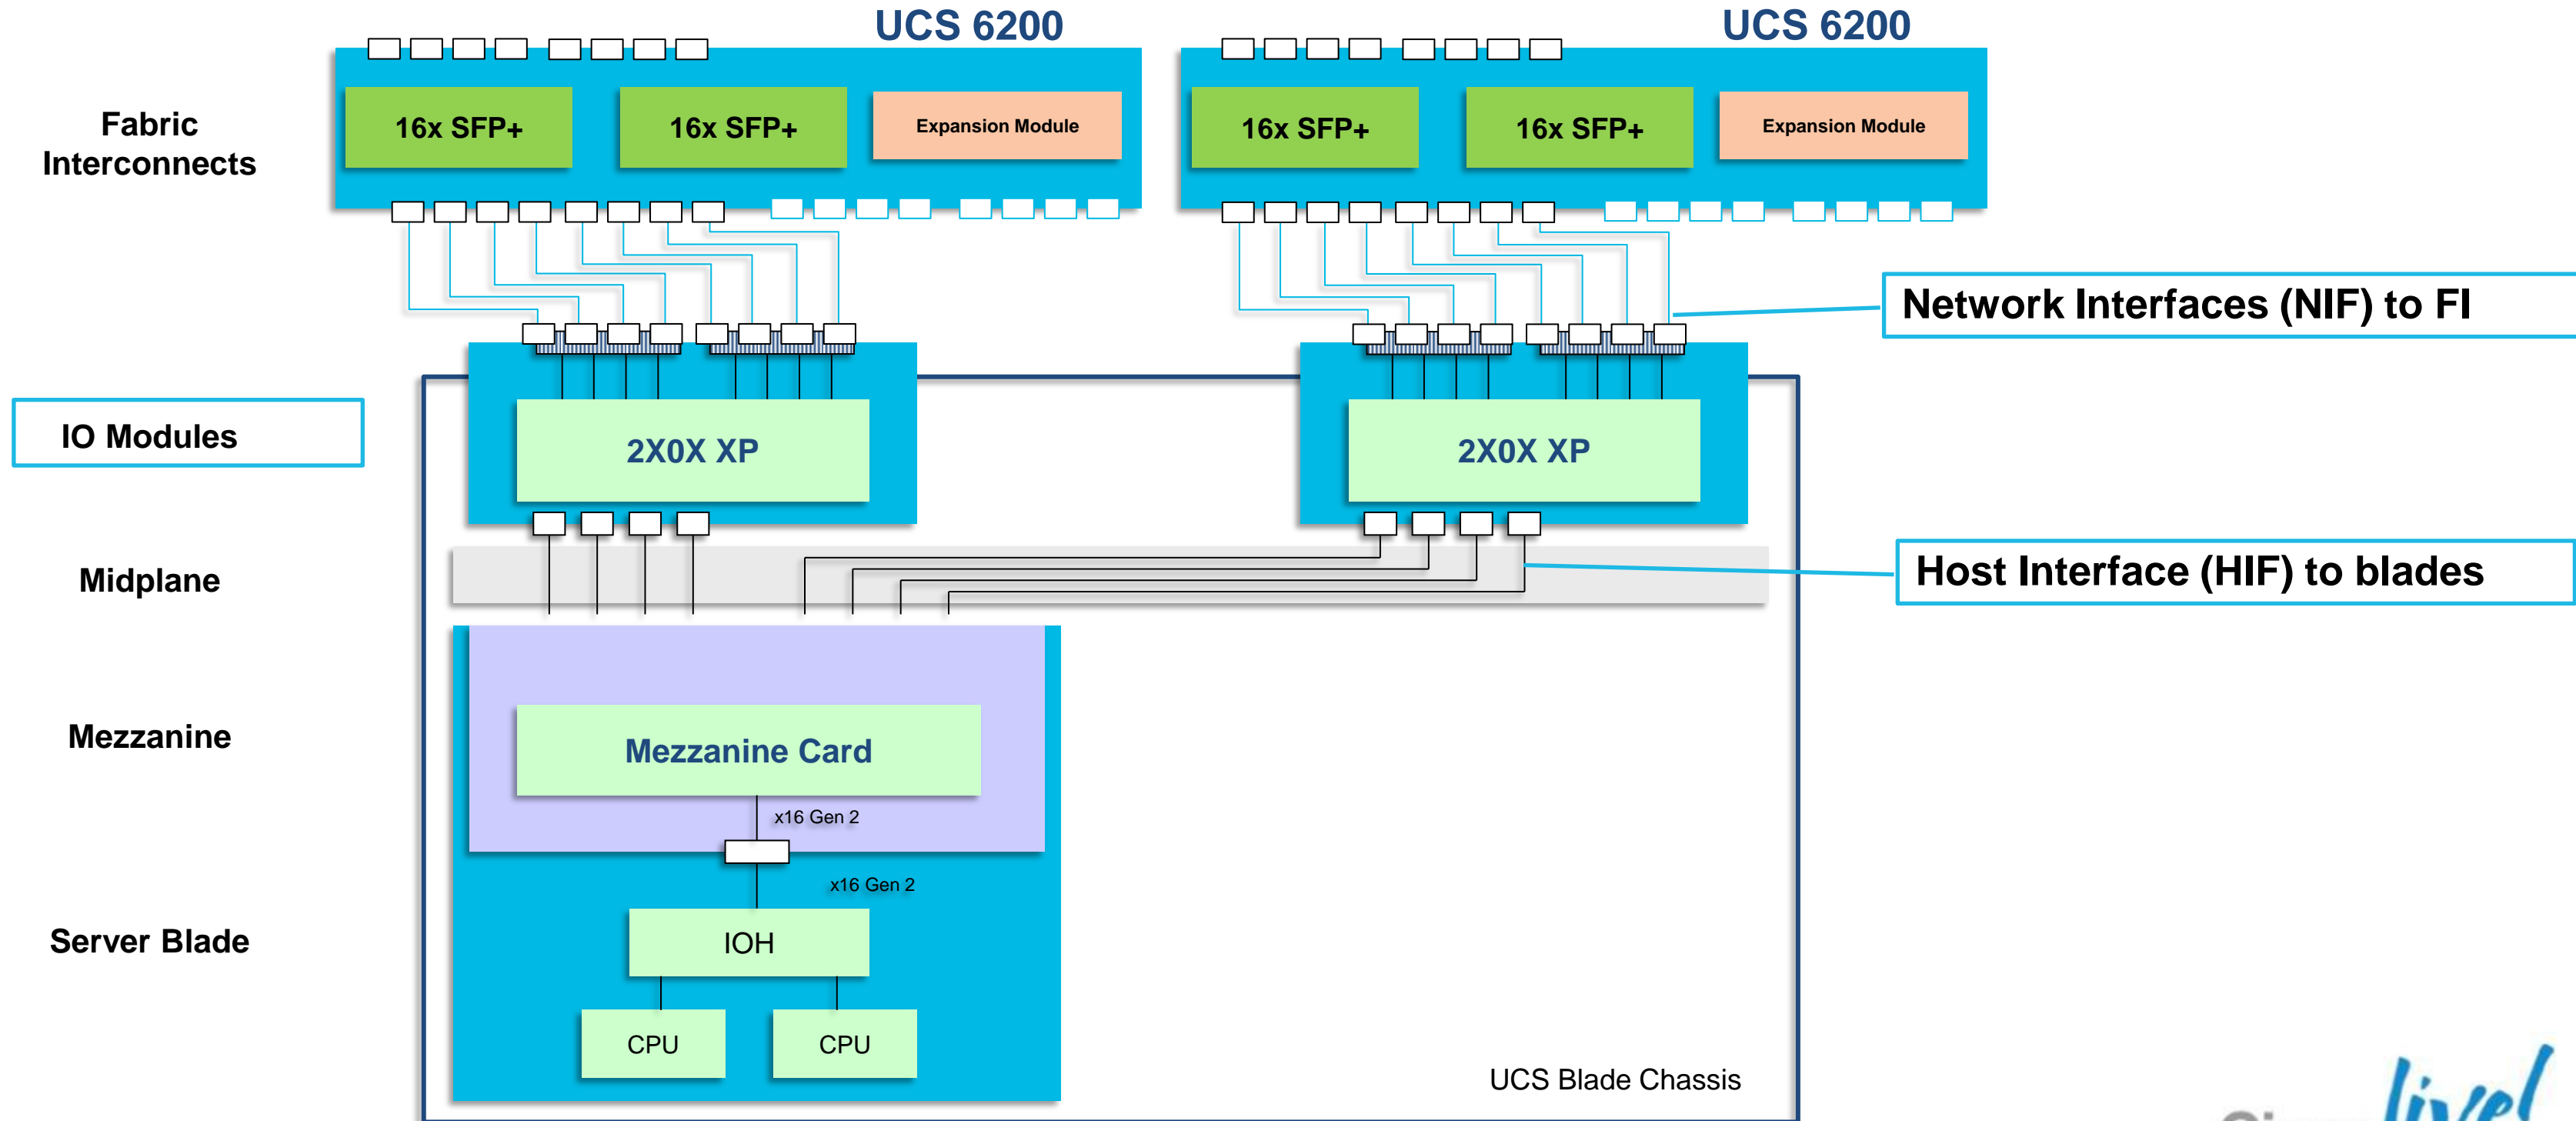


Switch Mode

Application Specific Scenarios

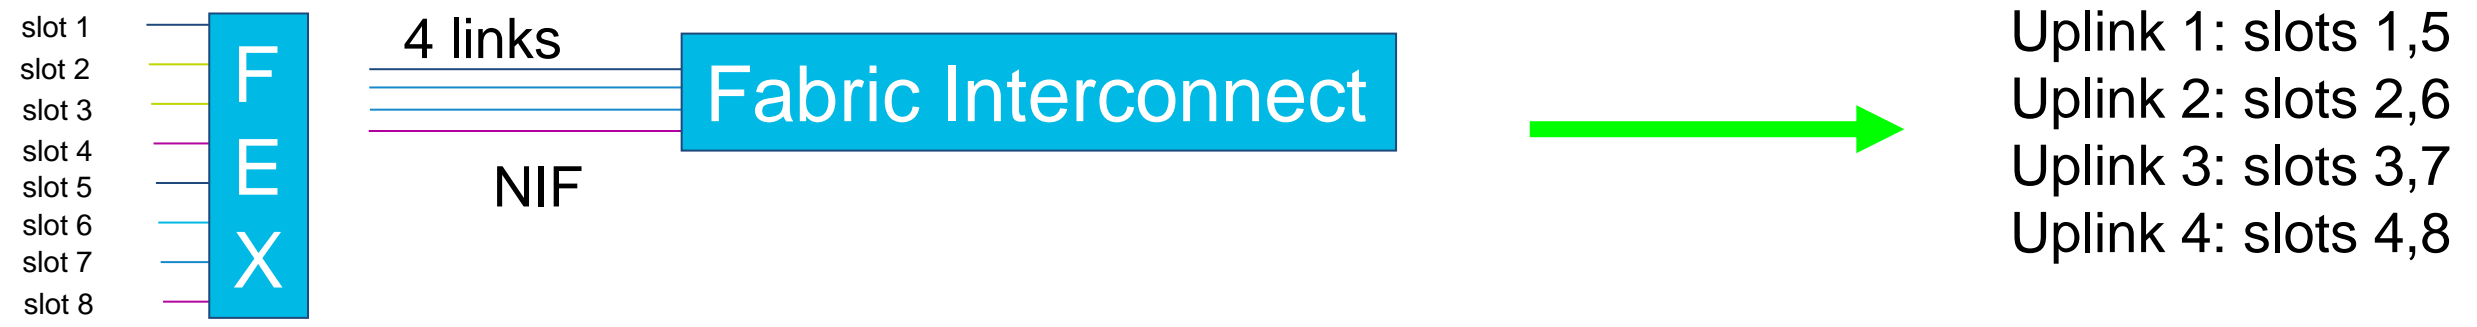
- Certain application like MS-NLB (unicast mode) have the need for unknown unicast flooding, which is not done in EHM(End Host Mode)
- Certain network topologies provide better network path out of the Fabric Interconnect due to STP root placement and HSRP L3 hop
- Switch Mode is “catch all” for different scenarios when EHM is not an option

Block Diagram: Host To Network Pinning




UCS 2X0X (Gen 1/2 FEX) : Host to Network Pinning

Server slots pinned to uplink



UCS 2208 (Gen 2 FEX) : Host to Network Pinning



Server slots pinned to uplink

- Uplink 1: slot 1
- Uplink 2: slot 2
- Uplink 3: slot 3
- Uplink 4: slot 4
- Uplink 5: slot 5
- Uplink 6: slot 6
- Uplink 7: slot 7
- Uplink 8: slot 8

