
cisco Live!
ALL IN

The bridge to possible

Cisco UCS 5th Generation Fabric

Connectivity and Best practices

Eldho Jacob,
@eljacob

BRKDCN-2587

cisco Live!

#CiscoLive

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 17, 2022.

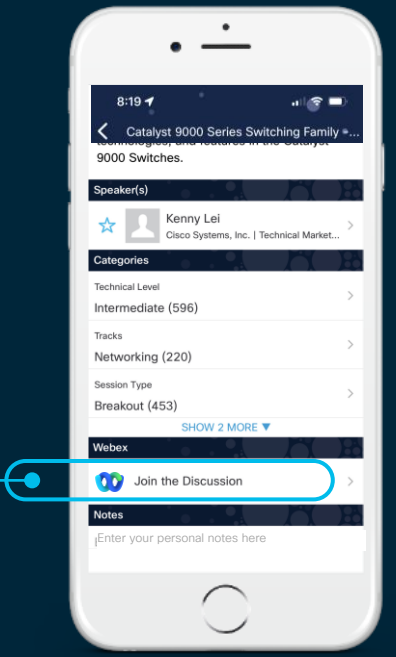

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKDCN-2587

Agenda

- Introduction
- UCS Fabric Simplified
- UCS 5th Gen Fabric Details
- Provision with Intersight
- 5th Gen Fabric connectivity
- Conclusion

UCS Fabric Simplified

UCS Fabric Simplified

Simplicity

Resiliency

TCO Reduction

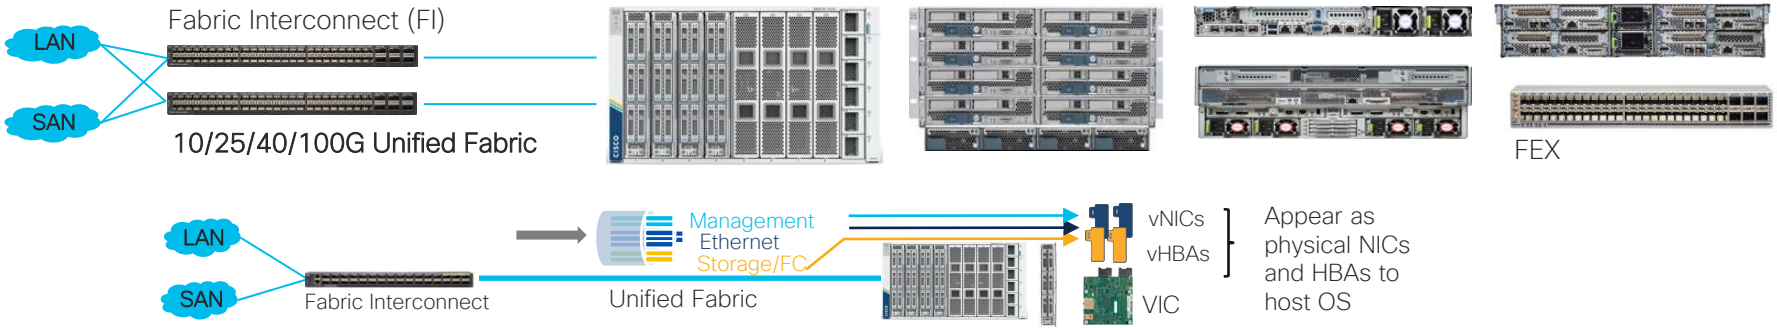

TRADITIONAL RACK

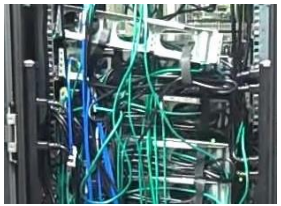

Ad Hoc and Inconsistent

TRADITIONAL BLADE

Structured, but Siloed and Complicated

CISCO UCS

Simplified, Optimized and Automated

UCS Fabric Innovation

UCS Fabric Innovation Cadence

	2009	2011	2016	2018	2019	2020	2021
	1st with.... Compute Converged Fabric powered by Dynamic IOV	**1st with...** 40G Adapter Single Wire Mgmt Common Fabric for Blade and Rack	**1st with...** End to End 40G converged Multi host adapter	Transition to 25G Fabric		**1st with** Cloud operated compute Fabric for blades	
Fabric Interconnect	1st Gen	2nd Gen	3rd Gen	4th Gen			
	 FI 6120 / 6140	 6248 / 6296	 6332 / 6332-16UP  6324	 6454		 64108	
IFM/ IOM/ FEX	 IOM 2104	 FEX 2232  IOM 2204/ 2208	 FEX 2348  IOM 2304	 IOM 2408		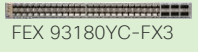 FEX 93180YC-FX3  IFM X-9108 25G	
VIC	 M81KR, P81E (10G)	 1200 Series (10/40G)	 1300 Series (10/40G)	 1400 / 14000 Series (10G/25G/40G/100G)		 14425 / 14825	
	Converged FC, Eth 10G Server/Uplink 8G FC uplink PCIe Gen1 128 PCIe device Hypervisor Bypass for ESX, KVM (VM-FEX)	10G Server/Uplink 8G FC uplink PCIe Gen2 256 PCIe device Flow Classification Single Wire Mgmt VM-FEX UsNIC, RSS,NetQueue	10/40G Server/Uplink 16G FC uplink UCS Mini (6324) PCIe Gen3 NVGRE/VxLAN VMQ, DPDK ROCEv1	40/100G uplink 10/25G Server 32G FC uplink 25G N9K FEX FC slow drain NVGRE/VXLAN/GENEVE NVMEoF (FC-NVMe, ROCEv2), VMMQ		X-Series Support 100G per VIC, 25G E2E 200G per X210c	

Currently Shipping

UCS 5th Generation Fabric

Introducing the 5th Generation UCS Fabric

Simplify

Cloud-operated infrastructure management
Reduces adapters and cables
Lowers server and network TCO
One VIC replaces multiple NICs, HBAs, and Mgmt adapters

Scalable and Resilient

Converged fabric for modern applications
Single-fabric for Mgmt, Data, Storage (FC, iSCSI, NFS/SMB, NVMeOF-ROCEv2, NVMEoF/FC, NVMeoF/TCP)
Flexible Blade/Rack connectivity at 10/25/40/100-Gbps

Future-ready

Industry-first End-to-End 100G fabric for blades
Enables High N/W workloads (BigData, SDS, HCI, CloudNative etc)
Single-flow 32G FC and 100G Ethernet
14.8 Tbps bandwidth per UCS domain
1.6 Tbps bandwidth per X9508 Chassis

6536 Fabric Interconnect

9108-100G IFM

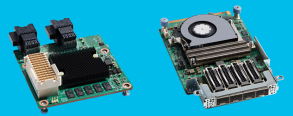

VIC 15000 Series

UCS 6536 Fabric interconnect

- 5th Generation UCS Fabric Interconnect (FI)
- 36x 100G Ethernet ports , 1 RU form-factor
  - 32x Ethernet ports (1/10/25/40/100 Gbps)
  - 4x Unified ports
    - 4x100G Ethernet (10/25/40/100) or
    - 16x 8/16/32G FC ports after breakout
- Support X9108-IFM-100G and X9108-IFM-25G
- Support IOM 2408 and FEX 93180YC-FX3
- Support UCS VIC 1400/14000 and 15000 Series
- Support M6 x210c, M5/M6 B- and C- Series
- Intersight Managed mode at FCS with 4.2(2)
- Support for UCSM, IOM 2304 (IMM/UCSM), VIC 1300 (UCSM) will be post 4.2(2) release

9108-100G IFM

- 2nd Generation UCS Intelligent Fabric Module (IFM)
- Supported with UCS 6536 Fabric Interconnect (FI) only
- 8 x 100G connectivity from IFM to FI (NIF Ports)
- 8 x 100G or 32 x 25G connectivity from IFM to VIC (HIF ports)
- Supports UCS VIC 14425/14825 and VIC 15231
- Connects X series compute node to external network.
- Hosts “Chassis Management Controller” (CMC) for chassis management
- Provide server CIMC connectivity via “Chassis Management Switch” (CMS)
- Enhanced Security – FPGA (Secure Boot), ACT2 (Anti-Counterfeit)
- Supported in Intersight Managed Mode (4.2.2)

IFM Comparison

	IFM-25G	IFM-100G
Fabric Interconnect	6454, 64108, 6536	6536
VIC	15231, 14425, 14825	15231, 14425, 14825
Network Interface (NIF) Ports	8 x 25G (port-channel)	8 x 100G (port-channel)
Host Interface (HIF) Ports	32 x 25G	8 x 100G or 32 x 25G
Oversubscription	4 : 1	1 : 1
Buffer	40 MB	40 MB
Latency	950ns	950ns

VIC 15000 Series

- 5th Gen VIC card for X- , B- *, C- Series
- Supports 10G/25G/40G/50G/100G/200G*
- CNA, Single Wire Mgmt
- Dynamic FC and Ethernet virtual interfaces
- x16 PCIe Gen 4
- NVMeoF: FC-NVMe, RoCEv2
- Overlays: NVGRE, VXLAN, Geneve
- RSS, NetQueue, VMQ, VMMQ, RSSv2*
- SR-IOV*, SIOV*, usNIC, DPDK
- PTPv2, L3ECN*, 16K Rx Ring Size
- 15000 series VIC's at FCS with 4.2(2)
  - VIC 15231 : 2x 100G MLOM for X210c
  - VIC 15428 : 4x 10/25/50G MLOM for M6 C-series

* Various VIC PIDs and features will be available post 4.2(2)

VIC 15231

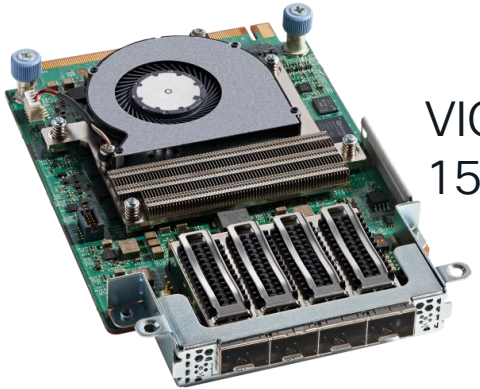

VIC 15428

VIC 15000 Series for X-Series and C-Series

	VIC 15231 (UCSX-ML-V5D200G)	VIC 15428 (UCSC-M-V5Q50G)
Speed	100G	10/25/50G
Max Ports	2	4
Form Factor	mLOM	mLOM
Server support	X210c M6	M6 C-Series
FI Series	6400/6536	6300/6400/6536
IOM / IFM / FEX	IFM-25G/ IFM-100G	93180YC-FX3
Chassis	X9508	-

Building Blocks of 5th Generation Fabric

Cisco Intersight
Policy-driven Infrastructure Provisioning

Cisco UCS Fabric Interconnect
100GE unified fabric switch
7.4 Tbps bandwidth

Cisco UCS Fabric Extender
IFM-100G, IFM-25G for X9508 Chassis
IOM-2408, FEX 93180YC-FX3

Cisco UCS I/O Adapters
VIC 1400/14000/15000 Series

Cisco UCS X-series Chassis
200G per x210c compute node
E2E 100G, 32G FC

Cisco UCS Blade and Rack Servers
Support for M5/M6 B-,C- series
E2E 25/40/100G, 32G FC

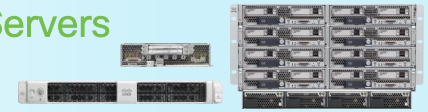

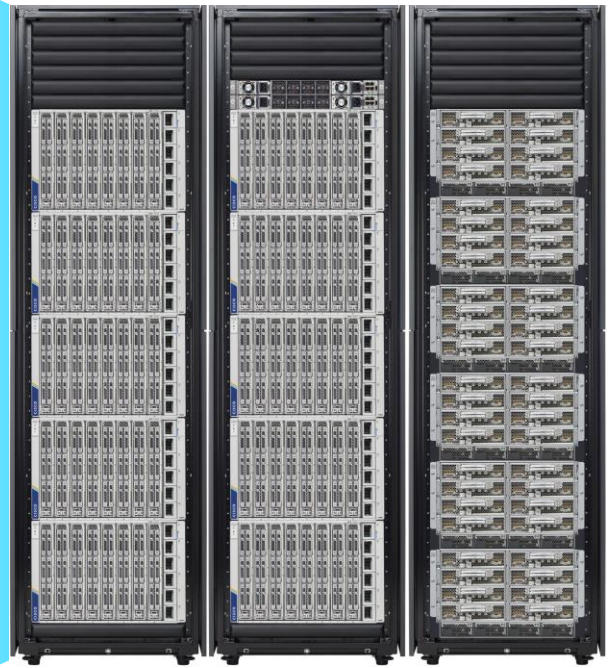

UCS 5th Generation Fabric Details

FI 6536 – Rear and front view

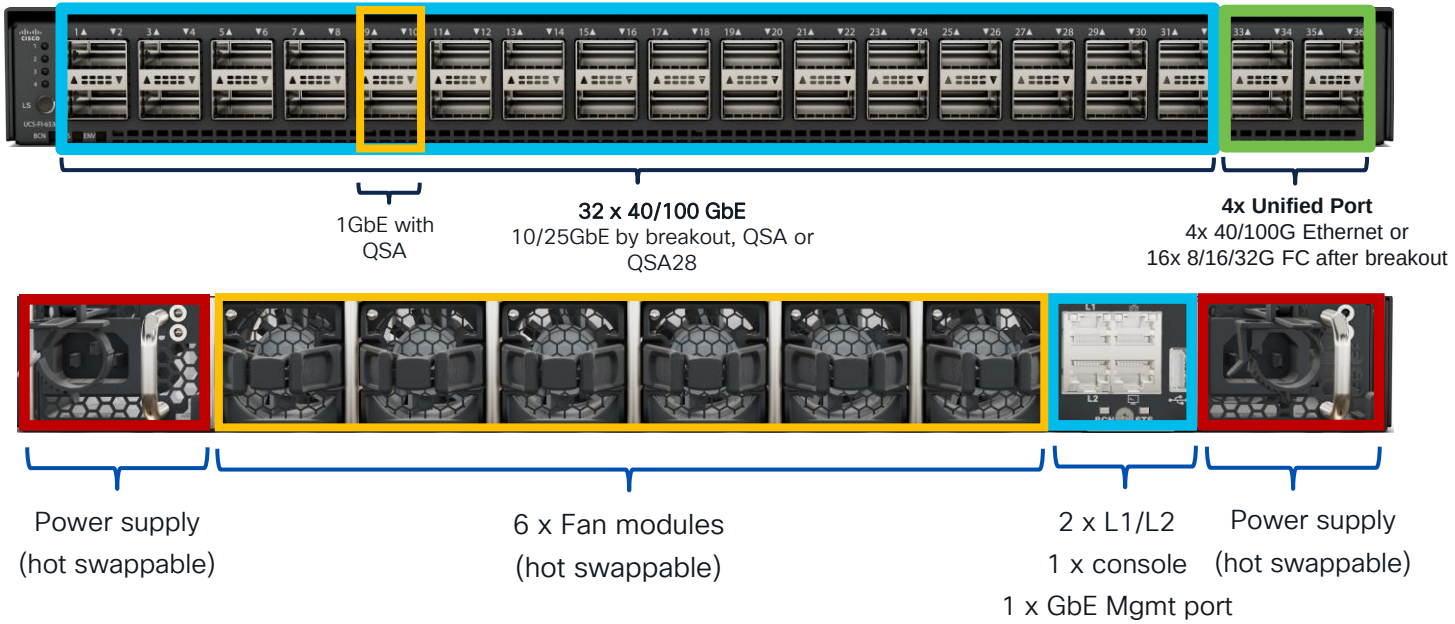

Fibre Channel Connectivity

FI 6536

Ports 33-36

128G FC QSFP
Cisco PID: DS-SFP-4X32G-SW

Multimode OM4 fiber, MTP/MPO female to 4x LC 8-fiber type-b breakout cable

8/16/32G FC SW SFP

SAN switch or Storage Array

UCS FI 6536 – Management Unit

- RJ-45 Management Port
- L1 and L2 Ports
- RS-232 Serial Port
- USB Port

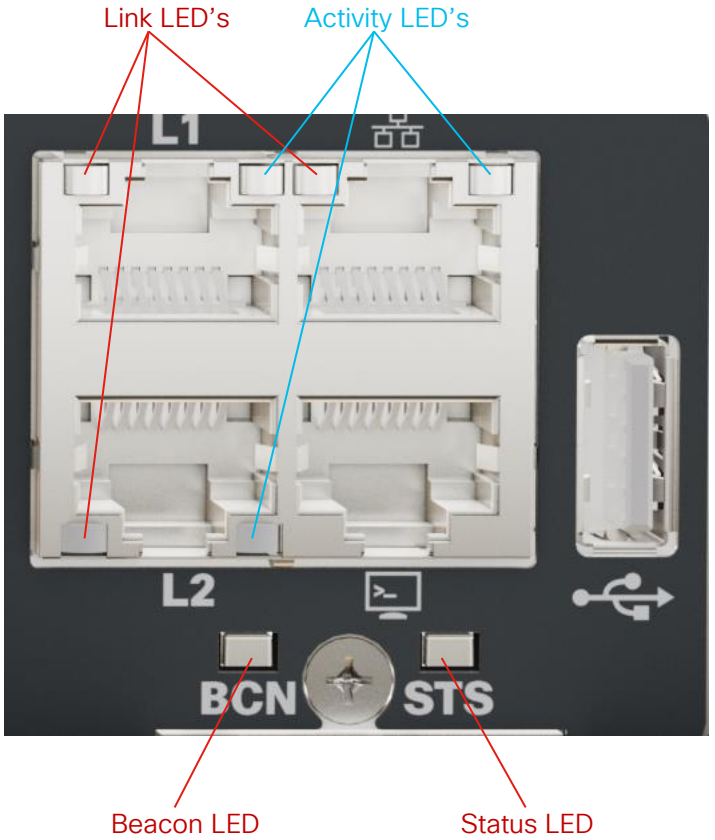

UCS FI 6536

Port no.	Interface type				Protocol support			Port role					
	1 Gigabit Ethernet QSA	10/25 Gigabit Ethernet Breakout, QSA or QSA28	40/100 Gigabit Ethernet	4 x 8/16/32-Gbps FC breakout	Ethernet	Fibre Channel	FCoE	Server 25/40/100G	Uplink: Ethernet 1/10/25/40/100G	Uplink: Fibre Channel 8/16/32G	Uplink: FCoE	Appliance port 10/25/40/100G	Storage port (Fibre Channel)
1 to 8		X	X		X		X	X	X		X	X	
9 to 10	X	X	X		X		X	X	X		X	X	
11 to 32		X	X		X		X	X	X		X	X	
33 to 36		X	X	X	X	X	X	X	X	X	X	X	X

FI 6536 Hardware Support

- X9508 Chassis
- X9108-IFM-25G, X9108-IFM-100G
- X210c M6 servers

- UCS 5108 rev 1 & 2 chassis
- IOM 2408
- B-Series M5 and M6 servers

- C-Series M5 and M6 servers
- VIC 1400/14000 , 15000 series
- Nexus 93180YC-FX3 in FEX mode
- Server support at 25/40G/100Gbps
- Supported only in Intersight Managed Mode at FCS.

