Best Practices for Compute Connectivity with the 6th Generation Fabric Interconnect

Joost van der Made
Fabric Computing Technologist & TME

BRKCOM-2658

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 13, 2025.

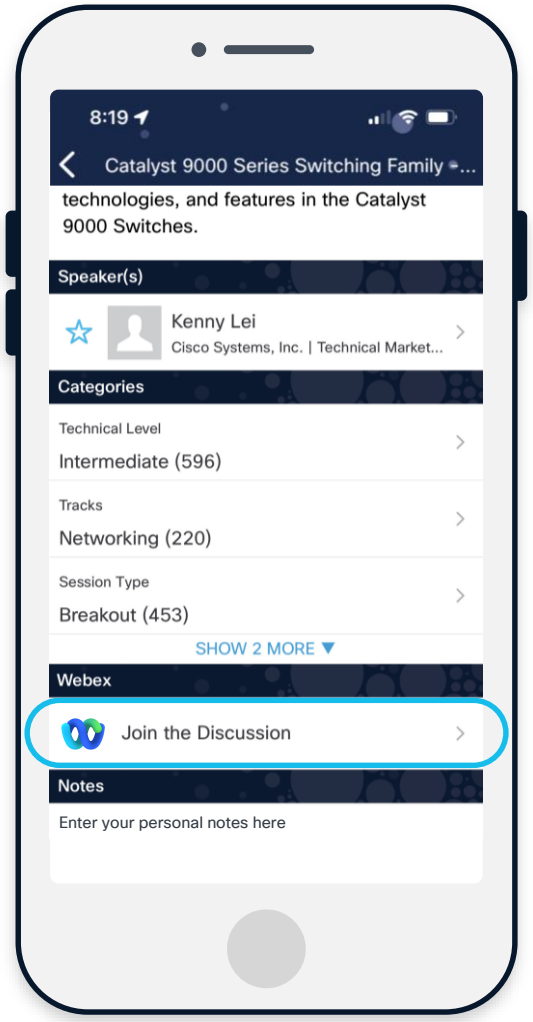

Introduction Joost van der Made
Fabric Computing Technologist & TME

https://www.linkedin.com/in/jvdmade/

UCS Fabric Value: A Look Back

Hardware Innovation | Unified Fabric | Business Benefit

UCS Fabric Architecture

10/25/40/100G, End-to-End NW Virtualization, QoS Guarantee, Redundant Fabric, Modular, Server-type/LAN/SAN independent

Unified/Converged Fabric

Unified fabric for Ethernet, Storage and Management, Dynamic I/O Virtualization, Cable/Adapter/Fabric-node Reduction, improving sustainability

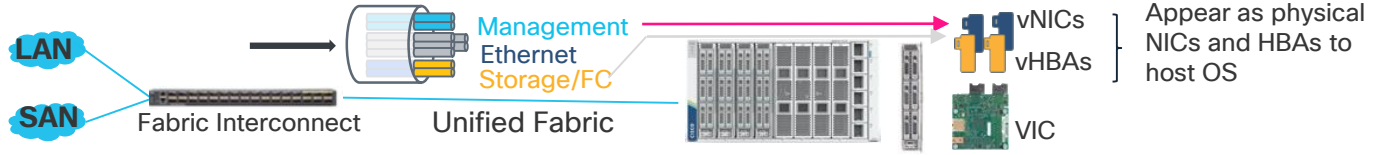

Stateless System, Unified Mgmt

Simplicity, Ease of deployment, Profile for Server/ Network/ Storage/ Mgmt, Upgrade/Replace/Operate at scale, CPU gen upgrade

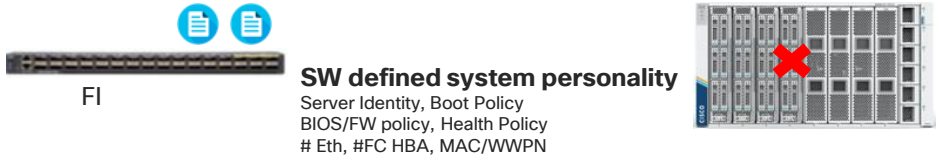

Port Consolidation, Local Switching (Traffic Seg)

Plug-n-play fabric, Better collab between NW/Storage/Server admin, Highly Scalable (160 servers), Flexible BW, Optimized E-W traffic

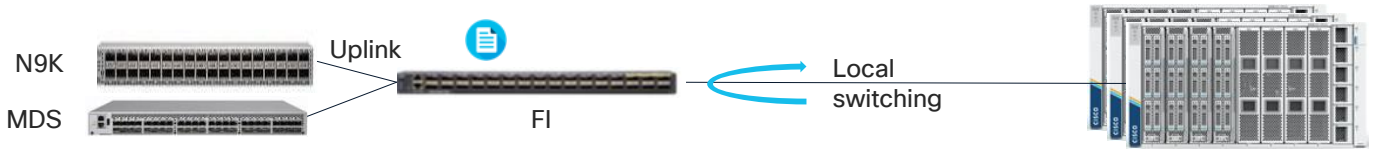

Resiliency and Performance For Mission Critical Apps

TCO savings, Consistent perf, No downtime – failures, upgrades
180+ Perf Benchmarks , Future proofed for compute evolution

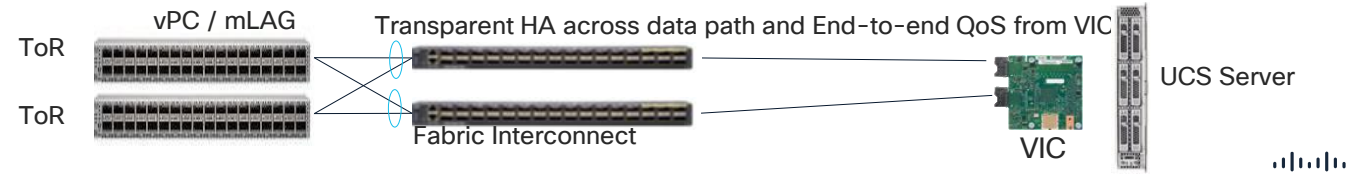

UCS Fabric Innovation Cadence & Leadership

FYI

UCS Fabric Innovation Cadence

2009	2011	2016	2018 / 2019	2020 / 2021	2022
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	**1st with..** E2E 100G Fabric for blades Pushing Performance Envelope

	1st Gen	2nd Gen	3rd Gen	4th Gen		5th Gen
Fabric Interconnect	 FI 6120 / 6140	 6248 / 6296	 6332 / 6332-16UP 6324	 6454	 64108	 6536
IOM / IFM / FEX	 IOM 2104	 FEX 2232 IOM 2204/ 2208	 FEX 2348 IOM 2304	 IOM 2408	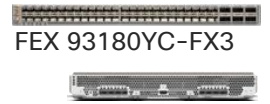 FEX 93180YC-FX3 IFM X-9108 25G	 IFM X-9108 100G
VIC	 M81KR, P81E (10G)	 1200 Series (10/40G)	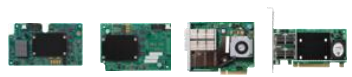 1300 Series (10/40G)	 1400 / 14000 Series (10G/25G/40G/100G)	 14425 / 14825	 15000 Series (10/25/50/100/200G)
	Converged FC, Eth 10G Server/Uplink 8G FC uplink PCIe Gen1 128 PCIe device Hypervisor Bypass for ESX, KVM (VM-FEX)	10G Server/Uplink 8G FC uplink PCIe Gen2 256 PCIe device Flow Classification 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 200G per X210c	100G Server and Uplink 32G FC uplink PCIe Gen4 Geneve perf RoCEv2 perf ECN, PTPv2 SR-IOV , SIOV, RSSv2 16K Rx ring-size

UCS Fabric SW Feature Roadmap

CY23

- X-series UCSM support
- Rx no–buffer counters
- FC/Eth congestion detection on FI ports

- RSSv2
- Netflow
- SR-IOV for ESXi in UCSM/IMM
- VIC QinQ Tunneling
- Cisco Patch Panel

CY24

- Windows-2025 support
- 50G/200G support on VIC
- Port-channel pinning
- FPIN

- NetQueue+ENS
- SR-IOV for Linux
- Windows poll mode
- MACsec on FI uplinks
- ERSPAN

CY25

- NXOS 10.x update
- VIC driver upstreaming
- PTP support in FI
- NFS w/ RDMA

- Additional FEC on FI for 25/100G
- Trunk in physical nic-mode
- IPv6 iSCSI boot
- FC/Eth congestion detection on VIC & IOM/IFM
- Enhanced Topology view

6th Generation
UCS Fabric Interconnects

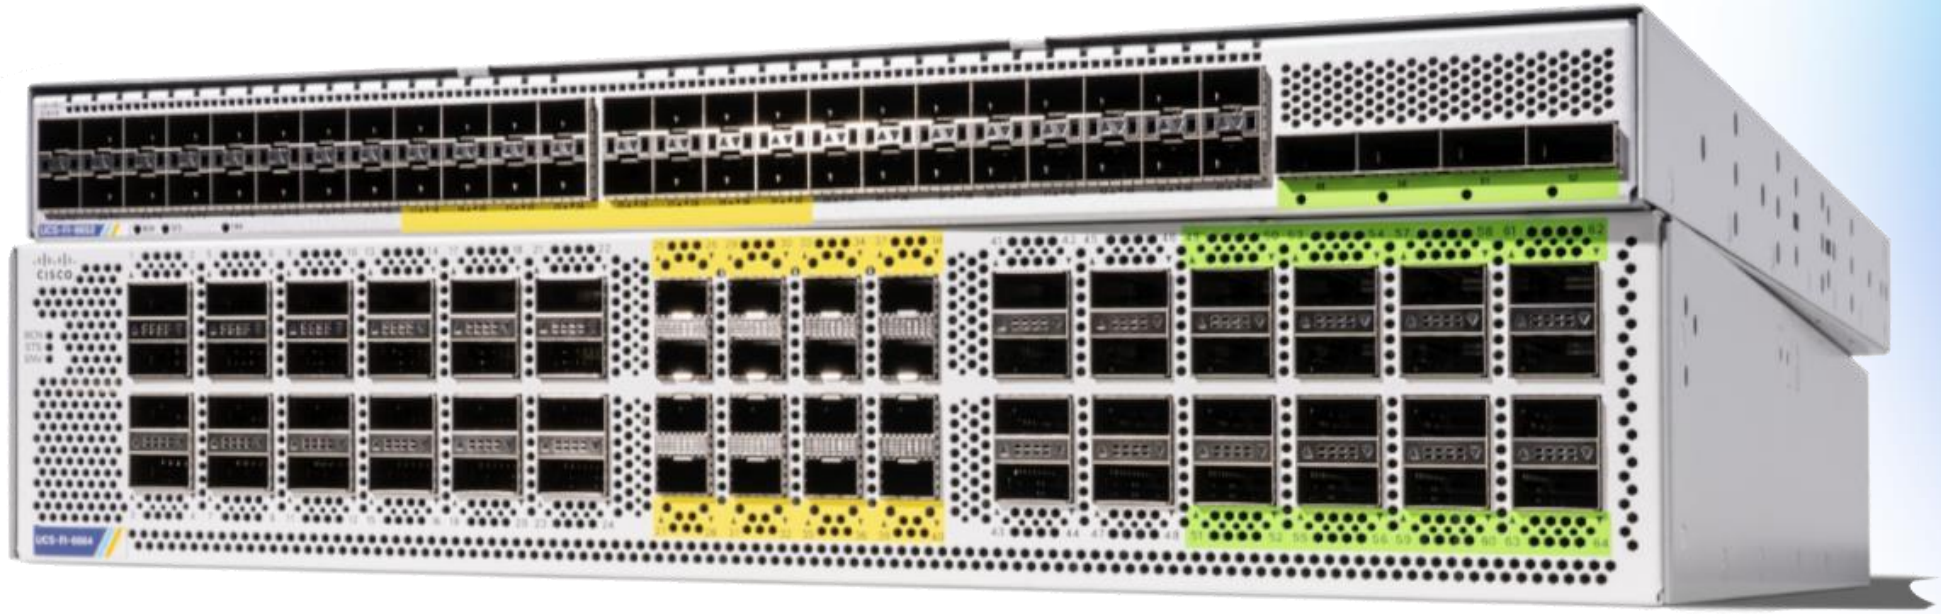

6th Gen UCS Fabric Interconnect 6664

11.65 T Bandwitdh

100G Ethernet

64G Fiber Channel

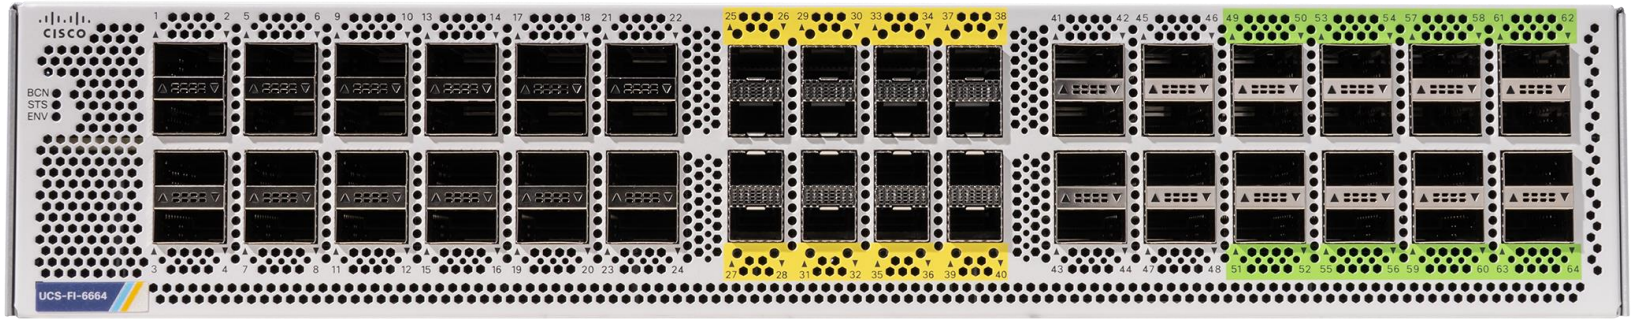

Features:

- 2 Rack Unit/High Density
- 64 Ports
  - 48 QSFP 40/100G Ethernet
  - 16 SFP 10/25G Ethernet or 16/32/64G Fiber Channel
- Supports both UCSM and IMM
- MACSec on Uplink Ports
- Perpetual Software License
- FC end-host and switch-mode

Interoperability:

- X series and C series server connectivity
- Support for M6, M7, M8+ generation Intel/AMD servers
- Support for Cisco MDS and Brocade SAN switch connectivity
- Support for IFM-100G, IFM-25G, Nexus 93180YC-FX3 FEX connectivity

Fabric Interconnect 6664 (Front)

Fabric Interconnect 6664 Port Layout (Back)

40/100G Eth/FCoE
10/25GbE by QSA or QSA28
QSFP Ports 1-24

10/25G Eth/FCoE or
16/32/64G FC
SFP Ports 25-40

40/100G Eth/FCoE
10/25GbE by QSA or QSA28
QSFP Ports 41-64

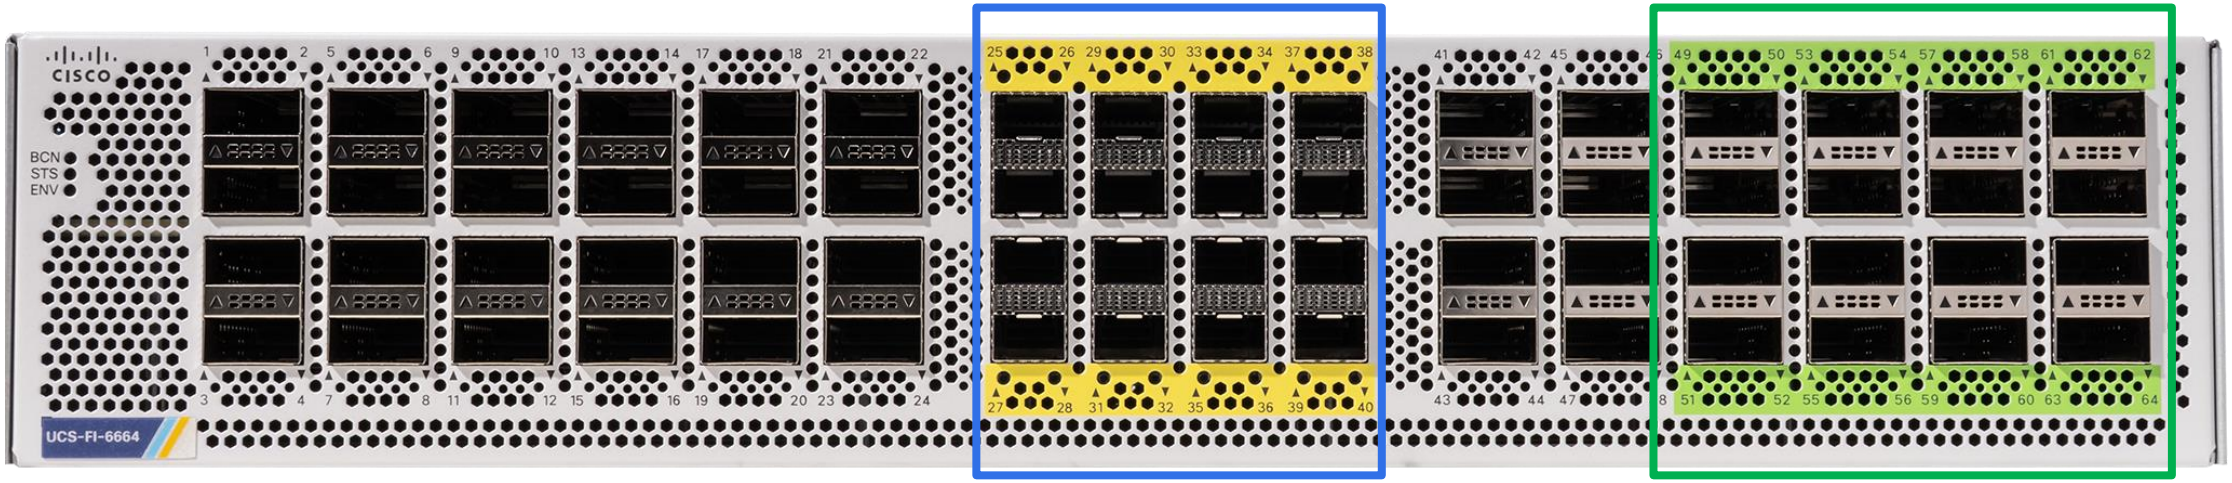

Unified Ports
(FC/Eth/FCoE)

MACsec capable
Recommended Uplinks

6th Gen UCS Fabric Interconnect 6652

8.44 Tbps Bandwidth

6454 Refresh

64G Fiber Channel

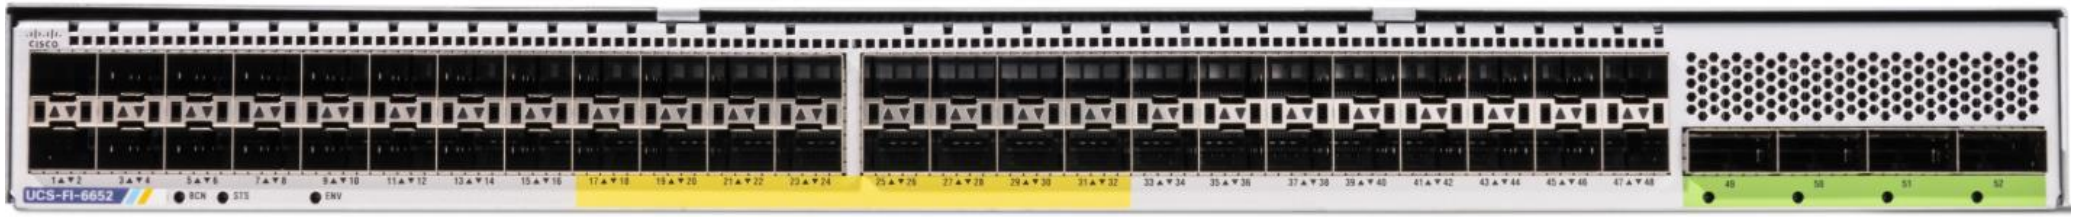

Features:

- 1 Rack Unit
- 52 Ports
  - 32 SFP 10/25G Ethernet
  - 16 SFP 10/25G Ethernet or 16/32/64G Fiber Channel
  - 4 QSFP-DD 40/100/200/400GbE Ethernet Uplinks
- Supports both UCSM and IMM
- MACSec on Uplink Ports
- Perpetual Software License
- FC end-host and switch-mode

Interoperability:

- X/B series (Blade) and C series (Rack) server connectivity
- Support for M5, M6, M7, M8+ generation Intel/AMD servers
- Support for Cisco MDS and Brocade SAN switch connectivity
- Support for IOM-2408, IFM-25G, Nexus 93180YC-FX3 FEX connectivity

Fabric Interconnect 6652 (Front)

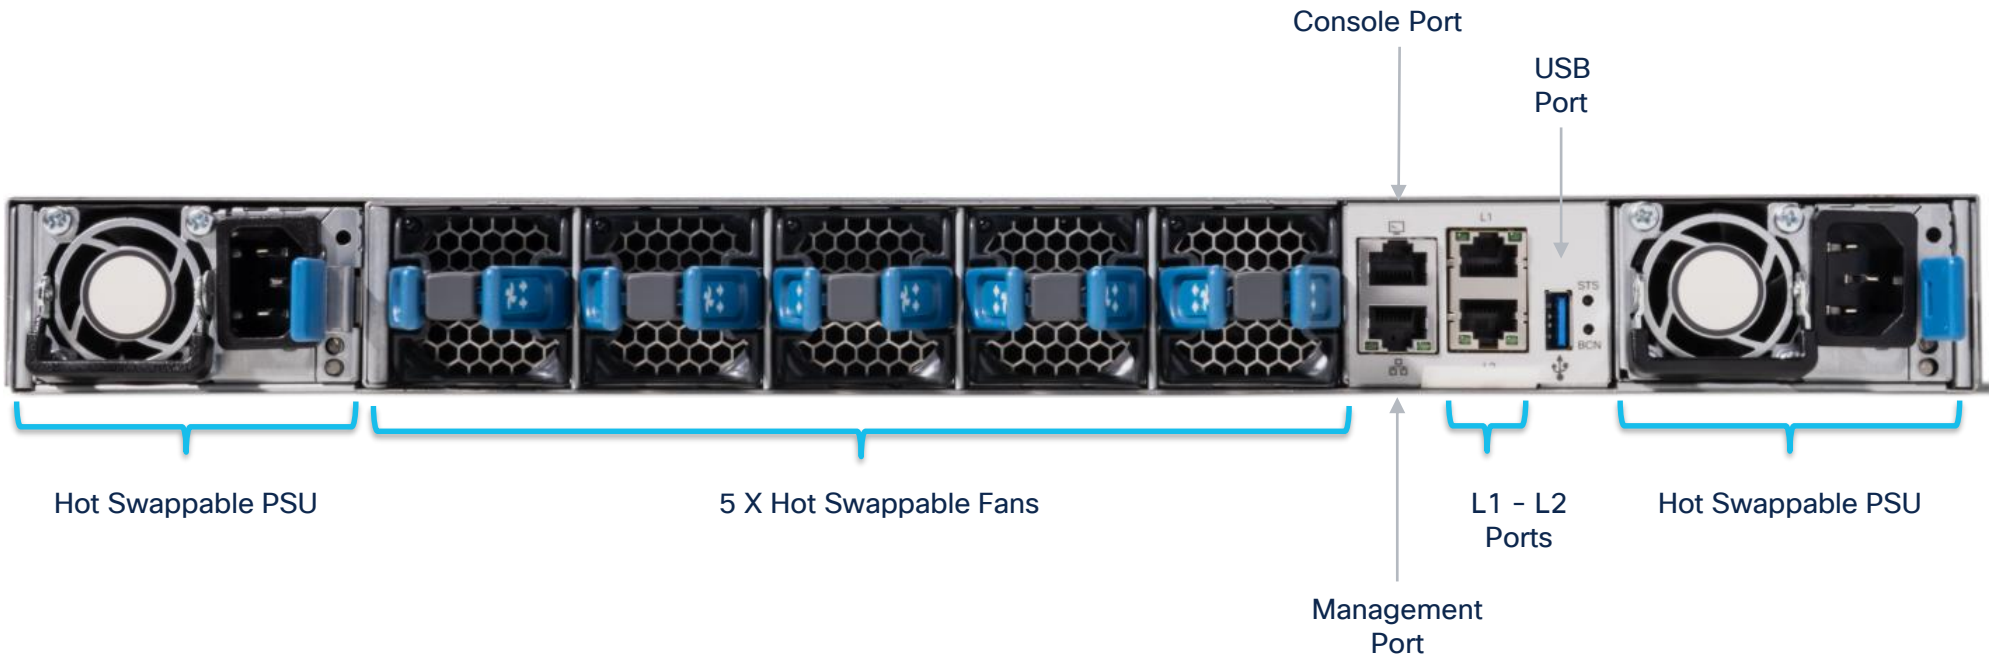

Fabric Interconnect 6652 Port Layout (Back)