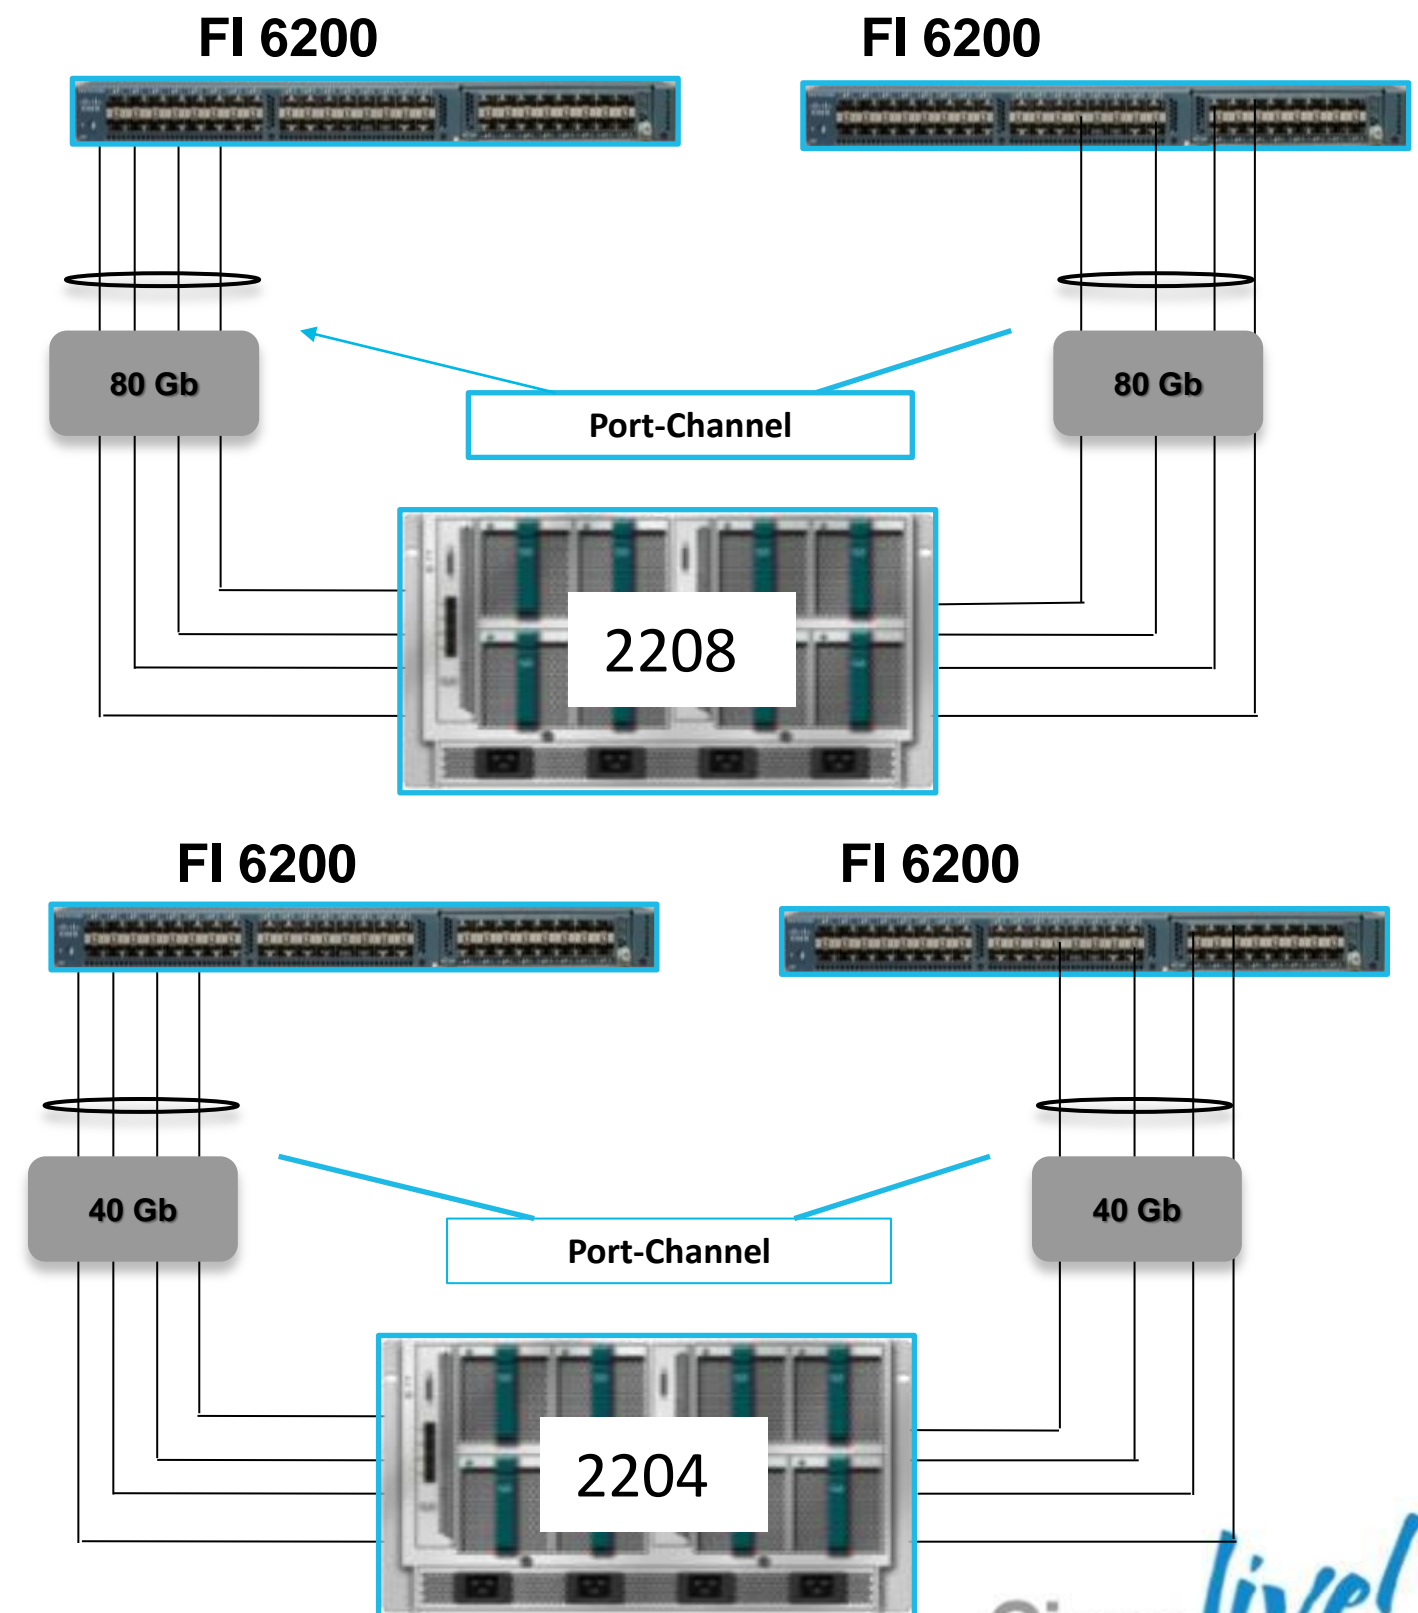
Fabric Port-Channel

Customer benefits

Resilience and flexibility


Feature details

- Up to 8 * 10 Gb aggregated bandwidth to each fabric in a chassis
- Port-channel is a user configurable choice, default mode discrete
- Addition of links does not require chassis re-acknowledge
- Supported number of links in a port-channel: 1,2,3,4,5,6,7 or 8
- Load balancing




UCS 2200 (Gen 2 FEX): Host to Network Pinning

Server slots channelled across all uplinks



- Uplink 1: slots 1-8
- Uplink 2: slots 1-8
- Uplink 3: slots 1-8
- Uplink 4: slots 1-8



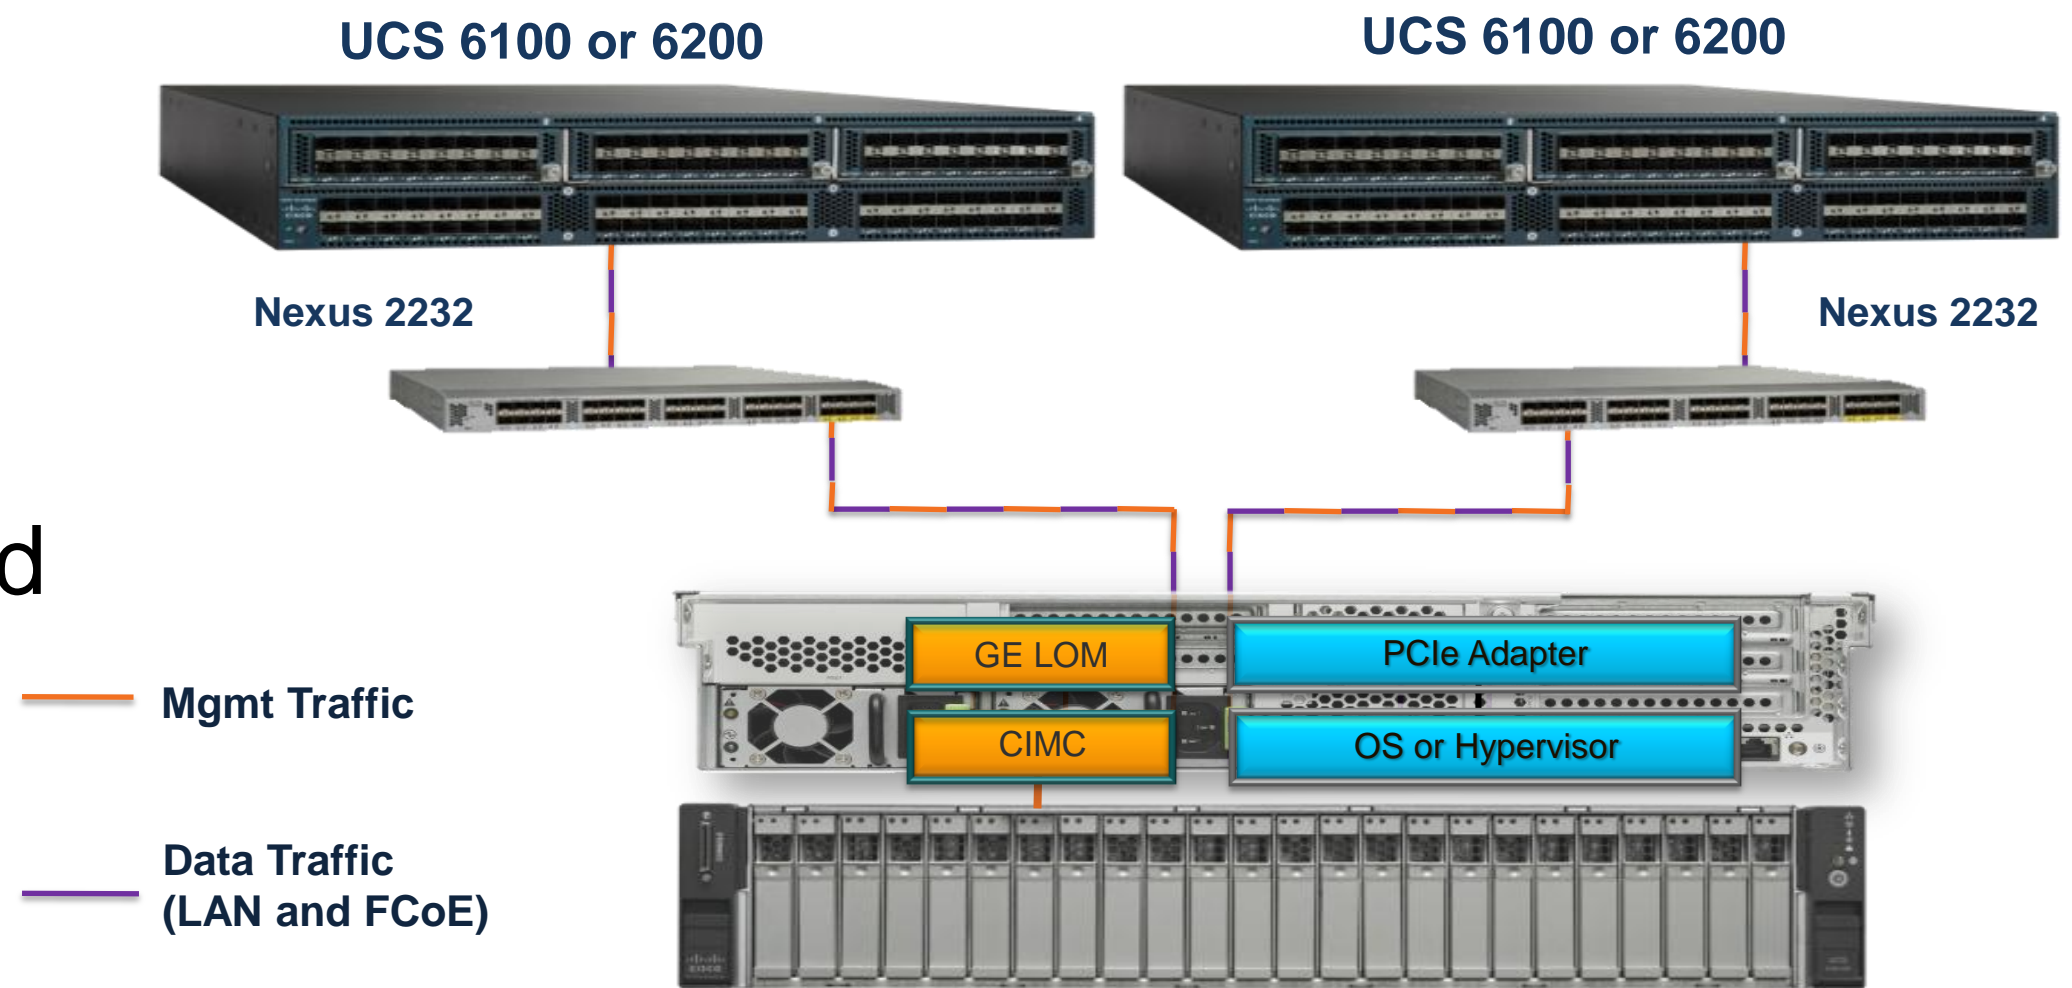
Server slots channelled across all uplinks

- Uplink 1: slots 1-8
- Uplink 2: slots 1-8
- Uplink 3: slots 1-8
- Uplink 4: slots 1-8
- Uplink 5: slots 1-8
- Uplink 6: slots 1-8
- Uplink 7: slots 1-8
- Uplink 8: slots 1-8