FI 6536 Series Support Matrix

UCS Servers, VIC, FEX, IOM

X-Series and B-Series

Server	M5	M6
B200	x	x
B480	x	-
X210c	-	x

VIC	IOM 2408 (B-Series)	IFM 25G (X-Series)	IFM 100G (X-Series)
1440	x	-	-
1440+PE	x	-	-
1480	x	-	-
14425	-	x	x
14425 + 14825	-	x	x
15231	-	x	x

C-Series

Server	M5	M6
C220	x	x
C240	x	x
C225	-	x
C245	-	x
C480	x	-

VIC
1455
1457
1467
1477
1495
1497
15428

FEX Support
Nexus 93180YC-FX3 in FEX mode

Rack-server connectivity
Direct-connect rack-server supported at 25G/40G/100G
93180YC-FX3 FEX uplink connectivity at 100G
93180YC-FX3 FEX port server connectivity at 25G

Feature Overview

FI 6200 series, 6300 series, 6400 series, 6536

Feature	FI 6200 Series	FI 6300 Series	FI 6400 Series	FI 6536
End host mode	Supported	Supported	Supported	Supported
Ethernet switch mode	Supported	Supported	Supported	Supported
FC / FCoE NPV mode	Supported	Supported	Supported	Supported
FC / FCoE switch mode	Supported	Supported	Supported	Supported
NetFlow	Supported	Supported	Not supported *	Not supported *
PVLAN	Supported	Supported	Supported	Not Supported *
Port-security	Supported	Supported	Supported	Supported
IOM	2200 series	2200 / 2300 series	220x / 2408	2408 / 2304*
IFM	-	-	IFM-25G	IFM-25G/ IFM-100G
Servers	Upto M5 B-/ C-/ S- series	Upto M6 B-/ C- /S- series	M4/M5/M6 B-/ C-/ S-/ X- series	M5/M6 B-/ C-/ X- series
VIC	1200/ 1300/ 1400	1200/ 1300/ 1400/15000	1300/ 1400/ 15000	1300* / 1400/ 15000
Management	UCSM	UCSM	UCSM/ IMM	UCSM*/ IMM

* Support post 4.2(2)

S-series is only supported in M4/M5 and X-series only in M6

FI Configuration Limits

Feature	FI 6100	FI 6200	FI 6300	FI 6400	FI 6536
Unicast MAC address per FI	13,800	20,000	32,000	32,000	32,000
Multicast MAC address per FI	6,000	7,000	7,000	7,000	7,000
Active VLAN per UCS domain	2,000	2,000	3,000	3,000	3,000
STP logical interfaces / VLAN port count	14,000	64,000	64,000	108,000	64,000*
IGMP groups	1,000	4,000	4,000	16,000	16,000
Virtual Interfaces (VIFs)	2,000	2,750	2,750	1,600	1,600
Host Bus Adapter (vHBAs)	320	320	320	320	320
Chassis per UCS domain	20	20	20	20	20
IP storage appliance per FI	4	16	16	16	16
VSANs	32	32	32	32	32

* 108K support post 4.2(2)

Fabric Interconnect Licenses

Fabric Interconnect	Licensed ports on base PID	License PIDs	Description
6248	12 x 10GbE	• UCS-LIC-10GE • UCS-L-6200-10G-C	• B-Series, C-Series, S-Series • C-Series only
6296	18 x 10GbE	• UCS-LIC-10GE • UCS-L-6200-10G-C	• B-Series, C-Series, S-Series • C-Series only
6332	8 x 40GbE	• UCS-LIC-6300-40G • UCS-LIC-6300-40GC	• B-Series, C-Series, S-Series • C-Series only
6332-16UP	4 x 40GbE + 8 x UP	• UCS-LIC-6300-40G • UCS-LIC-6300-40GC • UCS-LIC-6300-10G • UCS-LIC-6300-10GC	• B-Series, C-Series, S-Series • C-Series only • B-Series, C-Series, S-Series • C-Series only
6454	18 x 10/25GbE + 2 x 40/100GbE	• UCS-LIC-6400-25G • UCS-LIC-6400-25GC • UCS-LIC-6400-100G	• B-Series, C-Series, S-Series • C-Series only • Nexus, MDS , C-Series
64108	38 x 10/25GbE + 4 x 40/100GbE	• UCS-LIC-6400-25G • UCS-LIC-6400-25GC • UCS-LIC-6400-100G	• B-Series, C-Series, S-Series • C-Series only • Nexus, MDS , C-Series
6536	All ports enabled via a term-based subscription license	• DC-UFAB-SW-A	• Smart License for X-Series, B-Series, C-Series connectivity • Term based SW license for 36-60 month • Not a port-based license

Qualified Cabling

- Check the 6536 series data-sheet and the tmgmatrix links to get the supported cables/transceivers
- Some specific cable/transceiver PIDs have dependency on the IMM release, hence confirm against data-sheet and tmgmatrix.
- Server-port breakout supported only for 25G speed.
- Uplink-port breakout support for 1/10/25G Ethernet and 8/16/32G FC.
- At 25 and 100G speed the FI 6536 ports operates at RS-FEC by default for uplink as well as on server-port.
- FEC configuration change is possible only for uplink ports.
- https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs6536-fabric-interconnect-ds.html
- https://tmgmatrix.cisco.com

Provision with Intersight

Intersight Managed Mode: Configuration Model

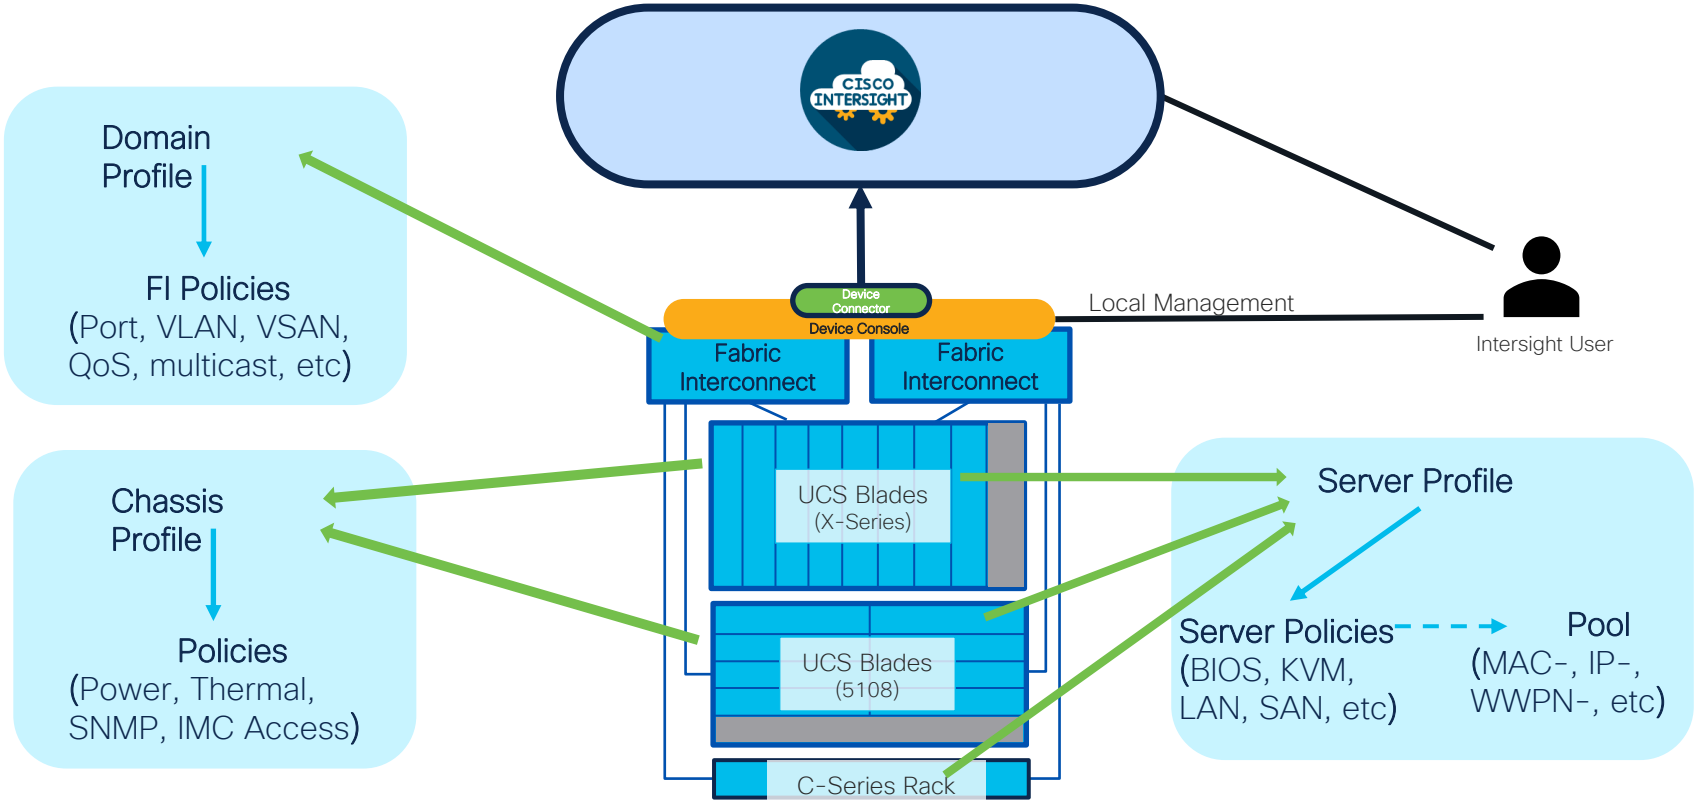

Intersight Domain profile

Minimum policies for FI Domain-profile

- Port policy
- VLAN Policy
- Multicast policy
- QoS Policy

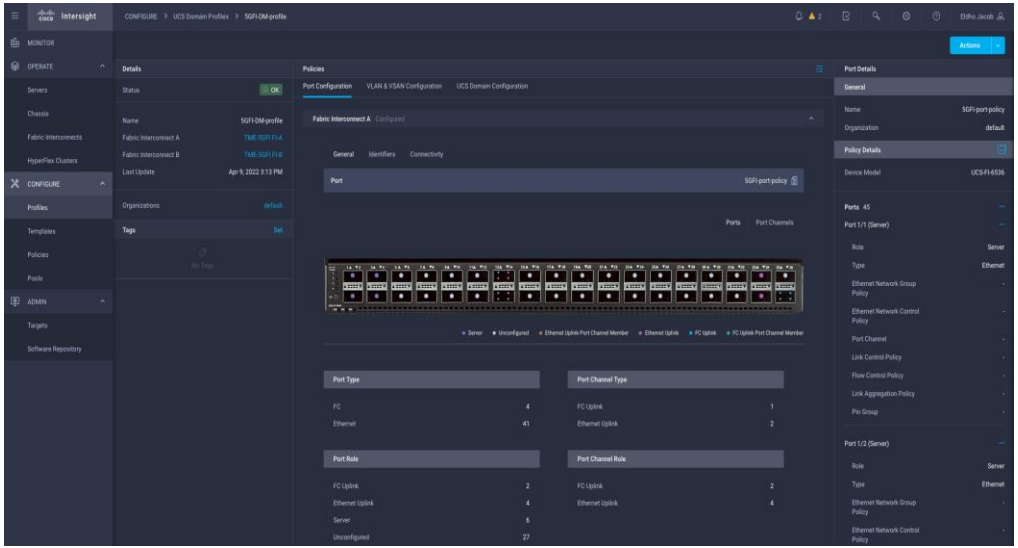

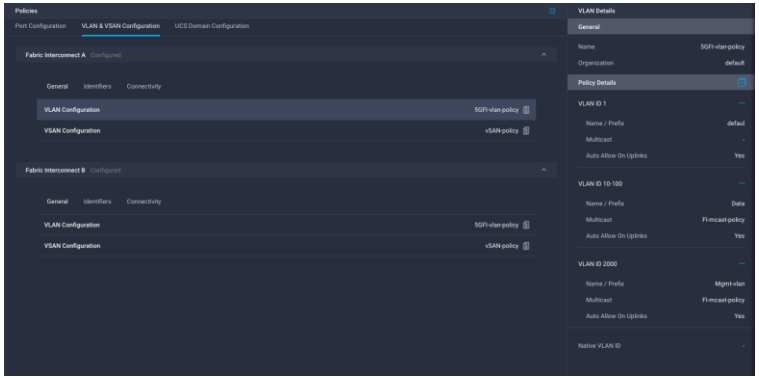

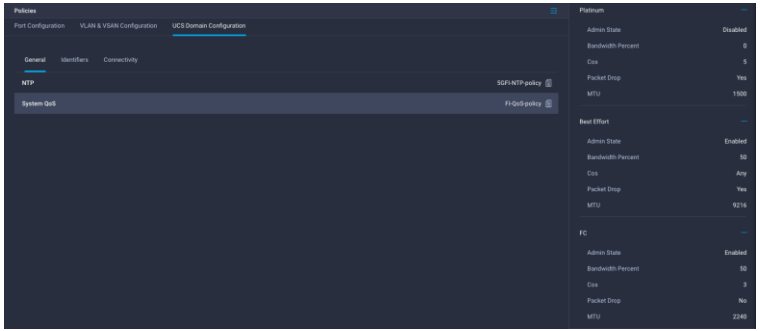

Ethernet Configuration

- Supports ethernet end-host and switching mode at FCS
- Ethernet ports : 1 to 32
  - Support speeds of 10/25/40/100G
  - 10/25G support via breakout or QSA
- Unified Ports : 33 to 36
  - Ethernet speeds of 10/25/40/100G
  - FC speed of 8/16/32G
- Dynamic ethernet breakout for 10G/25G
- Server ports at 25G/40G/100G
  - FEC is auto-determined for 25G/100G
- Ethernet uplink ports at 1/10/25/40/100G
  - FEC configuration possible for 25G/100G
- Appliance ports at 10/25/40/100G
- Monitor ports at 10/25/40/100G post 4.2(2)

Dynamic Ethernet Breakout

- Breakout allowed for all 36 ports of FI-6536
  - 4x 10/25G Ethernet appliance or uplink ports
  - 4x 25G Server ports
- Total of 144 ports after breakout
  - Max server ports limited to 128
- No reboot required with FI 6536 for ethernet breakout
- Breakout port members can connect multiple servers.

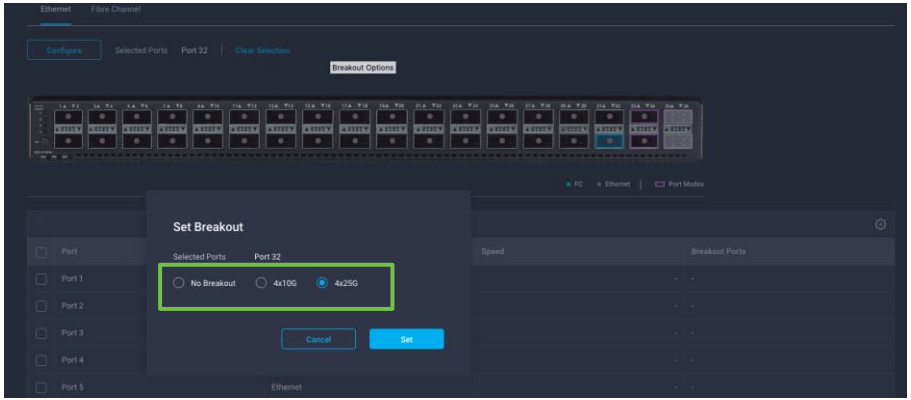

FC Configuration

- Supports FC end-host and FC switch mode.
- Configured via “Switch Control” policy, default without policy is FC end-host mode.
- Unified Ports : 33 to 36
  - FC support on unified ports using 128G FC QSFP and breakout cables
  - FC 33/1/1-4 to 36/1/1-4 after breakout
- FI reload required when changing from Eth <-> FC mode for a port and vice versa.
- FC speeds supported:
  - 8G, 16G and 32G speeds
  - 8G speed require IDLE fill pattern.
- Auto speed not supported for FC ports.
- FI reload not required when changing the port speed.