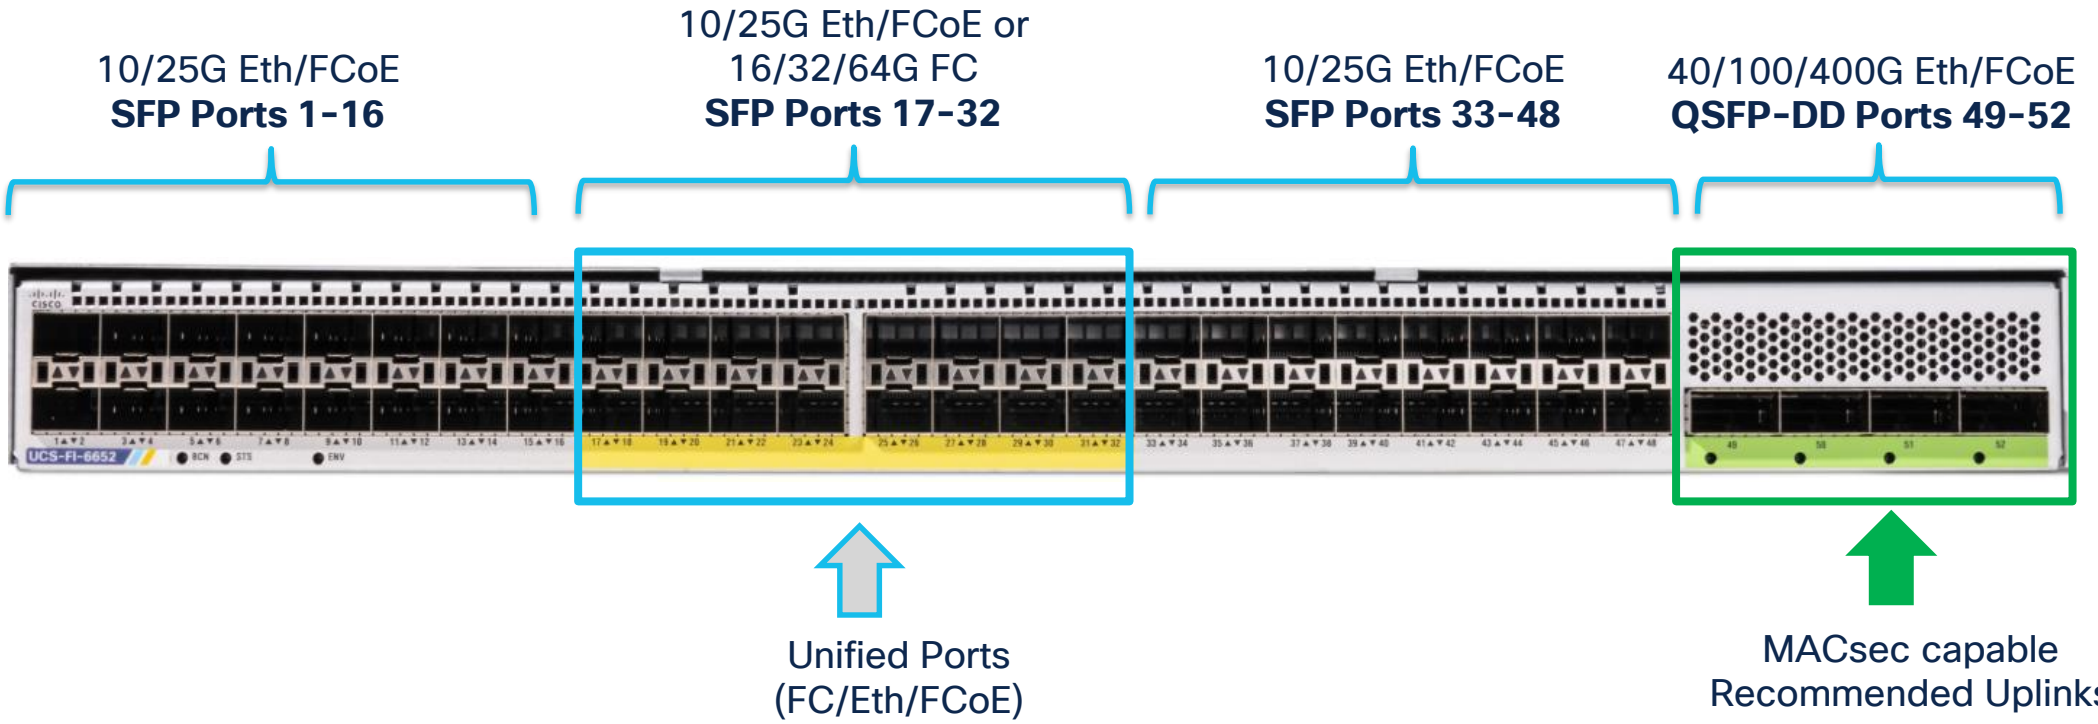

UCS FI 6664 Port Layout

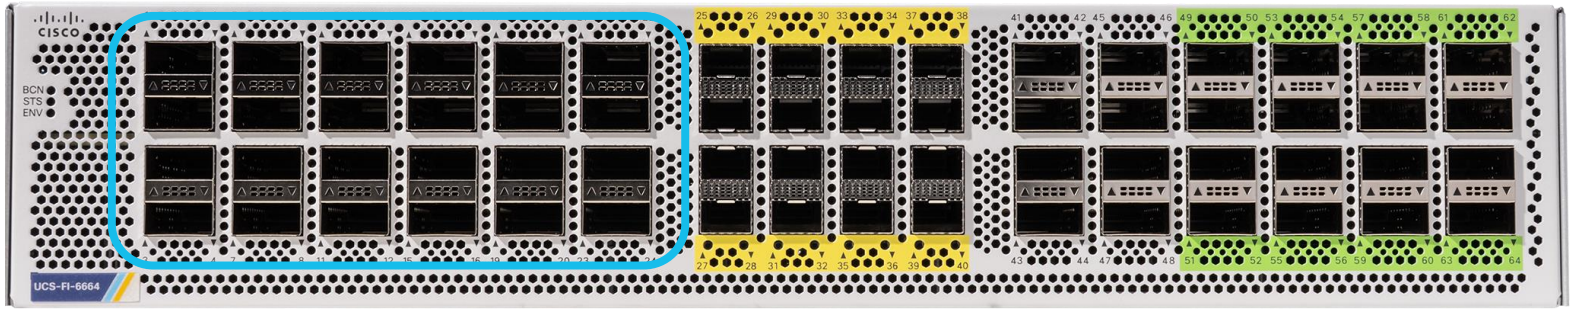

Port No.	Interface Type				Protocol Support			Port Role					
	10/25 Gigabit Ethernet	10/25 Gigabit Ethernet QSA or QSA28 only	40/100 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/40/ 100G	Uplink: Ethernet 10/25/4 0/100G	Uplink: Fibre Channel 16/32/6 4G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 24		X	X		X		X	X	X		X	X	
25 to 40	X			X	X	X	X	X	X	X	X	X	X
41 to 48		X	X		X		X	X	X		X	X	
49 to 64		X	X		X		X	X	X		X	X	

Note: 1G connectivity and Port Breakouts are not supported

UCS FI 6664 Port Layout

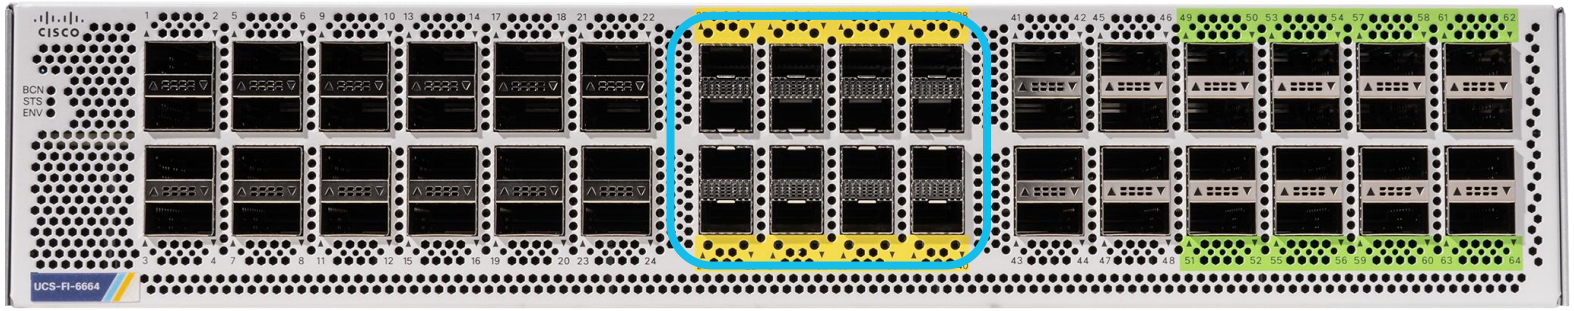

Port No.	Interface Type				Protocol Support			Port Role					
	10/25 Gigabit Ethernet	10/25 Gigabit Ethernet QSA or QSA28 only	40/100 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/40/ 100G	Uplink: Ethernet 10/25/40/100G	Uplink: Fibre Channel 16/32/64G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 24		X	X		X		X	X	X		X	X	
25 to 40	X			X	X	X	X	X	X	X	X	X	X
41 to 48		X	X		X		X	X	X		X	X	
49 to 64		X	X		X		X	X	X		X	X	

Note: 1G connectivity and Port Breakouts are not supported

UCS FI 6664 Port Layout

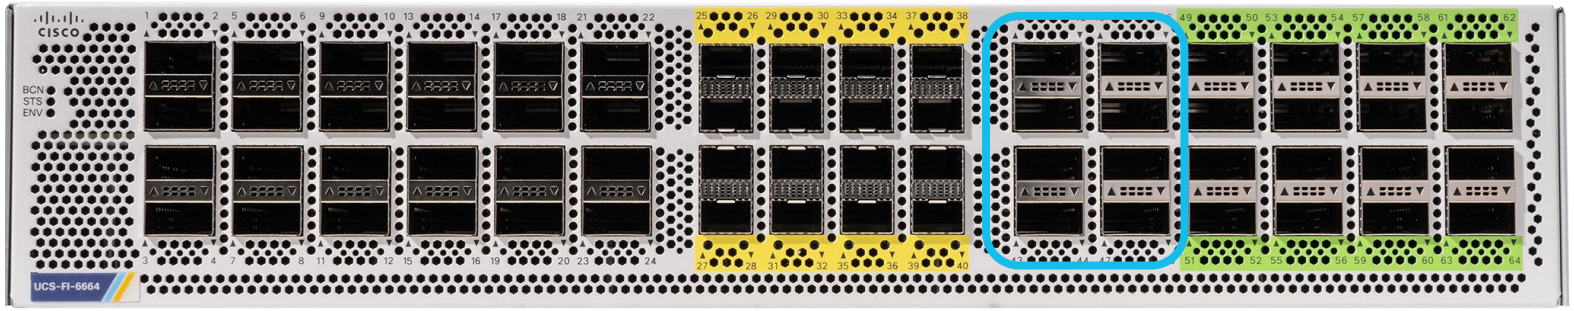

Port No.	Interface Type				Protocol Support			Port Role					
	10/25 Gigabit Ethernet	10/25 Gigabit Ethernet QSA or QSA28 only	40/100 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/40/ 100G	Uplink: Ethernet 10/25/40/100G	Uplink: Fibre Channel 16/32/64G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 24		X	X		X		X	X	X		X	X	
25 to 40	X			X	X	X	X	X	X	X	X	X	X
41 to 48		X	X		X		X	X	X		X	X	
49 to 64		X	X		X		X	X	X		X	X	

Note: 1G connectivity and Port Breakouts are not supported

UCS FI 6664 Port Layout

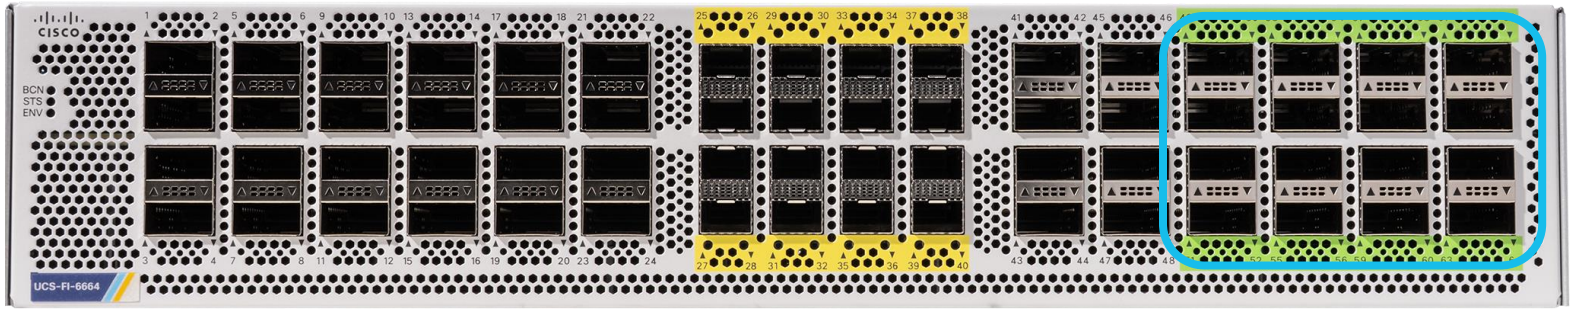

Port No.	Interface Type				Protocol Support			Port Role					
	10/25 Gigabit Ethernet	10/25 Gigabit Ethernet QSA or QSA28 only	40/100 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/40/ 100G	Uplink: Ethernet 10/25/4 0/100G	Uplink: Fibre Channel 16/32/6 4G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 24		X	X		X		X	X	X		X	X	
25 to 40	X			X	X	X	X	X	X	X	X	X	X
41 to 48		X	X		X		X	X	X		X	X	
49 to 64		X	X		X		X	X	X		X	X	

Note: 1G connectivity and Port Breakouts are not supported

FI 6664 Support Matrix in IMM/UCSM

X-Series

Server	M6	M7	M8
X210c	X	X	X
X215c	-	-	X
X410c	-	x	-

VIC	IFM 25G (X-Series)	IFM 100G (X-Series)
15420	x	x
15422	x	x
15230	x	x
15231	x	x
14425	x	x
14825	x	x

Note: B-Series server support is not available with FI 6664

C-Series

Server	M6	M7	M8
C220	x	x	x
C240	x	x	x
C225	x	-	x
C245	x	-	x

FEX Support
Nexus 93180YC-FX3 in FEX mode

VIC
1455
1467
1477
1495
15425
15427
15428
15235
15237
15238

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

UCS FI 6652 Port Layout

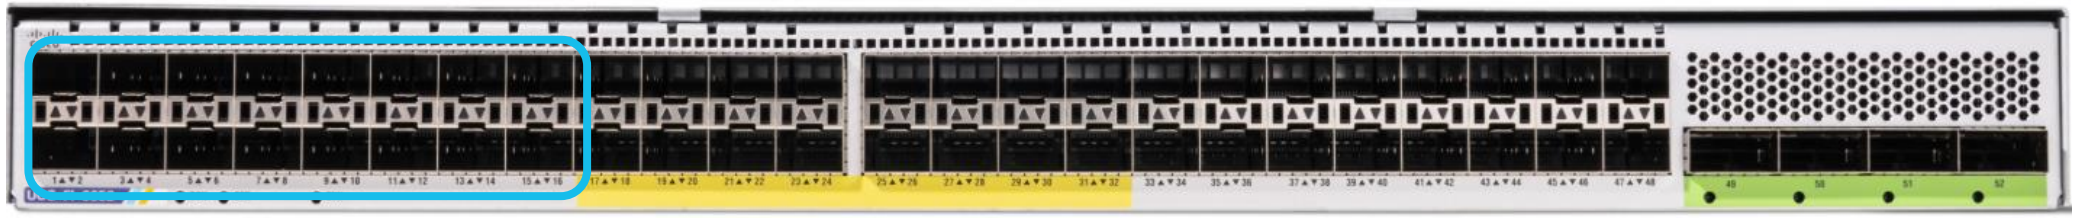

Port No.	Interface Type				Protocol Support			Port Role						
	1 Gigabit Ethernet QSA	10/25 Gigabit Ethernet	40/100/200/400 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/50 G	Uplink: Ethernet 10/25G	Uplink: Ethernet 40/100/200/400G	Uplink: Fibre Channel 16/32/64G	Uplink: FCoE	Appliance port 10/25/40/100G	Storage port (Fibre Channel)
1 to 16		X			X		X	X	X			X	X	
17 to 32		X		X	X	X	X	X	X		X	X	X	X
33 to 48		X			X		X	X	X			X	X	
49 to 52		X	X		X		X		X	X		X	X	

UCS FI 6652 Port Layout

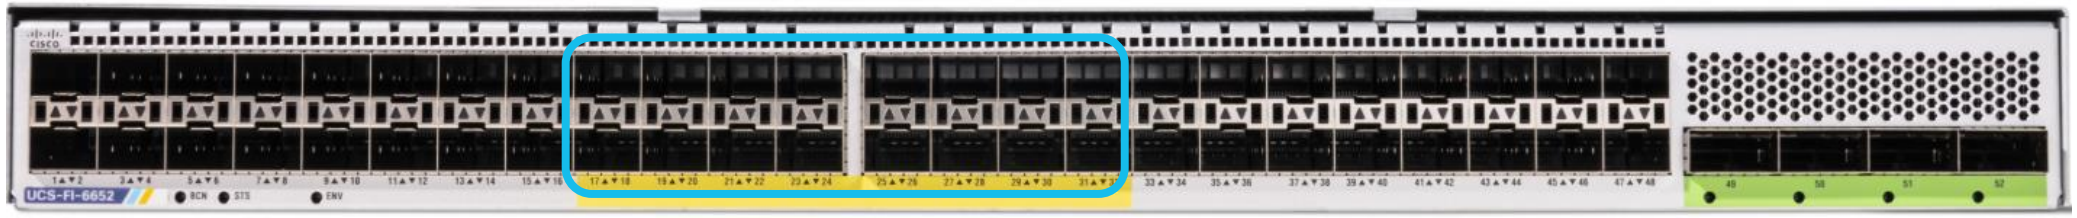

Port No.	Interface Type				Protocol Support			Port Role						
	1 Gigabit Ethernet QSA	10/25 Gigabit Ethernet	40/100/ 200/400 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/50 G	Uplink: Ethernet 10/25G	Uplink: Ethernet 40/100/ 200/ 400G	Uplink: Fibre Channel 16/32/ 64G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 16		X			X		X	X	X			X	X	
17 to 32		X		X	X	X	X	X	X		X	X	X	X
33 to 48		X			X		X	X	X			X	X	
49 to 52		X	X		X		X		X	X		X	X	

UCS FI 6652 Port Layout

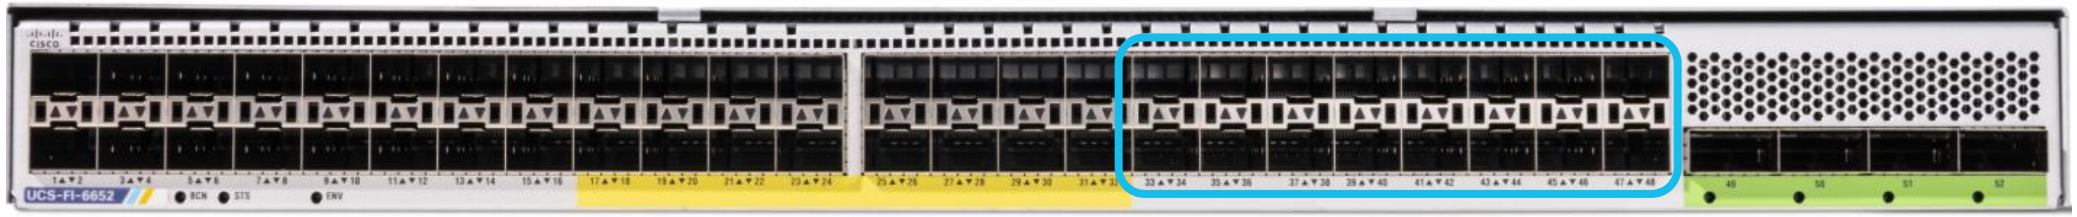

Port No.	Interface Type				Protocol Support			Port Role						
	1 Gigabit Ethernet QSA	10/25 Gigabit Ethernet	40/100/ 200/400 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/50 G	Uplink: Ethernet 10/25G	Uplink: Ethernet 40/100/ 200/ 400G	Uplink: Fibre Channel 16/32/ 64G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 16		X			X		X	X	X			X	X	
17 to 32		X		X	X	X	X	X	X		X	X	X	X
33 to 48		X			X		X	X	X			X	X	
49 to 52		X	X		X		X		X	X		X	X	