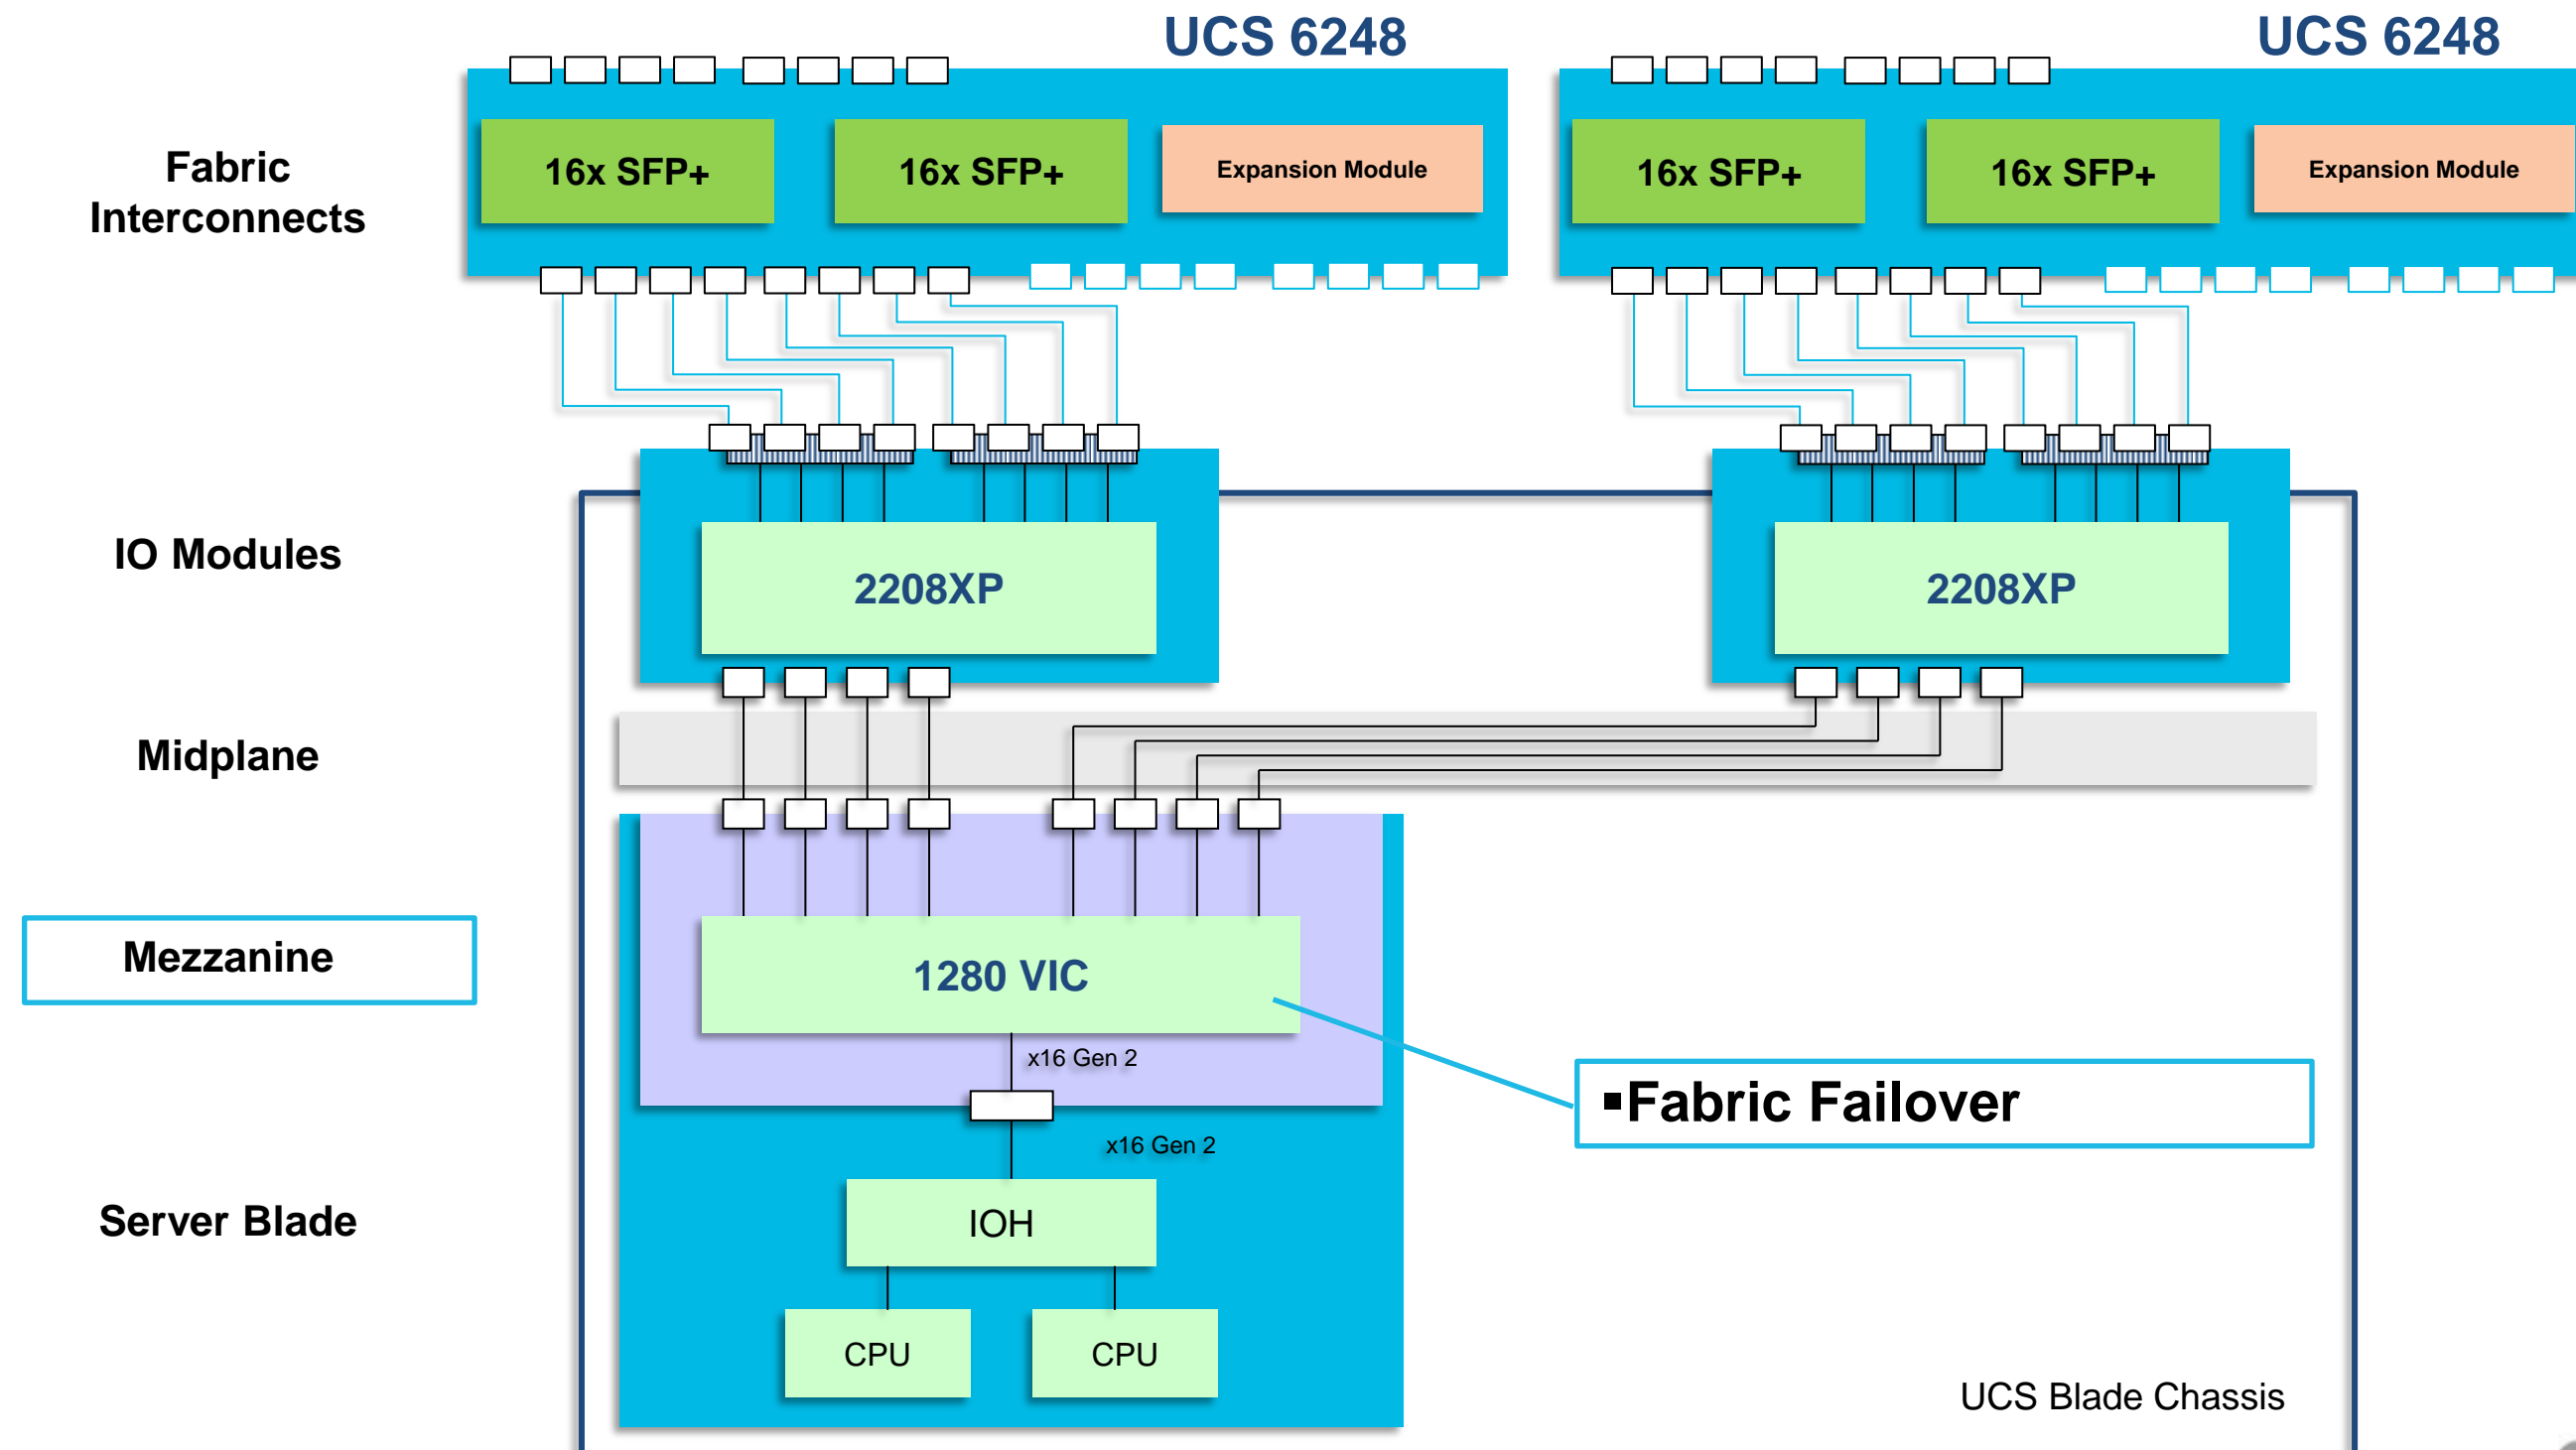


UCSM C-series Integration (Nexus 2232)

- DelMar and C-Peak MR6 Required for VIC 1225
- Default CIMC Settings
 - Shared-LOM-EXT
- Dual Wire Mode still supported
- Multiple VIC 1225 supported
 - C240, C260, C420, C460
- NO Direct FI SUPPORT