FC Port Configurations

- Slider bar for FC port selection
  - Contiguous ports from right to left
  - Select port 36 thru 33 in increments of 1
  - Require system Reboot on deploy

- Unified ports selected
  - Configure unique FC break-out mode of “4x 8/16/32G” per unified port
  - All breakout port members operate at same FC speed
  - Each breakout member operates as individual FC port like in previous FI generations

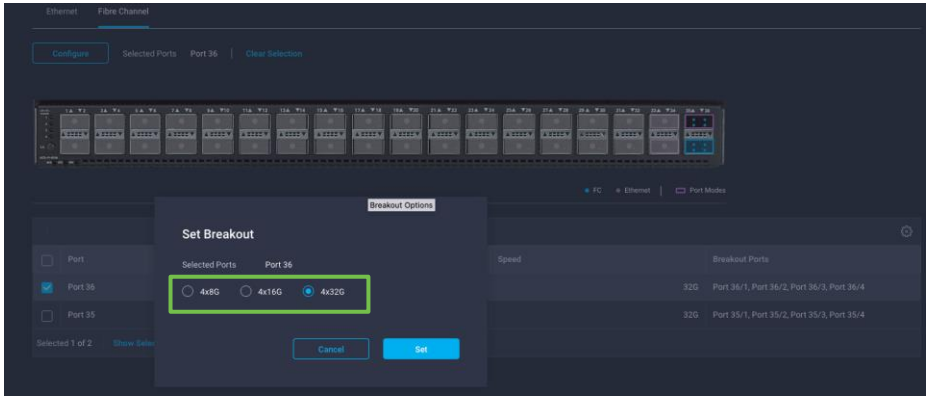

QoS Configuration

FI 6536 supports 6 system classes

- Enabled via FI QoS policy
- 2 default classes : Best-effort and FC
- Recommendation is to use the 2 default-classes
- Additional 4 user selectable classes.
- max of 2 no-drop class including fibre-channel

QoS at vNIC

- Enabled via “Ethernet QoS policy” attached to vNIC
- Sets CoS, MTU and Rate per vNIC

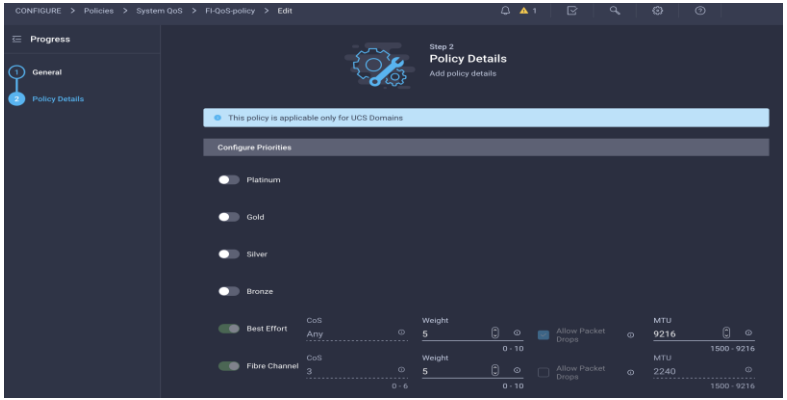

5th Gen Fabric Connectivity

IFM-25G/100G to 6536 Fabric Interconnect Support

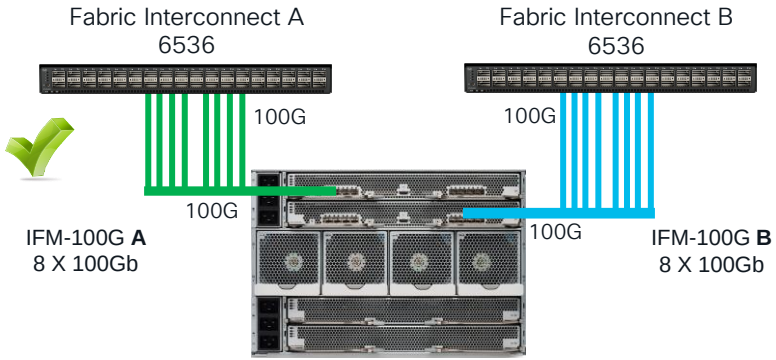

X9108-IFM-100G to 6536 FIs :
Supported

- 1600G per X9508 chassis
- 100G E2E single-flow per x210c
- 32G E2E FC I/O
- 200G per x210c with 1:1 oversubscription

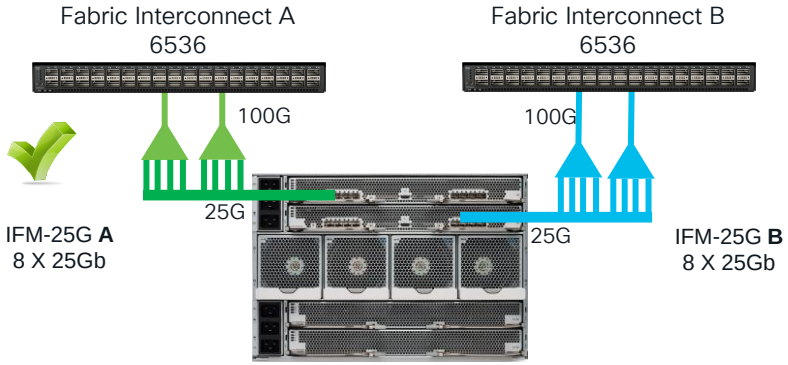

X9108-IFM-25G to 6536 FIs :
Supported

- 400G per X9508 chassis
- 25G E2E single-flow per x210c
- 200G per x210c with 4:1 oversubscription

Throughput per UCS x210c compute node

x210c Compute Node	A FI-6536 + X9108-IFM-100G	B FI-6536 + X9108-IFM-25G	C FI-6536 + X9108-IFM-25G/100G	D FI-6536 + X9108-IFM-25G/100G	
x210c configuration	VIC 15231	VIC 15231	VIC 14425	VIC 14425+ VIC 14825	
Throughput per node	200G (100G per IFM)	100G (50G per IFM)	100G (50G per IFM)	200G (100G per IFM)	
vNICs needed for max BW	2	2	2	4	
KR connectivity per IFM	1x 100GKR	2x 25GKR	2x 25GKR	4x 25GKR	
Single vNIC throughput on VIC	100G	50G (2x25G KR)	50G (2x25G KR)	50G (2x25G KR)	50G (2x25G KR)
Max Single flow BW per vNIC	100G	25G	25G	25G	25G
Single vHBA throughput on VIC	100G	50G	50G	50G	50G

2 x 100Gb – 100G End-to-end

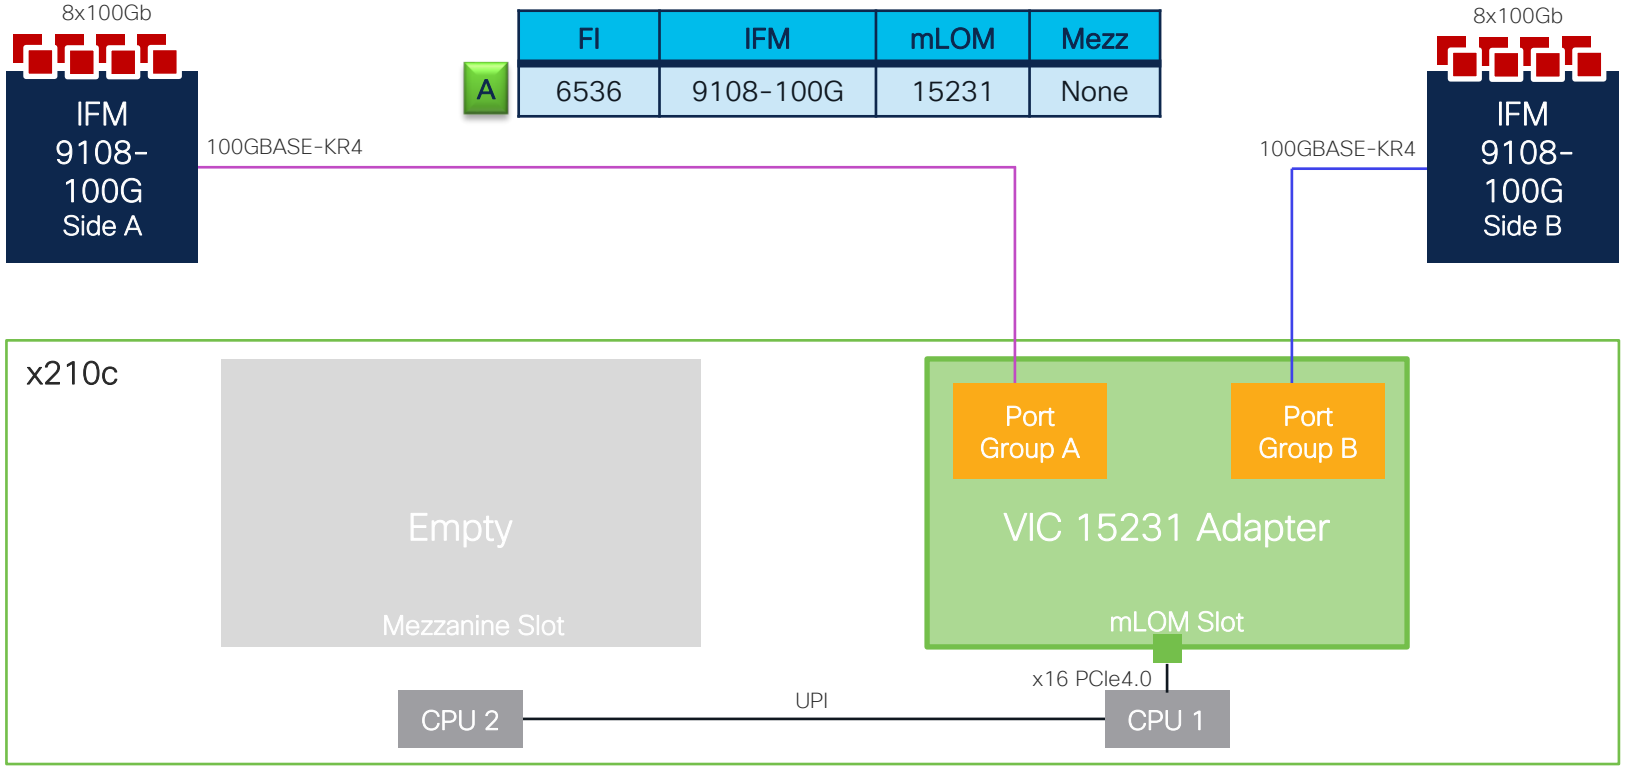

All vNICs and vHBAs are 100Gb

100G vNIC/vHBA in X210c with VIC 15231

100G vNIC/vHBA as seen on RHEL and ESXi host

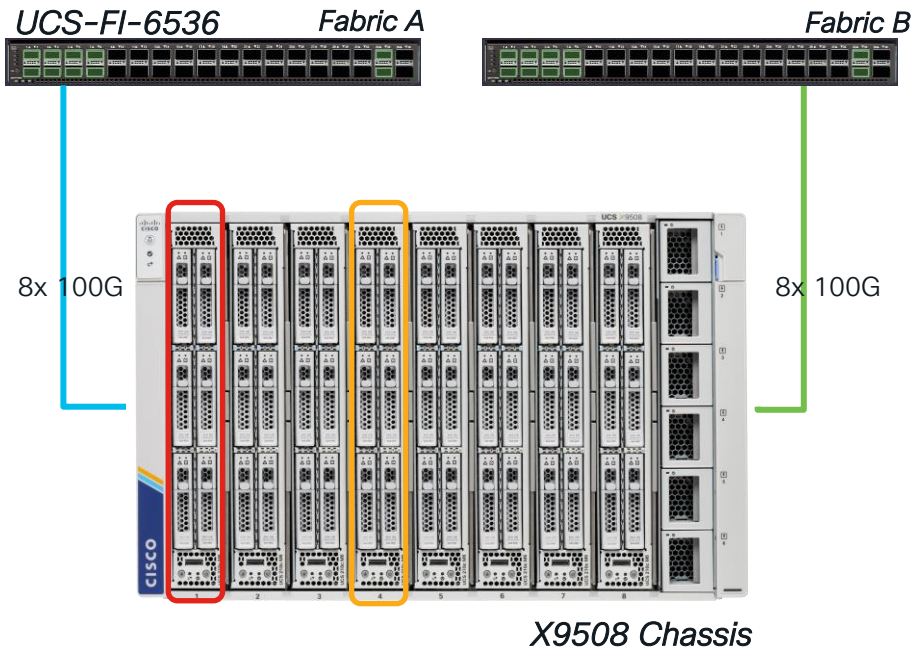


```
[root@ora5g-node1 ~]#
[root@ora5g-node1 ~]# lshw -short | grep -i eno7
/0/101/2/0/0/0/0.2        eno7        network        VIC Ethernet NIC
[root@ora5g-node1 ~]# ethtool eno7 | grep -i speed
        Speed: 100000Mb/s                                   100G vNIC
[root@ora5g-node1 ~]# ifconfig eno7
eno7: flags=4163<UP,BROADCAST,RUNNING,MULTICAST>  mtu 9000
        inet 10.10.21.61  netmask 255.255.255.0  broadcast 10.10.21.255
        inet6 fe80::225:b5ff:fe01:2710  prefixlen 64  scopeid 0x20<link>
        ether 00:25:b5:01:27:10  txqueuelen 1000  (Ethernet)
        RX packets 83  bytes 12616 (12.3 KiB)
        RX errors 0  dropped 83  overruns 0  frame 0
        TX packets 84  bytes 10010 (9.7 KiB)
        TX errors 0  dropped 0 overruns 0  carrier 0  collisions 0

[root@ora5g-node1 ~]# cat /sys/class/fc_host/host3/speed
100 Gbit                                                    100G vHBA
[root@ora5g-node1 ~]#
```

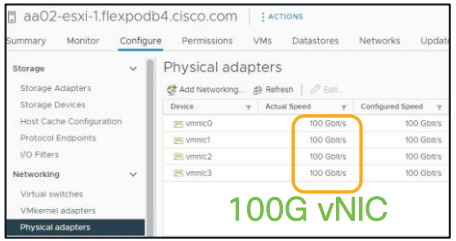


```
[root@aa02-esxi-1:~] esxcli storage san fc list
   Adapter: vmhba0
   Port ID: 5600A6
   Node Name: 20:00:00:25:b5:17:00:0a
   Port Name: 20:00:00:25:b5:17:0a:13
   Speed: 100 Gbps                          100G vHBA
   Port Type: NPort
   Port State: ONLINE
   Model Description: none
   Hardware Version: none
   OptionROM Version: 1
   Firmware Version: 1
   Driver Name: nfnic
   DriverVersion: 1
```

100G Single Flow : VIC 15231 performance

Single Flow 100G:

Blade6 vnic1 (client) --> Blade1 vnic1 (Server)

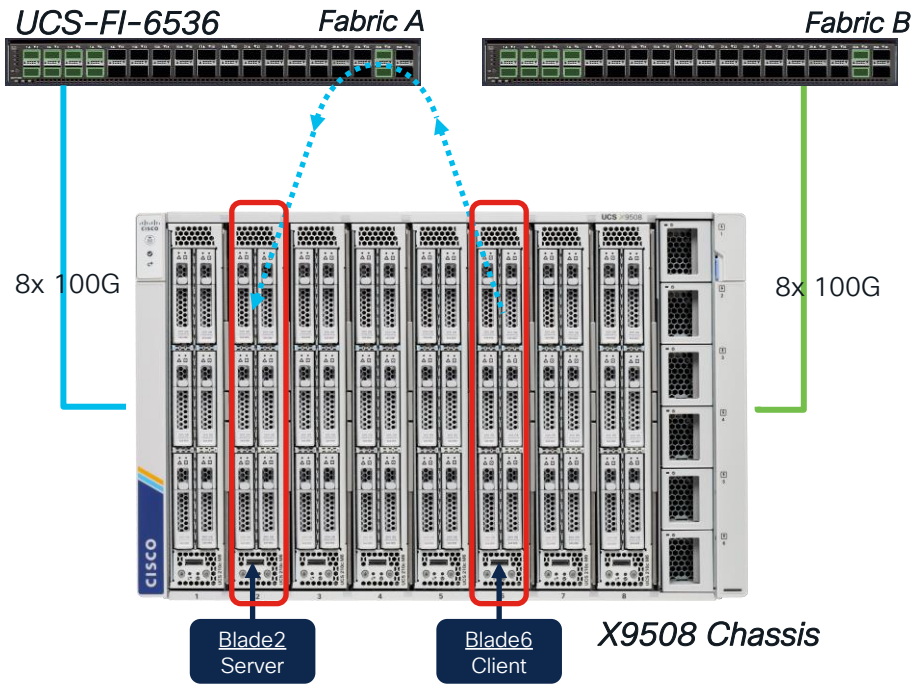

iPerf results on Fabric

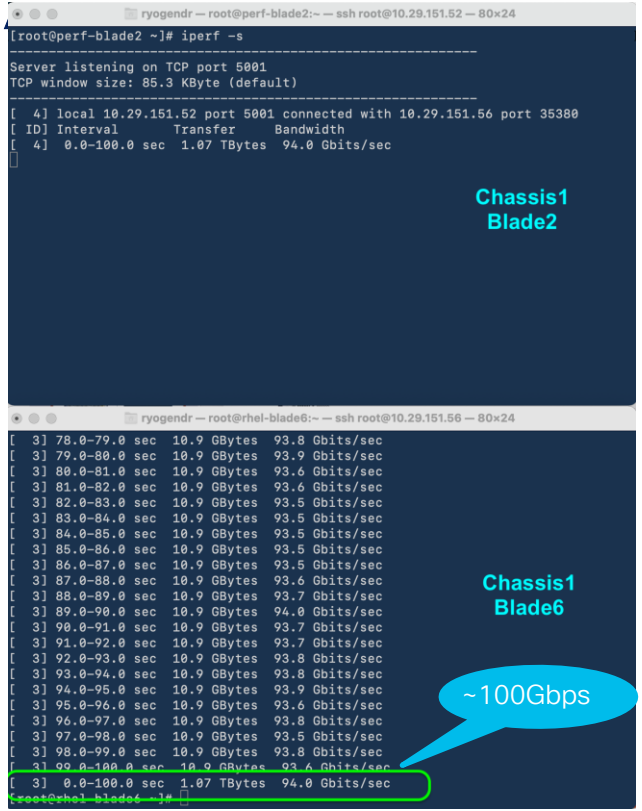

200G per X210c: VIC 15231 performance

200G per X210c

Fabric A : Blade6 vnic1 (client) -> Blade2 vnic1 (Server)

Fabric B : Blade6 vnic2 (client) -> Blade2 vnic2 (Server)

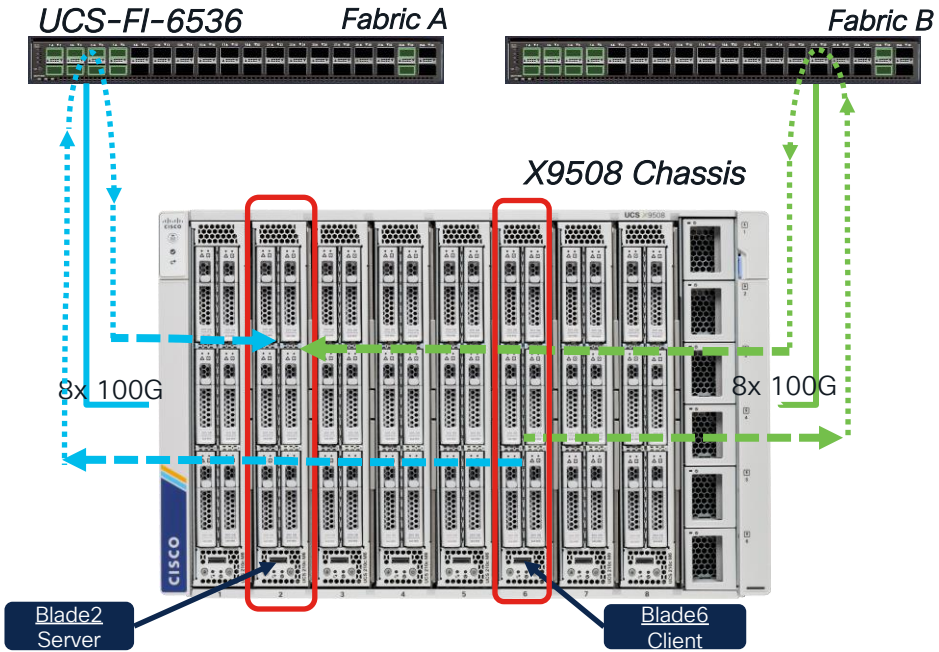

iPerf results

```
[root@blade2 ~]# taskset -c 16-24 iperf -s -B 10.0.11.41
------------------------------------------------------------
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
------------------------------------------------------------
[  1] local 10.0.11.41 port 5001 connected with 10.0.11.42 port 55864
[  2] local 10.0.11.41 port 5001 connected with 10.0.11.42 port 55870
[  3] local 10.0.11.41 port 5001 connected with 10.0.11.42 port 55868
[  4] local 10.0.11.41 port 5001 connected with 10.0.11.42 port 55866
```