UCS FI 6652 Port Layout

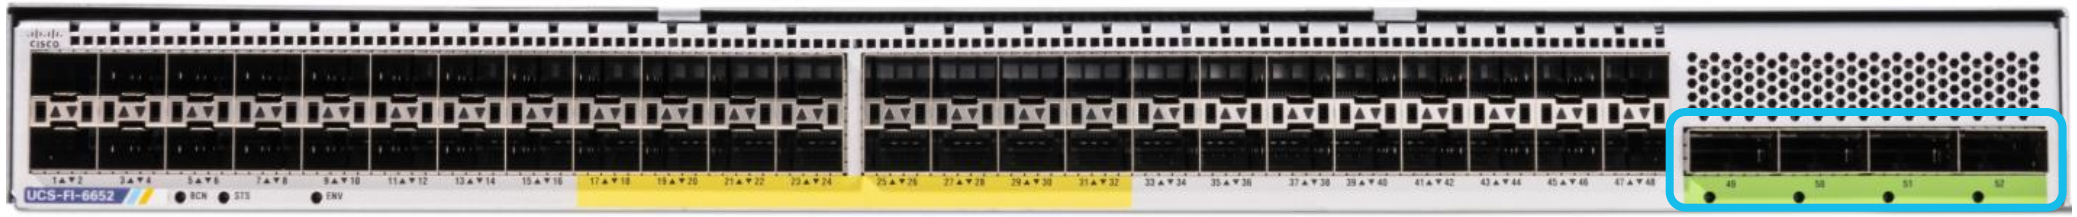

Port No.	Interface Type				Protocol Support			Port Role						
	1 Gigabit Ethernet QSA	10/25 Gigabit Ethernet	40/100/ 200/400 Gigabit Ethernet	16/32/64 Gbps FC	Ethernet	Fibre Channel	FCoE	Server 25/50 G	Uplink: Ethernet 10/25G	Uplink: Ethernet 40/100/ 200/ 400G	Uplink: Fibre Channel 16/32/ 64G	Uplink: FCoE	Appliance port 10/25/40/ 100G	Storage port (Fibre Channel)
1 to 16		X			X		X	X	X			X	X	
17 to 32		X		X	X	X	X	X	X		X	X	X	X
33 to 48		X			X		X	X	X			X	X	
49 to 52		X	X		X		X		X	X		X	X	

FI 6652 Support Matrix in IMM/UCSM

X-Series & B-Series

Server	M5	M6	M7	M8
B200	X	x	-	-
B480	X	-	-	-
X210c	-	X	x	x
X215c	-	-	-	x
X410c	-	-	x	-

VIC	IOM 2408 (B-Series)	IFM 25G (X-Series)
15420	-	x
15422	-	x
15230	-	x
15231	-	x
15411	x	-
1440	x	-
1440+PE	x	-
1480	x	-
14425	-	x
14825	-	x

C-Series

Server	M5	M6	M7	M8
C220	x	x	x	x
C240	x	x	x	x
C225	-	x	-	x
C245	-	x	-	x

VIC
1455
1457
1467
1477
1495
1497
15235
15237
15238
15425
15427
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 w/ 4x25G
93180YC-FX3 FEX port server connectivity at 25G

Hardware Comparison

	6300		6400		6536	6600
	6332	6332-16	6454	64108	6536	6664
ASIC	Trident-2* Donner	Trident-2* Tiburon Donner	Homewood	Heavenly	Heavenly	Ararat
Bandwith	2.56 Tbps	2.82 Tbps	3.82 Tbps	7.42 Tbps	7.42 Tbps	11.65 Tbps
Switching Type	Cut-through	Cut-through	Cut-through	Cut-through	Cut-through	Cut-through
IFM	-	-	X9108-IFM-25G	X9108-IFM-25G	X9108-IFM-25G / X9108-IFM-100G	X9108-IFM-25G / X9108-IFM-100G
IOM	2204/2208/ 2304	2204/2208/ 2304	2204/2208/ 2408	2204/2208/ 2408	2408 / 2304	-
FEX	2232/2348	2232/2348	2232/ 93180YC-FX3	2232/ 93180YC-FX3	93180YC-FX3	93180YC-FX3
Eth Port	32x40G	32x40G +16x10G	48x25G + 6x40/100G	96x25G + 12x40/100G	36x 40/100G	48x 40/100G, 16x 10/25G
FC Port	-	4/8/16G	8/16/32G	8/16/32G	8/16/32G	16/32/64G
UCS Unified Management	UCSM	UCSM	Intersight / UCSM	Intersight / UCSM	Intersight /UCSM	Intersight /UCSM

Feature Comparison

Feature	FI 6300 Series	FI 6400 Series	FI 6536	FI 6664
NXOS Release	7.x	9.x , 10.x*	9.x, 10.x*	10.x
End host & switch mode	Supported	Supported	Supported	Supported
FC/FCoE NPV & switch mode	Supported	Supported	Supported	Supported
NetFlow	Supported	Supported	Supported	Supported
PVLAN	Supported	Supported	Supported	Supported
Port-security	Supported	Supported	Supported	Supported
IOM	2200 / 2300 series	220x / 2408	2408 / 2304	-
IFM	-	IFM-25G	IFM-25G/ IFM-100G	IFM-25G/ IFM-100G
Servers	Upto M6 B-/ C- /S- series	M4/M5/M6/M7/M8 B-/ C-/ S-/ X- series	M4/M5/M6/M7/M8 B-/ C-/ X- / S- series	M6/M7/M8 C-/ X- series
VIC	1200/ 1300/ 1400/15000	1300/ 1400/ 15000	1300/ 1400/ 15000	1400/ 15000
Management	UCSM	UCSM / IMM	UCSM / IMM	UCSM / IMM
NVMeoF, SR-IOV, VIC QinQ Tunneling	Supported*	Supported	Supported	Supported
Congestion monitoring, MACSEC, ERSPAN		Supported	Supported	Supported

FI Configuration Limits

Feature	FI 6300	FI 6400	FI 6536	FI 6600
Unicast MAC address per FI	32,000	92,000	92,000	92,000
Multicast MAC address per FI	7,000	7,000	7,000	7,000
Active VLAN per UCS domain	3,000	3,000	3,000	3,000
STP logical interfaces / VLAN port count	64,000	108,000	108,000	108,000
IGMP groups	4,000	16,000	16,000	16,000
Virtual Interfaces (VIFs)	2,750	1,600	1,600	1,600
Host Bus Adapter (vHBAs)	320	320	320	320
Chassis per UCS domain	20	20	20	20
IP storage appliance per FI	16	16	16	16
VSANs	32	32	32	32

Fabric Interconnect Migration to 6600 Series

FI 6400/6536

Migrate to FI 6600 in UCSM/ IMM

FI 6454

FI 64108

FI 6536

FI 6454

FI 6664

FI 6652

FI 6664 Advantages

- 6664 supports 48 x 40/100G + 16 x 10/25G or 16x 16/32/64G FC
- MACSec PHY on uplink ports
- 10/25/40/100G server-port support
- 10/25/40/100G ethernet uplink ports
- IFM-100G/IFM-25G, N9K-C93180YC-FX3 in FEX mode

FI 6652 Advantages

- 6652 supports 32 x 10/25G + 16 x 10/25G or 16x 16/32/64G FC + 4 X 40/100/200/400G
- MACSec PHY on uplink ports
- 10/25G server-port support
- 40/100/200/400G ethernet uplink ports
- IFM-25G, IOM-2408, N9K-C93180YC-FX3 in FEX mode

Fiber Channel Connectivity

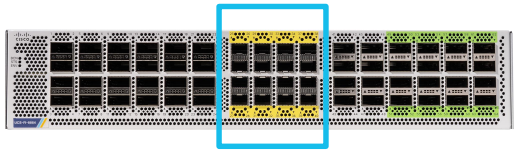

FI 6664 or FI6652

Ports 25-40

+

16/32/64G FC SFP
Cisco PIDs (64G)
- DS-SFP-FC64G-SW
- DS-SFP-FC64G-LW

+

or

MMF or SMF LC cable

+

16/32/64G FC SFP

SAN switch or Storage Array

FI 6664 Supported Forward Error Correction (FEC)

Port	Supported FEC
1-24	100F: RS-FEC, FC-FEC 40G: No FEC QSA: No FEC, FC-FEC, KR-FEC, RS-FEC
25-40	RS-FEC, FC-FEC
41-64	100F: RS-FEC, FC-FEC 40G: No FEC QSA: No FEC, FC-FEC, KR-FEC, RS-FEC

FI 6664 Auto-Negotiate

Auto Negotiate

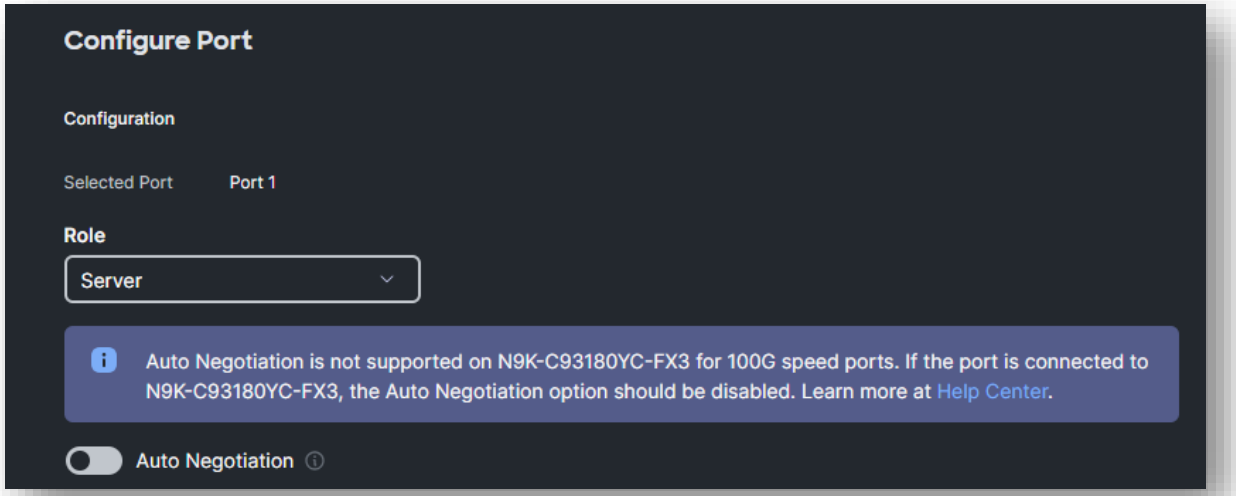

Disable Auto Negotiation when Connected to IFM-100G with Copper Cables.

FI-6664 Breakout Cable

- FI-6664: Does not support ANY Breakout cable
- 4x25G or 4x10G on the QSFP ports is NOT supported
- 1x10G and 1x 25G with QSA/QSA28 with SFP+/SFP28 in 100G QSFP Ports
- 400G/800G Nexus ToR can connect to FI with following Breakout cables.
  - QDD-400G-SR8-S*
  - QDD-400G-SR4.2-BD
  - QDD-4X100G-FR-S
  - QDD-4X100G-LR-S
  - QDD-2X100-SR4-S
  - QDD-4ZQ100-CUxM 400G

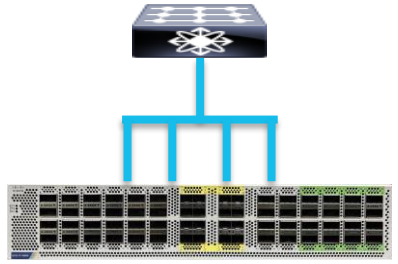

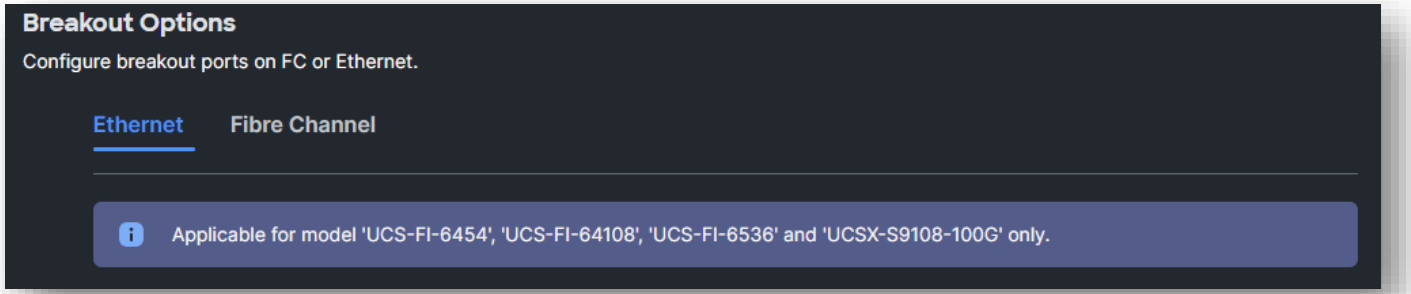

FI-6664 100G End-to-End

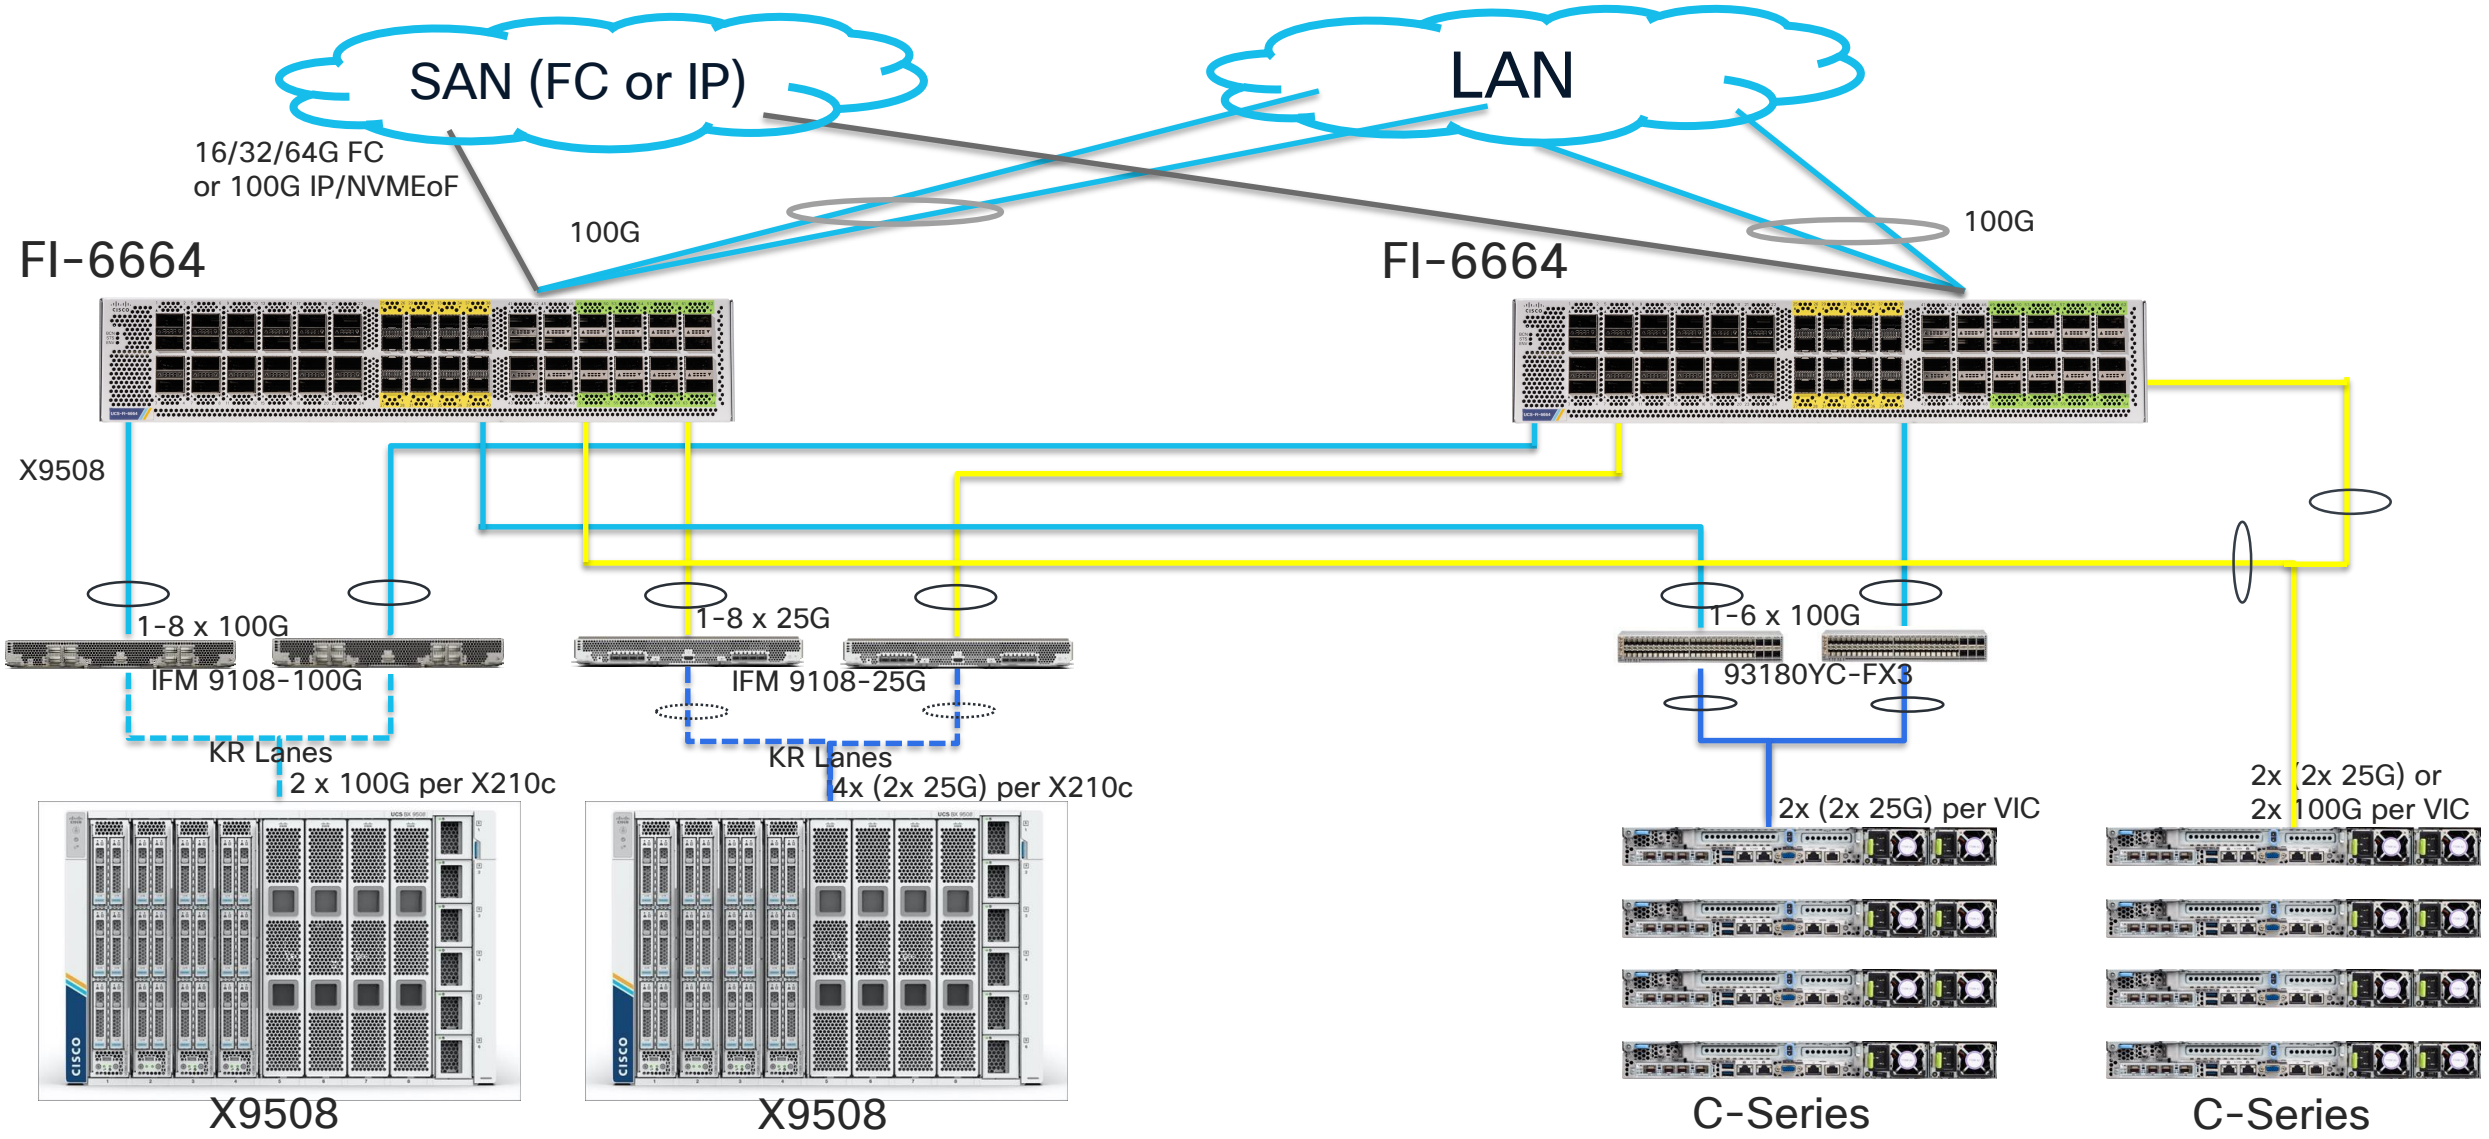

FI 6664 to IFM-100G/25G Connectivity

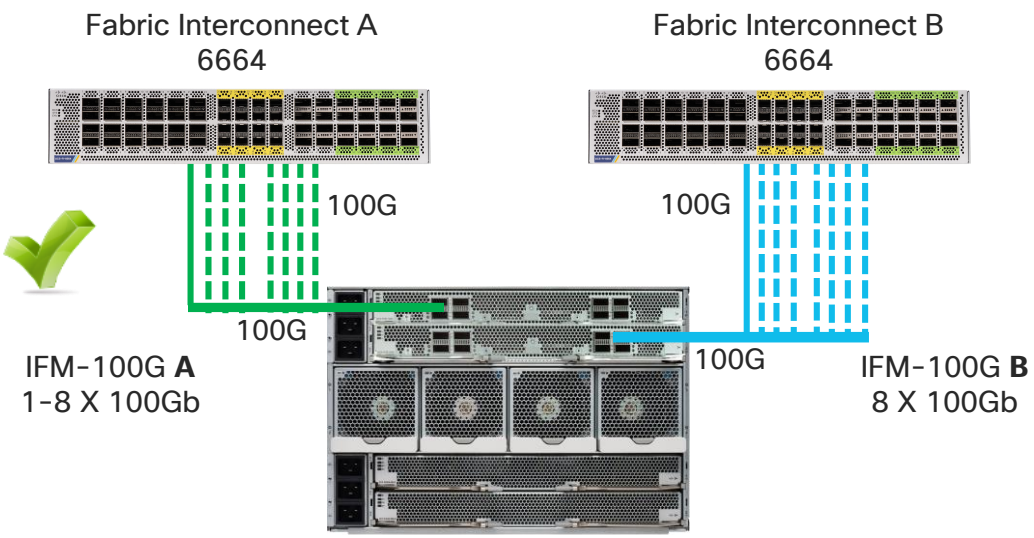

X9108-IFM-100G to 6664 FIs : **Supported**

- 1600G per X9508 chassis
- 100G E2E single-flow per x210c
- 64G E2E FC I/O
- 200G per x210c with 1:1 oversubscription