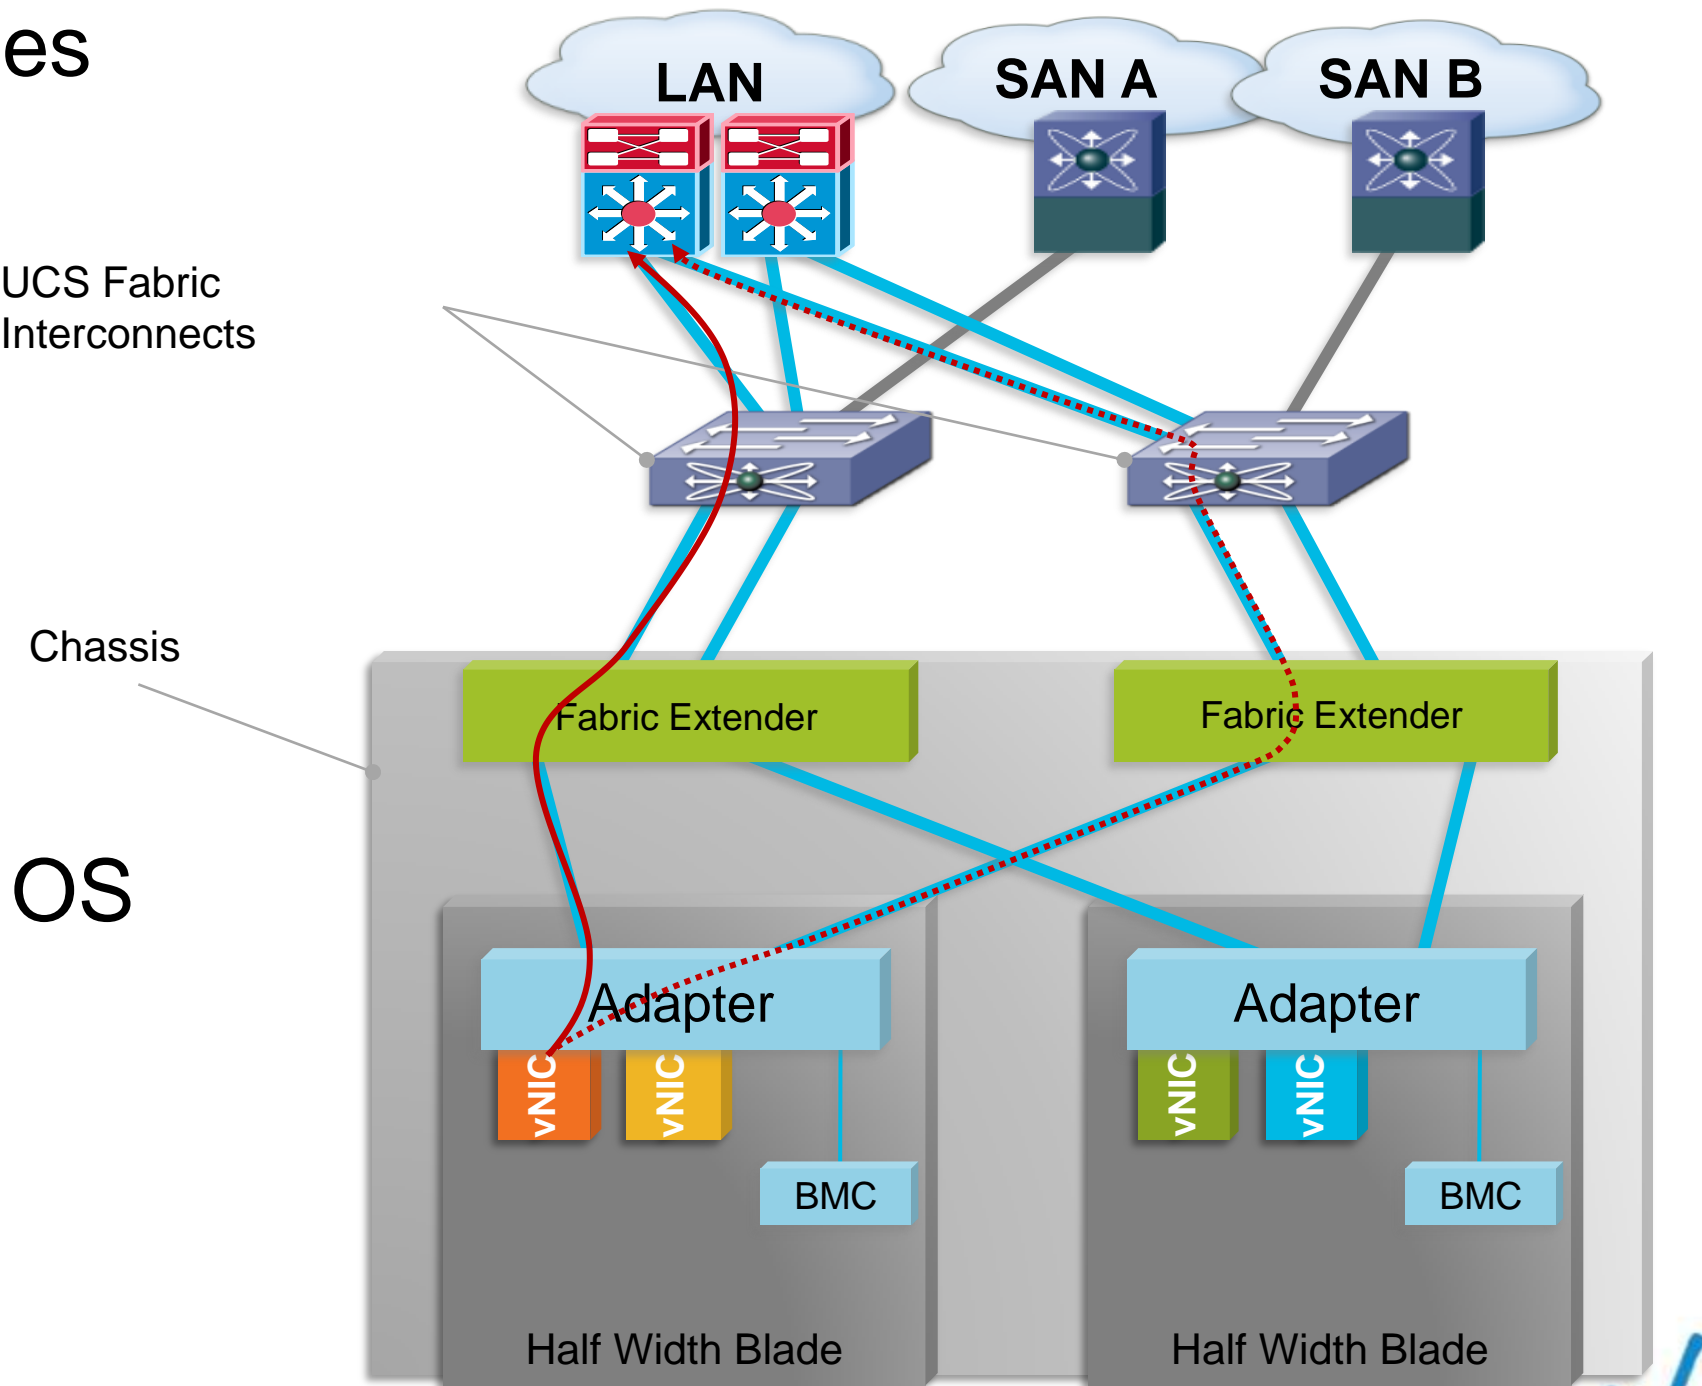


Block Diagram: Fabric Failover

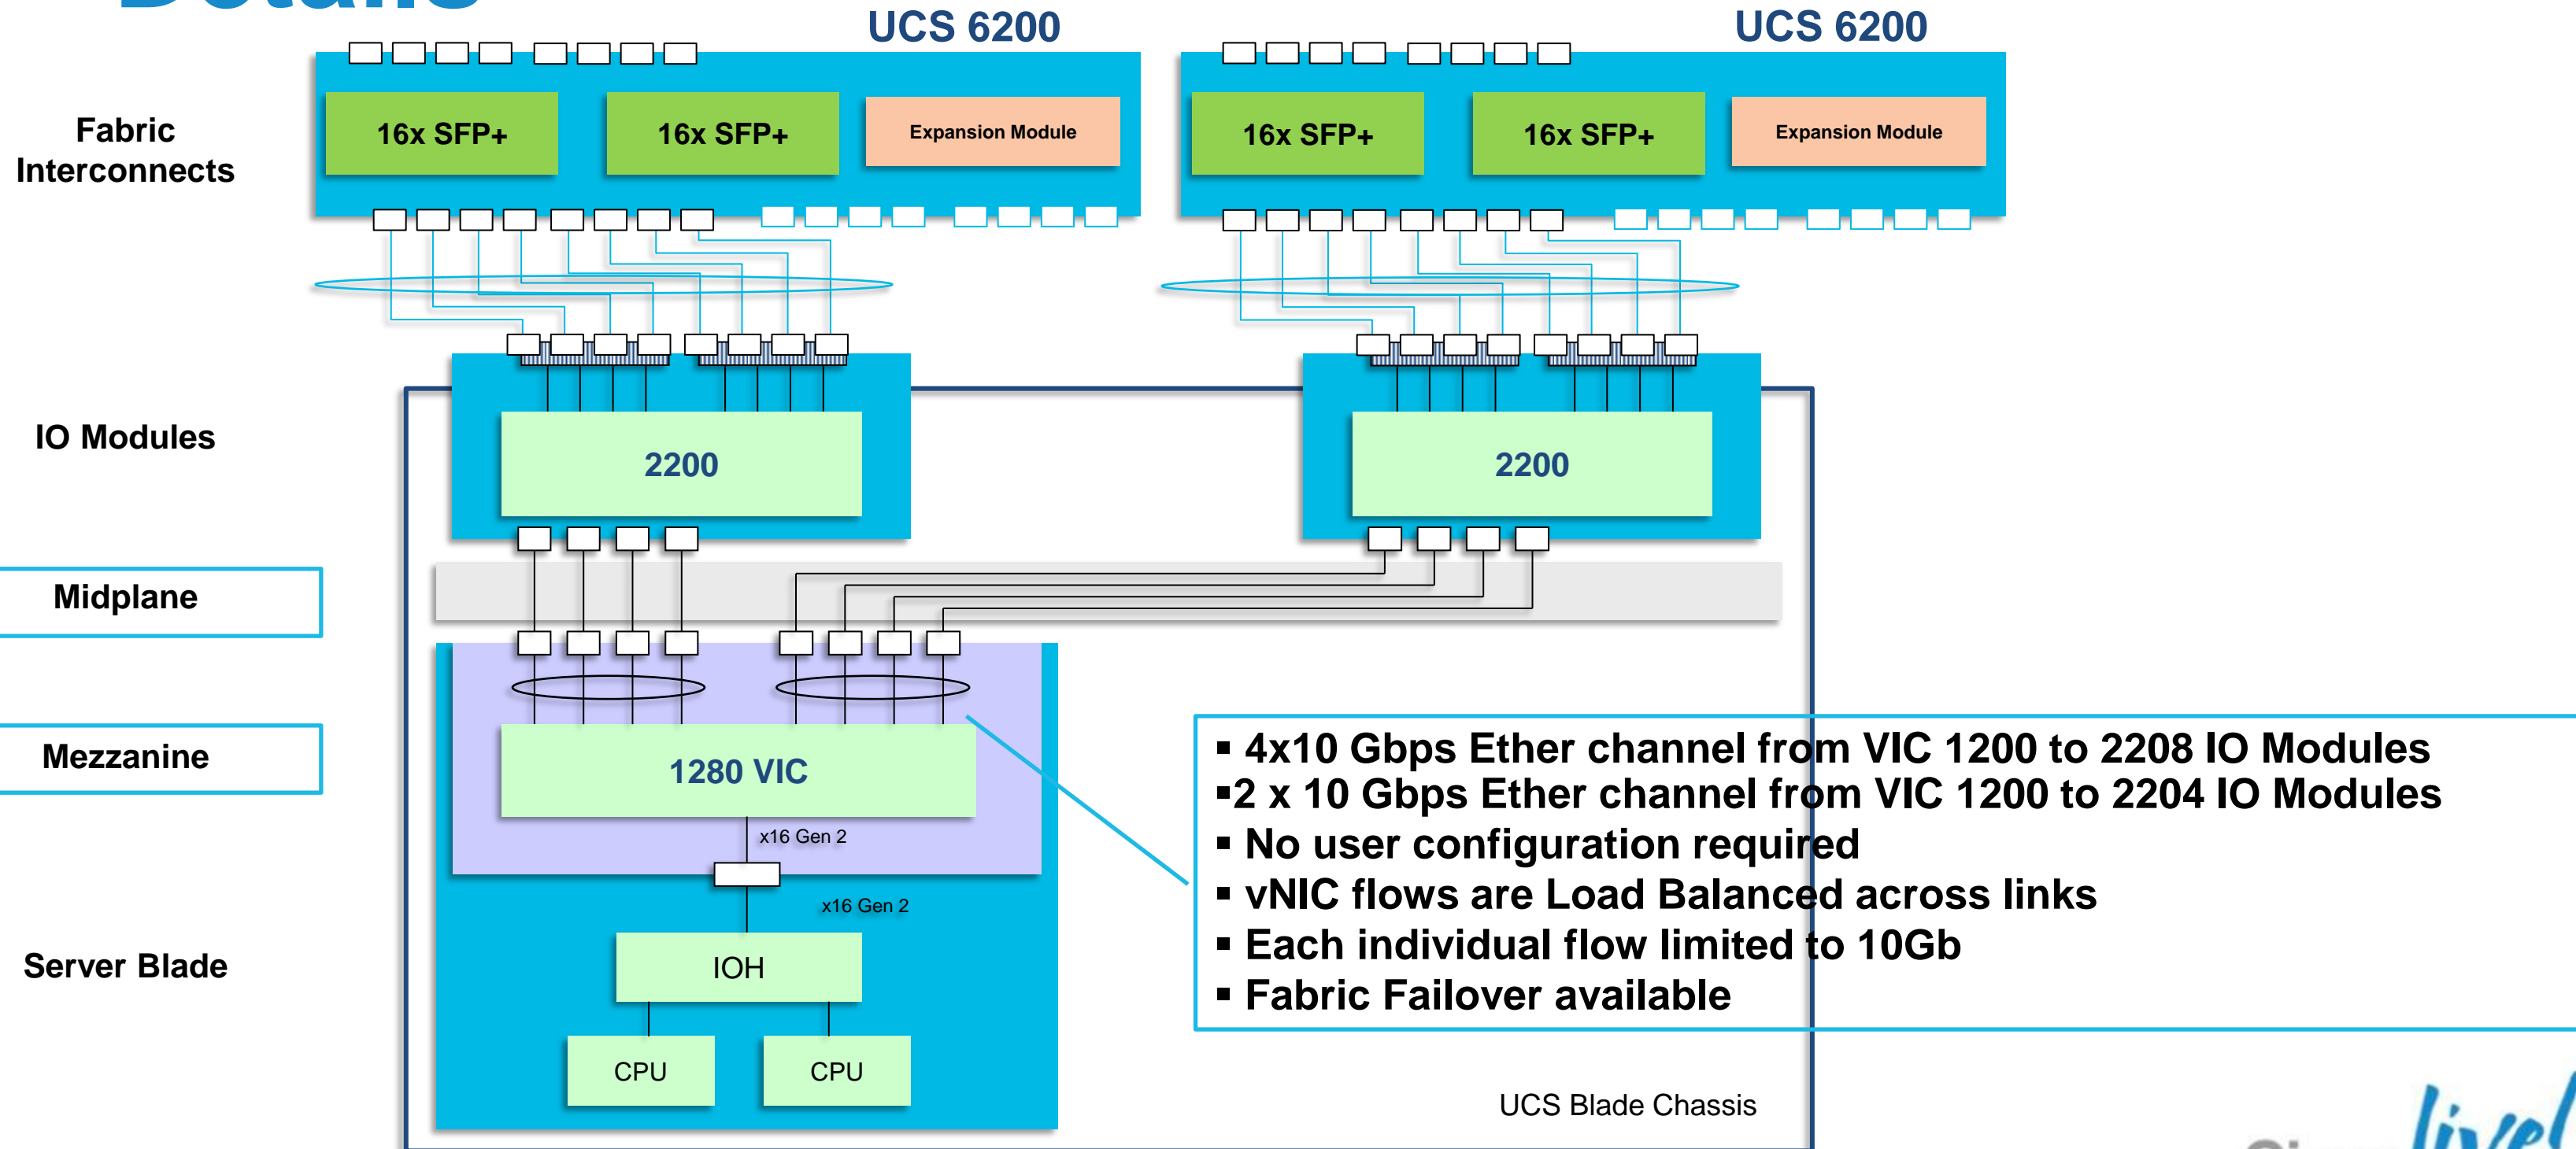


End Host Mode: Fabric Failover for Ethernet

- Fabric provides NIC failover capabilities 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
- Recommended in case of bare metal OS
- Hyper - V