Disable this terminal from "MultiExec" mode

```
[  4] 21.00-24.00 sec  5.75 GBytes  16.5 Gbits/sec
[  3] 21.00-24.00 sec  11.5 GBytes  33.0 Gbits/sec
[  2] 21.00-24.00 sec  11.5 GBytes  33.0 Gbits/sec
[  1] 21.00-24.00 sec  5.76 GBytes  16.5 Gbits/sec
[SUM] 21.00-24.00 sec  34.5 GBytes  98.9 Gbits/sec
[  4] 24.00-27.00 sec  5.76 GBytes  16.5 Gbits/sec
[  3] 24.00-27.00 sec  11.5 GBytes  33.0 Gbits/sec
[  2] 24.00-27.00 sec  11.5 GBytes  33.0 Gbits/sec
[  1] 24.00-27.00 sec  5.76 GBytes  16.5 Gbits/sec
[SUM] 24.00-27.00 sec  34.5 GBytes  98.9 Gbits/sec
```

~100Gbps Fabric A

Disable this terminal from "MultiExec" mode

```
[root@blade2 ~]# taskset -c 32-40 iperf -s -B 10.0.21.41
------------------------------------------------------------
Server listening on TCP port 5001
TCP window size: 85.3 KByte (default)
------------------------------------------------------------
[  1] local 10.0.21.41 port 5001 connected with 10.0.21.42 port 36552
[  2] local 10.0.21.41 port 5001 connected with 10.0.21.42 port 36548
[  3] local 10.0.21.41 port 5001 connected with 10.0.21.42 port 36550
[  4] local 10.0.21.41 port 5001 connected with 10.0.21.42 port 36558
```

Disable this terminal from "MultiExec" mode

```
[  2] 6.00-9.00 sec  7.64 GBytes  21.9 Gbits/sec
[  1] 6.00-9.00 sec  9.95 GBytes  28.5 Gbits/sec
[  4] 6.00-9.00 sec  7.65 GBytes  21.9 Gbits/sec
[  3] 6.00-9.00 sec  8.49 GBytes  24.3 Gbits/sec
[SUM] 6.00-9.00 sec  33.7 GBytes  96.6 Gbits/sec
[  2] 9.00-12.00 sec  8.70 GBytes  24.9 Gbits/sec
[  1] 9.00-12.00 sec  8.72 GBytes  25.0 Gbits/sec
[  4] 9.00-12.00 sec  8.71 GBytes  24.9 Gbits/sec
[  3] 9.00-12.00 sec  8.39 GBytes  24.0 Gbits/sec
[SUM] 9.00-12.00 sec  34.5 GBytes  98.9 Gbits/sec
```

~100Gbps Fabric B

End to End 32G : VIC 15231 performance

32G FC connectivity to a Storage Target:

Chassis-Blade2 vHBA1 <-- HBA1 Storage (100% read)
Chassis-Blade2 vHBA1 <-- HBA5 Storage (100% read)

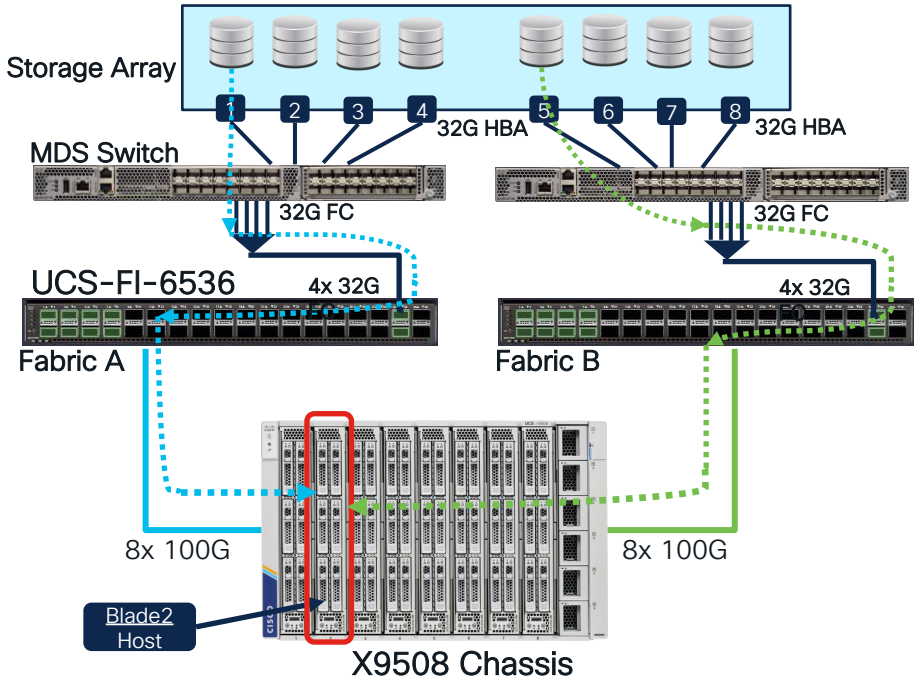

FIO results per x210c

```
[root@ora5g-node1 ~]# cat /sys/class/fc_host/host3/speed
100 Gbit
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol1:/dev/mapper/5g_vol2 --direct=1 --rw=r
w --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=8 --numjobs=8 --runtime=120 --group_re
porting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB,                    512KiB-512KiB, ioengi
ne=libaio, iodepth=8
...
fio-3.19
Starting 8 processes
Jobs: 8 (f=16): [R(8)][11.6%][r=3141MiB/s]r=6281 IOPS][eta 01m:47s]
```

~32G FC Fabric A

```
[root@ora5g-node1 ~]# cat /sys/class/fc_host/host7/speed
100 Gbit
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol9:/dev/mapper/5g_vol10 --direct=1 --rw=
rw --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=8 --numjobs=8 --runtime=120 --group_r
eporting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB                    T) 512KiB-512KiB, ioengi
ne=libaio, iodepth=8
...
fio-3.19
Starting 8 processes
Jobs: 8 (f=16): [R(8)][12.4%][r=3142MiB/s]r=6284 IOPS][eta 01m:46s]
```

~32G FC Fabric B

Result from Storage Array

```
FlexPod-A800::>
FlexPod-A800::> qos statistics performance show
Policy Group             IOPS      Throughput     Latency Is Adaptive? Is Shared?
-------------------- -------- ---------------- ---------- ------------ ----------
-total-                108565     6342.20MB/s     4.68ms - -
User-Best-Effort       101469     6342.12MB/s     5.01
_System-Work             7096       74.24KB/s
-total-                101496     6341.45MB/s
User-Best-Effort       101455     6341.44MB/s     5.01
_System-Work               41       15.06KB/s
-total-                101503     6342.33MB/s     5.00ms - -
User-Best-Effort       101465     6342.31MB/s     5.01ms false true
_System-Work               38       19.36KB/s   105.00us false true
```

Total 64G across Fabric A & B

100G FC per Fabric : VIC 15231 performance

100G vHBA performance in a x210c:

Chassis Blade2 <-- 4x 32G FC storage target (100% read)

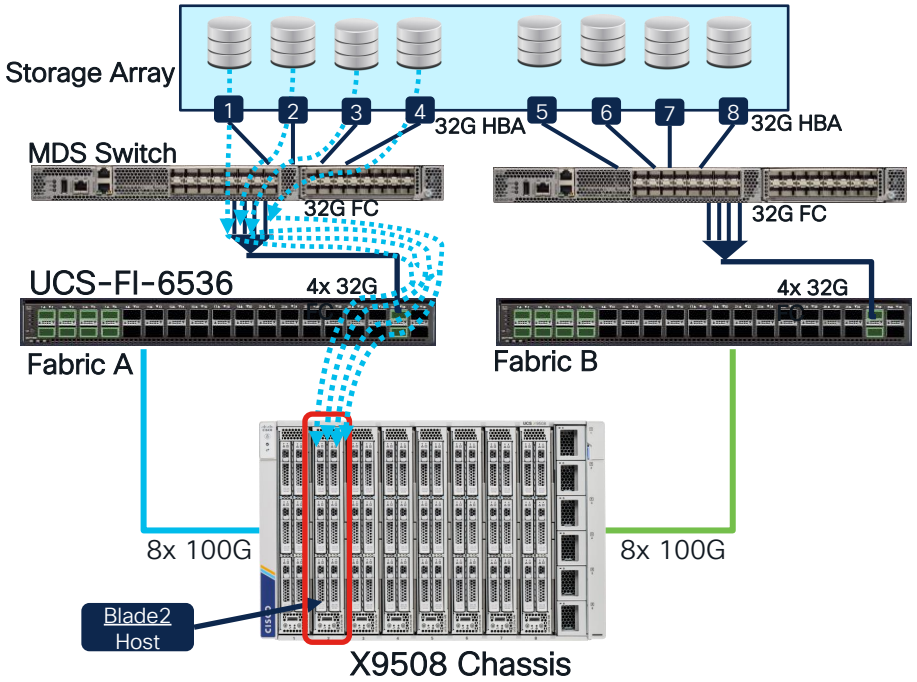

FIO results per x210c

```
[root@ora5g-node1 ~]# fio --filename=/dev/mapper/5g_vol1:/dev/mapper/5g_vol2:/dev/mapper/5g_vol3:/dev/
mapper/5g_vol4:/dev/mapper/5g_vol5:/dev/mapper/5g_vol6:/dev/mapper/5g_vol7:/dev/mapper/5g_vol8 --direc
t=1 --rw=rw --ioengine=libaio --bs=512k --rwmixread=100 --iodepth=32 --numjobs=64 --runtime=30 --group
_reporting --name=seqreadwrite
seqreadwrite: (g=0): rw=rw, bs=(R) 512KiB-512KiB, (W) 512KiB-512KiB, (T) 512KiB-512KiB, ioengine=libai
o, iodepth=32
...
fio-3.19
Starting 64 processes
Jobs: 64 (f=512): [R(64)][64.5%][r=11.1GiB/s]r=22.7k IOPS][eta 00m:11s]
```

```
FlexPod-A800::> q s p s
  (qos statistics performance show)
Policy Group             IOPS      Throughput    Latency Is Adaptive? Is Shared?
------------------- -------- ---------------- ---------- ------------ ----------
-total-                183591    11479.40MB/s       4.
User-Best-Effort       183529    11479.38MB/s       4.
_System-Work               62       22.08KB/s          0ms
-total-                183511    11474.38MB/s       4.
User-Best-Effort       183489    11474.38MB/s          0ms false true
_System-Work               22          0KB/s     45.00us false true
-total-                183569    11478.01MB/s       5.09ms - -
User-Best-Effort       183529    11478.00MB/s       5.09ms false true
_System-Work               40       1.34KB/s     200.00us false true
```

92G FC (~100G) per Fabric

100G NFS: VIC 15231 performance

100G NFS connectivity to a Storage Filer:

Chassis-Blade2 vNIC1 <-- 100G NIC on Storage (100% read)

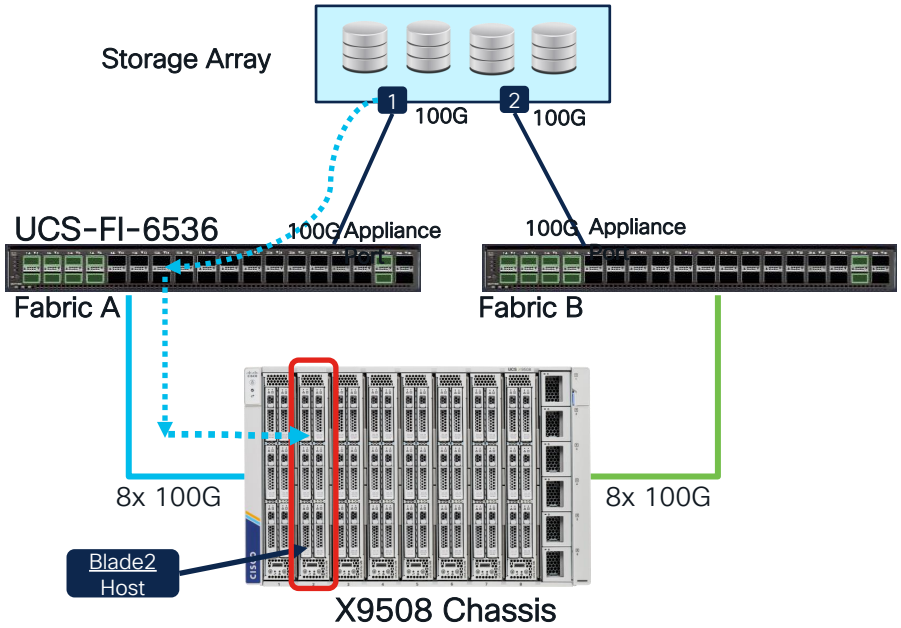

NFS results per x210c

```
[root@ora5g-node1 ~]# fio --name=test1 --filename=/fiovol1/data1:/fiovol1/data2:/fiovol1/data3:/fiovol1/data4:/fiovol2
/fiovol2/data22:/fiovol2/data23:/fiovol2/data24:/fiovol3/data31:/fiovol3/data32:/fiovol3/data33:/fiovol3/data34:/fiovo
l:/fiovol4/data42:/fiovol4/data43:/fiovol4/data44 --rw=read --direct=1 --ioengine=libaio --bs=512k --numjobs=16 --iode
-runtime=300 --time_based --group_reporting
test1: (g=0): rw=read, bs=(R) 512KiB-512KiB, (W) 512KiB-512KiB, (T) 512KiB-512KiB, ioengine=libaio, iodepth=64
...
fio-3.19
Starting 16 processes
Jobs: 16 (f=256): [R(16)][100.0%][r=11.3GiB/s][r=23.2k IOPS][eta 00m:00s]
test1: (groupid=0, jobs=16): err= 0: pid=82192: Fri Apr 22 13:53:52 2022
  read: IOPS=23.3k, BW=11.4GiB/s (12.2GB/s)(3408GiB/300028msec)
    slat (usec): min=26, max=5329, avg=56.33, stdev=25.13
    clat (usec): min=1957, max=531941, avg=43939.16, stdev=9451.20
     lat (usec): min=1997, max=531987, avg=43995.59, stdev=9450.83
    clat percentiles (msec):
     |  1.00th=[   25],  5.00th=[   28], 10.00th=[   31], 20.00th=[   38],
     | 30.00th=[   41], 40.00th=[   44], 50.00th=[   46], 60.00th=[   47],
     | 70.00th=[   49], 80.00th=[   51], 90.00th=[   53], 95.00th=[   56],
     | 99.00th=[   62], 99.50th=[   66], 99.90th=[  103], 99.95th=[  131],
     | 99.99th=[  205]
   bw (  MiB/s): min= 5977, max=12232, per=100.00%, avg=11646.89, stdev=19.82, samples=9584
   iops        : min=11947, max=24465, avg=23293.77, stdev=39.65, samples=9584
  lat (msec)   : 2=0.01%, 4=0.01%, 10=0.01%, 20=0.06%, 50=78.43%
  lat (msec)   : 100=21.39%, 250=0.10%, 500=0.01%, 750=0.01%
  cpu          : usr=0.22%, sys=8.45%, ctx=6244500, majf=0, m
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%,
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64
     issued rwts: total=6980099,0,0,0 short=0,0,0,0 dropped=0
     latency   : target=0, window=0, percentile=100.00%, d
Run status group 0 (all jobs):
   READ: bw=11.4GiB/s (12.2GB/s), 11.4GiB/s-11.4GiB/s (12.2GB/s-12.2GB/s), io=3408GiB (3660GB), run=300028-300028msec
```

~100G NFS

```
FlexPod-A800: cluster.cluster: 4/22/2022 13:50:07
 cpu  cpu    total                    fcache   total    total data     data     data cluster  cluster  cluster    disk    disk    pkts    pkts
 avg busy      ops  nfs-ops cifs-ops     ops     recv     sent busy     recv     sent    busy     recv     sent    read   write    recv    sent
---- ---- -------- -------- -------- -------- -------- -------- ---- -------- -------- ------- -------- -------- ------- ------- ------- -------
 18%  43%   187117   187117        0        0   38.1MB   11.4GB  97%   37.7MB   11.4GB      0%    338KB    336KB   478MB   1.71MB  230535 1371110
 18%  42%   186775   186775        0        0   38.4MB   11.5GB  98%   38.1MB   11.5GB      0%    285KB    286KB   481MB   3.91MB  236859 1375341
 18%  43%   187603   187603        0        0   39.0MB   11.5GB  98%   38.7MB   11.5GB      0%    286KB    285KB   480MB   23.9KB  245848 1376171
 18%  43%   186820   186820        0        0   38.0MB   11.4GB  98%   37.8MB   11.4GB      0%    266KB    266KB   609MB   1.16MB  231085 1371855
 18%  43%   186775   186775        0        0   38.5MB   11.5GB  98%   38.2MB   11.5GB      0%    301KB    301KB   643MB   17.9KB  237274 1374307
 18%  43%   186634   186634        0        0   38.4MB   11.4GB  97%   38.1MB   11.4GB      0%    251KB    251KB   613MB   1.46MB  239374 1370362
 18%  43%   185568   185568        0        0   37.7MB   11.4GB  97%   37.4MB   11.4GB      0%    278KB    278KB   616MB   3.91MB  226770 1369212
 18%  43%   187085   187085        0        0   38.3MB   11.4GB  97%   38.0MB   11.4GB      0%    307KB    307KB   582MB   23.9KB  235274 1365179
 18%  43%   186116   186116        0        0   38.4MB   11.4GB  98%   38.2MB   11.4GB      0%    269KB    269KB   577MB   1.17MB  237298 1369754
 18%  43%   186856   186856        0        0   38.1MB   11.4GB  98%   37.8MB   11.4GB      0%    270KB    270KB   575MB   29.9KB  234293 1372264
 18%  43%   186380   186380        0        0   38.5MB   11.4GB  97%   38.2MB   11.4GB      0%    281KB    281KB   590MB   646KB   238783 1371946
 18%  43%   186480   186480        0        0   38.6MB   11.4GB  97%   38.3MB   11.4GB      0%    309KB    309KB   577MB   2.25MB  241873 1366584
 18%  43%   185636   185636        0        0   38.2MB   11.4GB  97%   37.9MB   11.4GB      0%    284KB    282KB   568MB   23.9KB  235542 1368954
 18%  43%   185936   185936        0        0   37.6MB   11.3GB  97%   37.3MB   11.3GB      0%    280KB    282KB   594MB   1.18MB  225957 1360859
 18%  43%   187402   187402        0        0   38.3MB   11.5GB  98%   38.0MB   11.5GB      0%    268KB    268KB   587MB   29.9KB  234891 1375560
```