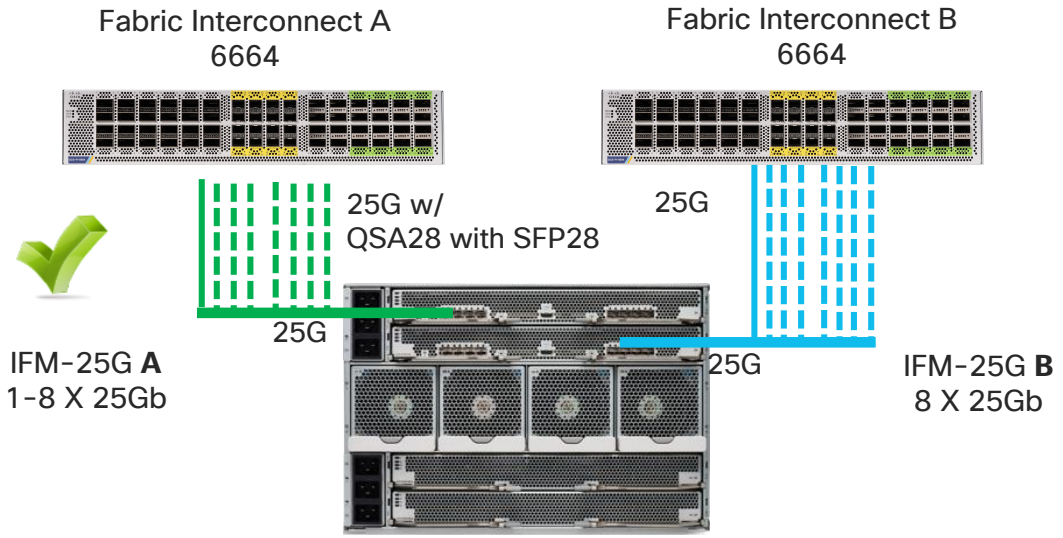

X9108-IFM-25G to 6664 FIs : **Supported**

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

FI 6652- Connection

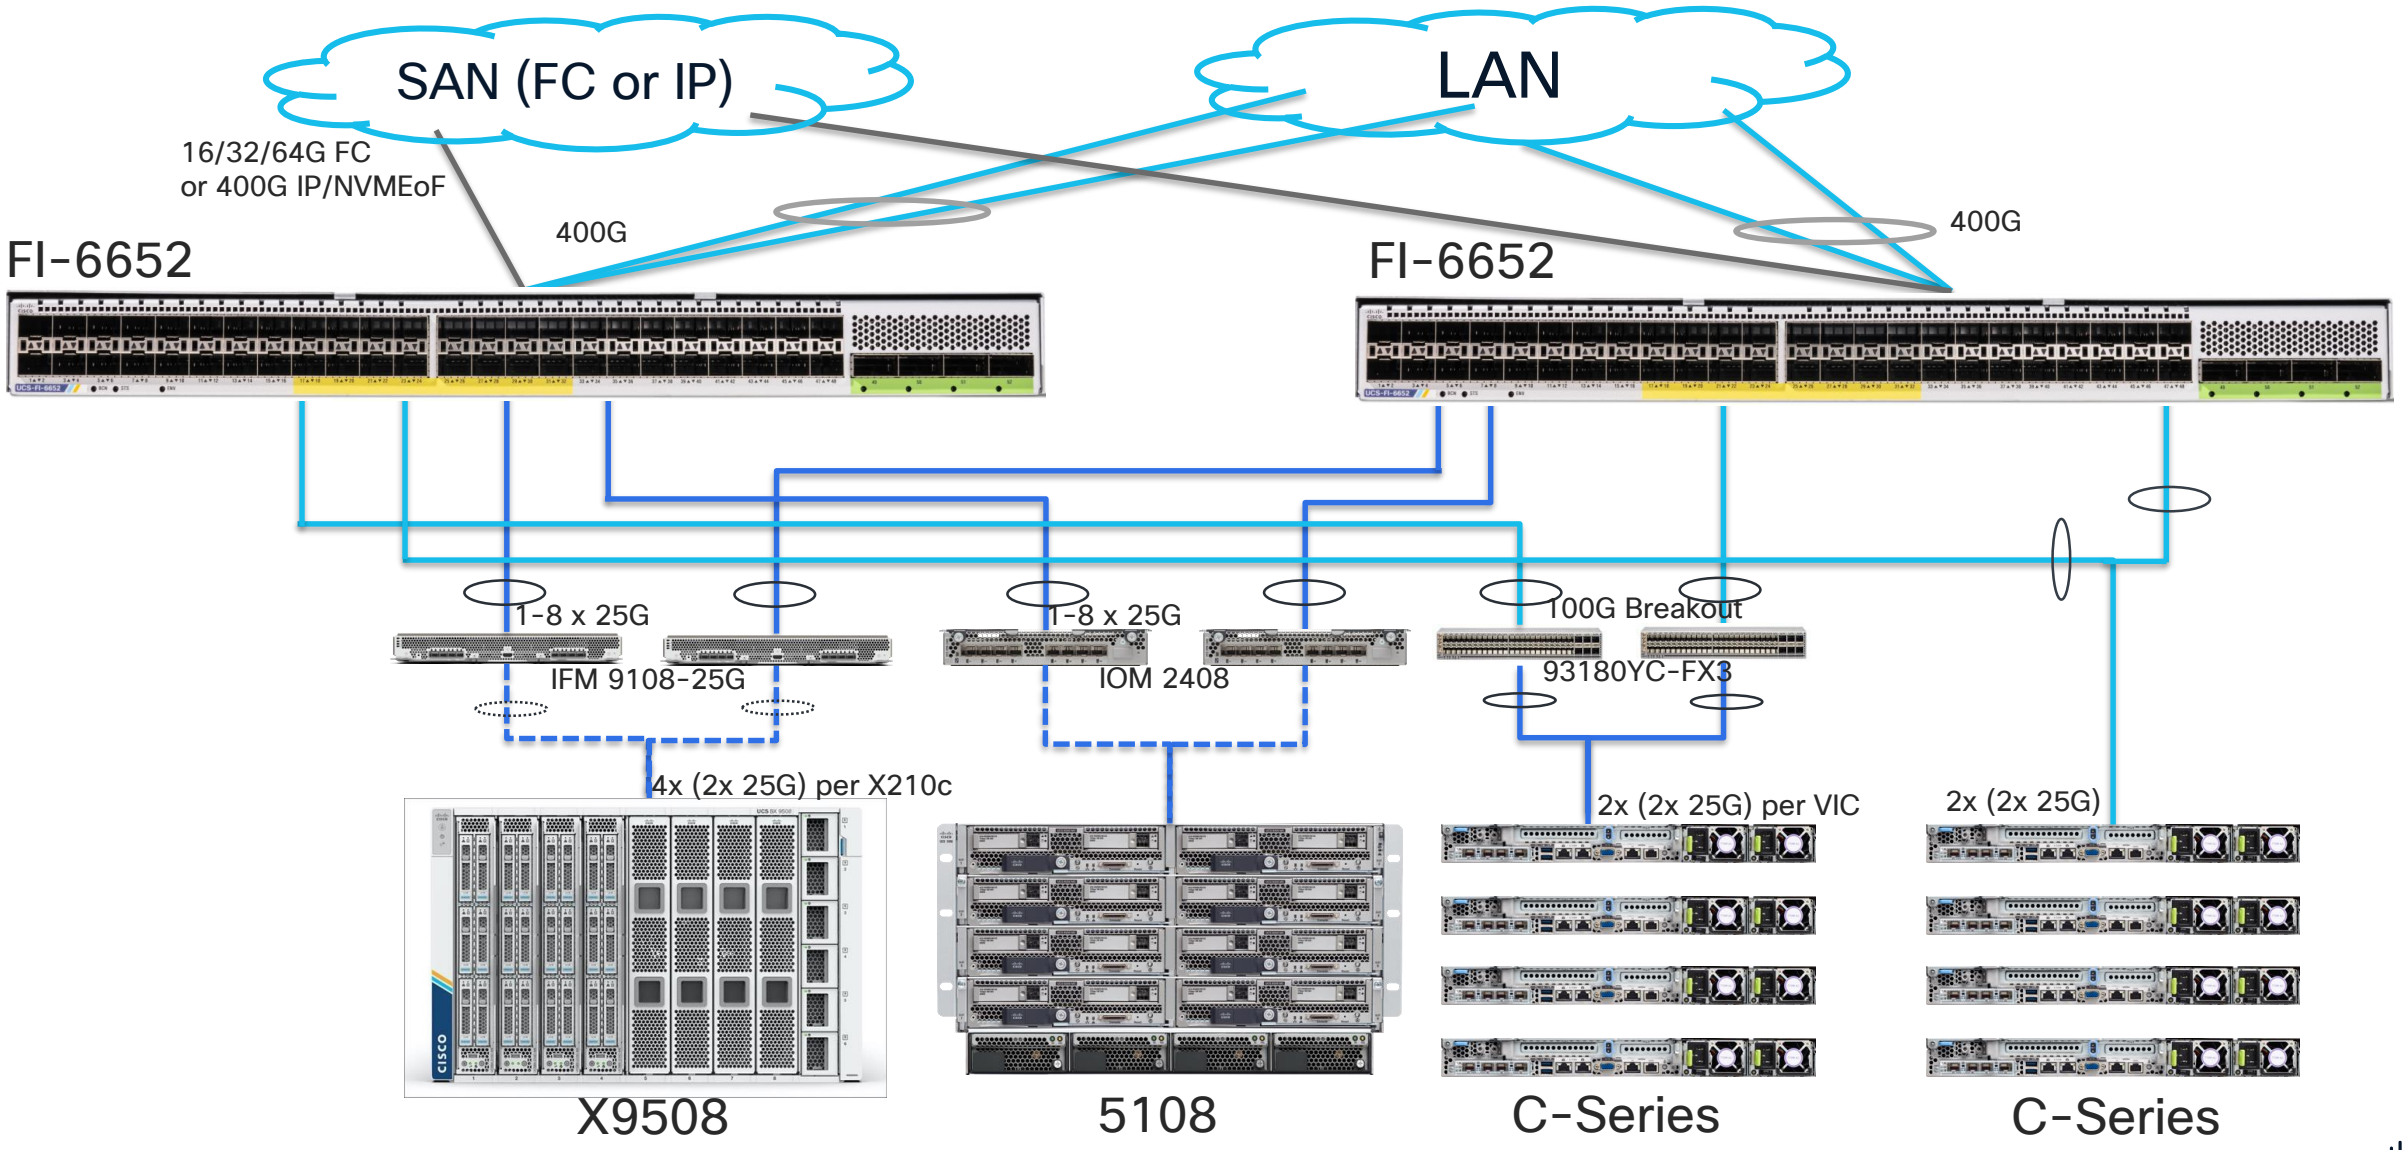

IO Modules to FI 6652 Support

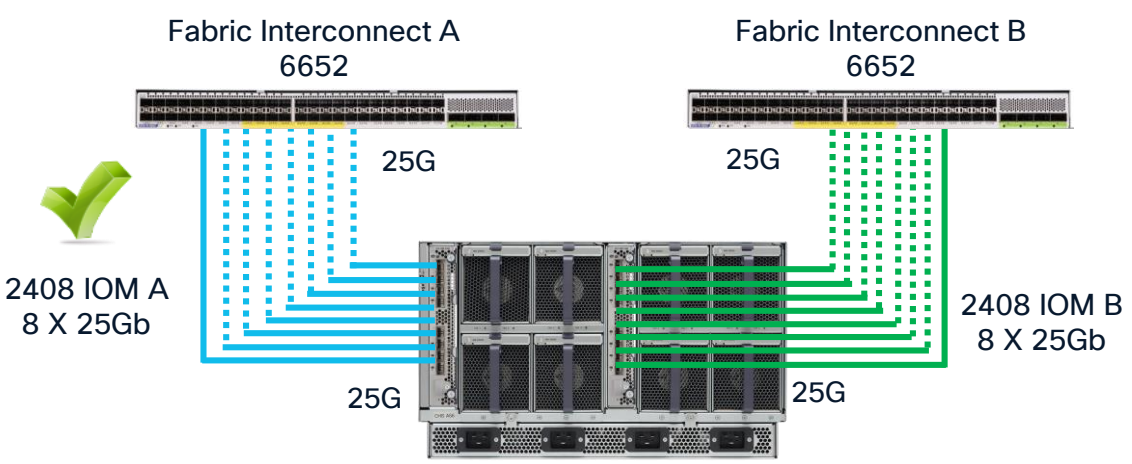

2408 IOMs to 6652 FIs : **Supported**

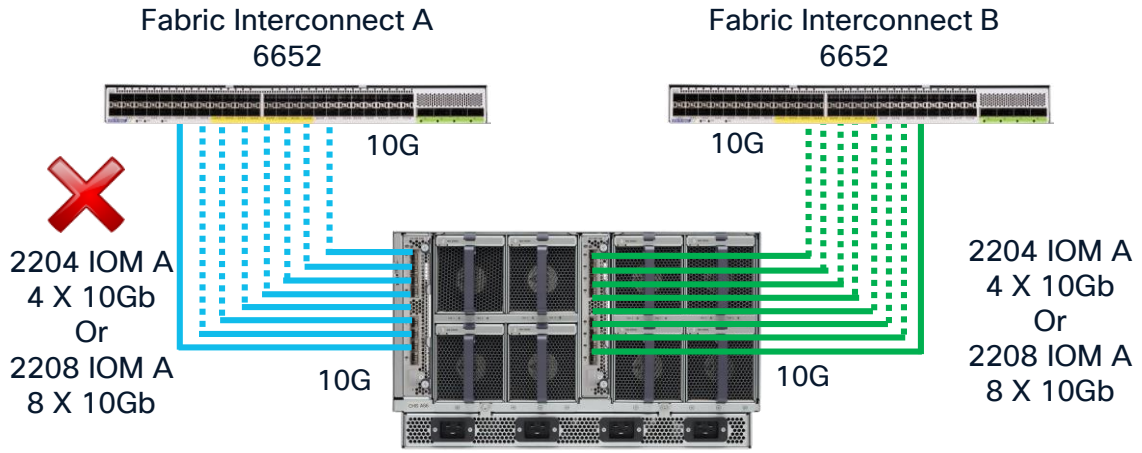

2204 or 2208 IOMs to 6652 FI : **Not Supported**

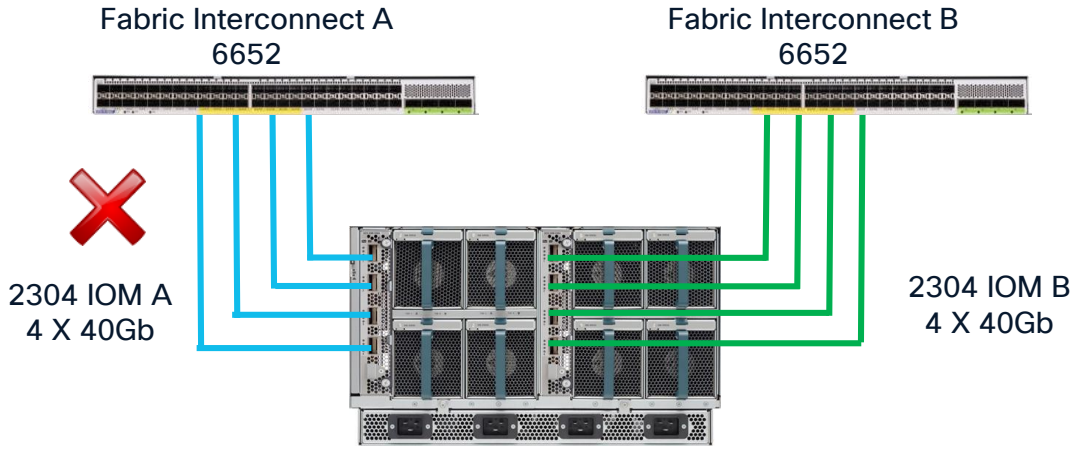

2304 IOMs to 6652 FIs : **Not Supported**

LAN connectivity with 6th Gen FI in end-host mode

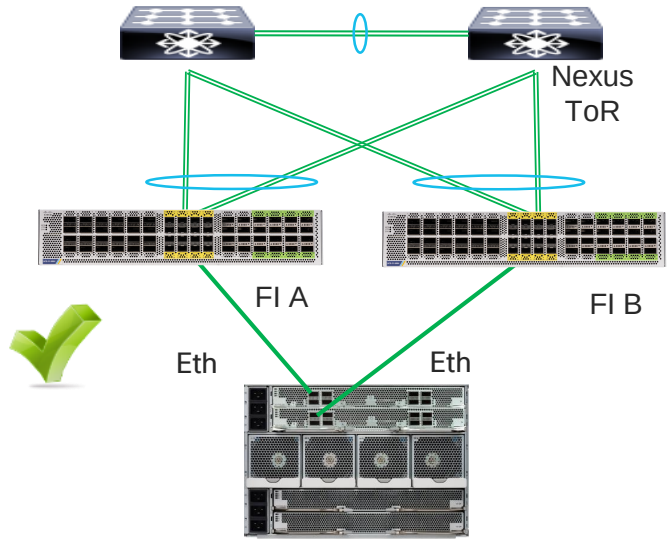

- **Recommended configuration**
- Port-channel from FI to Nexus-9K in VPC or a ToR in multi-chassis trunking (MCT)
- ToR switch ports should be STP edge-port
- Provides redundancy for FI, ToR and uplinks
- FI uplink BW aggregation via port-channel
- Avoids ToR to ToR L2 switching
- ToR fabric could be 3-tier, EVPN-VxLAN, IP-fabric, CLOS

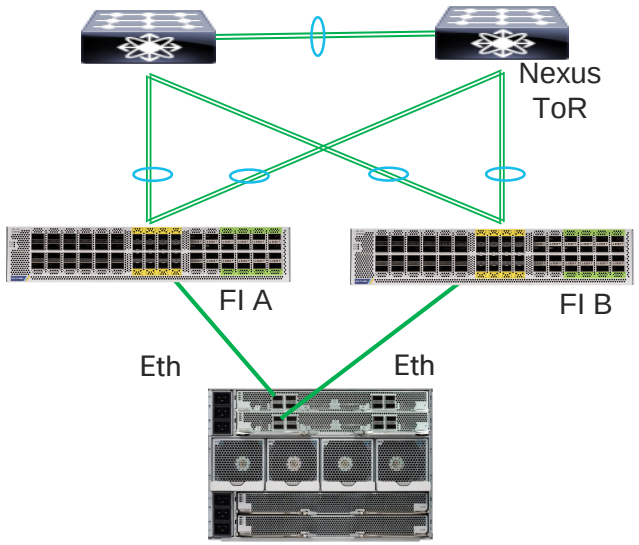

- Recommended with non-VPC/non-MCT
- Supported but not recommended with VPC/MCT
- vNICs gets pinned to one uplink & are distributed across multiple uplinks.
- ToR & uplink redundancy is available but will have re-pinning failover time

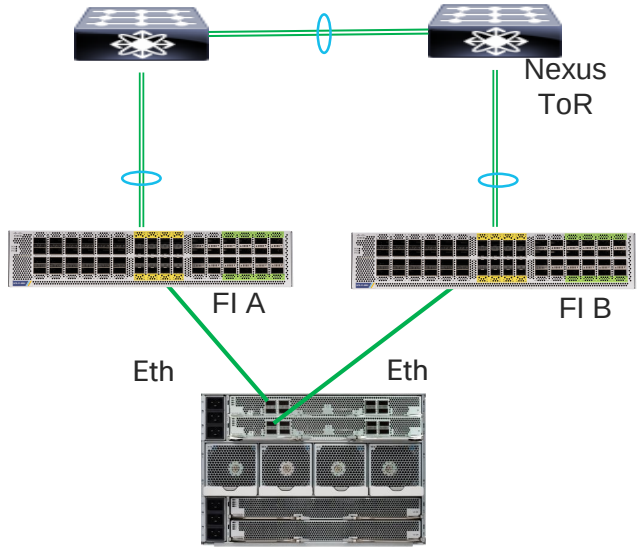

- Supported and **not recommended**
- No ToR level redundancy

FC-SAN connectivity with 6th Gen FI

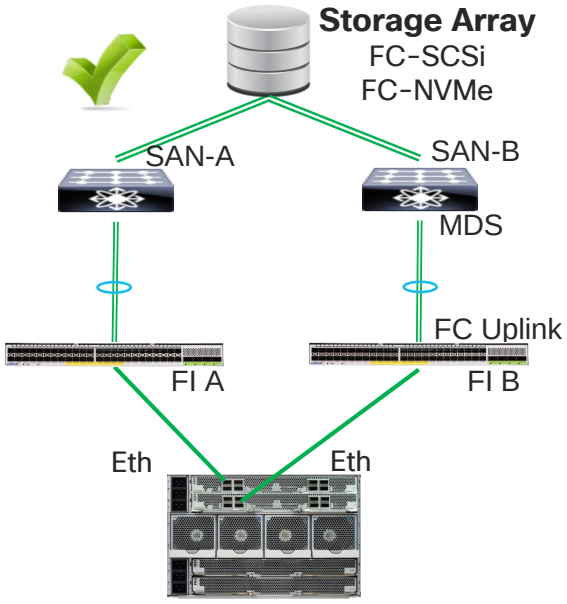

- FC end-host mode (N-port on FI & F-port on MDS)
- Port-channel from FI to MDS
- Port-channel for HA & BW aggregation
- VSAN is carried into MDS SAN with VSAN trunking
- 4 vHBA per server for higher redundancy
- Bigger SAN domain

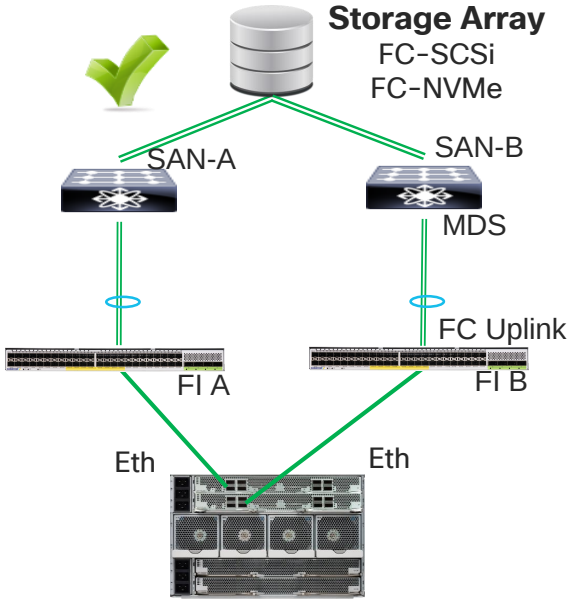

- FC switch-mode (E-port on both)
- Port-channel from FI to MDS
- Provides HA & BW aggregation
- VSAN is carried into MDS SAN
- 4 vHBA per server for higher redundancy
- SAN domain is limited to 255
- Can have storage array connected to FI along with MDS SAN connectivity

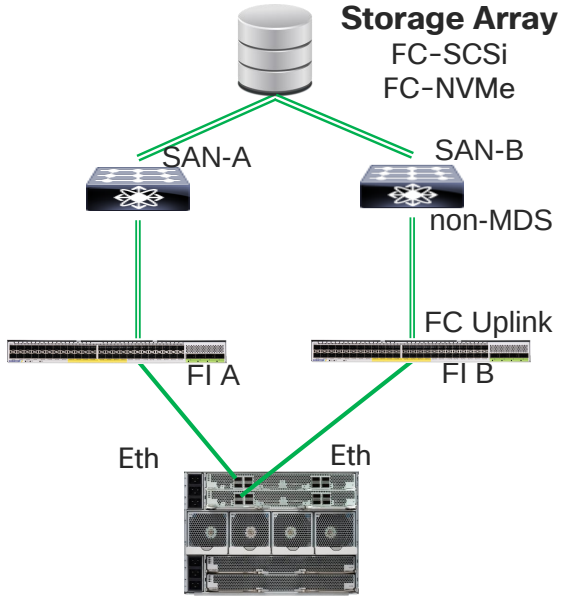

- FC end-host mode
- Recommended for non-MDS
- No port-channel with non-MDS
- VSAN virtualization is not available on non-MDS
- 4 vHBA per server for higher redundancy