Block Diagram: Next Gen UCS Fabric Details



Adapter Port Channel

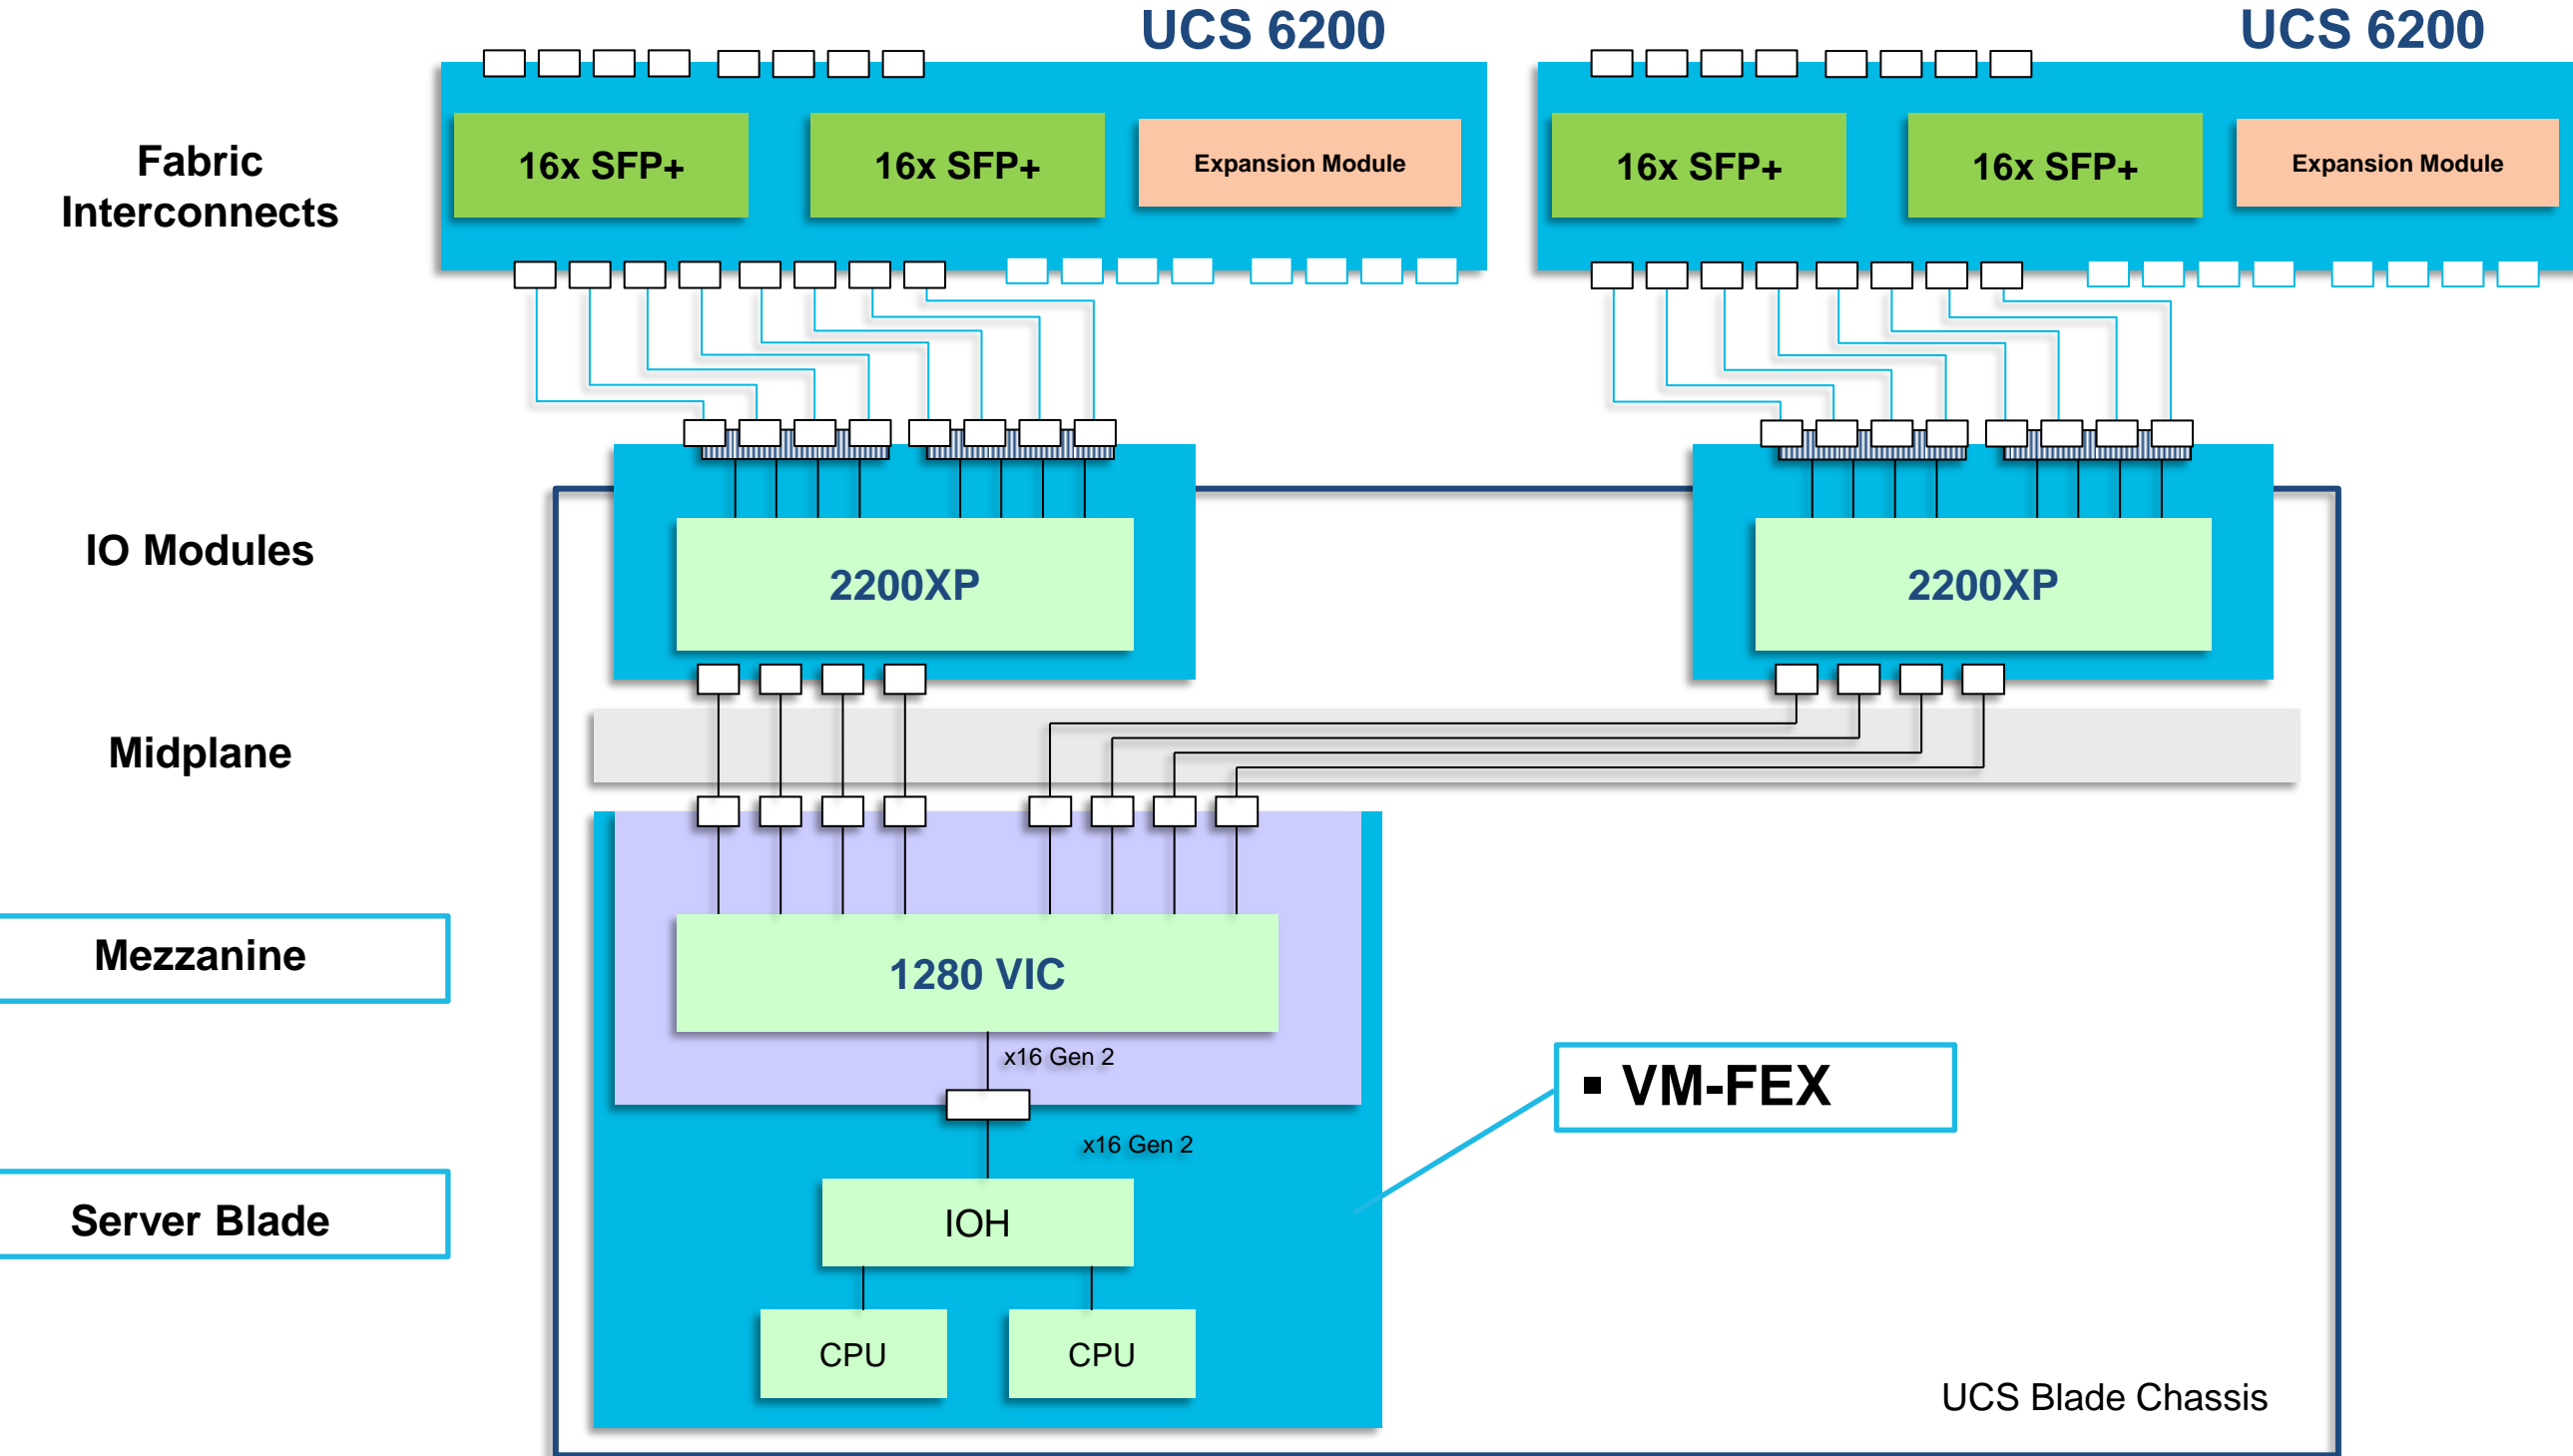
Up to 32 Gbps throughput per vNIC using flow based port-channel hash



- Implicit Port-channel between UCS 1280 VIC adapter and UCS 2200 IOM
- Flow based hash
- A vNIC is active on side A or B
- A vNIC have access to up to 32 Gbps throughput