x210c, IFM 9108-25G and VIC 15231

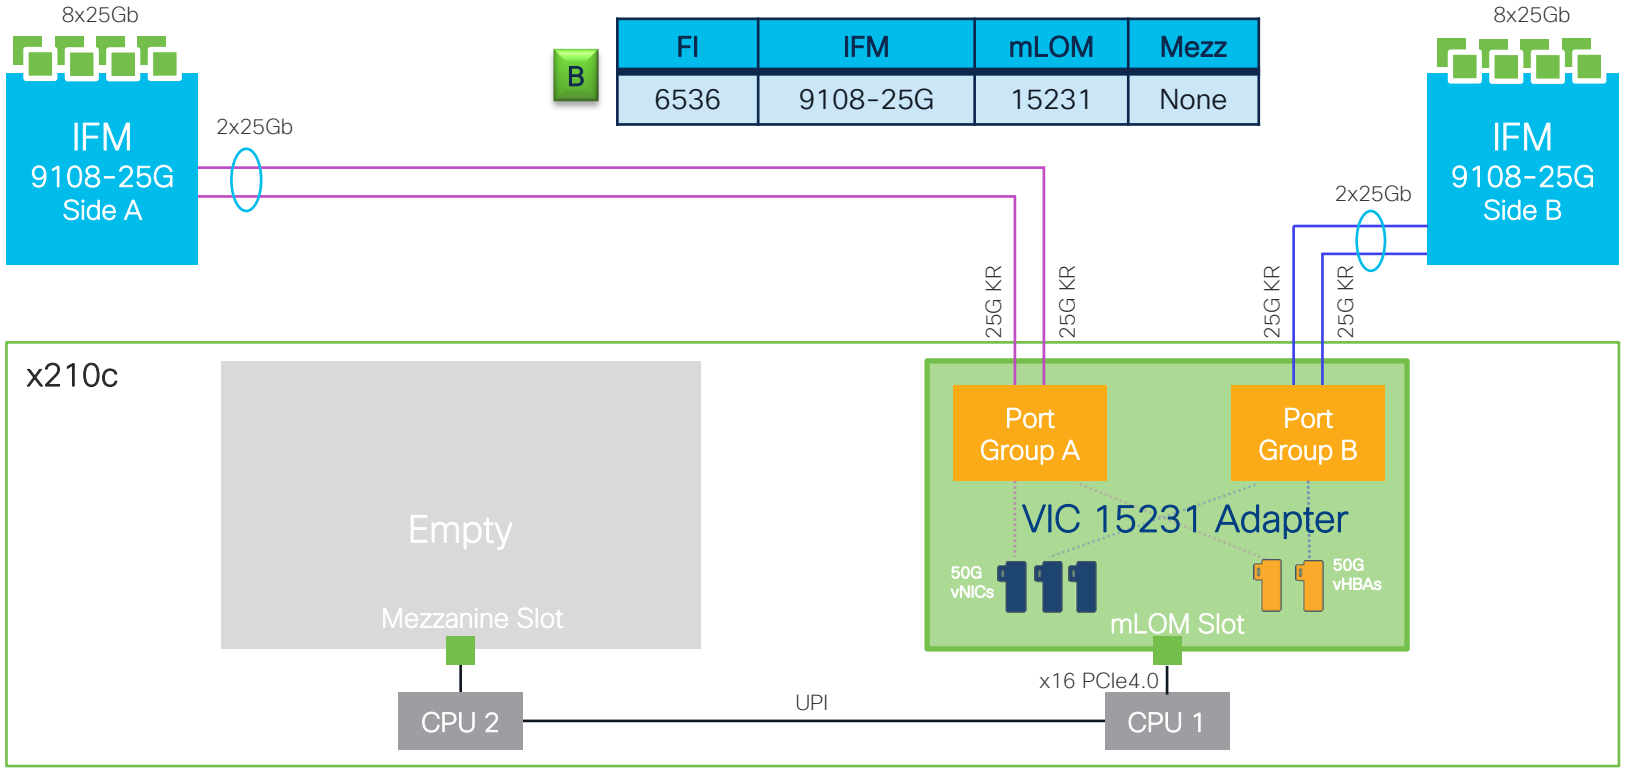

x210c, IFM 9108-100G and VIC 14425

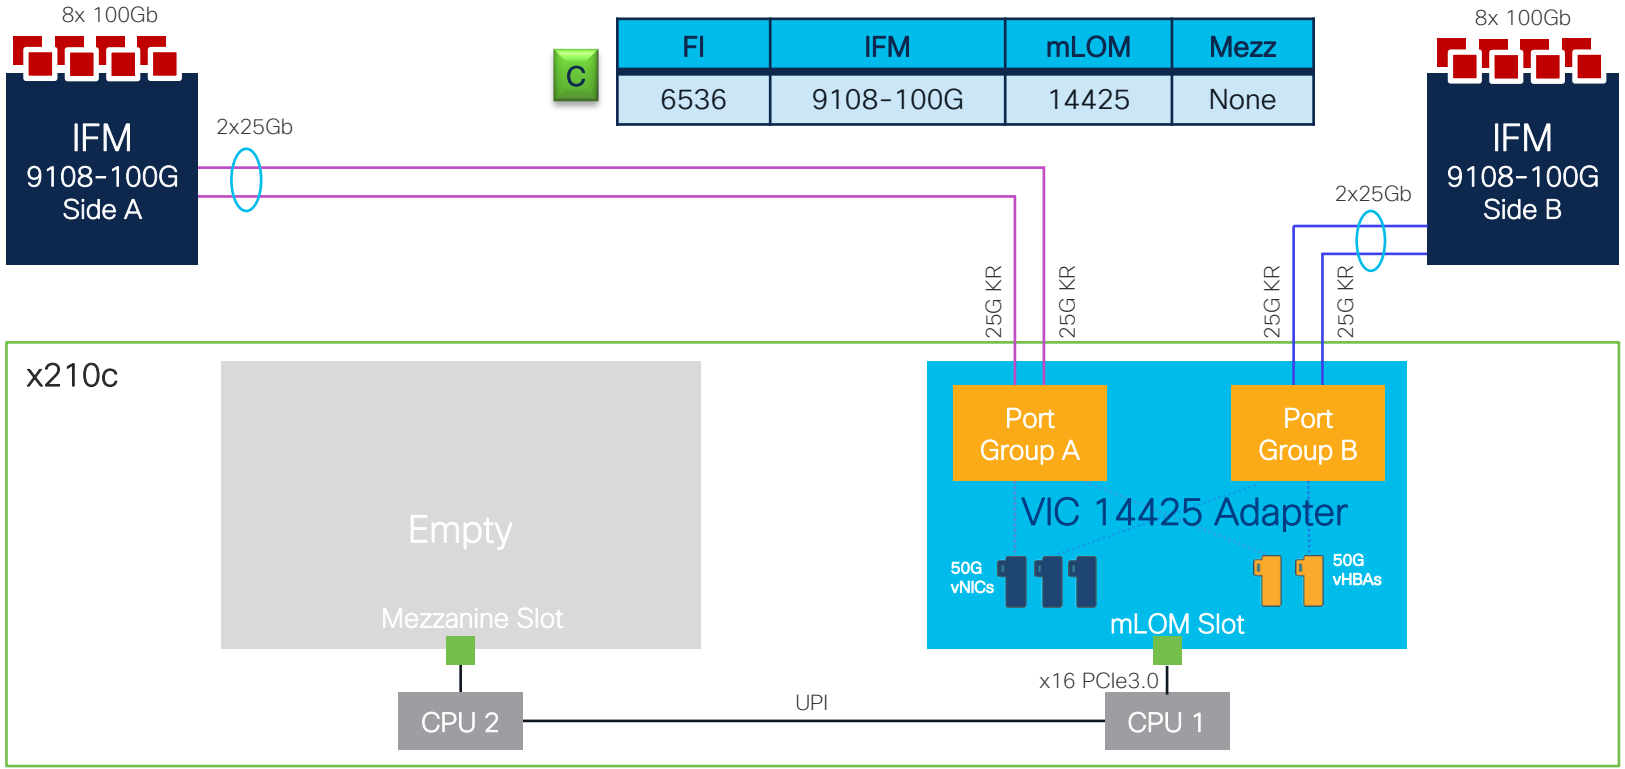

x210c, IFM 9108-100G, VIC 14425 and VIC 14825

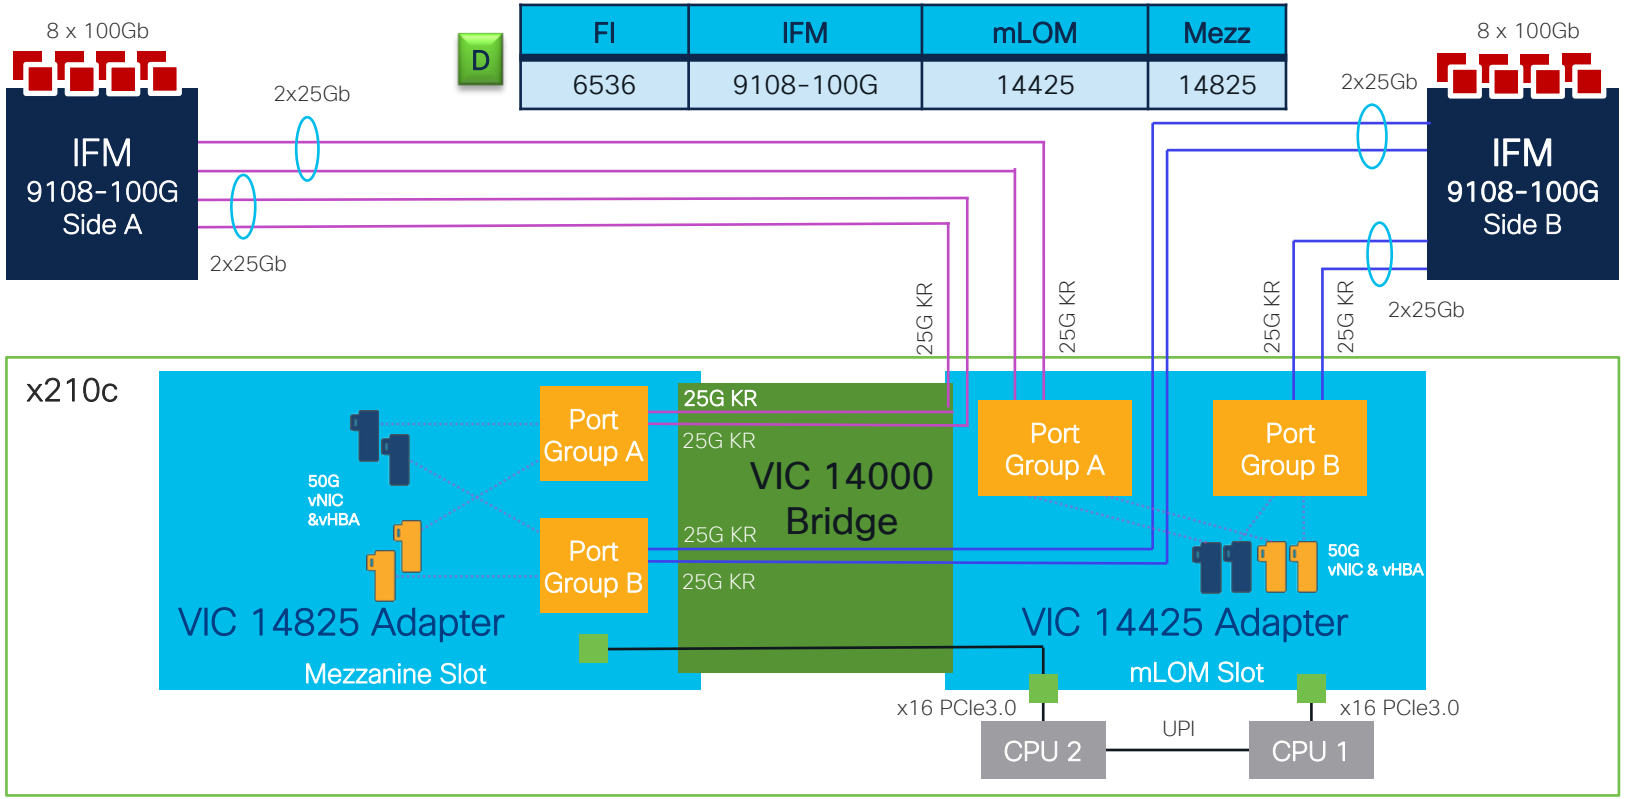

x210c, IFM 9108-25G and VIC 14425

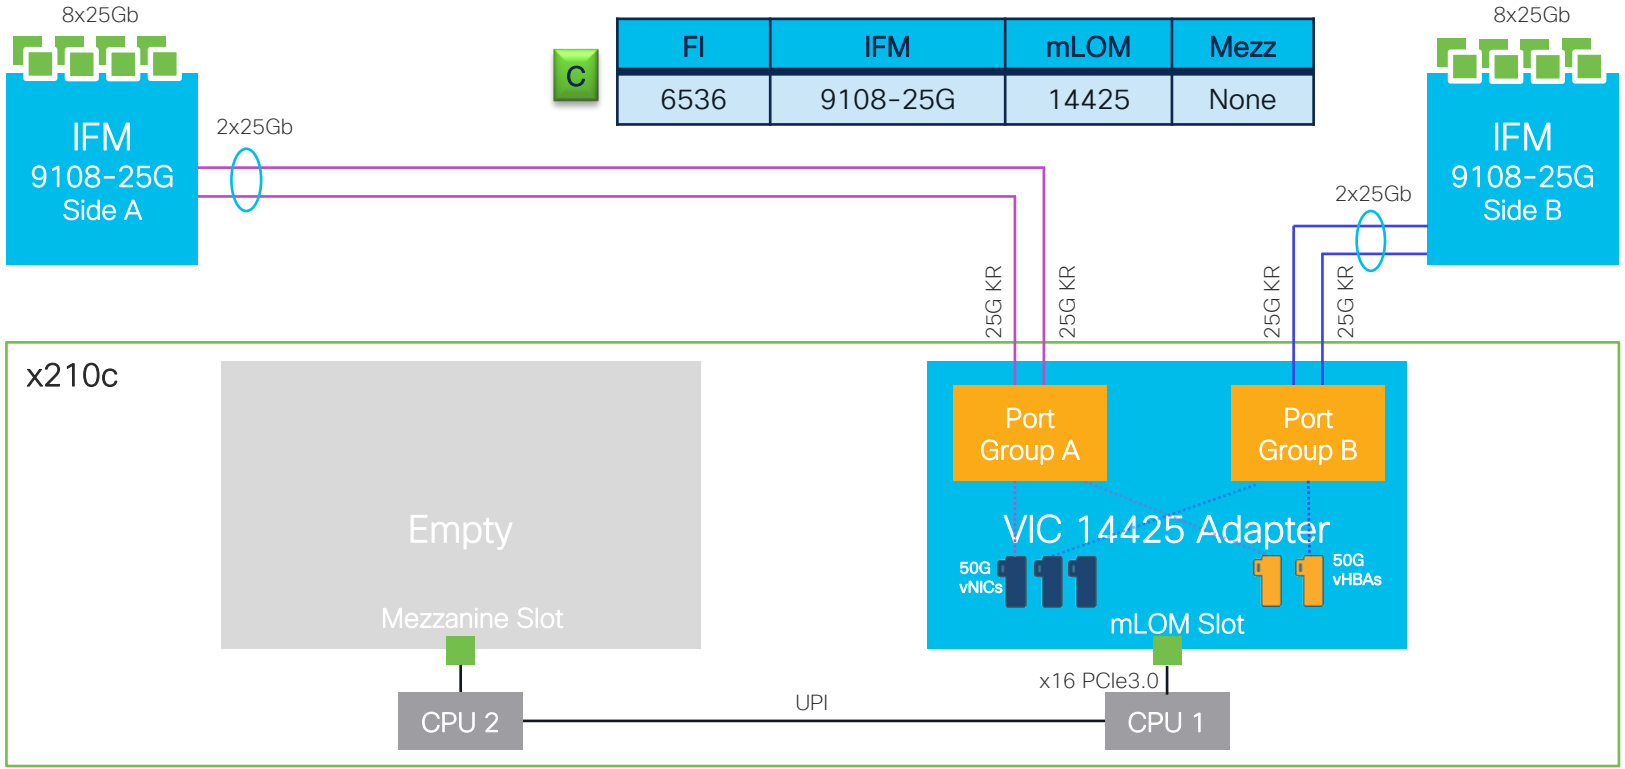

x210c, IFM 9108-25G, VIC 14425 and VIC 14825

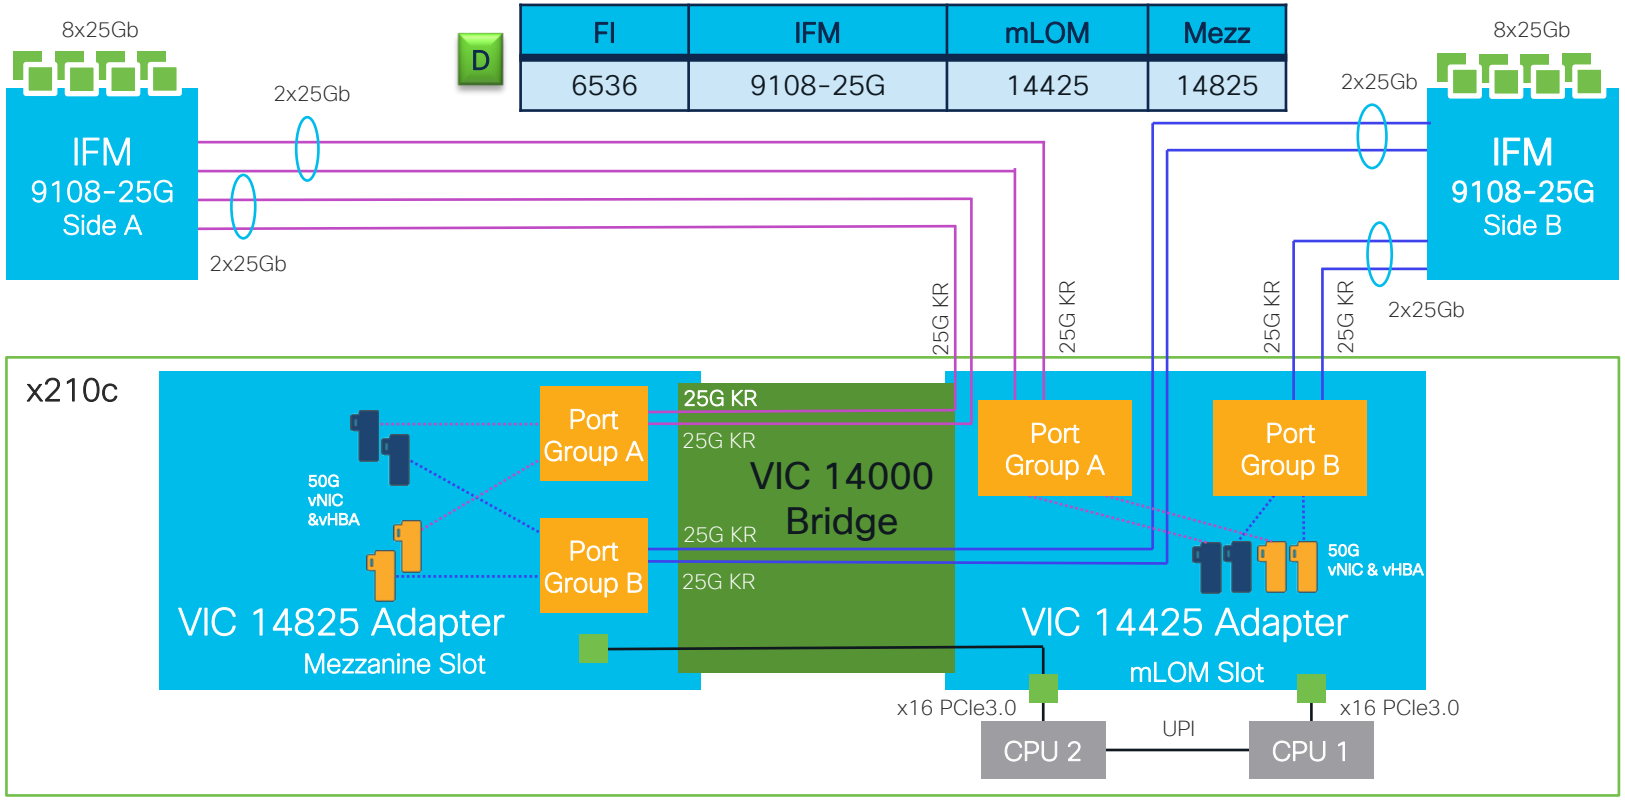

FI 6536 to IOM-2xxx : B-Series Connectivity

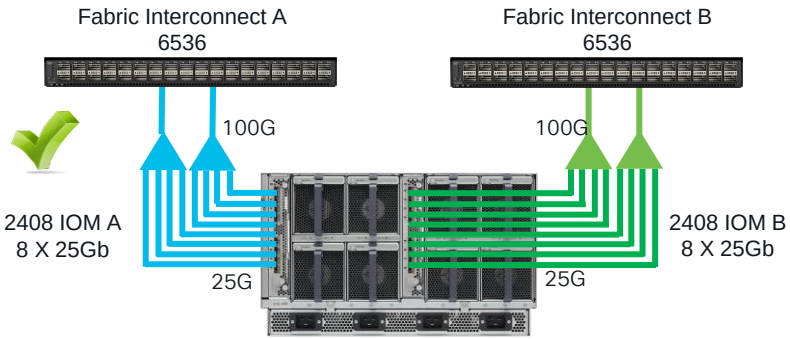

2408 IOMs to 6536 FIs : Supported

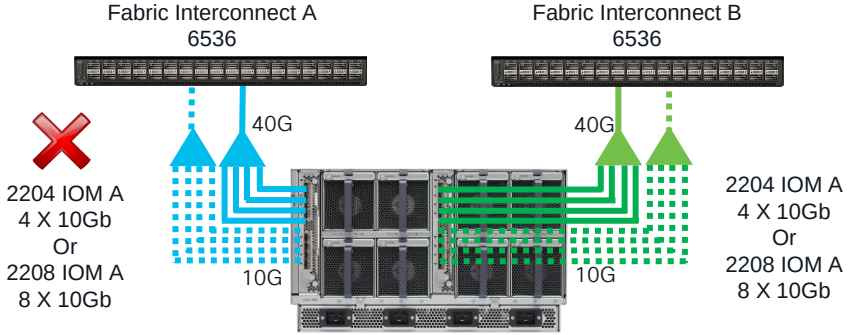

2204 or 2208 IOMs to 6536 FI : Not Supported

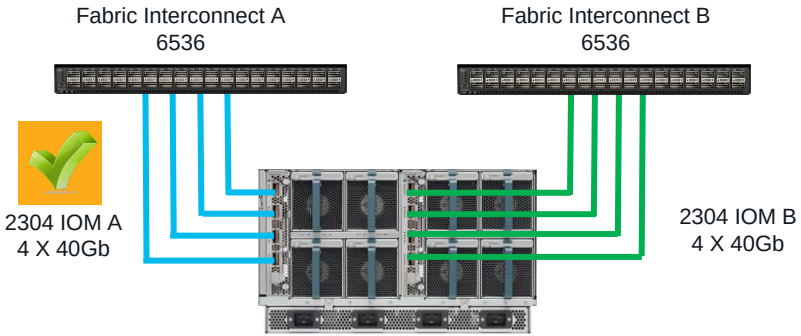

2304 IOMs to 6536 FIs : Targeted post 4.2.2

Throughput per B-series blade server with 6536 FI

	6536 + 2408	
	B200-M5 / M6	B480-M5
1440	40G	40G
1440 + PE	80G*	80G*
1440 + 1480	80G	120G
1440 + PE + 1480	N/A	160G*

* 1440+PE enables 40G-KR4 path between the VIC and IOM.

Note that aggregate per vNIC would be 40G and the max single-flow from blade-server will be 25G since the IOM-2408 to FI-6536 connectivity is 25Gbps.

B200M5/M6, IOM 2408, & VIC 1440 with PE

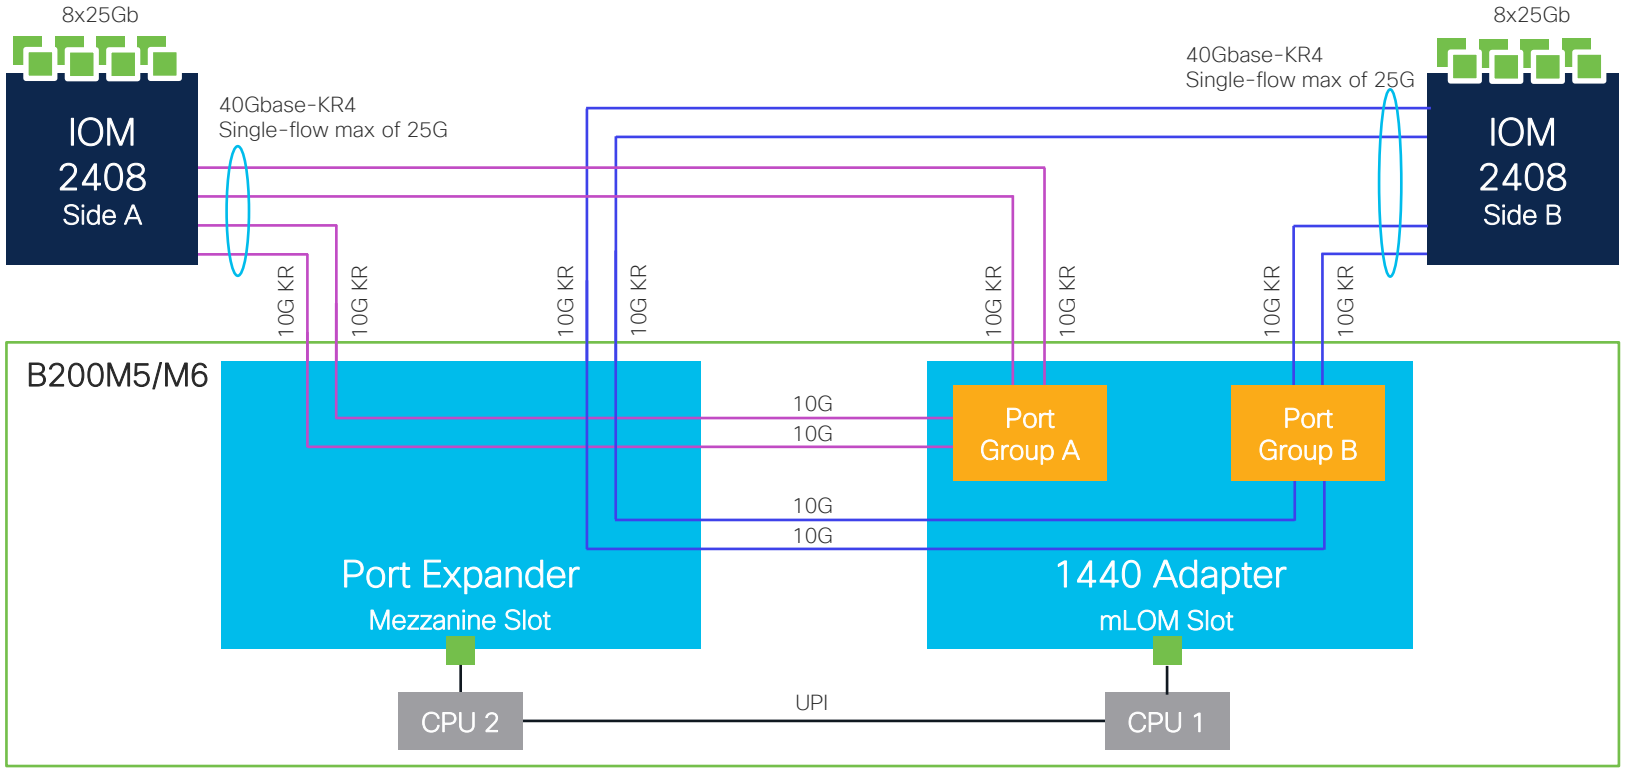

FI 6536 Series Use Cases with C-Series

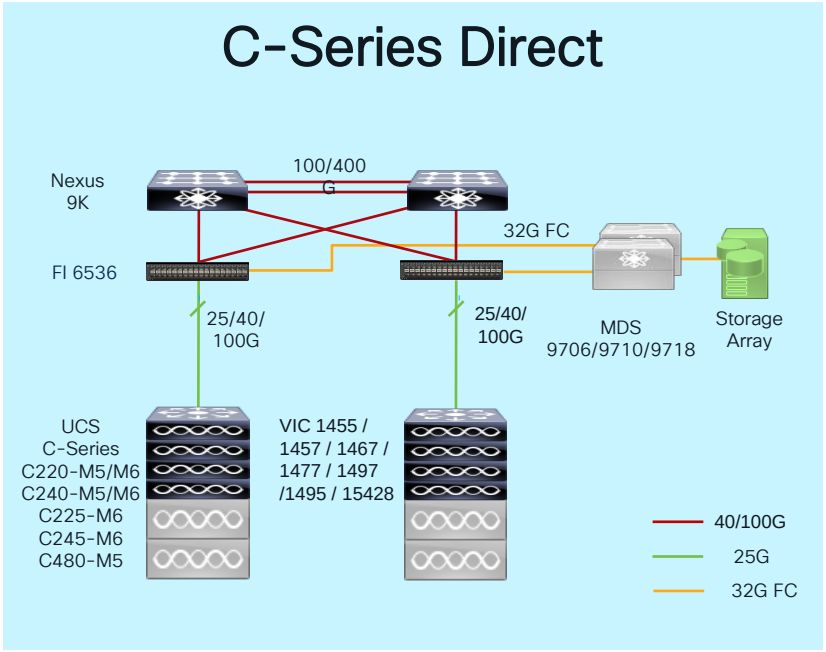