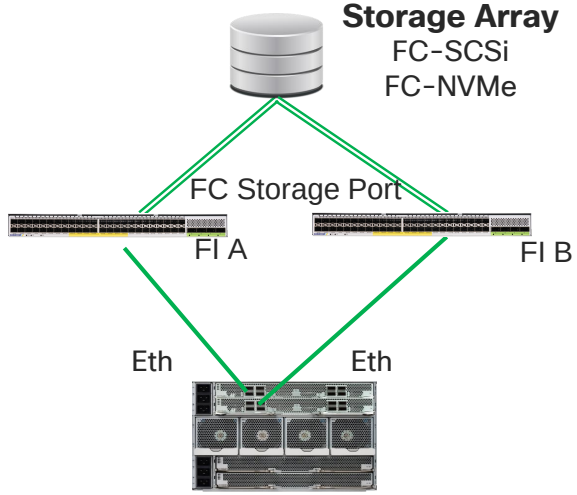

- FC switch mode
- 4 vHBA per server for higher redundancy

IP-SAN connectivity with 6th Gen FI

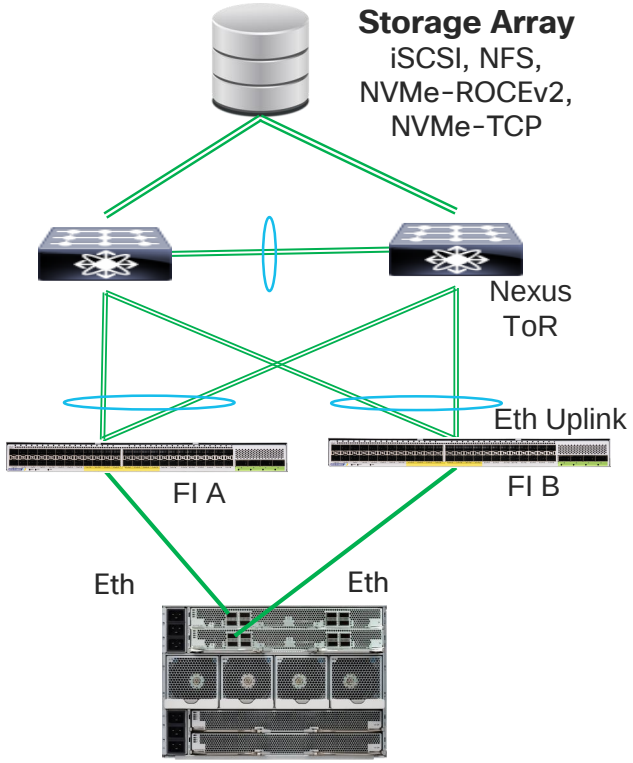

- VPC/MCT port-channel is recommended
- MTU 9216 should be enabled in system-qos
- TCP based storage can use best-effort class but if required no-drop could be enabled across FI & TOR
- ROCEv2 would require no-drop QoS-class along with PFC enabled on FI & ToR
- Multiple vNICs on server for redundancy
- Best-practice connectivity from Nexus to Storage is different for each vendor

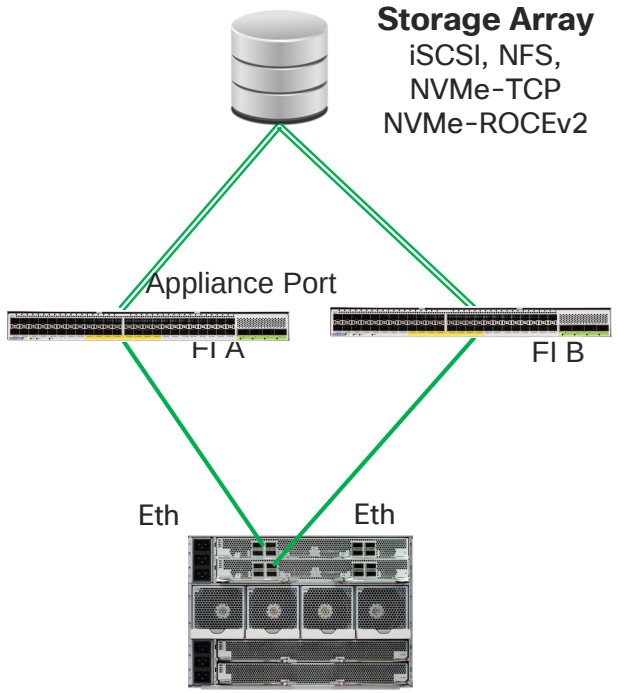

- For small to medium deployments
- MTU 9216 should be enabled in system-qos
- Can avoid ToR for storage access
- Direct port-channel from an FI to a Storage controller is possible.
- No VPC like port-channel towards Storage Array