1. 10 Gb FTP traffic
2. 10 Gb UDP traffic

Block Diagram: VM-FEX



UCS Blade Chassis



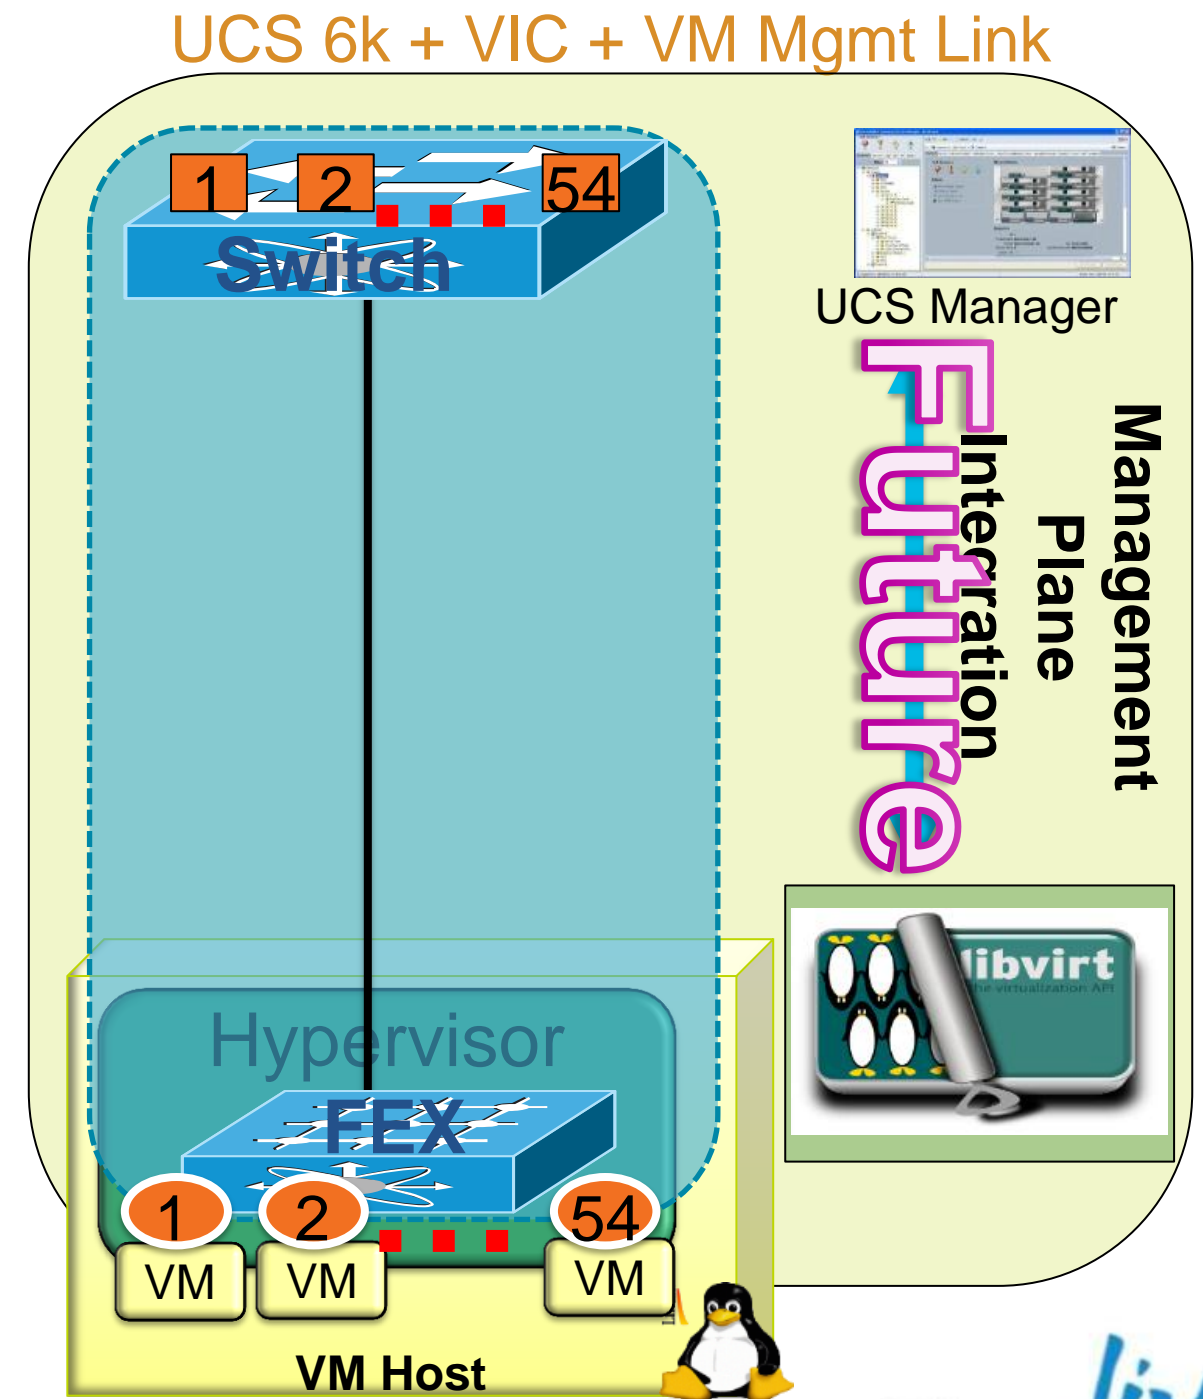
VM-FEX with RedHat KVM

Customer benefits

Ability to unify virtual and physical network infrastructure for Red Hat KVM

Feature details

- Available with RHEL 6.1
- Default cluster for standard based port extenders
- No UCSM and RHEV-M interactions in UCS 2.0
- Virtual machine management via Libvirt tool

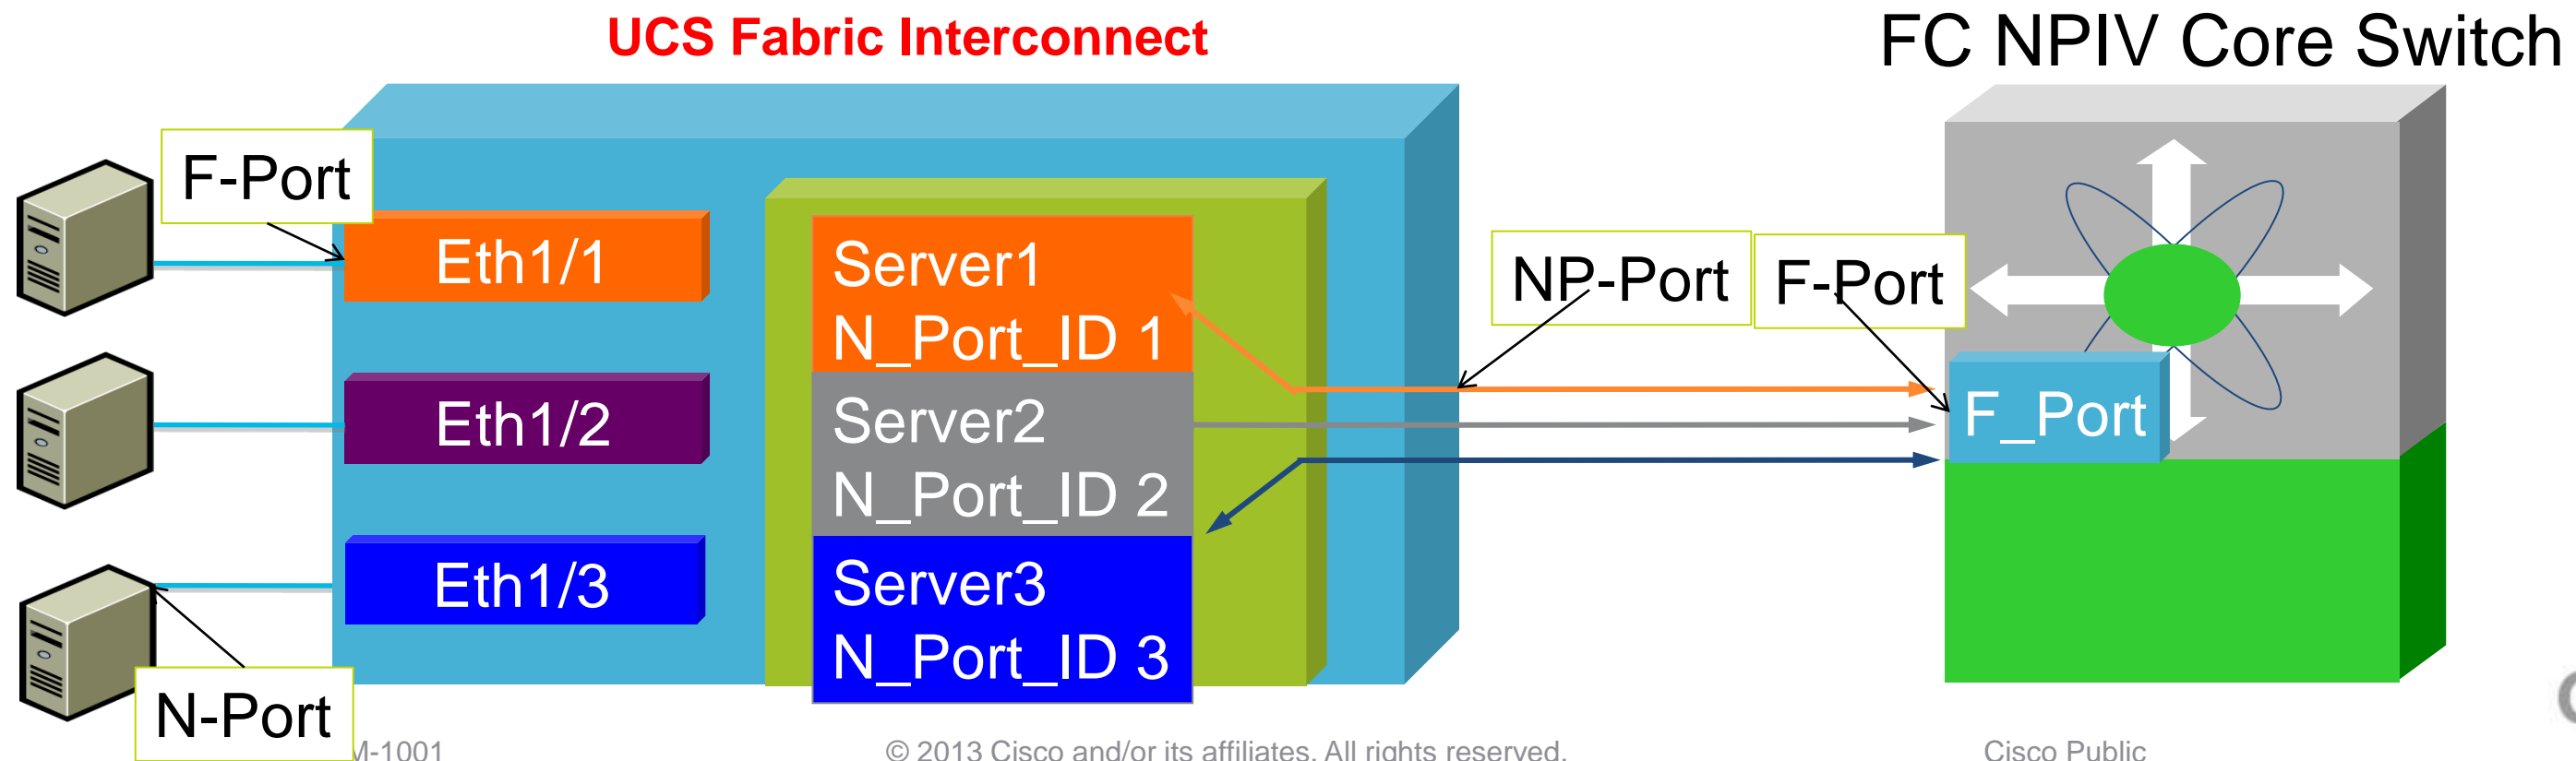


Storage Connectivity



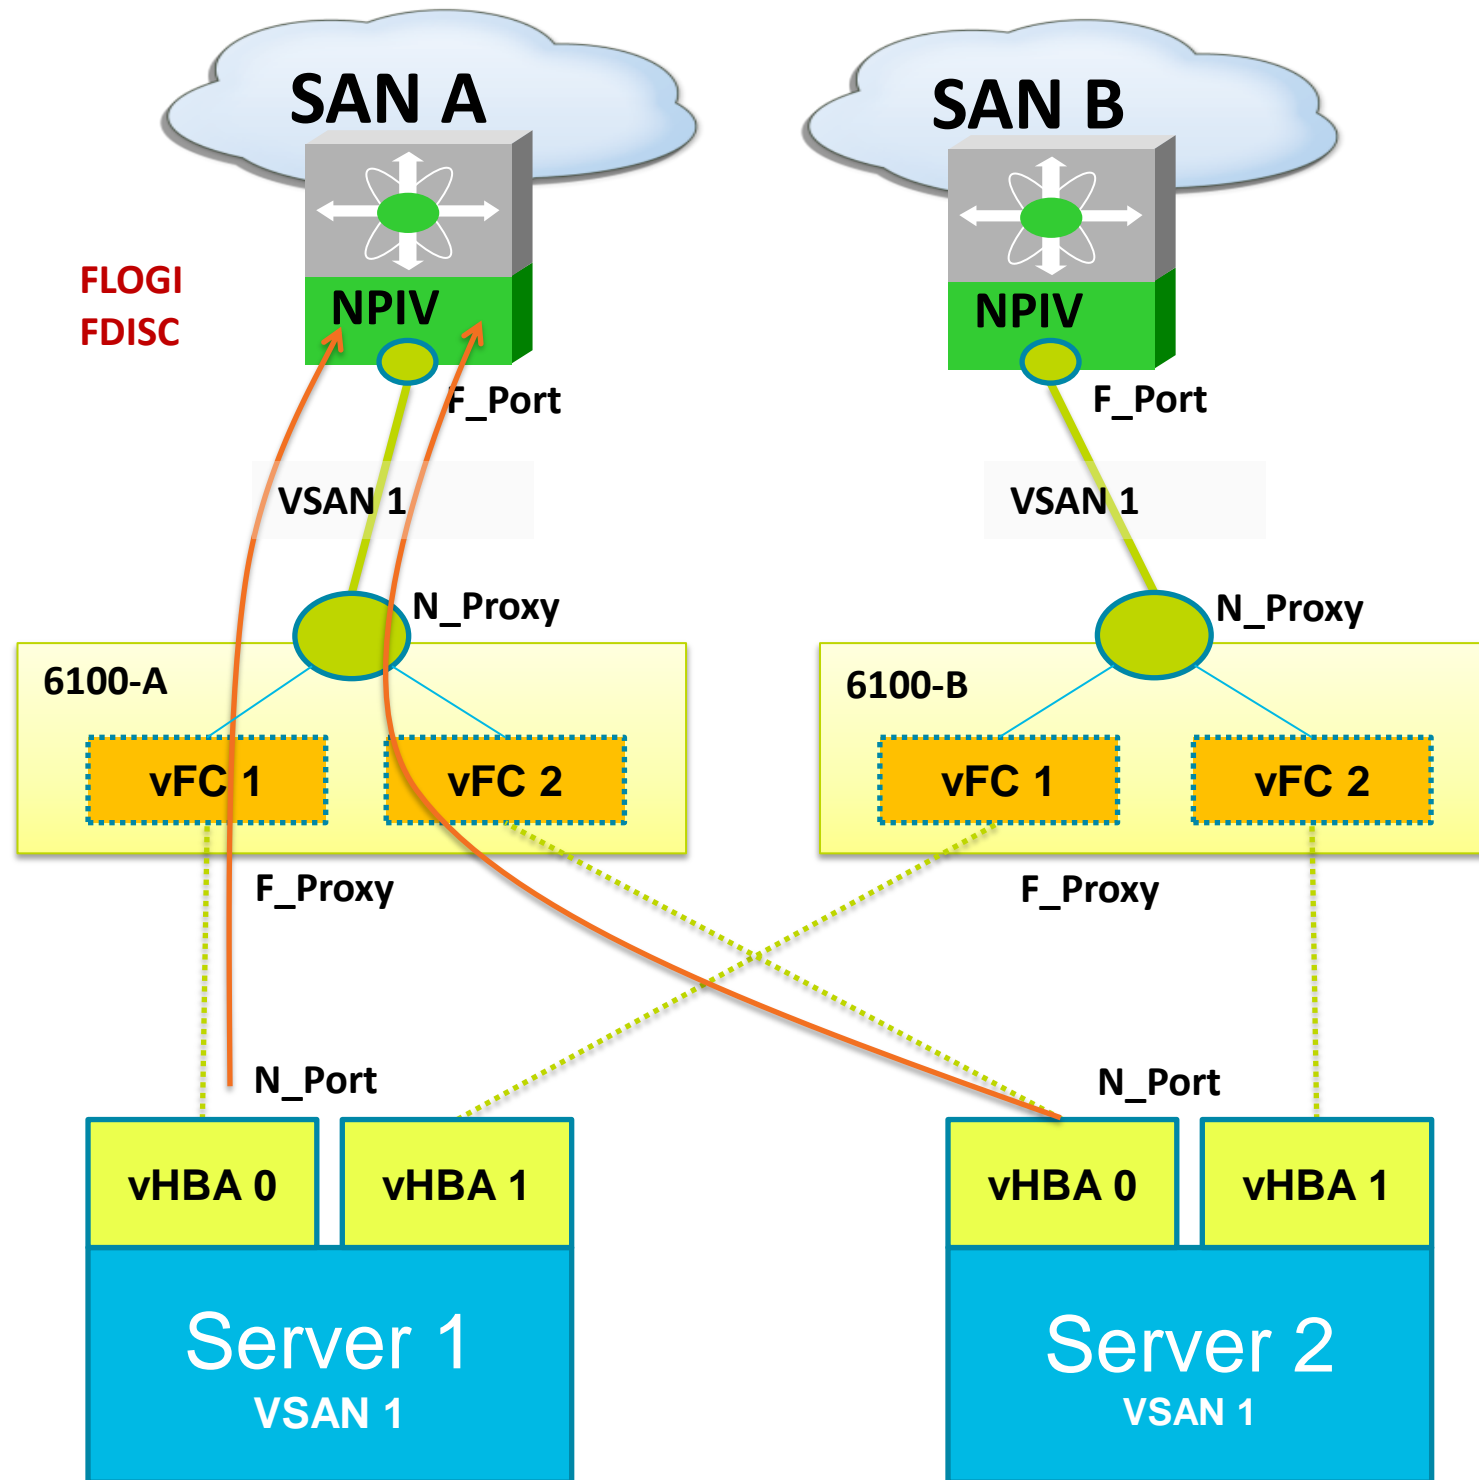
NPV

- N-Port Virtualiser (NPV) utilises NPIV functionality to allow a “switch” to act like a server performing multiple logins through a single physical link
- Physical servers connected to the NPV switch login to the upstream NPIV core switch
 - Physical uplink from NPV switch to FC NPIV core switch does actual “FLOGI”
 - Subsequent logins are converted (proxy) to “FDISC” to login to upstream FC switch
- No local switching is done on an FC switch in NPV mode
- FC edge switch in NPV mode does not take up a domain ID