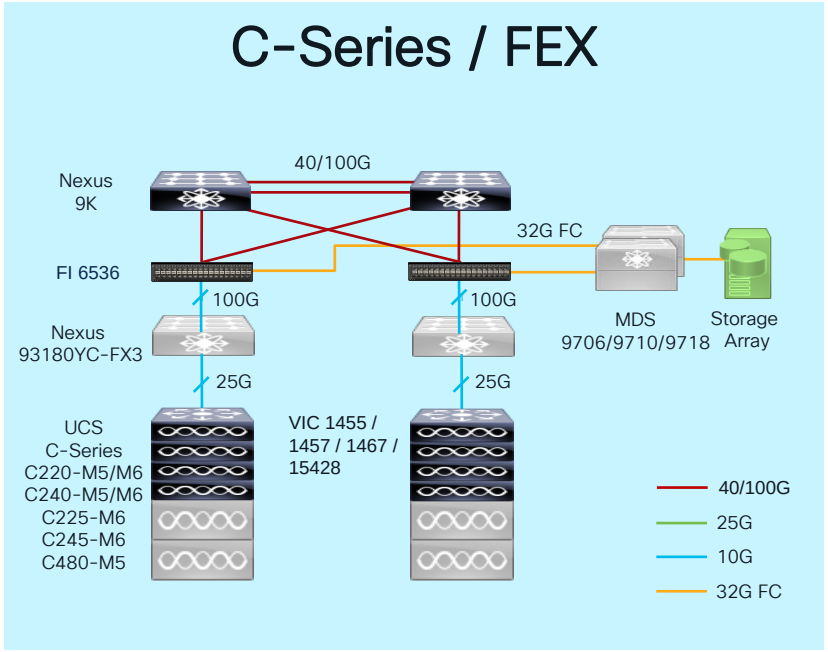

Migrate to 5th Gen Fabric

Cisco UCS 9508 Chassis fabric upgrade for 100G E2E

Migrate from IFM-25G to IFM-100G

- Backplane less X-series chassis design
- Hot-swappable IFMs

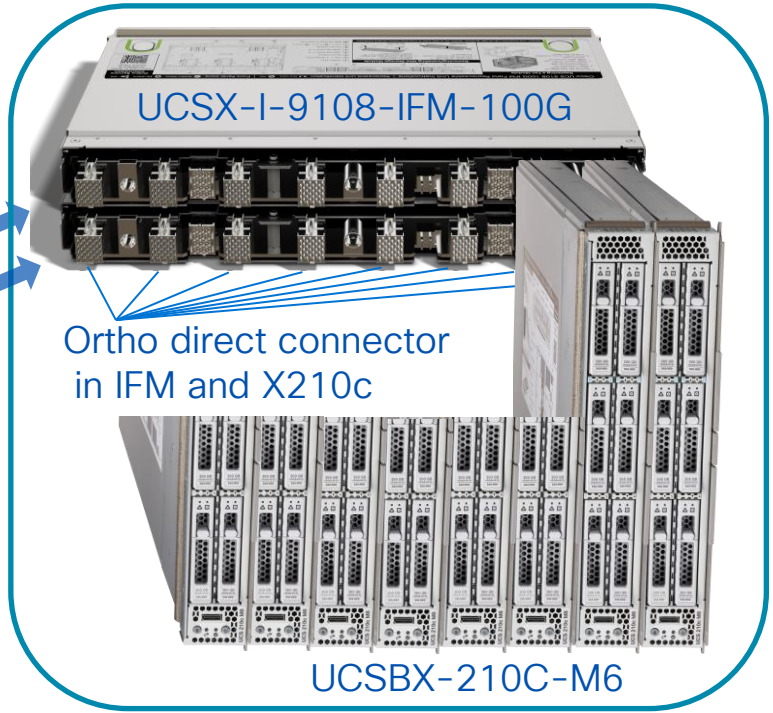

Fabric Interconnect Migration Options in IMM

FI 6400 Series in IMM

FI 6454

FI 64108

Migrate to FI 6536 in IMM

FI 6536

FI 6400 series

- 6454 supports 48 x 10/25G + 6 x 40/100G
- 64108 supports 96 x 10/25G (16xUP)+ 12 x 40/100G
- 8/16/32G FC uplink/storage port
- 10/25G server ports
- 1/10/25/40/100G ethernet uplink ports
- IFM-IFM-25G, IOM-220x/2408

FI 6536 Advantages

- 6536 supports 32 x 40/100G + 4 x 40/100G or 16x 8/16/32G FC
- 8/16/32G FC uplink/storage port
- All 36 ports can breakout into 10/25G
- 25/40/100G server-port support
- 1/10/25/40/100G ethernet uplink ports
- IFM-100G/IFM-25G, IOM-2408 , IOM-2304(post-FCS)

IMM Transition Tool

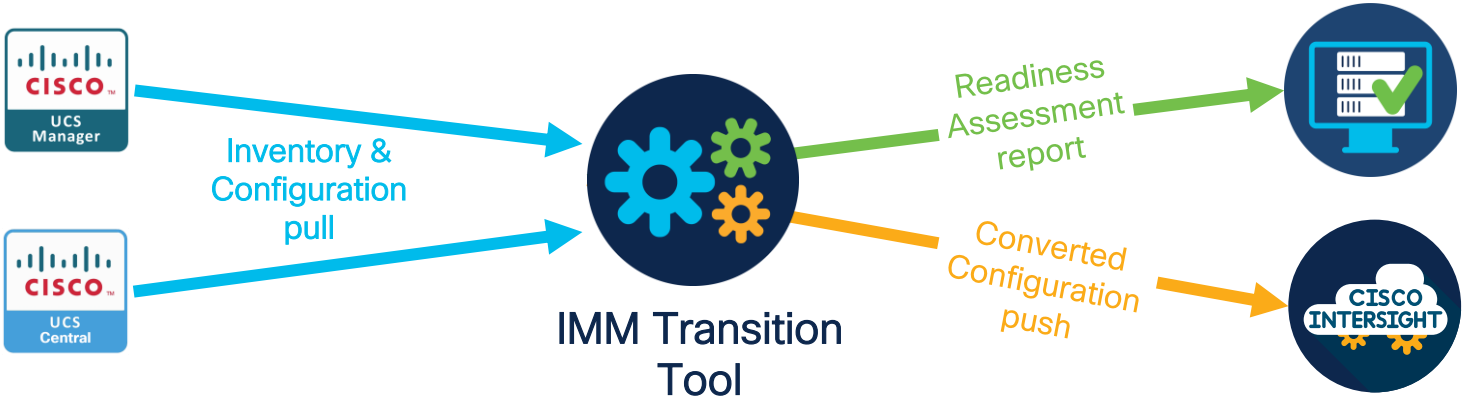

IMM Transition Tool

- Assesses **hardware & firmware compatibility**
- **Extends** existing **Service Profile Templates** to Intersight
- Automatically converts related **server policies** (boot, BIOS, LAN/SAN connectivity, etc.)
- Converts **fabric configuration** (VLANs/VSANs, port configuration, etc.)