6th Gen Domain Profile

IMM Configuration Model

FYI

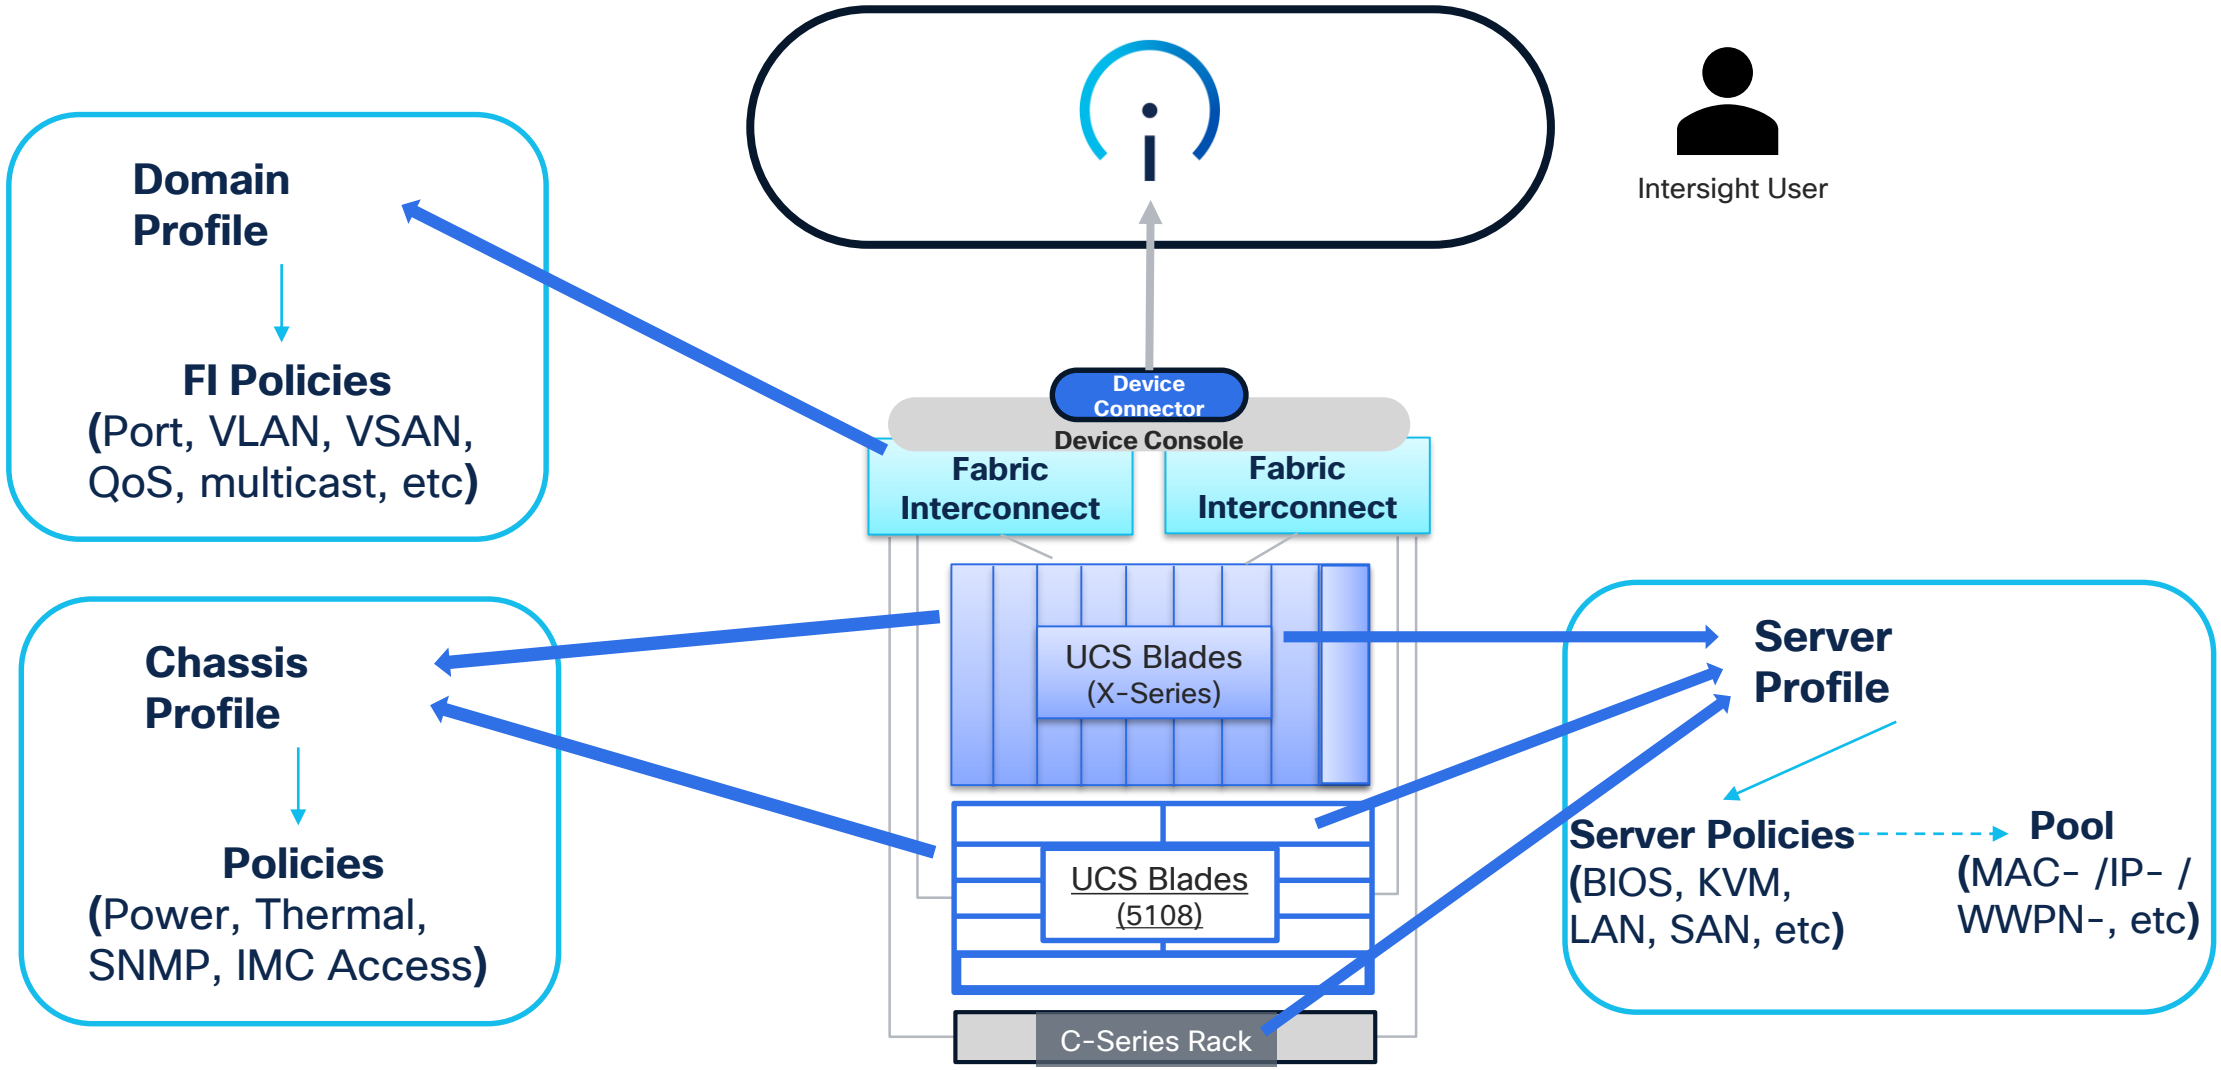

Intersight Domain profile

Best Practice

Minimum policies for FI Domain-profile

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

VLAN policy

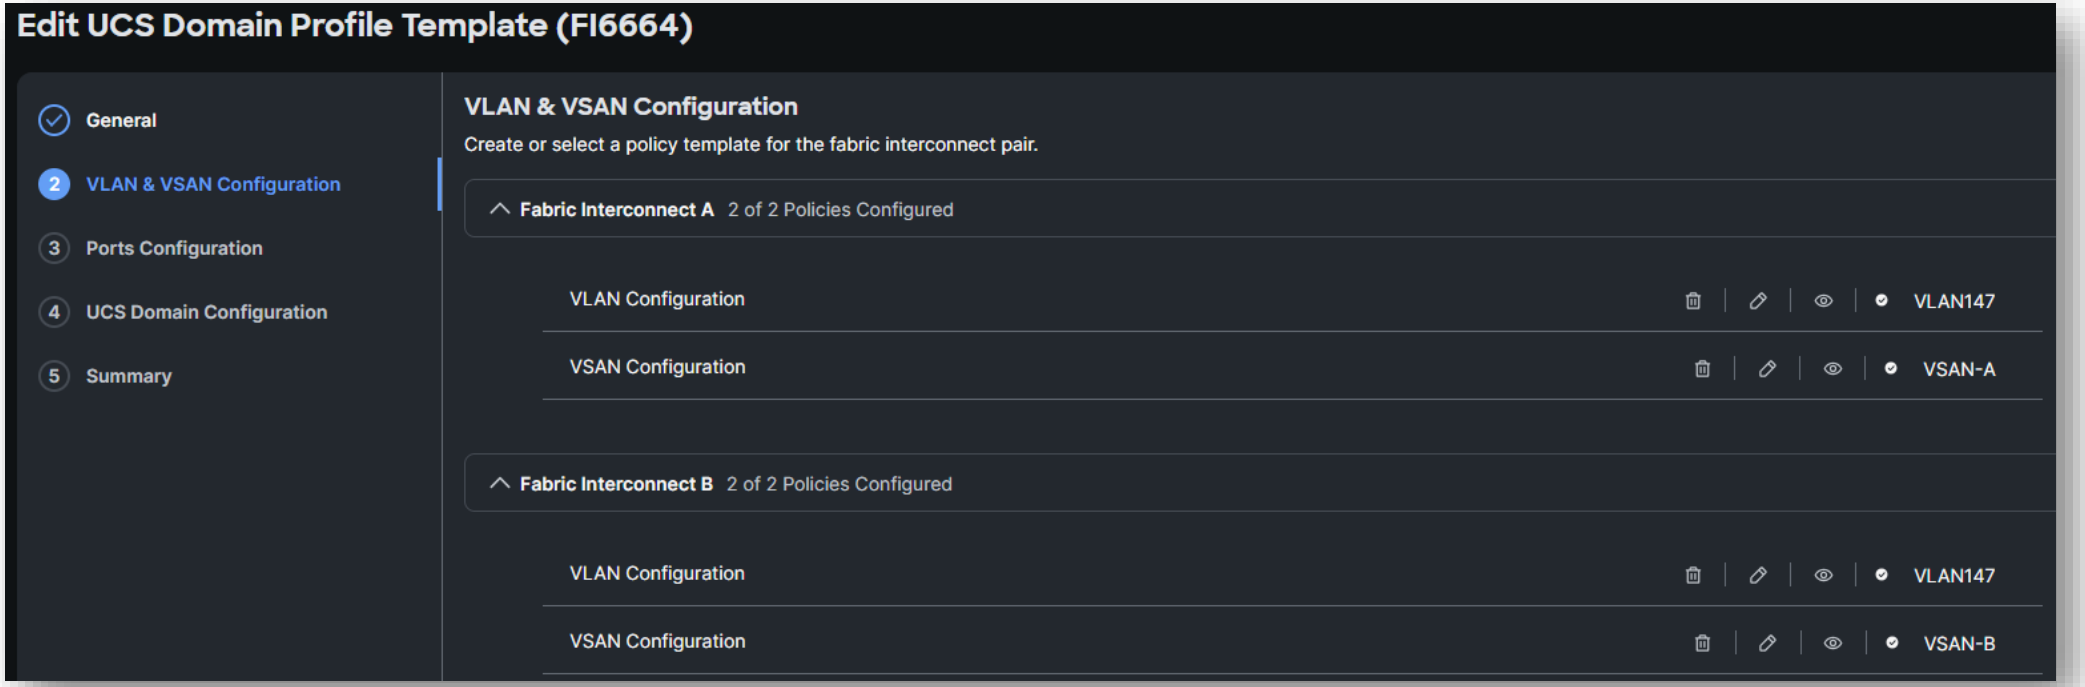

Intersight Domain profile

Minimum policies for FI Domain-profile

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

Port policy

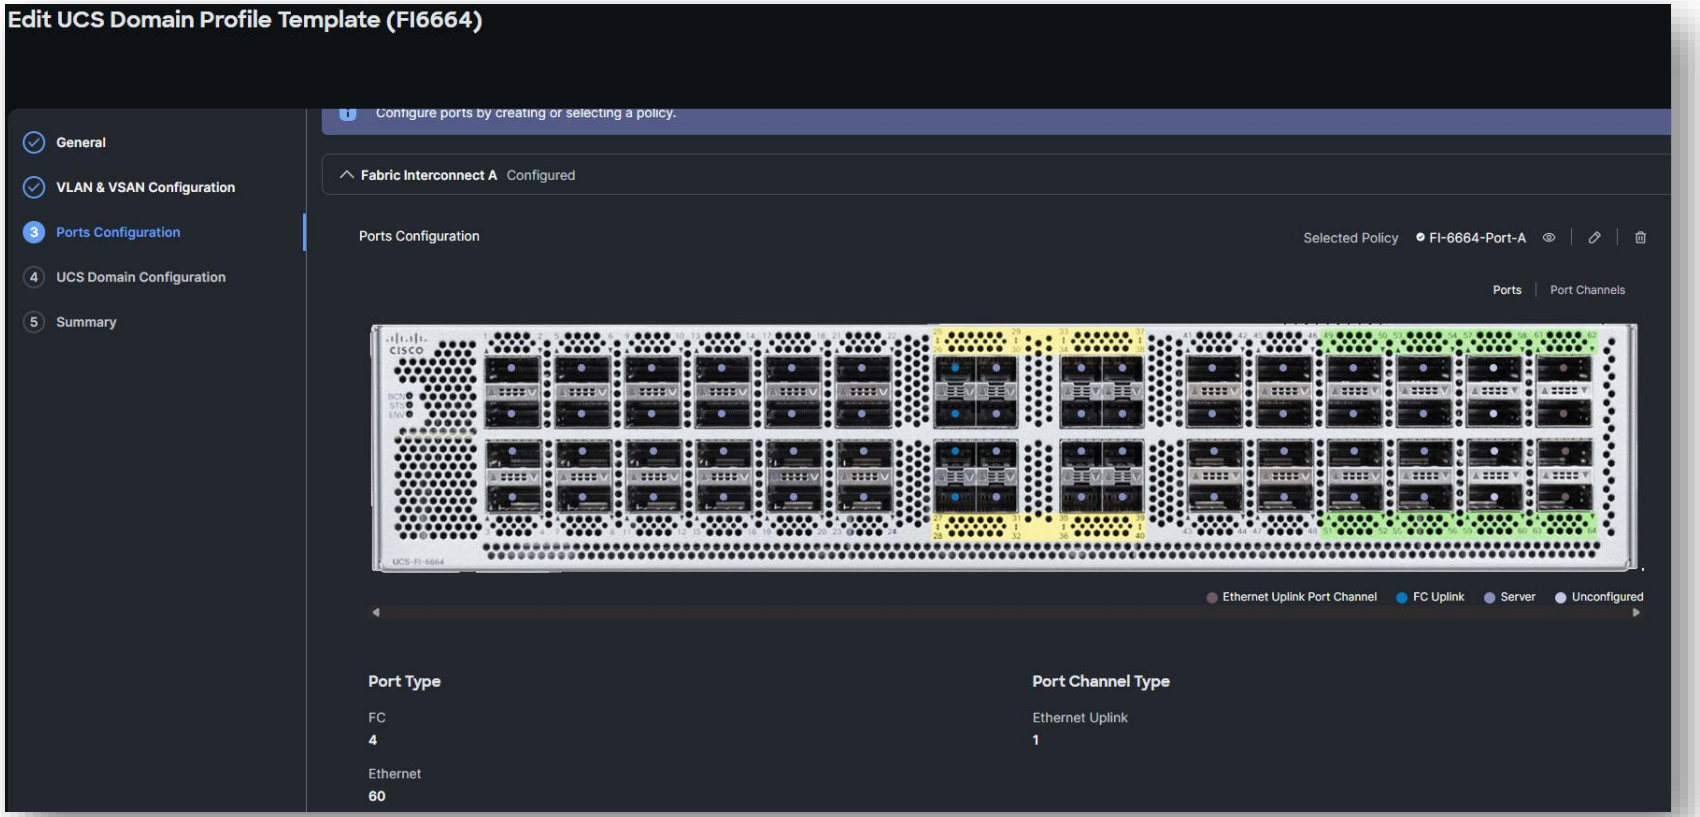

Intersight Domain profile

Minimum policies for FI Domain-profile

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

QoS policy

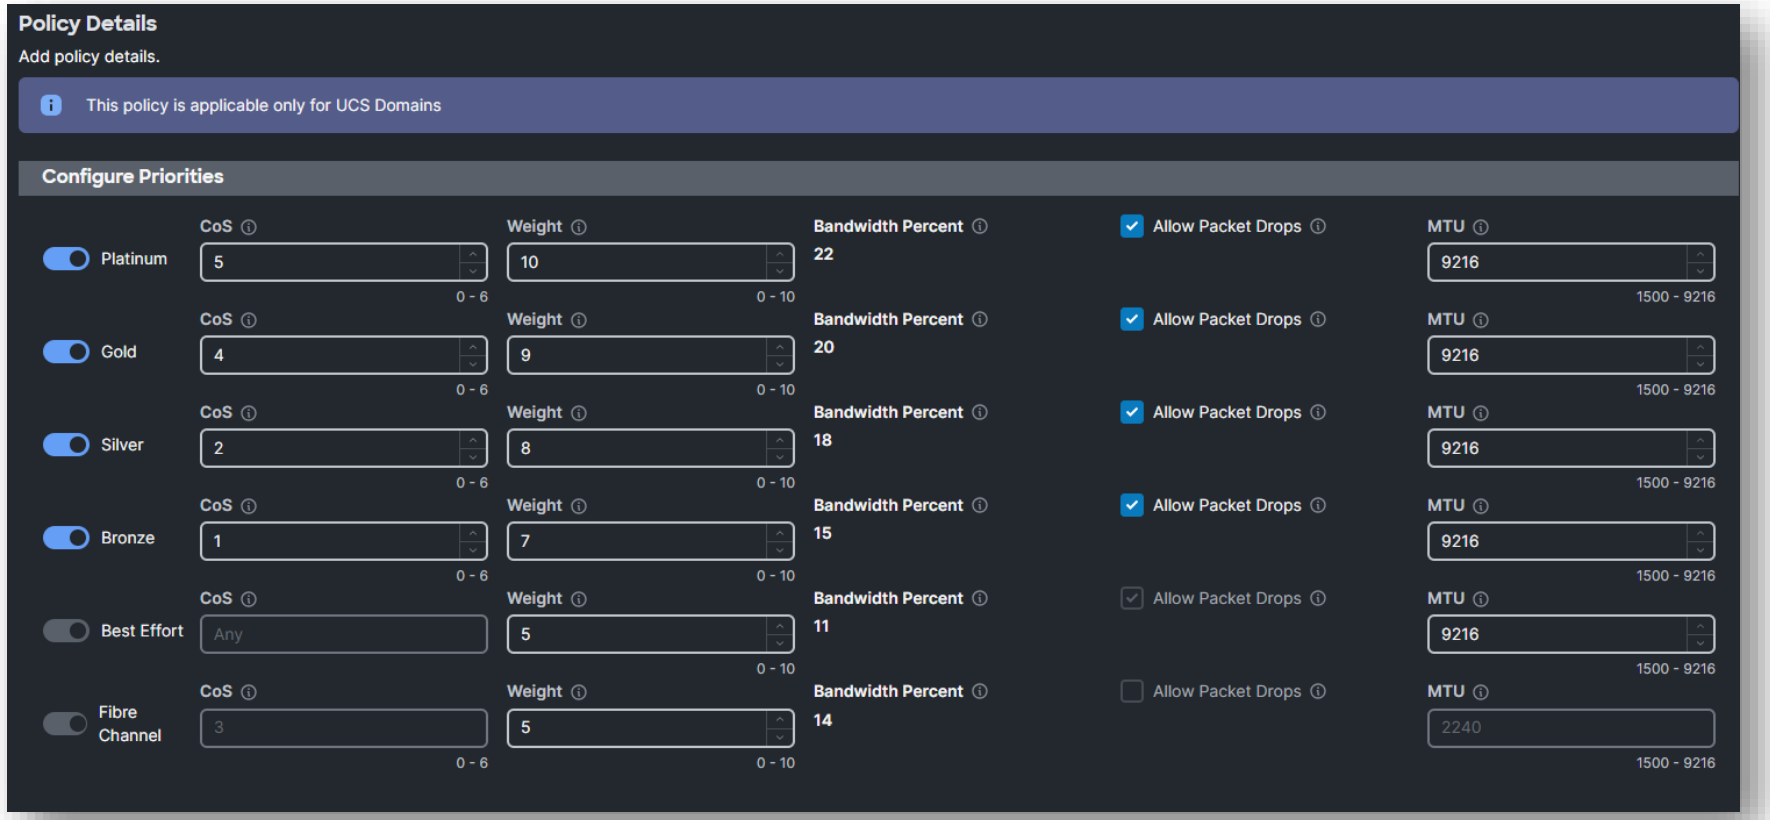

X-Series Out-of-Band Management

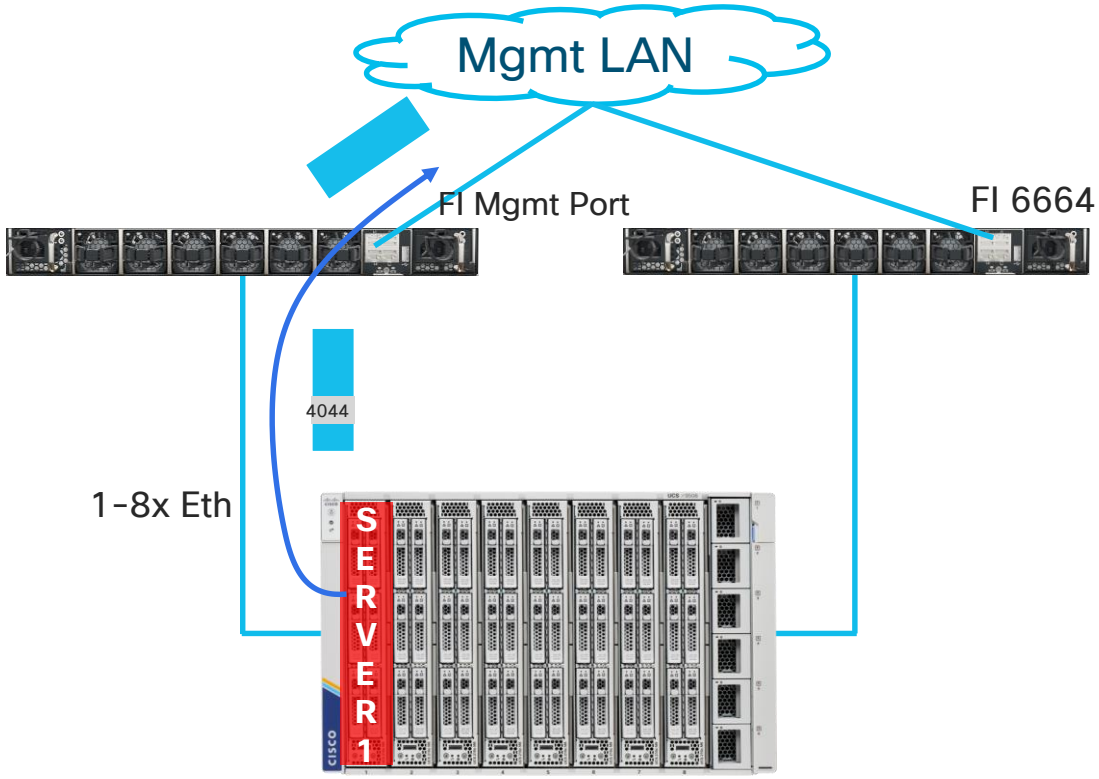

UCS X-Series Chassis with X9108-IFM-100G

Server's vKVM access will be via the mgmt port on the FI with Out-of-Band configuration done with **IMC Access Policy**

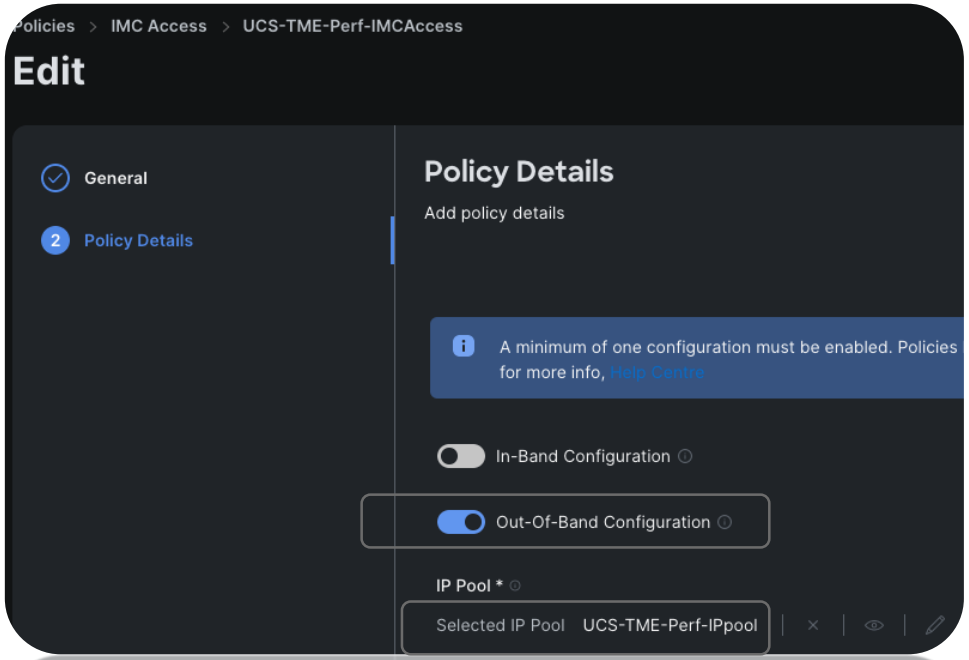

FC Storage

FC Configuration

- Supports FC end host and switch mode at FCS
- Unified Ports
  - 25 to 40 (FI-6664)
  - 17 to 32 (FI-6652)
- FI reload required when changing from Eth <-> FC mode for a port and vice versa.
- FC speeds supported:
- 16G, 32G and 64G speeds
- Auto speed not supported for FC ports.
- FI reload not required when changing the port speed.
- 64 Buffer Credits per Port

FI 6664 Intersight FC Configuration

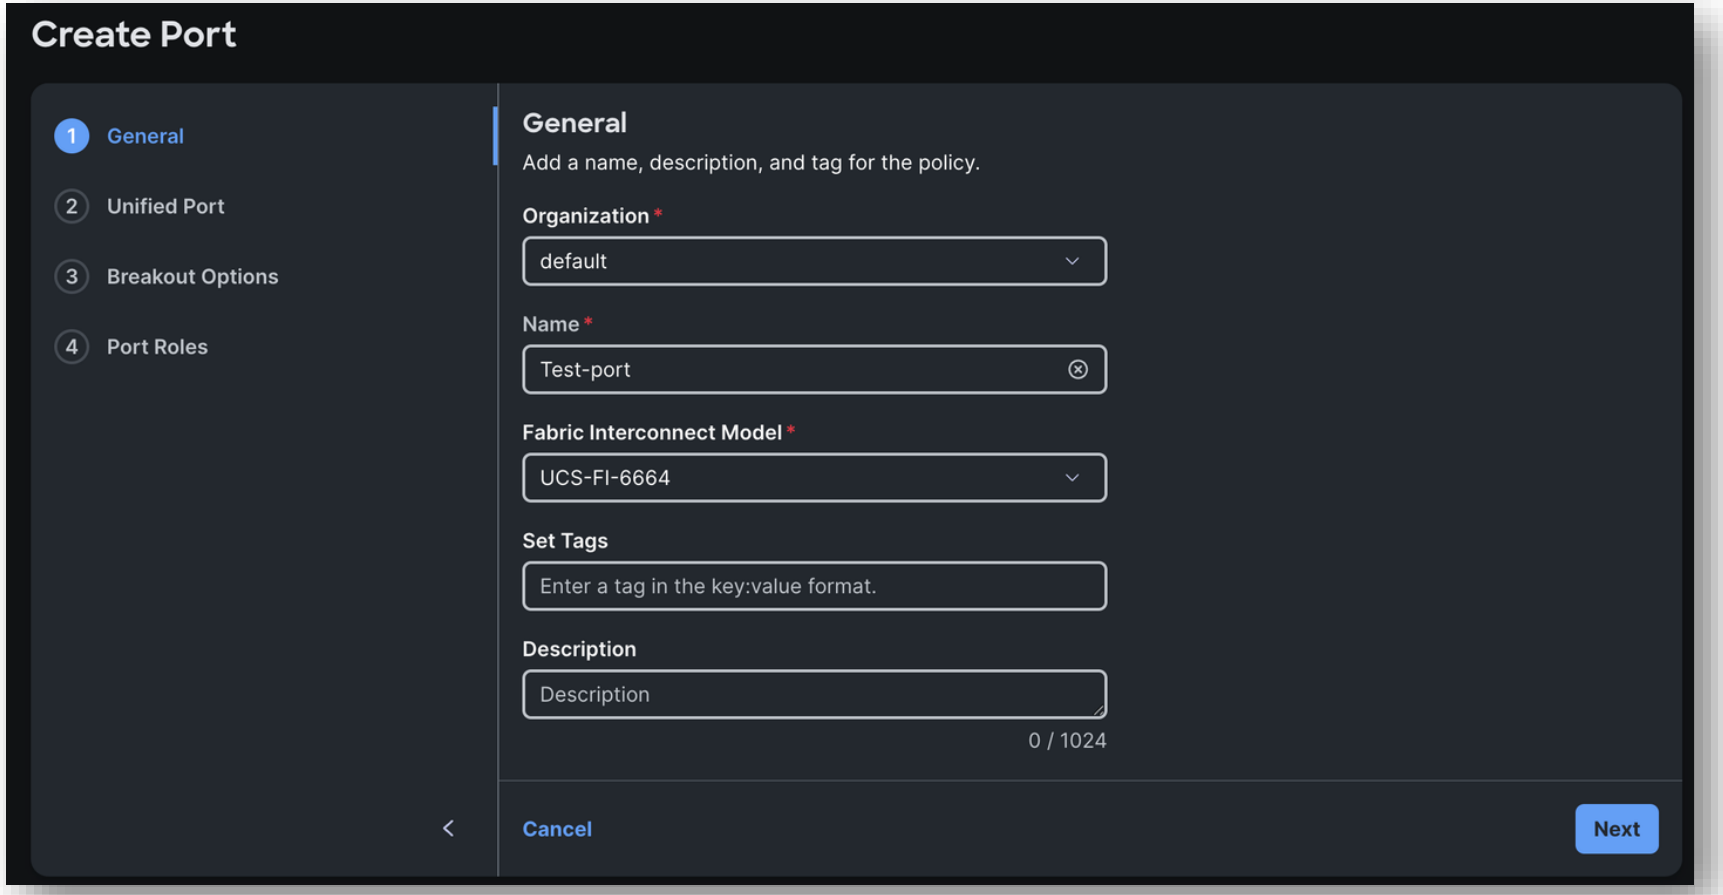

FI 6664 Intersight FC Configuration

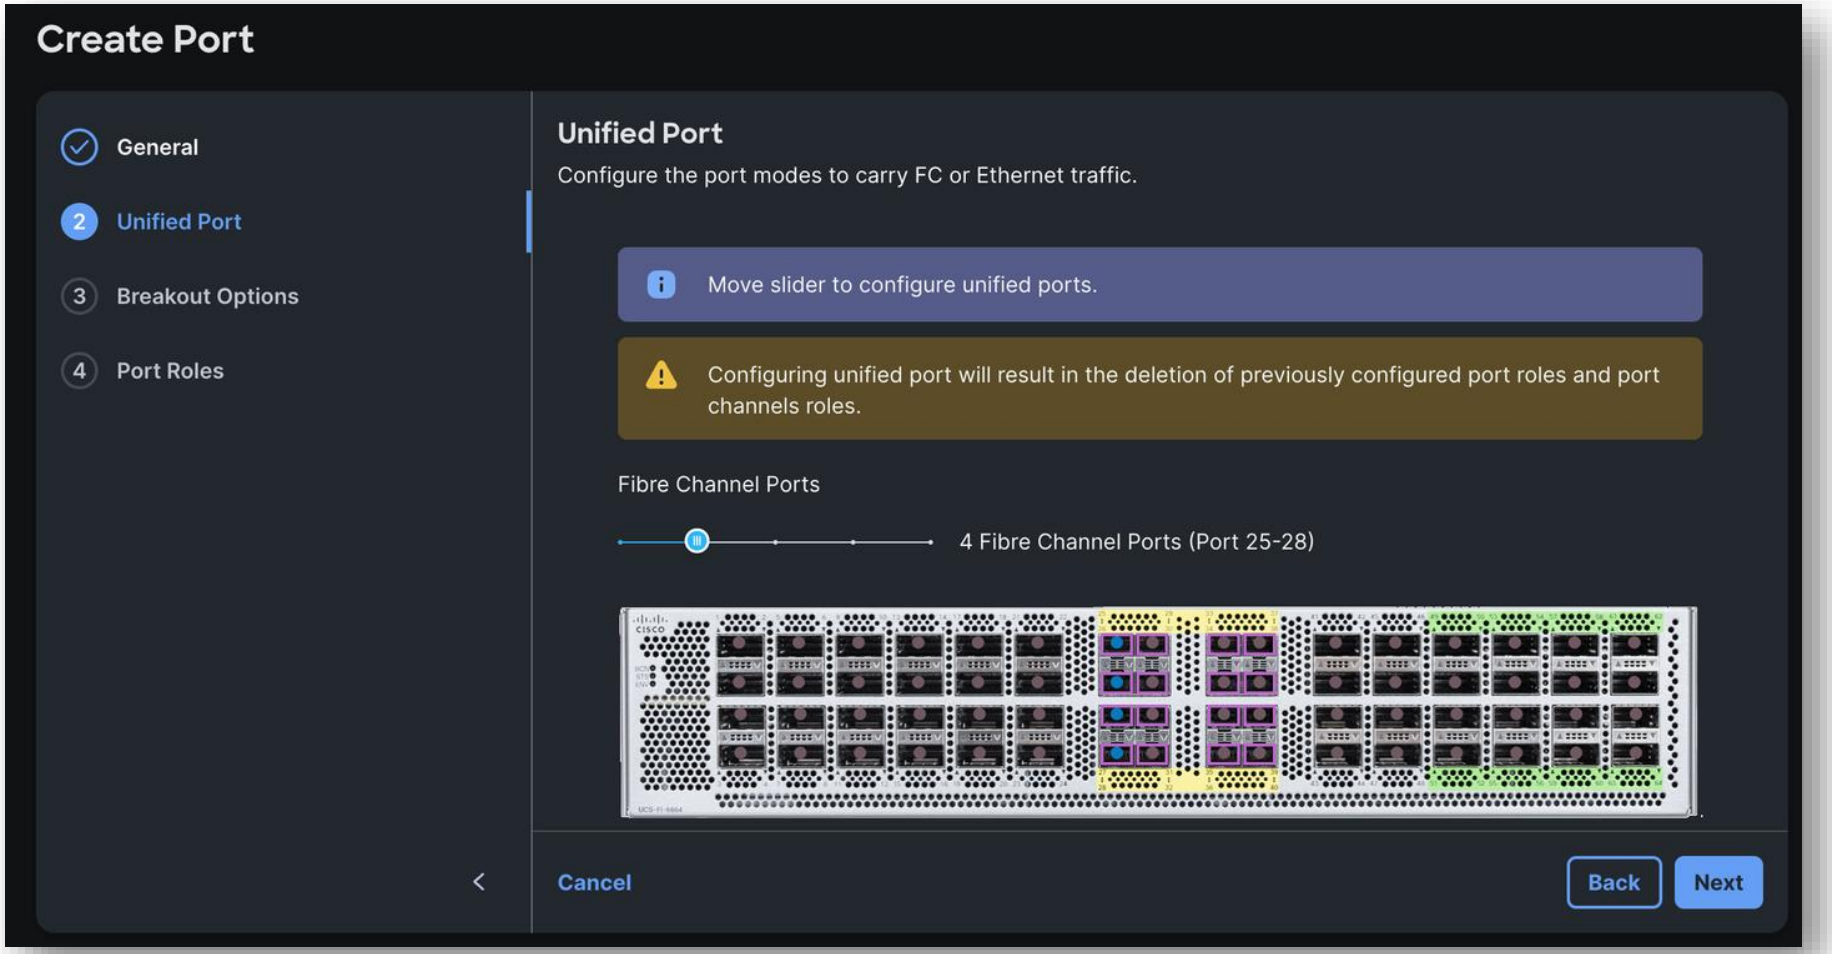

FI 6664 Intersight FC Configuration

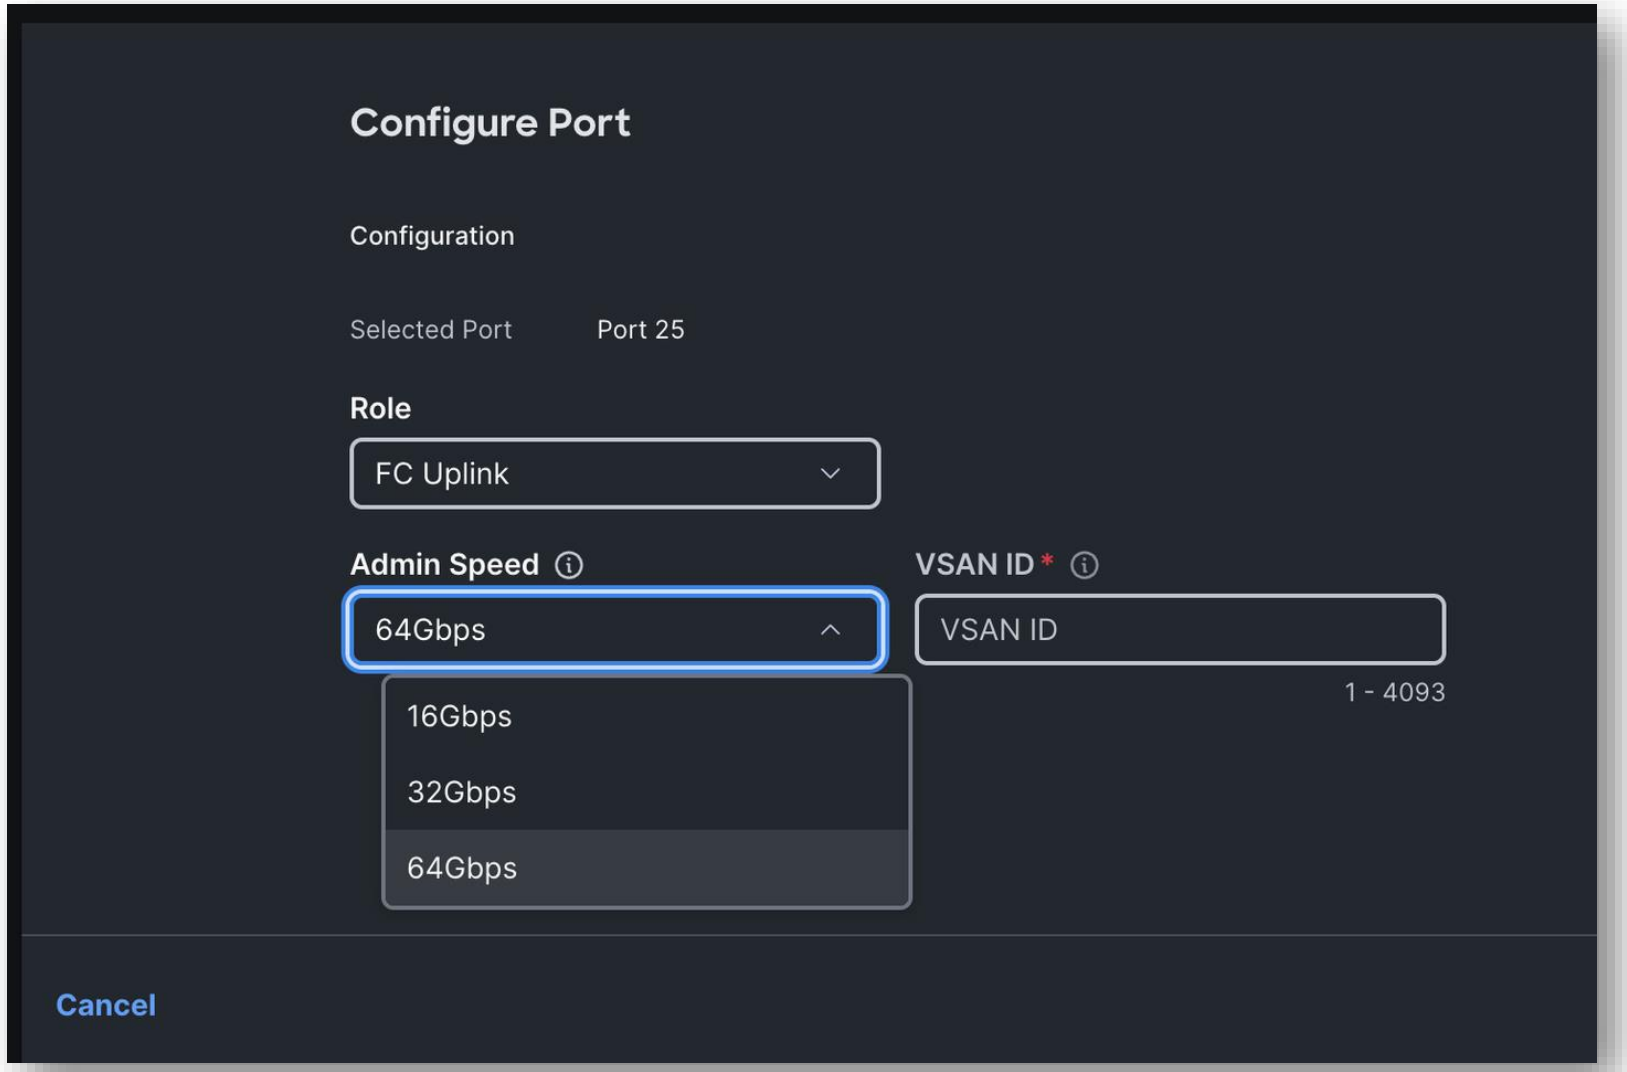

FI 6664 Unified Ports in UCS Manager

Configure Unified Ports

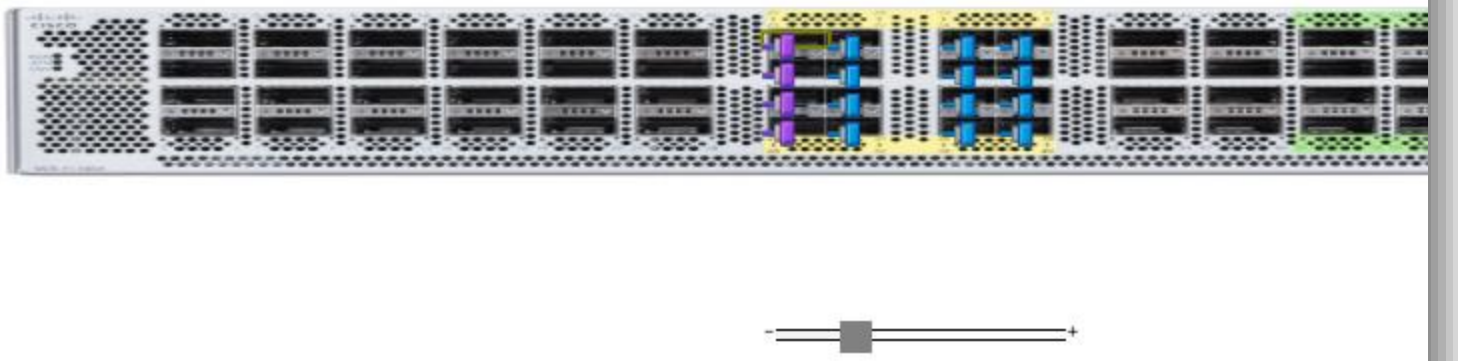

Instructions

The position of the slider determines the type of the ports.
All the ports to the left of the slider are Ethernet ports (Blue), while the ports to the right are Fibre Channel ports (Purple).