SAN “End Host” NPV Mode


N-Port Virtualisation Forwarding



- Fabric Interconnect operates in N_Port Proxy mode
 - Simplifies multi-vendor interoperoperation
 - Simplifies management
- SAN switch sees Fabric Interconnect as an FC End Host with many N_Ports and many FC IDs assigned
- Server facing ports function as F-proxy ports
- Server vHBA pinned to an FC uplink in the same VSAN. Round Robin selection
- Provides multiple FC end nodes to one F_Port off an FC Switch
- Eliminates the FC domain on UCS Fabric Interconnect
- One VSAN per F_port (multi-vendor)
- F_Port Trunking and Channelling with MDS, 5K

SAN "End Host" NPV Mode

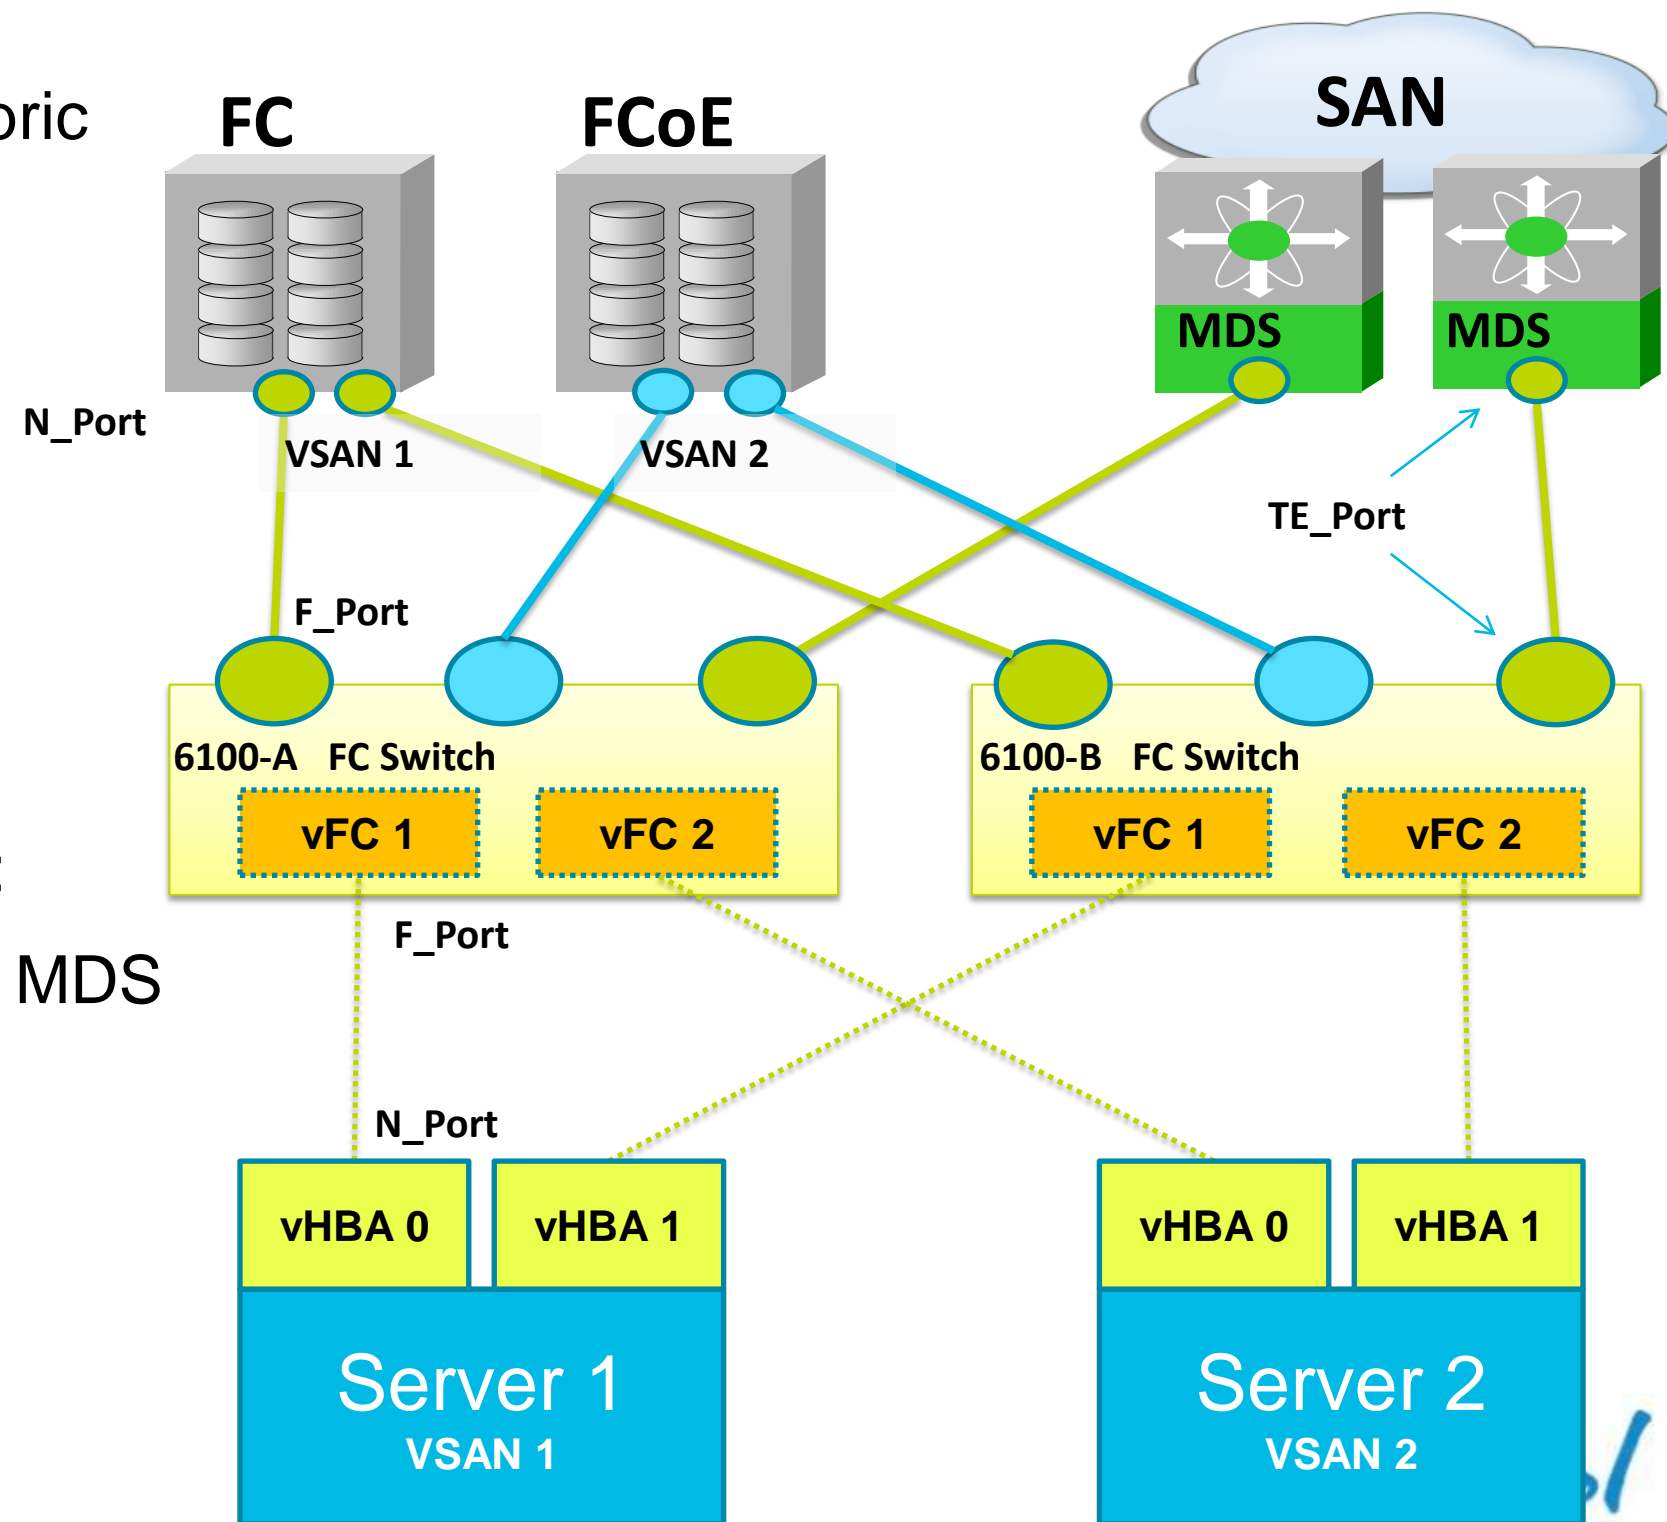
N-Port Virtualisation Forwarding with MDS, Nexus 5000



- F_Port Channelling and Trunking from MDS or Nexus 5000 to UCS
- FC Port Channel behaves as one logical uplink
- FC Port Channel can carry all VSANs (Trunk)
- UCS Fabric Interconnects remains in NPV end host mode
- Server vHBA pinned to an FC Port Channel
- Server vHBA has access to bandwidth on any link member of the FC Port Channel
- Load balancing based on FC Exchange_ID
 - per flow