5th Gen Fabric – 100G End-to-End

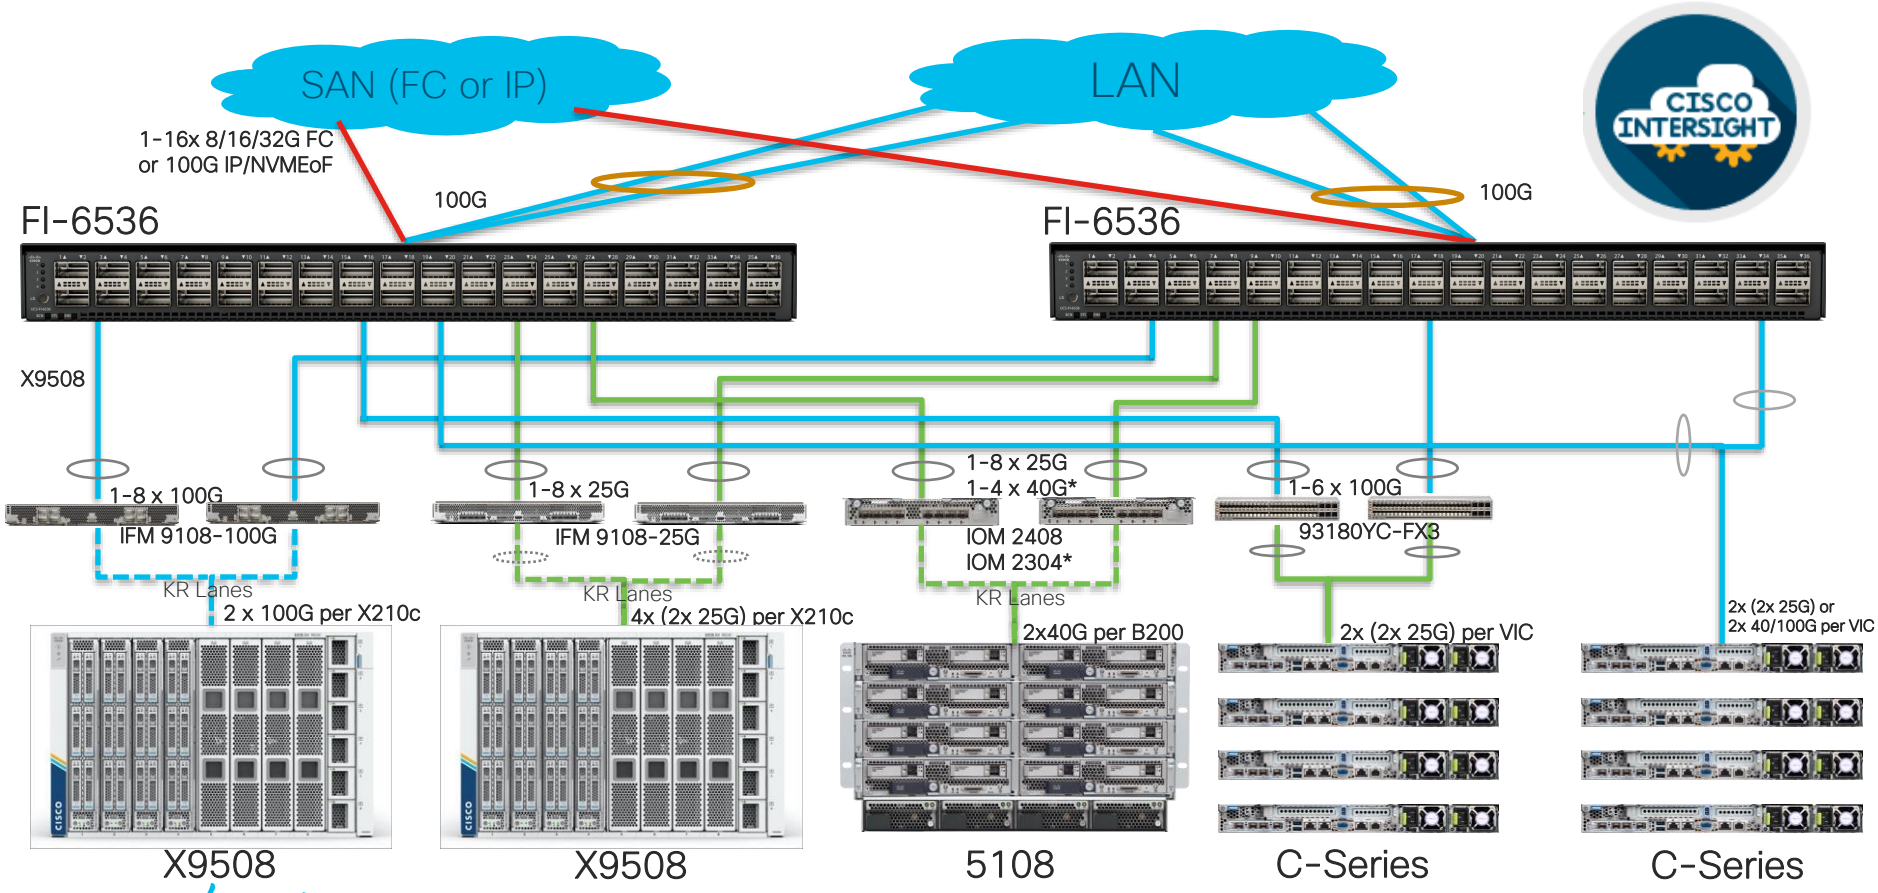

* Post 4.2(2) FCS

Addendum:

UCS VIC 15000 Series

VIC 15231

- MLOM VIC for X210c-M6 with FI 6400/6536
- Backplane connectivity of 25G or 100G depending on IFM
  - 1 port of 100G-KR4 with an IFM-100G
  - 2 ports of 25G-KR with an IFM-25G
- With IFM-100G the logical uplink is a 100G interface
  - vNIC/vHBA bound to logical uplink will be 100G
  - vNIC/vHBA can do single-flow of 100Gbps
- With IFM-25G the logical uplink is a 50G port-channel
  - vNIC/vHBA bound to logical uplink will be 50G
  - vNIC/vHBA can do single-flow of 25Gbps
- vNIC/vHBA defined by LAN/SAN connectivity policy
  - No default vNICs/vHBAs

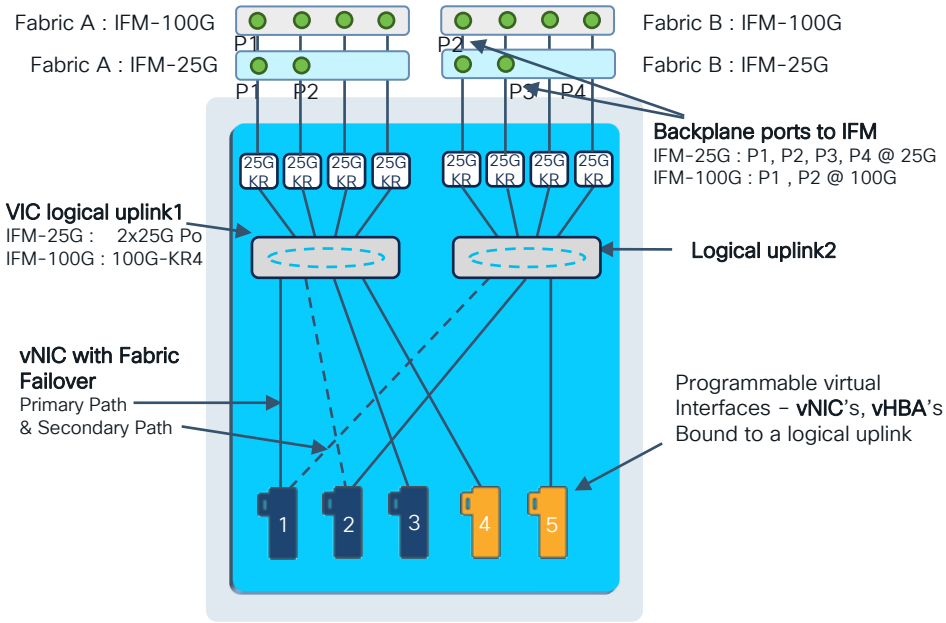

VIC 15428 (FI managed)

- MLOM VIC for M6 C-series with FI 6300/6400/6536
- Four physical ports which can run at 10G/25G
- Port speed is determined by inserted transceiver
- Physical ports (P1, P2) are statically HW port-channeled as logical uplink 1 and ports (P3, P4) are bundled as logical uplink 2
- Connectivity to Fabric Interconnects
  - Supports one link or two links to each FI.
  - The links connected to each FI will be in a port-channel
  - VIC HW port-channel cannot be disabled when FI managed
  - vNIC speed determined by the transceiver type and the number of active links
  - vNIC/vHBA speed of 10G, 20G, 25G or 50G per fabric.
  - No FEC or auto-negotiation configuration required, link settings are auto-determined.
- vNIC/vHBA defined by LAN/SAN connectivity policy
  - No default vNICs/vHBAs

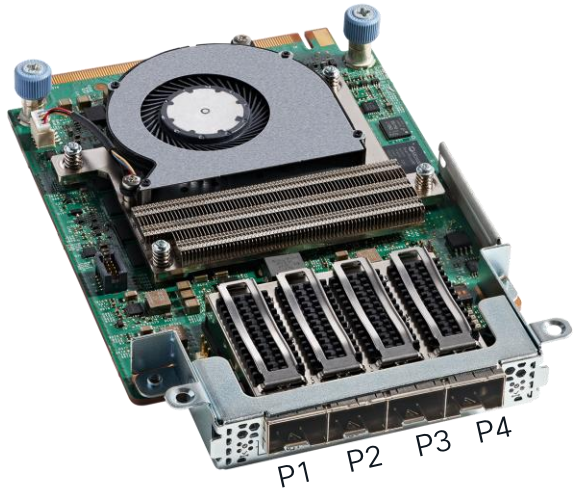

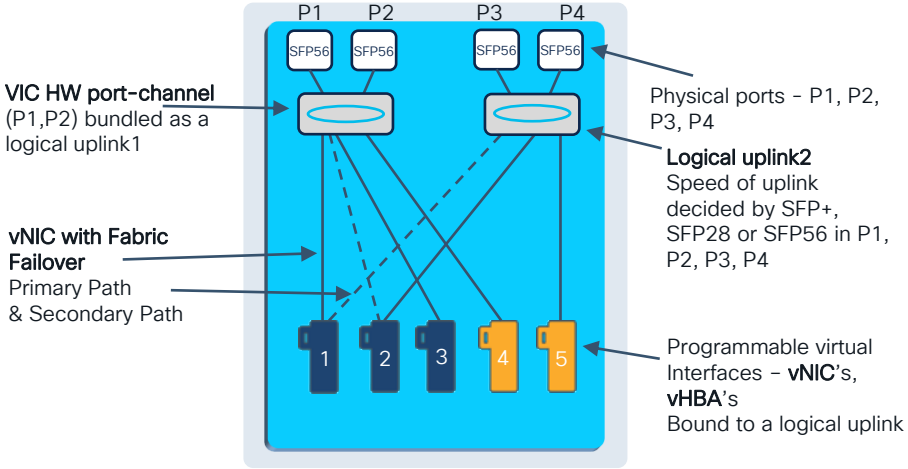

VIC 15428 connectivity to FI 6536

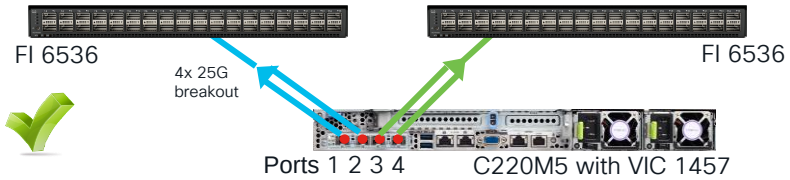

VIC ports 1 & 2 to FI-A and ports 3 & 4 to FI-B - Supported

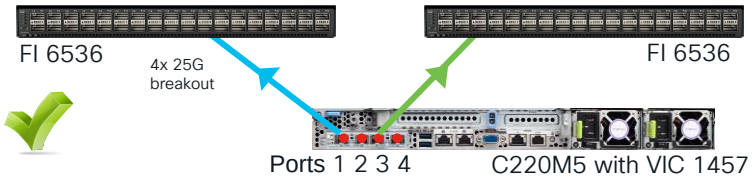

VIC port 1 to FI-A and port 3 to FI-B - Supported

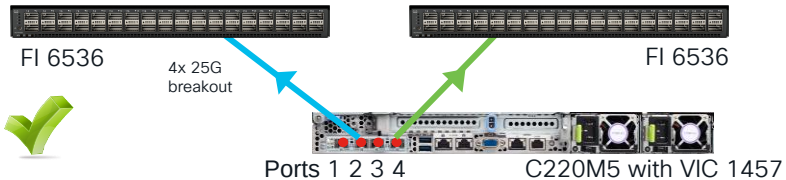

VIC ports 2 to FI-A and ports 4 to FI-B - Supported
VIC ports 1 or 2 to FI-A and ports 3 or 4 to FI-B - Supported

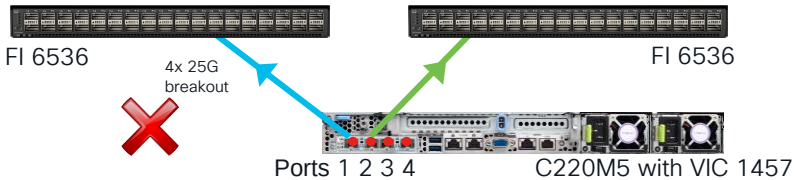

Not Supported
1. VIC port 1 to FI-A and port 2 to FI-B
2. VIC port 3 to FI-A and port 4 to FI-B

Intersight Server Profile

Common Policies for Server Profile

- BIOS
- Boot Order
- IMC Access
- LAN Connectivity
- Storage
- Virtual KVM
- Virtual Media
- SAN Connectivity

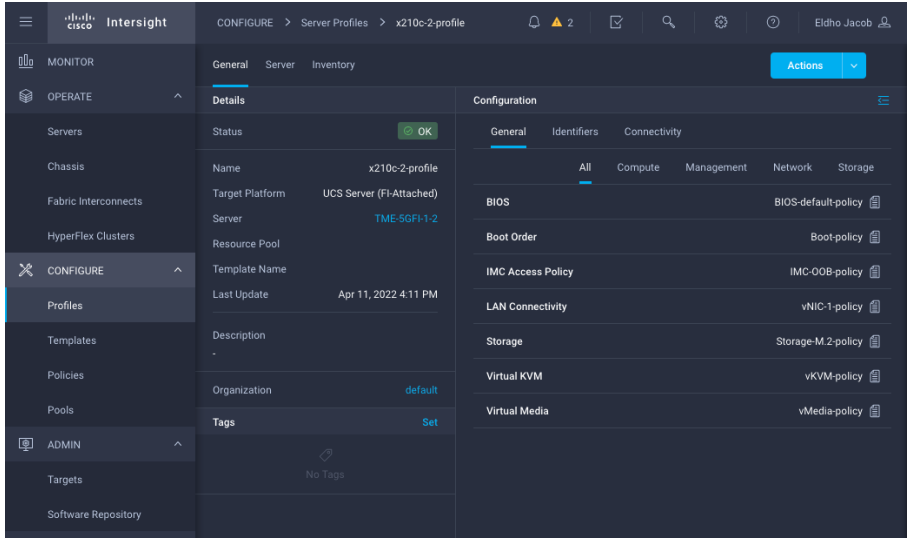

LAN Connectivity Policy

LAN Connectivity policy

- Creates multiple vNICs and decide adapter placement
  - Ethernet policies per vNIC
    - **Network Group** (Allowed VLANs, Native VLAN)
    - **Network Control** (CDP/LLDP, uplink failure, MAC register etc)
    - **QoS** (CoS, rate-limit, MTU)
    - **Adapter** (Rx/Tx queues, ring-size, HW acceleration/Offload features)

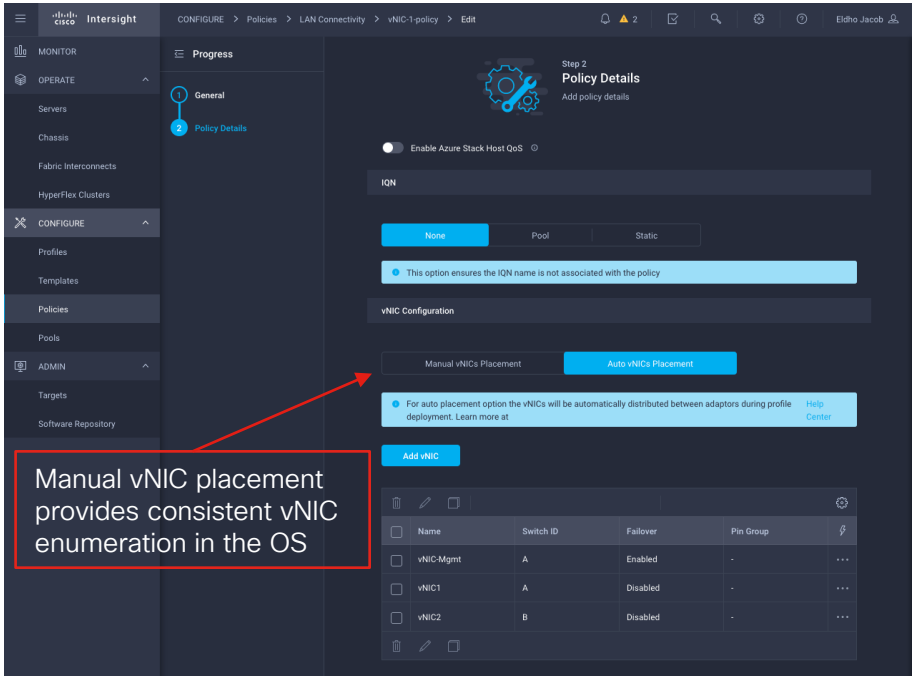

RSS Support

- VIC HW feature supported for ESXi, Linux, and Windows
- RSS provides better server CPU utilization, higher throughput and handles bursty traffic
- Achieved by Rx traffic distribution across multiple queues/cpu-cores based on L2/L3/L4 packet header fields
- VIC 15000 series support 16K Tx and Rx ring size, while VIC 1400 series supports up to 4K Tx and Rx ring size

Adapter policy for performance with RSS

Parameter	ESXi	Linux	Windows
TX queue	1	1	1
TX ring size	4K /16K	4K /16K	4K /16K
RX queue	8	8	8
RX ring size	4K /16K	4K /16K	4K /16K
CQ	9	9	9
Interrupt	11	11 or 10	512
Interrupt Calculation	CQ + 2	“CQ + 2” or “Rx-Queue + 2”	512 or “2 x CPU-cores + 4”
RSS	Enabled	Enabled	Enabled

NetQueue Support

- NetQueue is an integrated hardware and software solution from both Cisco and VMware
- NetQueue achieves higher throughput and performance by having dedicated TX/RX queue per VM
- VIC RSS enables multiple RX queues across multiple VM's, while NetQueue dedicates a TX/RX queue per VM.
- NetQueue on the vNIC is enabled through the VMQ connection policy.
- Interrupt for NetQueue is calculated as "2 x VMQ + 2."