Port	**Transport**	**If Role** or **Port Channel Membership**	Desired **If Role**
Port 25	ether	Unconfigured	FC Uplink
Port 26	ether	Unconfigured	FC Uplink
Port 27	ether	Unconfigured	FC Uplink
Port 28	ether	Unconfigured	FC Uplink

End to End 64G : VIC 15231 performance

64G FC connectivity to a Storage Target:

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

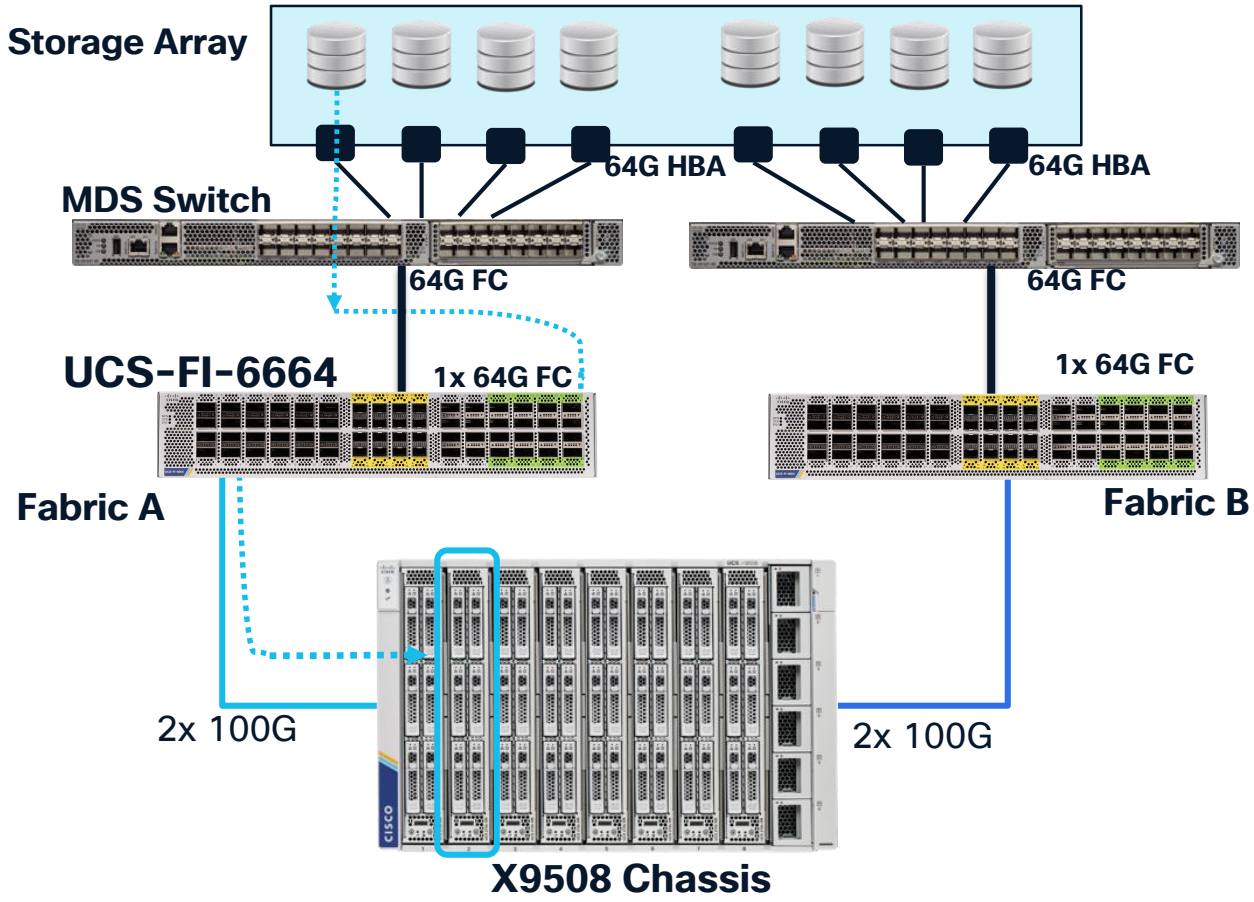

End to End 64G : VIC 15231 performance

```
jvdmade@localhost:~> cat /sys/class/fc_host/host3/speed
100 Gbit
```

```
jvdmade@localhost:~> cat ciscolive.ini
[global]
bs=256k  # Block size
iodepth=256 # I/O depth
direct=1  # Direct I/O
ioengine=libaio # I/O engine (Linux)
numjobs=8 # Number of jobs
name=my-test-job
rw=rw
rwmixread=100
group_reporting


[job1]
filename=/mnt/fc1/testfile.dat # Path to the test file
size=10G # 10GB of I/O
runtime=60 # Run for 60 seconds
```

End to End 64G : VIC 15231 performance

```
jvdmade@localhost:~> sudo fio ciscolive.ini
job1: (g=0): rw=rw, bs=(R) 256KiB-256KiB, (W) 256KiB-256KiB, (T) 256KiB-256KiB, ioengine=libaio, iodepth=256
...
fio-3.23
Starting 8 processes
Jobs: 8 (f=8): [R(8)][92.9%][r=6082MiB/s][r=24.3k IOPS][eta 00m:01s]
job1: (groupid=0, jobs=8): err= 0: pid=50105: Thu Jun  5 12:59:35 2025
  read: IOPS=24.3k, BW=6075MiB/s (6370MB/s)(80.0GiB/13485msec)
    slat (nsec): min=1385, max=4299.8k, avg=326399.90, stdev=143799.28
    clat (msec): min=10, max=119, avg=83.20, stdev= 7.12
     lat (msec): min=10, max=120, avg=83.53, stdev= 7.15
    clat percentiles (msec):
     |  1.00th=[   44],  5.00th=[   79], 10.00th=[   81], 20.00th=[   82],
     | 30.00th=[   83], 40.00th=[   84], 50.00th=[   84], 60.00th=[   85],
     | 70.00th=[   86], 80.00th=[   87], 90.00th=[   88], 95.00th=[   90],
     | 99.00th=[   93], 99.50th=[   96], 99.90th=[  103], 99.95th=[  113],
     | 99.99th=[  118]
   bw (  MiB/s): min= 5141, max= 6278, per=99.50%, avg=6044.50, stdev=26.86, samples=208
   iops        : min=20566, max=25111, avg=24174.65, stdev=107.36, samples=208
  lat (msec)   : 20=0.36%, 50=0.71%, 100=98.79%, 250=0.14%
  cpu          : usr=0.14%, sys=2.82%, ctx=326196, majf=0, minf=91
  IO depths    : 1=0.1%, 2=0.1%, 4=0.1%, 8=0.1%, 16=0.1%, 32=0.1%, >=64=99.8%
     submit    : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.0%
     complete  : 0=0.0%, 4=100.0%, 8=0.0%, 16=0.0%, 32=0.0%, 64=0.0%, >=64=0.1%
     issued rwts: total=327680,0,0,0 short=0,0,0,0 dropped=0,0,0,0
     latency   : target=0, window=0, percentile=100.00%, depth=256

Run status group 0 (all jobs):
   READ: bw=6075MiB/s (6370MB/s), 6075MiB/s-6075MiB/s (6370MB/s-6370MB/s), io=80.0GiB (85.9GB), run=13485-13485msec
```

NVMe Over Fabric

NVMe Protocol: How does it work?

- NVMe communicates directly with the CPU
- Multi Queue (up to 64k queues) involves many communication between host and Target.
- Queue can have up to 64k commands per queue, which makes it designed to handle ultra-fast CPU and storage.

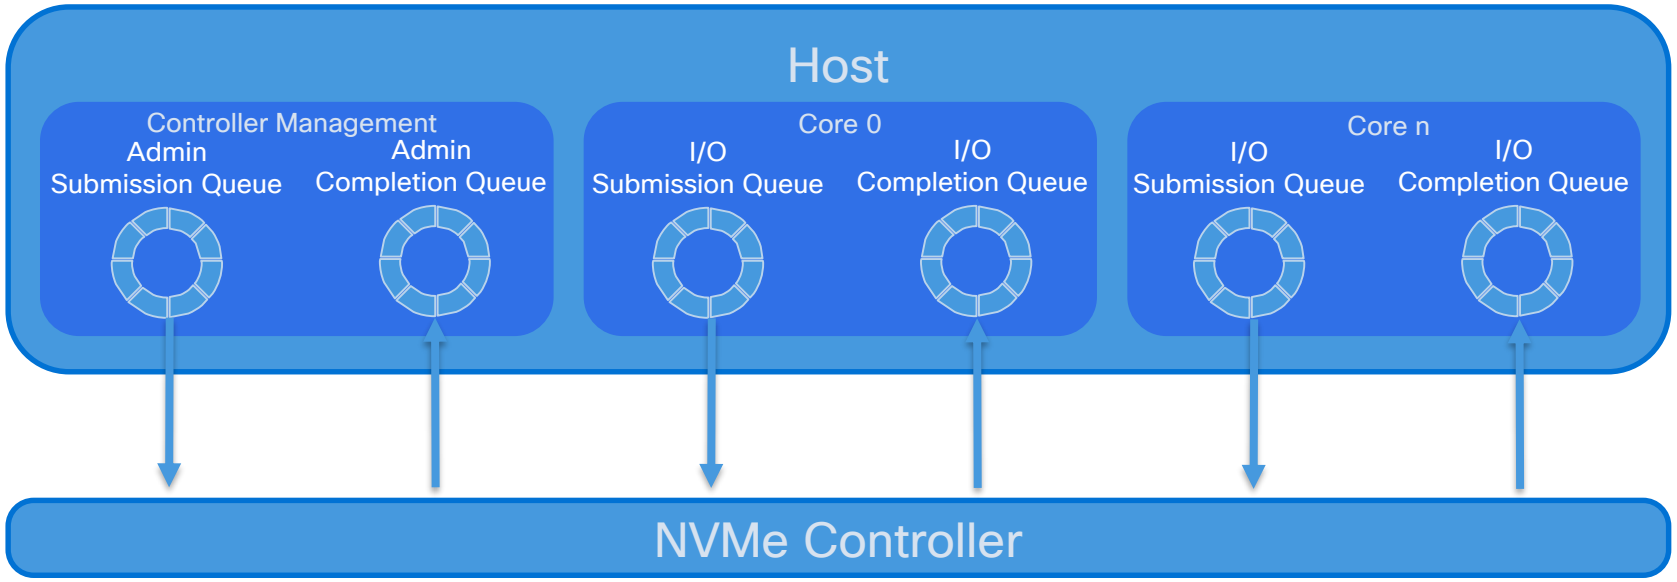

What is NVMe Over Fabric?

- Protocol that allows hosts to connect to storage across a network fabric using non-volatile memory express (NVMe) commands.

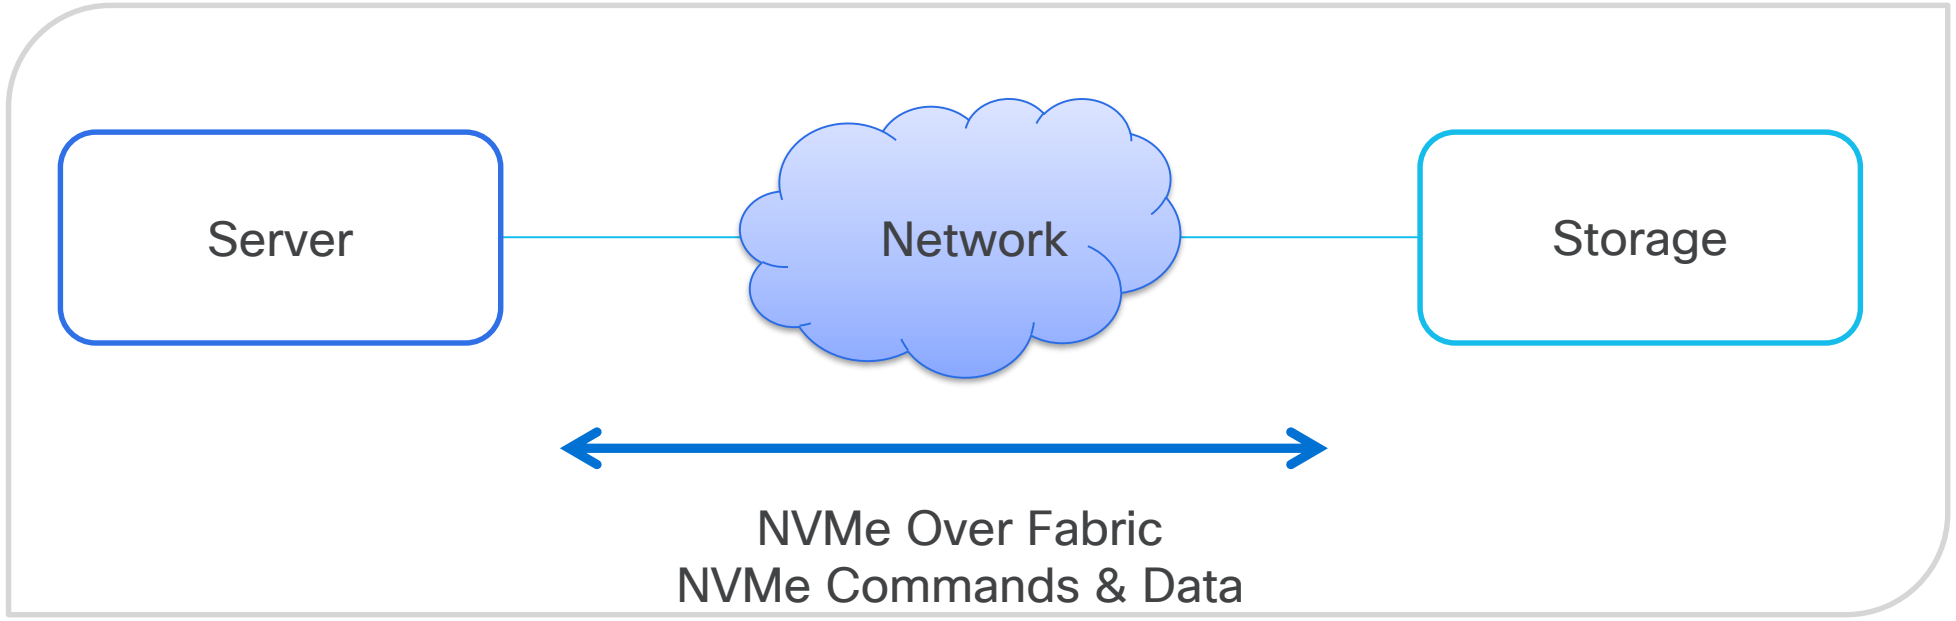

NVMe Over TCP

NVMe Over TCP Network Best Practices

QoS	• Classification • Queuing • Queue management • Scheduling
Traffic Prioritization	• Storage can have 100% capacity when other classes don't have traffic.
Avoid Policing or Shaping	• Can lead to mediocre performance.
Dedicated Storage Network	• Avoid Traffic congestion • Troubleshoot • Change Management
Load Balancing	• One TCP Connection per I/O Submission and Completion Queue Pair.

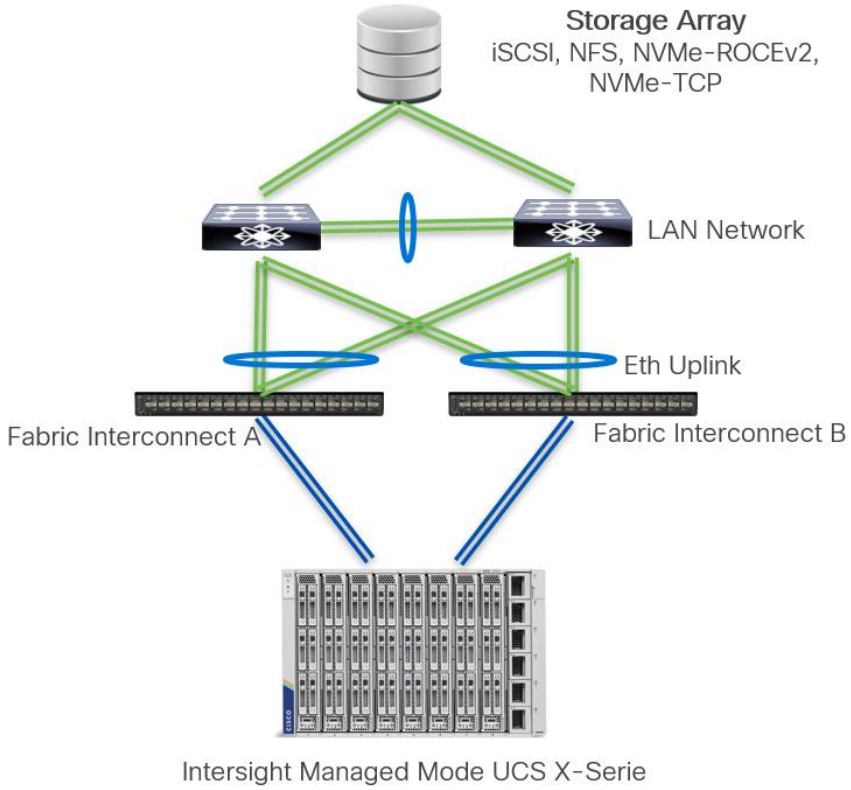

NVMe Over TCP

Best Practice

- Change Ethernet Adapter Policies for ESXi

  RX=16
  Receive and Transmit Ring = 16384

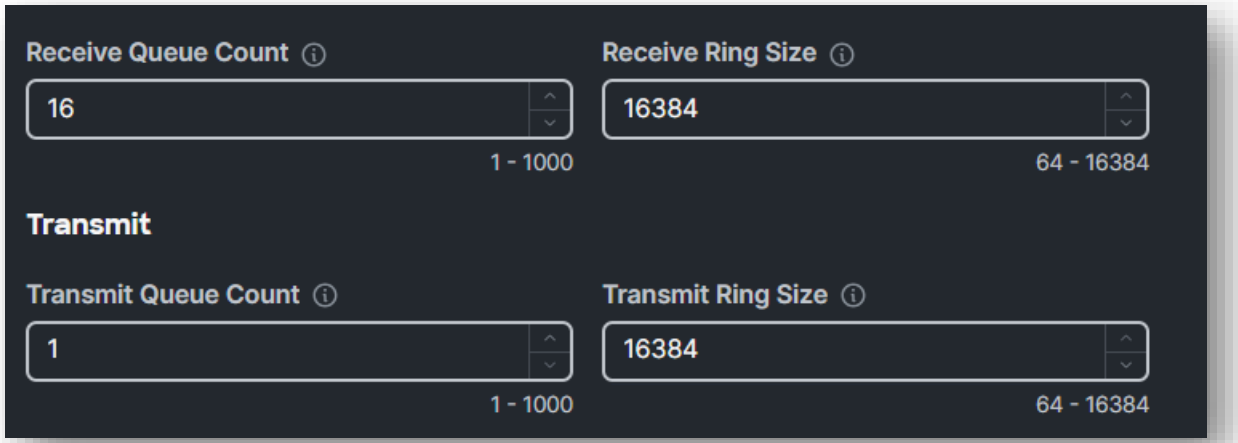

*

Increasing the ring size can increase the latency, but on higher speed (100GbE) interfaces the higher speeds mean less time the data sits in the buffers, thereby minimize the latency impact.

NVMe Over RoCEv2

NVMe Over RoCE (RDMA over Converged Ethernet)

- Lossless Traffic
  - Use Data Center Bridging (DCB) Tools, like Priority Flow Control (PFC) etc.