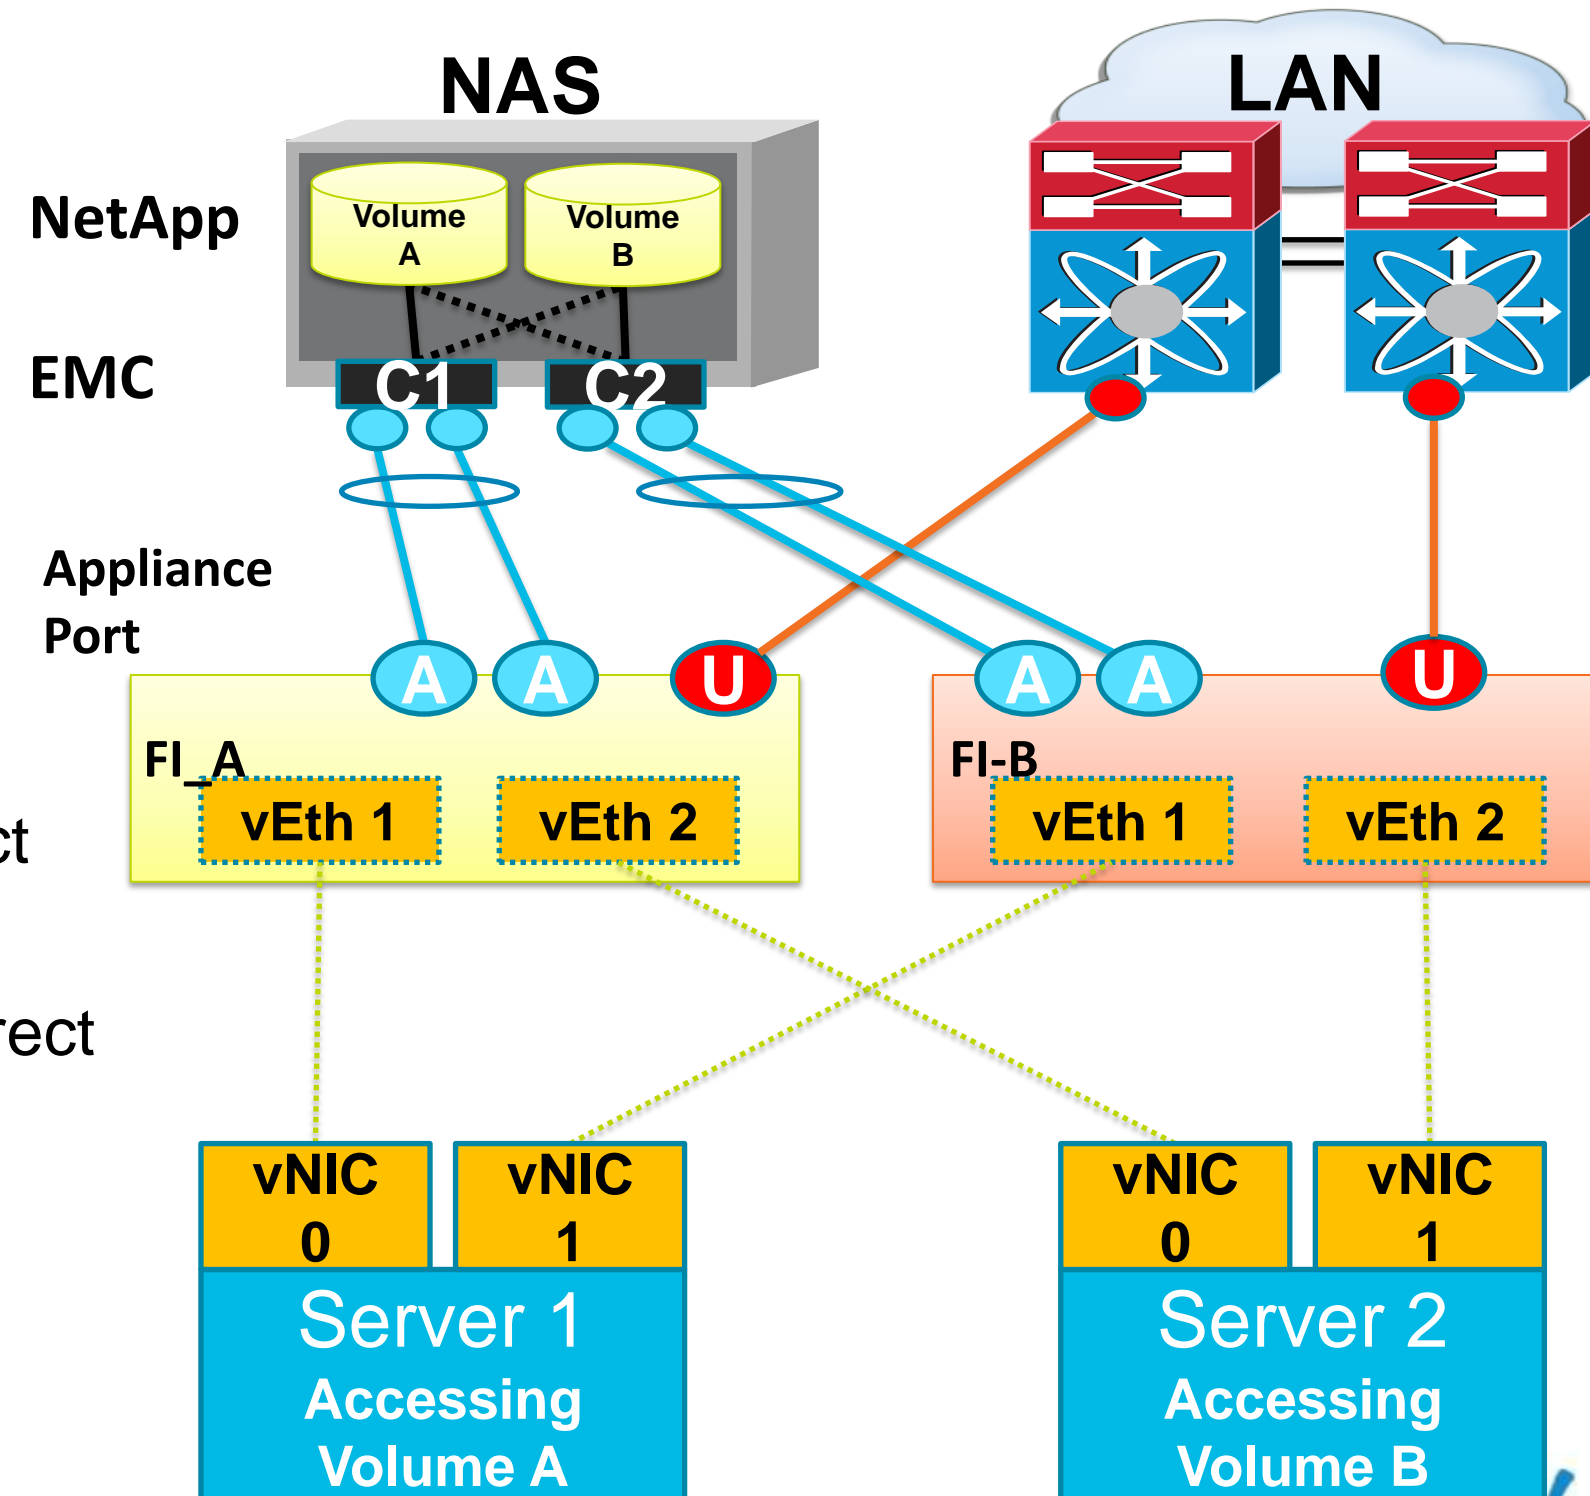
SAN FC Switch Mode

- UCS Fabric Interconnect behaves like an FC fabric switch
- Storage ports can be FC or FCoE
- Light subset of FC Switching features
 - Select storage ports
 - Set VSAN on storage ports
- No zoning configuration inputs in UCSM
- Connection to an external FC switch is required:
 - Zoning configured and pushed to UCS from MDS
- Fabric Interconnect uses a FC Domain ID



NAS Direct Attach

- Default (recommended) - End Host Mode
 - Superior traffic engineering
 - Easier integration into network
 - 1.4 Introduced Appliance Ports which allow direct connect NAS filers
- Options - Ethernet Switching Mode
 - As of 1.4, no need to use this mode for NAS direct connect
 - Previous releases required switching mode for direct connect NAS



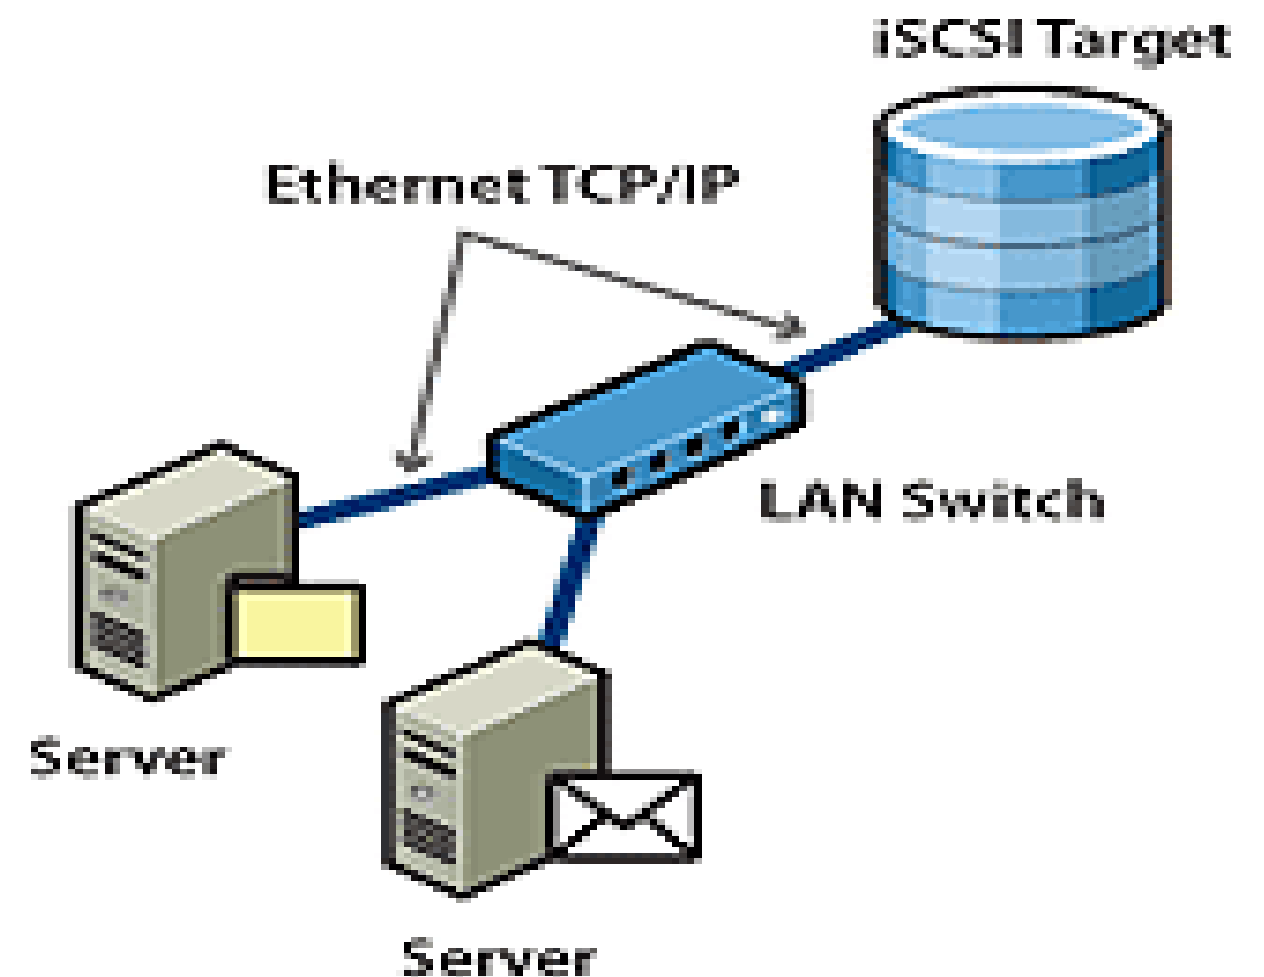
iSCSI Boot Support in UCS Manager

Customer benefits

Comprehensive storage protocol support

Feature details

- Adapters
 - Virtual Interface Card
 - iBFT, no iSCSI offloads
 - Broadcom 57112
 - Full offload, iSCSI HBA
- Operating System support
 - VMware 4.1U1
 - Windows 2008 R2
 - Linux Red Hat 5.6, 6.0, 6.1



Direct Connection of FC/FCoE Storage

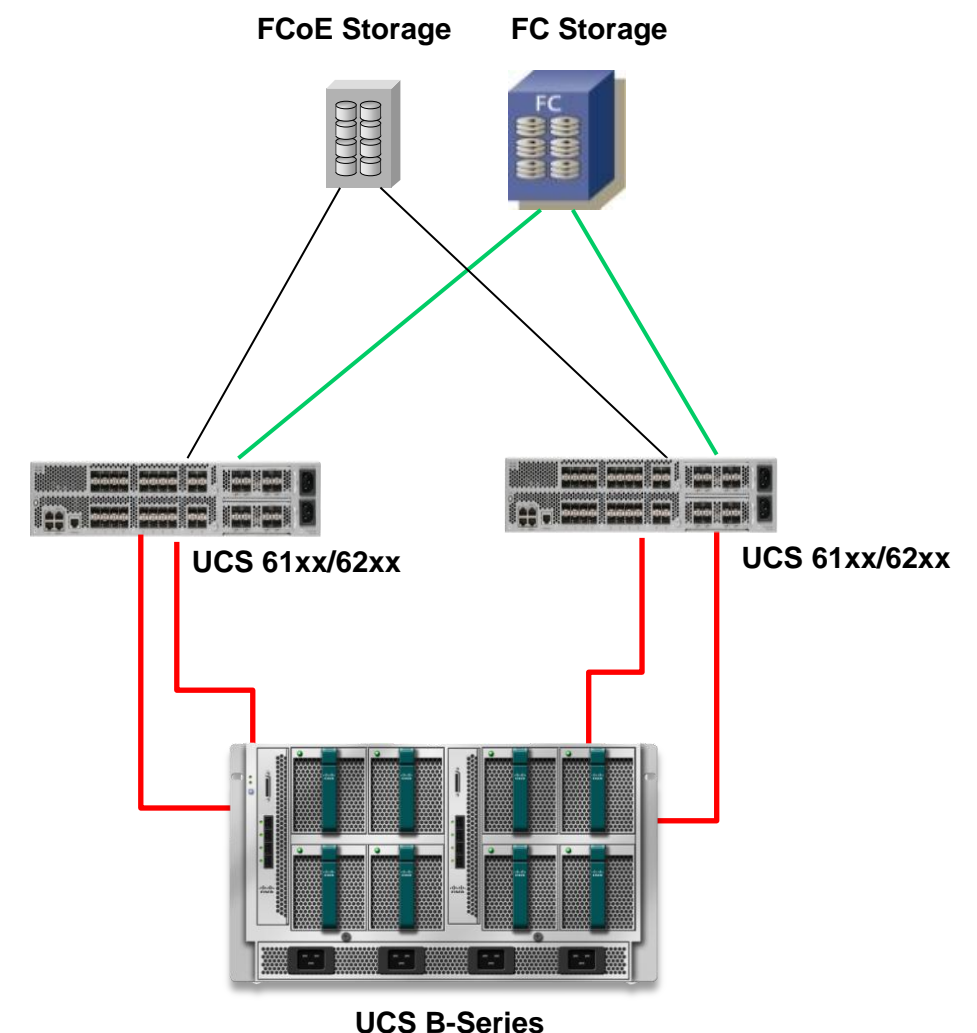
Customer benefits

Support to directly connect FC/ FCoE

Lower cost point for small deployments

Feature details

- Support for EMC and NetApp direct attached FC/FCoE storage (other vendors planned)
- Local Zoning OR Upstream Zoning
- UCS Configured in FC Switch Mode
- Ethernet and FC switching modes are independent




Key Considerations of FC/FCoE Direct Connect

New In 2.1 !!!

- As of April 2011 Default Zoning **Not** Supported
 - Default Zoning Explicitly set to **DENY** in 2.1.1a
 - Default Zoning GUI/CLI Controls Removed
 - Do **Not** Upgrade to 2.1.1a if Default Zoning is being Used
- Local Zoning **OR** Upstream Zoning
 - Parallel Local and Upstream Zoning Currently **NOT** Supported
 - Upstream Zoning Provided by MDS/N5k
 - Migrating from Upstream to Local Zoning
 - CLI Commands to Remove Residual Upstream Zoning
- Supported FC/FCoE Direct Connect Arrays
 - Check Note 5 on HCL for Updated List

Unified Appliance Support

- File and block data over a single port/cable
 - FCoE, iSCSI, NFS, CIFS
- Port and cable consolidation
- New port type: Unified Appliance Port
 - Appliance port of today + FCoE
- Initial support for NetApp storage and their Unified Target Adapter



Demo



Q & A



Complete Your Online Session Evaluation

Give us your feedback and receive a Cisco Live 2013 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 8 March 12:00pm-2:00pm



Cisco *live!* 365

Don't forget to activate your Cisco Live 365 account for access to all session material,

communities, and on-demand and live activities throughout the year. Log into your Cisco Live portal and click the "Enter Cisco Live 365" button.

www.ciscoliveaustralia.com/portal/login.wv

Cisco *live!*