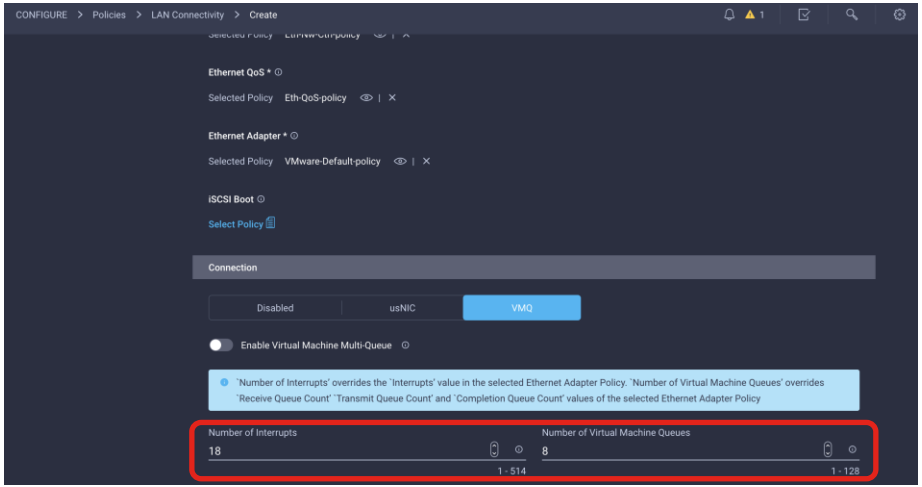

VMMQ Support

- Virtual Machine Multi-Queue, allows allocating multiple RX queues per vPort in a Windows Hyper-V host.
- Thus, providing higher throughput and distributes traffic load across multiple CPU cores.
- Supported in IMM, UCSM, and CIMC
- VMMQ is recommended over VMQ, or RSS for Windows Hyper-V. Both VMQ and RSS are supported by VIC 15000 for Windows.
- Use default adapter policy values in Intersight of “Win-HPN” and “MQ” to enable VMMQ. And the policy definition is good for 64 vPorts. Sample example for IMM is attached

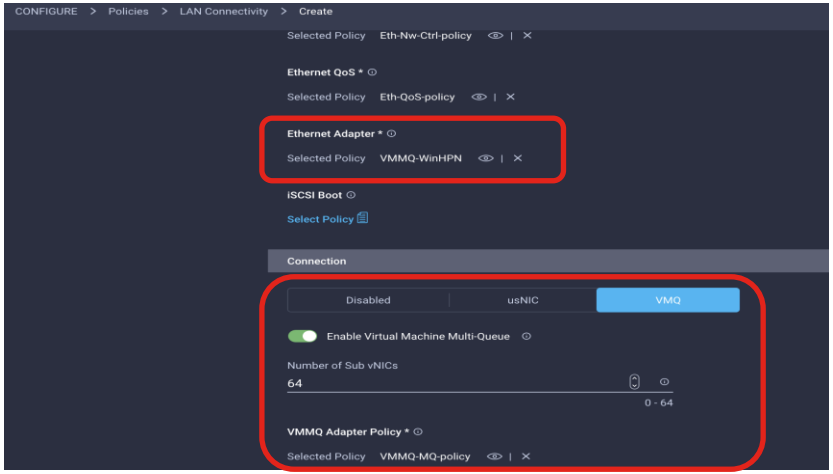

NVMeoF Support

VIC 1400/15000	ESXi	RHEL	SLES
FC-NVMe	Yes	Yes	Yes
NVMe-RoCEv2	Yes*	Yes	-
NVMe-TCP	Yes*	Yes*	-

* Targeted post 4.2.2

- ROCEv2 is supported with Windows SMBdirect.
- Refer UCS RoCEv2 configuration guide for details : https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/ucs-manager/GUI-User-Guides/RoCEv2-Configuration/4-1/b-RoCE-Configuration-Guide-4-1.html

PTP Support

- Supported on all VIC 15000 series
- < 100ns precision
- Supports PTPv1 and PTPv2
- Supports multicast and unicast PTP
- Supported with RHEL
- Only one vNIC per VIC should be enabled with PTP
- Check with “ethtool -T <eth-id>
- Supported in IMM, CIMC, and UCSM

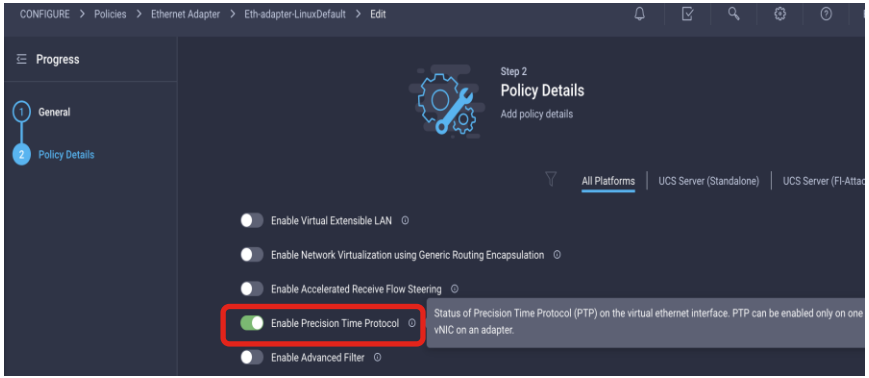

VIC 15428 (in standalone)

- VIC 15428 connected to a Nexus switch (non-FI)
- Ports can run at 10G/25G/50G speed
- Speed is determined by inserted transceivers.
- By default, ports (P1, P2) and (P3, P4) are in an HW port-channel
  - **2x vNICs** and **2x vHBAs** by default
- With HW port-channel disabled from CIMC, 4 uplink ports corresponding to each physical port gets enabled
  - **4x vNICs** and **4x vHBAs** by default
- Auto-Negotiation Mode (enabled by Admin link- training config)
  - Enabled for 50G Copper Transceivers by default
  - Disabled for 25G Copper Transceivers by default
  - Auto Negotiation can be enabled/disabled using CIMC CLI or WebUI in standalone mode.
- Admin Link Training Configuration per port
  - **Auto**, **On**, **Off** (Auto means VIC firmware decides the correct mode)
  - Auto-fec supported with 25G-CUx cables when link-training is on.
  - Link training ensures greater link reliability
- FEC (Forward Error Correction) Configuration per port
  - FEC Configuration for 25G (**cl74**, **cl91-cons16**, **cl91**, **cl108**)
  - FEC Configuration for 50G (**cl91**)
  - FEC Configuration ignored for 10G speed
- Physical NIC mode supported to disable priority-tagging

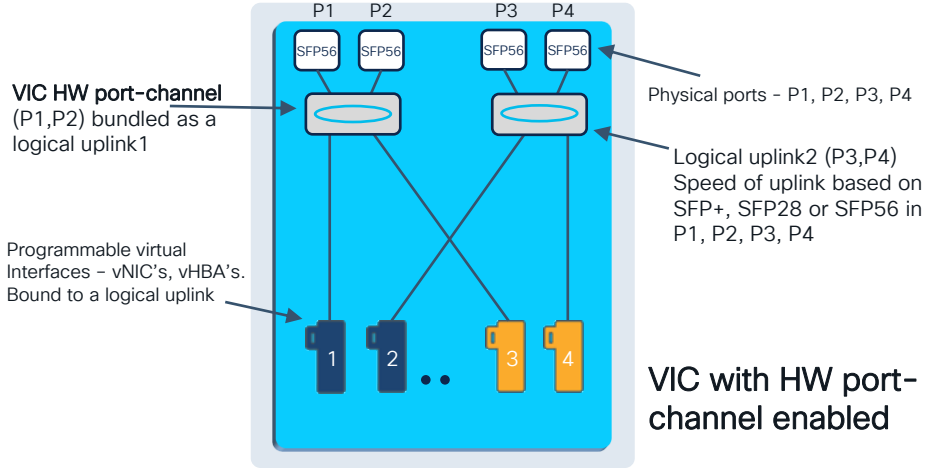

VIC with HW port-channel enabled

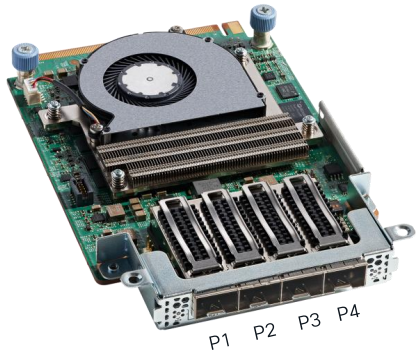

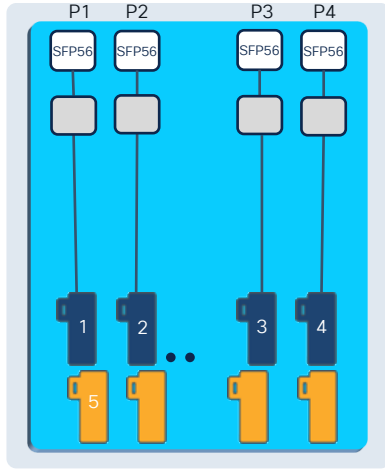

VIC with HW port-channel disabled

VIC 15428 connectivity to Nexus switch in standalone mode

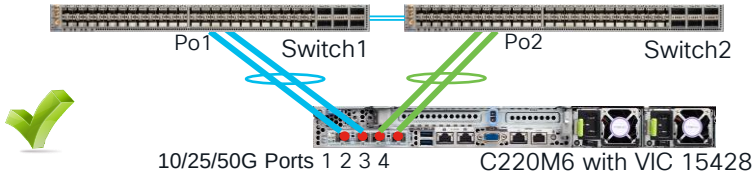

VIC ports 1 & 2 to SW1 and ports 3 & 4 to SW2 – Supported

- With Default VIC port-channel (PO) **enabled**
- Requires PO config on switch with switch dependent bonding
- Cannot support MCT/VPC at ToR switch and OS IP-hash kind of load-balancing
- MAC-hash or port-ID load-balancing in OS should be used to avoid mac-move on ToR

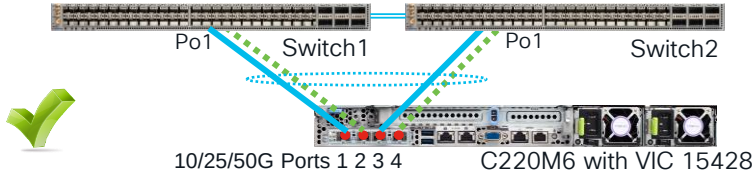

VIC ports 1 or 2 to SW1 and ports 3 or 4 to SW2 – Supported

- With VIC port-channel **enabled** and using one link in (1,2) & (3,4) port pair
- Supports switch dependent & switch-independent OS teaming/bonding
- Supports MCT/VPC at ToR switch and all OS teaming load-balancing options

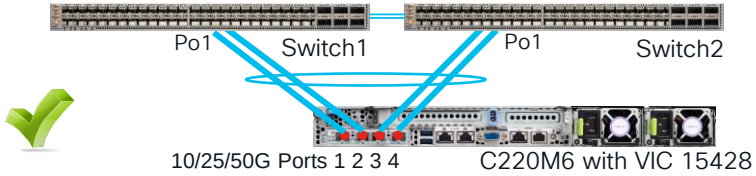

VIC ports 1, 2, 3, 4 to SW1 & SW2 – Supported

- With VIC port-channel **disabled**
- Supports MCT/VPC at ToR switch and all OS teaming/load-balancing options

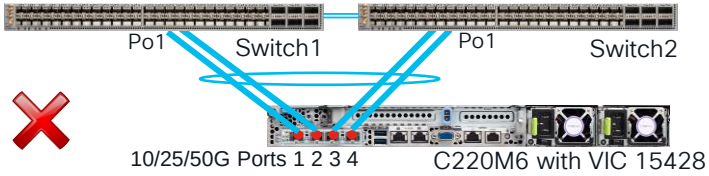

VIC ports (1,2) to SW1 & (3,4) to SW2 – Not Supported

- With VIC port-channel **enabled**

These connectivity options at 10G/25G are applicable for VIC 1400 series as well

Physical NIC mode (in Standalone Server)

- Supported on VIC 1400/15000 series with standalone rack servers from 4.2.2 IMC release
- Physical NIC mode disables default priority tagging on VIC vNICs
  - Allows interoperability with switches that don't support priority-tagging
- Only default vNICs are supported, no additional vNICs can be created
  - 2 vNICs for dual-port MLOM or PCIe rack VIC
  - 4 vNICs for quad-port MLOM or PCIe rack VIC
- Disable FIP and LLDP on VIC to enable physical NIC mode

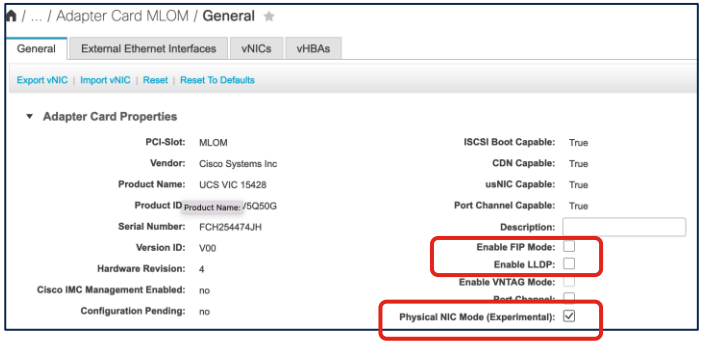

Technical Session Surveys

- Attendees who fill out a minimum of four session surveys and the overall event survey will get Cisco Live branded socks!
- Attendees will also earn 100 points in the Cisco Live Game for every survey completed.
- These points help you get on the leaderboard and increase your chances of winning daily and grand prizes.

Continue your education

- BRKCLD-2028 - Build Your Hybrid Cloud with Cisco
- BRKCLD-2807 - Real Life Deployment of Hybrid Cloud Infrastructure using Intersight
- BRKCLD-3010 - UCS X-Series: Blurring the Line Between Rack and Blade for Modern Applications
- BRKCLD-3011 - Learn Why Cisco UCS Builds the Best AMD Servers in the Market
- BRKDCN-2310 - Unleash the Power of Converged Infrastructure Using Cisco Intersight and UCS X-Series
- BRKDCN-2587 - Best Practices for Cloud and Compute Connectivity with the 5th Gen UCS Fabric and Intersight
- BRKDCN-2794 - How Will the UCS X-Series Hardware Bring Your Applications Together?
- CCP-1210 - Roadmap: Hyperflex, Intersight, and Unified Computing Systems Update
- DEVNET-3018 - Provision Infrastructure at Cloud scale with Terraform Provider for Cisco Intersight
- DEVNET-3026 - Simplify IAC for On Premises Environment using Cisco Intersight and Hashicorp Terraform Cloud
- DEVWKS-2060 - Automating Cisco UCS Server Provisioning with Terraform and Intersight
- DEVWKS-2942 - DevNet Workshop - Programmability with Intersight Managed Mode
- IBODCN-1551 - An Interactive Conversation on the Evolution of Compute Hardware in the Cloud-First Era
- IBODCN-2301 - An interactive conversations on IMM Transition – How do I get there?
- LABDCN-1115 - Cisco Intersight Cloud-based Infrastructure management
- TSCDCN-2012 - Troubleshooting New Installation of UCS X-Series System in Intersight Managed Mode

World of Solutions – X-Series & Fabric Demos

- Product innovation on display
  - X-Series, X-Fabric, 5th Gen Fabric Interconnect, 5th Gen VIC, GPU Node
- On Demo
  - Experience simplicity with X-Series
    - Deploy X-Series live with just a few clicks - Power of Intersight
    - Define FlexPod with X-Series in Intersight
    - Pervasive visibility across X-Series FlashStack with Intersight
    - Deploy Red Hat OpenShift Container Platform with confidence (Ceph)
    - Automation with X-Series leveraging Intersight Ansible
  - Experience innovative X-Fabric – Blurring line between rack and blade
    - Simplicity of adding GPUs to compute nodes…in few minutes!
  - Experience amazing performances on X-Series
    - Performance out of the box
    - Oracle Swingbench performance
  - Experience new 5th Gen Unified Fabric – Simplify, accelerate and reduce components
    - 100G/200G Ethernet and NFS with a single VIC 15231
    - 100G FC aggregate and 32G E2E FC per vHBA on a single VIC 15231

X-Series white papers

- Cisco UCS X-Series Quick Start Guide
- Cisco UCS X210c M6 Compute Node Disk I/O Characterization
- Deploy Cisco UCS X210c Compute Node with Cisco Intersight Management Mode for VDI
- FlashStack with Cisco UCS X-Series and Cisco Intersight
- FlexPod Datacenter with Cisco UCS X-Series and Cisco Intersight
- Power SAP HANA with the Cisco UCS X-Series Certified by SAP
- Deploy SAP HANA Scale-Up Appliance with UCS X-Series
- Cisco UCS and Intel SGX with Fortanix Confidential Computing Manager
- Deploy a High-Performance Standalone Oracle Database Solution: Oracle 19c on Cisco UCS X-Series
- FlexPod Datacenter with Citrix VDI and VMware vSphere 7 for up to 2500 Seats
- FlexPod XCS Solution with Cisco Intersight Platform Tech Preview
- Red Hat OpenShift Container Platform with OpenShift Data Foundation on Cisco UCS X-Series
- Cisco UCS X-Series Servers with Intel Optane Persistent Memory for Virtual Desktop Infrastructure White Paper
- Get Answers from Your Data with Cisco UCS Integrated Infrastructure for Splunk Enterprise

Cisco Learning and Certifications

From technology training and team development to Cisco certifications and learning plans, let us help you empower your business and career. www.cisco.com/go/certs

Pay for Learning with Cisco Learning Credits

(CLCs) are prepaid training vouchers redeemed directly with Cisco.

Learn

Cisco U.
IT learning hub that guides teams and learners toward their goals

Cisco Digital Learning
Subscription-based product, technology, and certification training

Cisco Modeling Labs
Network simulation platform for design, testing, and troubleshooting

Cisco Learning Network
Resource community portal for certifications and learning

Train

Cisco Training Bootcamps
Intensive team & individual automation and technology training programs

Cisco Learning Partner Program
Authorized training partners supporting Cisco technology and career certifications

Cisco Instructor-led and Virtual Instructor-led training
Accelerated curriculum of product, technology, and certification courses

Certify

Cisco Certifications and Specialist Certifications
Award-winning certification program empowers students and IT Professionals to advance their technical careers

Cisco Guided Study Groups
180-day certification prep program with learning and support

Cisco Continuing Education Program
Recertification training options for Cisco certified individuals

Here at the event? Visit us at **The Learning and Certifications lounge at the World of Solutions**

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

cisco Live!

Thank you


cisco Live!
ALL IN