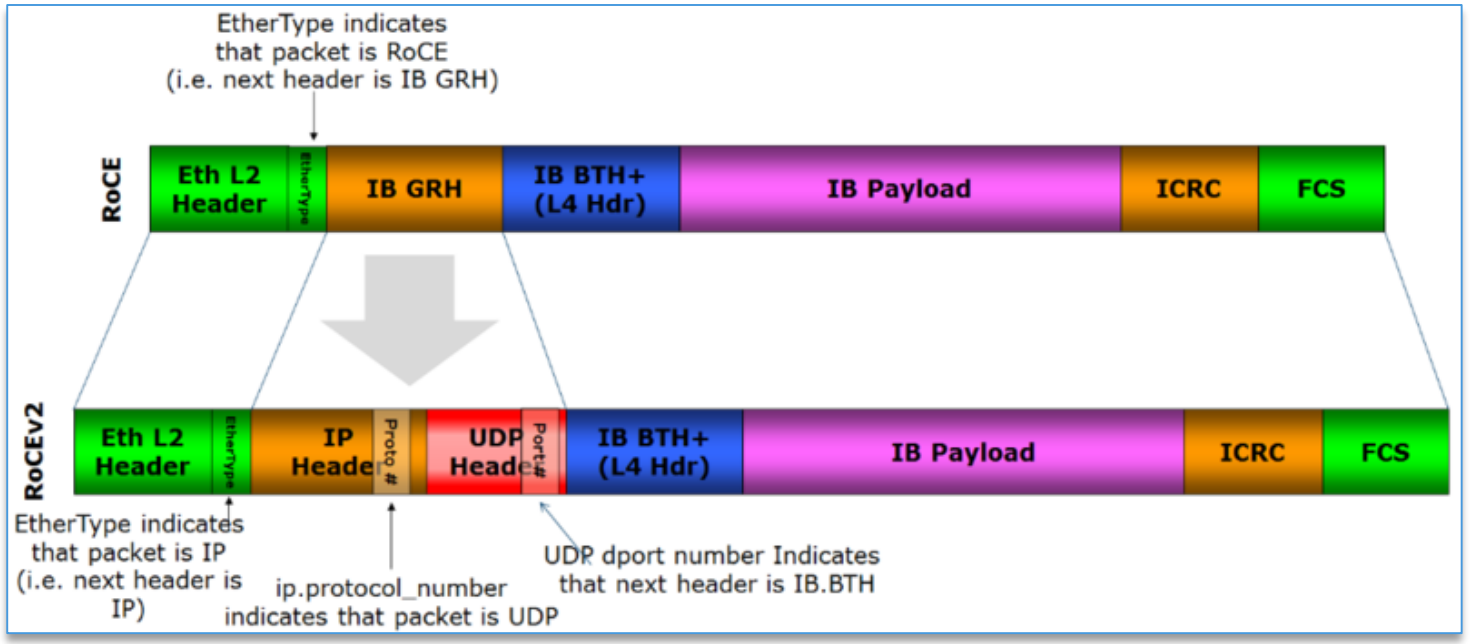

Source:
https://en.wikipedia.org/wiki/RDMA_over_Converged_Ethernet

NVMe Over RoCEv2 Guidelines

FYI

- NVMe over RDMA with RoCE v2 is currently **supported only 4th and 5th gen VIC** Series adapters.
- RoCE v2 supports **maximum two** RoCE v2 enabled **interfaces** per adapter.
- NVMe over RoCEv2 **cannot** be used with usNIC, VxLAN, VMQ, VMMQ, NVGRE, GENEVE Offload, ENS, and DPDK features.
- Cisco **Intersight** does **not support fabric failover for vNICs** with RoCE v2 enabled.
- During the failover or failback event, the Spanning Tree Protocol (STP) can result temporary loss of network connectivity. To prevent this connectivity issue, **disable STP** on **uplink switches**.
- **NVMe Over RoCE** does **not support boot** from **SAN** at this moment of writing.
- When creating **RoCEv2 interfaces**, use Intersight provided **Linux-NVMe-RoCE adapter policy**. **Do not use the default Linux Adapter policy with RoCEv2**; RoCEv2 interfaces will not be created in the OS.
- **RoCEv2 does not support bonding**.

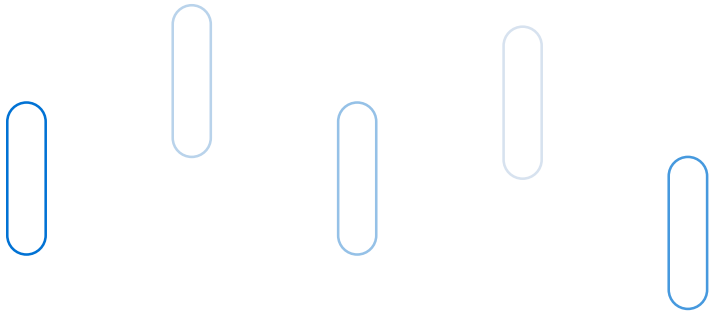

NVMe Over RoCEv2 Intersight

Best Practice

FYI

- Create VLANs for NVMe Over RoCE.
- Enable Platinum QoS CoS5 with Jumbo MTU.
- Two adapter policies are the best practice as there are diverse types of traffics each with different priories.
  One traffic is the **public traffic** for managing servers and **other traffic** is the **Storage traffic** using NVMe protocol Over RoCE which needs be classified as high.

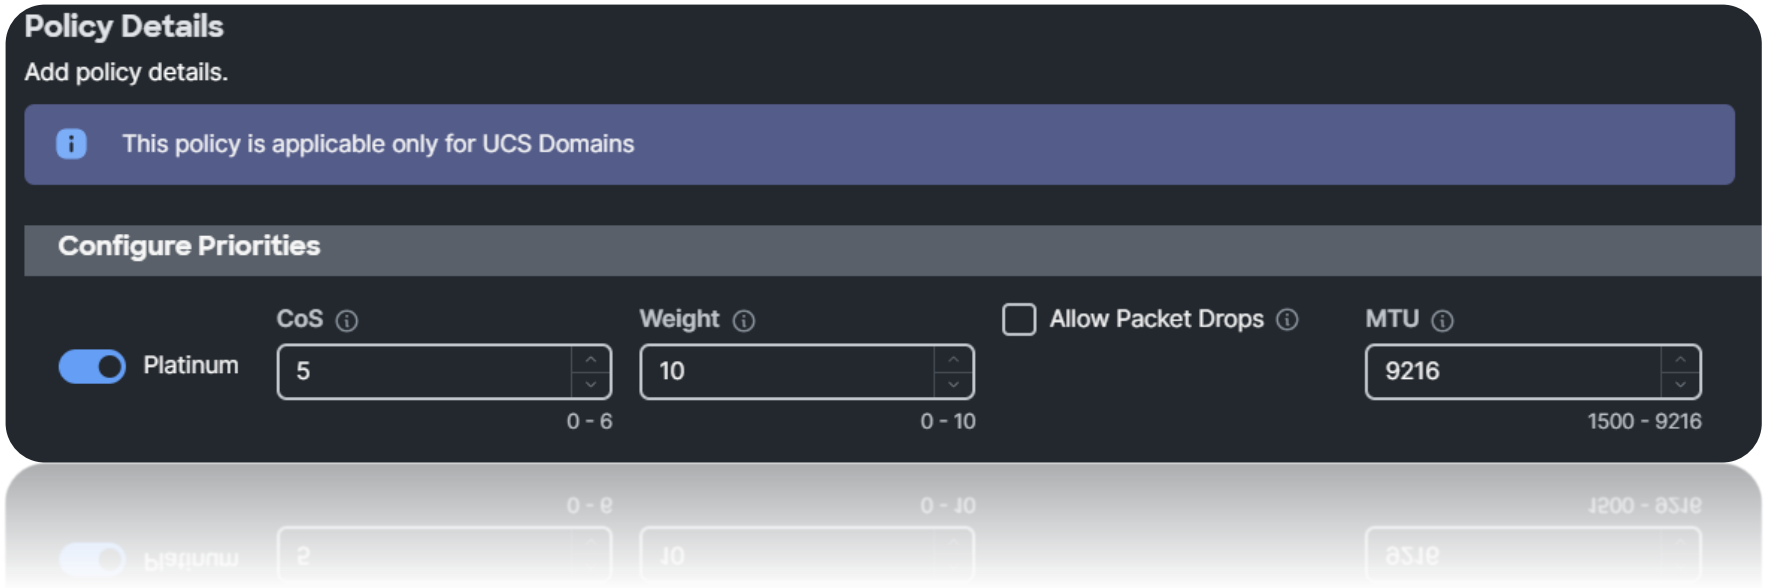

NVMe Over RoCEv2 Intersight BIOS Policy

Best Practice

- Create a BIOS policy that with the following options enabled:
- For **Intel** based servers, **Intel VT for directed IO** under **Intel Directed IO** tab.
- For **AMD** based servers **SVM Mode** under **Processor** tab.
- Enable **IOMMU** under **Memory** tab.

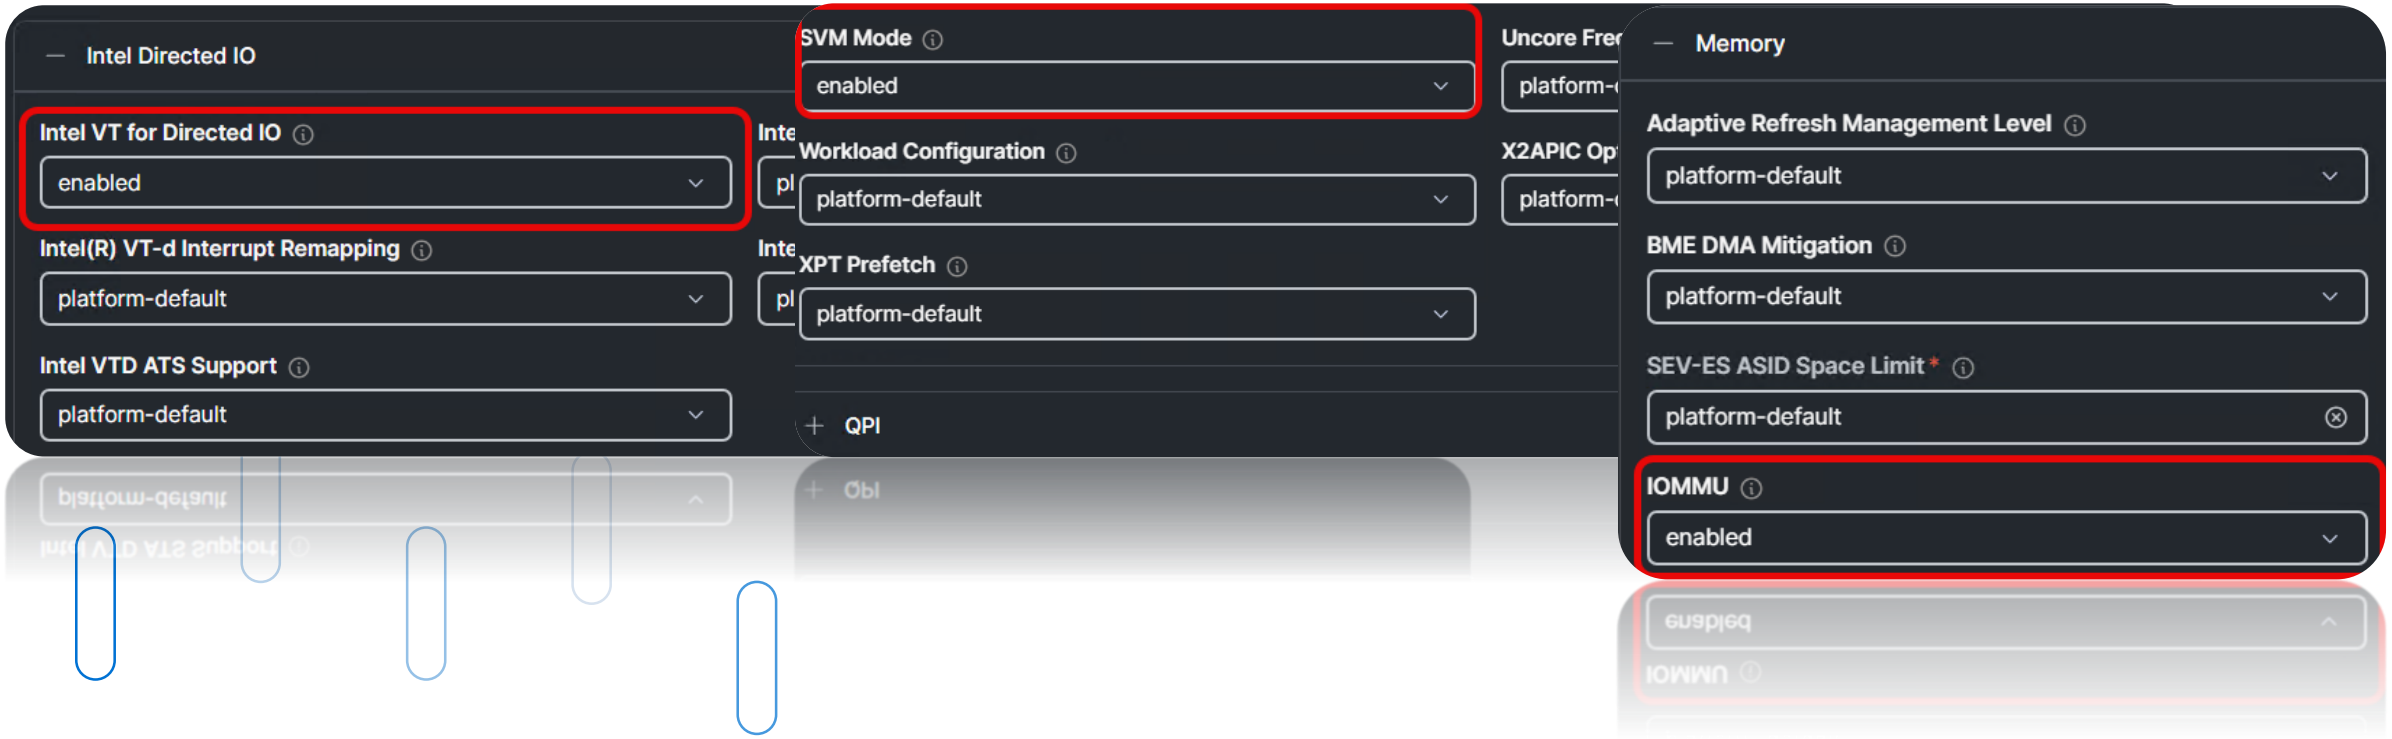

NVMe Over RoCEv2 Intersight LAN Connectivity Policy

FYI

- Do not enable Failover because RoCE will not work is enabled.
- Create an Ethernet Network Group policy and fill in the RoCE VLAN.
- Create an Ethernet Network Control policy and leave the defaults.
- Create an Ethernet QoS Policy and set the MTU Bytes to 9000
  The Priority should be set to Platinum.

Best Practice

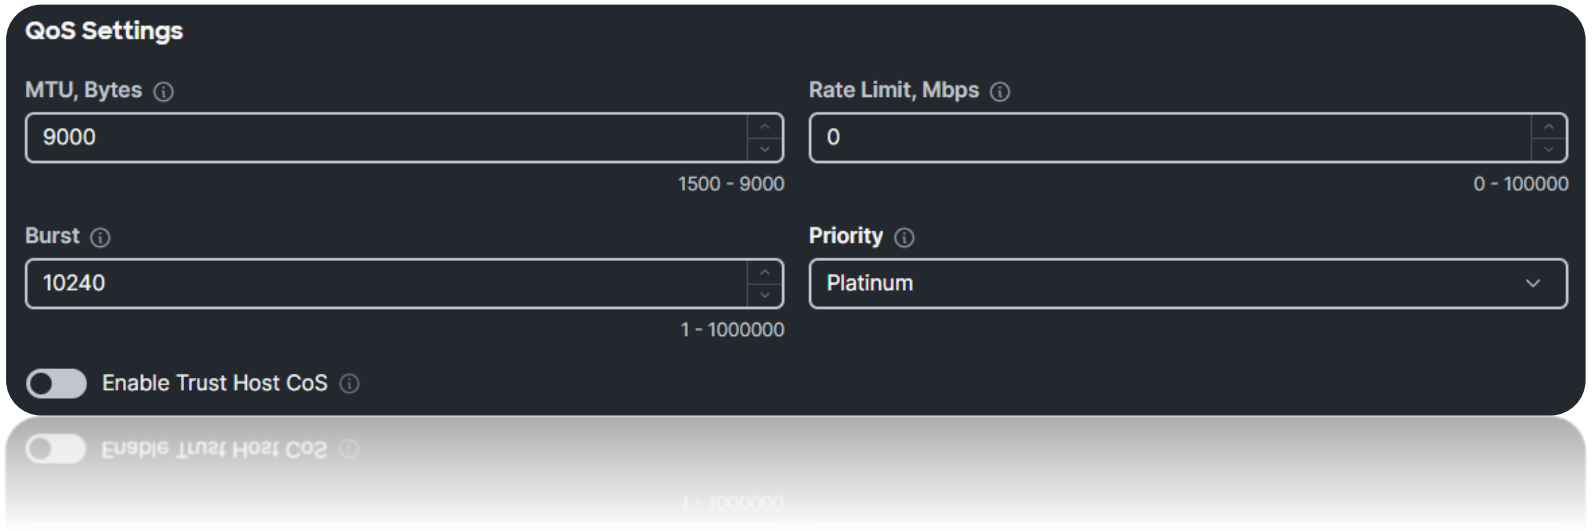

NVMe Over RoCEv2 Intersight Adapter Policy

Best Practice

- Create an Ethernet Adapter Policy and Select Cisco Provided Configuration. Choose **Linux-NVMe-RoCE** because this predefined configuration has the best setting for NVMe Over RoCE.
- Set Queue Pairs to 1024
- Set Memory Regions to 131072
- Set Resource Groups to 8
- Select Version 2
- Set Class of Service to 5
- Set Interrupts to 256
- Select Interrupt mode MSIx
- Set Interrupt Timer to 125
- Select Interrupt Coalescing Type Min
- Set Receive Queue Count to 1
- Set Receiving Ring Size to 512
- Set Transmit Queue Count to 1
- Set Transmit Ring Size to 256
- Set Completion Queue Count to 2
- Set Completion Ring Size to 1
- Set Uplink Failback Timeout to 5

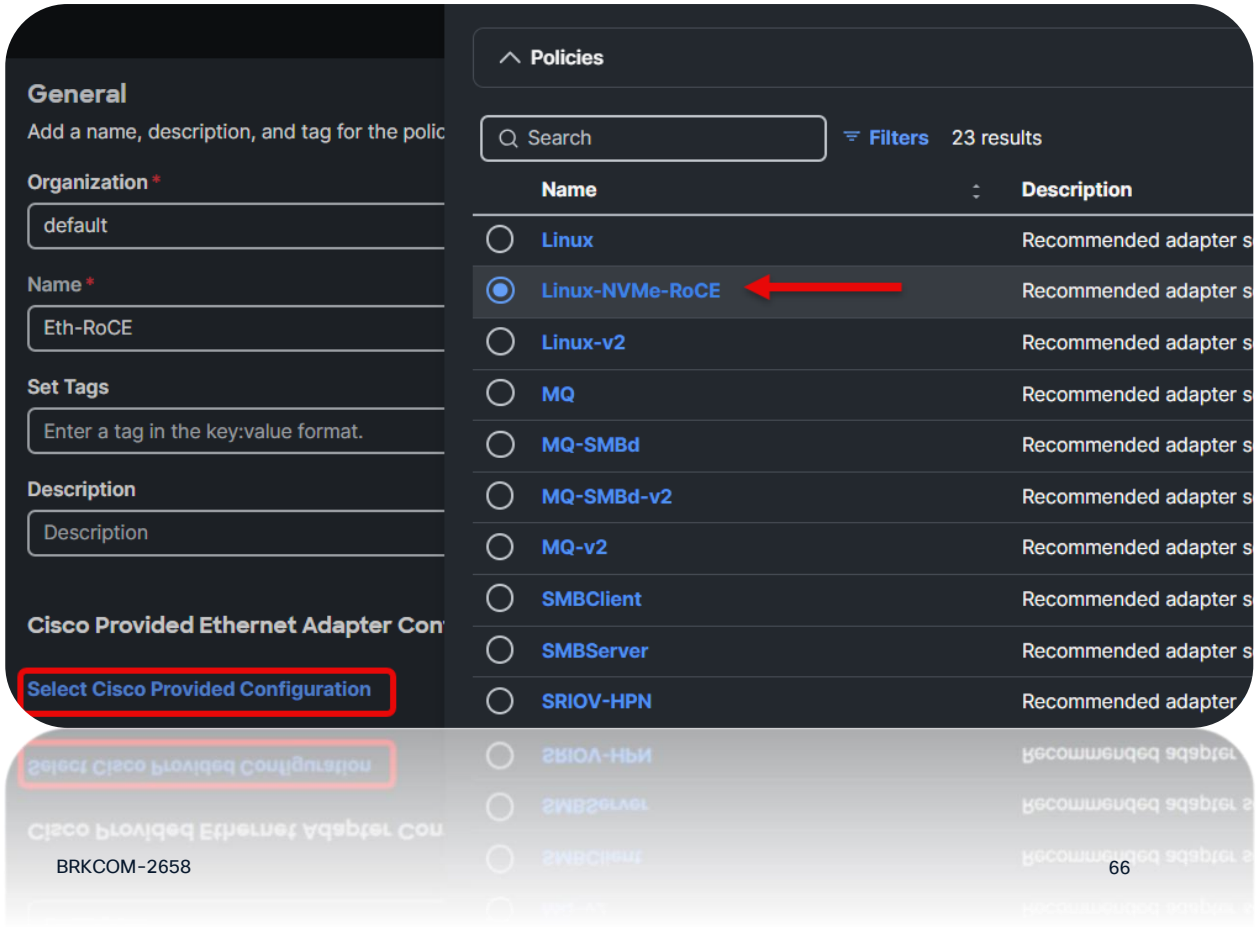

NVMe Over RoCE Verify in Fabric Interconnect

FYI

- Connect to the Fabric Interconnect via SSH.
- Connect to the adapter of the server.
  Use the command **connect adapter x/y/z**
- Type the command **attach-mcp**
- Run the command **vnicl**

```
AC06-FI-6536-A# connect adapter 1/3/1

Entering character mode
Escape character is '^]'.

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

```
adapter (mcp):2# vnicl 17
              vnicid : 17
                name : Eth0
                type : enet
               state : UP
             adminst : UP
               flags : OPEN, INIT, LINKUP, NOTIFY_INIT, ENABLE, USING_DEVCMD2
           ucsm name : Eth0
          spec_flags : MULTIFUNC, TRUNK, ROCEV2
       mq_spec_flags :
                slot : 0
               h:bdf : 0:03:00.0
              vs.mac : 00:25:b5:06:00:0a
                 mac : 00:25:b5:06:00:0a
               vifid : 822
            vifcookie : 822
                 uif : 1
  portchannel_bypass : 0x0
                 cos : 0
                vlan : 0
          rate_limit : unlimited
            cur_rate : unlimited
           stby_vifid : 0
       stby_vifcookie : 0
 stby_recovery_delay : 0
             channel : 0
       stdby_channel : 0
             profile :
       stdby_profile :
          init_errno : 0
                 cdn : Eth0
       devspec_flags : TSO, LRO, RXCSUM, TXCSUM, RSS, RSSHASH_IPV4, RSSHASH_TCPIPV4
                 lif : 18
               vmode : STATIC
         encap mode : NONE
            host wq : [45] (n=1)
            host rq : [2003-2018] (n=16)
            host cq : [2020-2021] (n=2)
          host intr : [2752-3007] (n=256)
             notify : pa=0x141e97000/44 intr=2
          rdma_resgrp : 8
             rdma_qp : 1024
             rdma_mr : 131072
           rdma_cos : 5
        preferences : 0x0 []
       rdma dcmd wq : [28-35] (n=8)
       rdma rcmd wq : [36-43] (n=8)
       rdma dcmd cq : [2012-2019] (n=8)
        rdma pkt wq : [12-27] (n=16)
         devcmd2 wq : [44] (n=1)
```

NVMe Over FC

NVMe Over FC

Increased Performance | Seamless insertion | Phased Transition | Multiprotocol flexibility | Ecosystem support

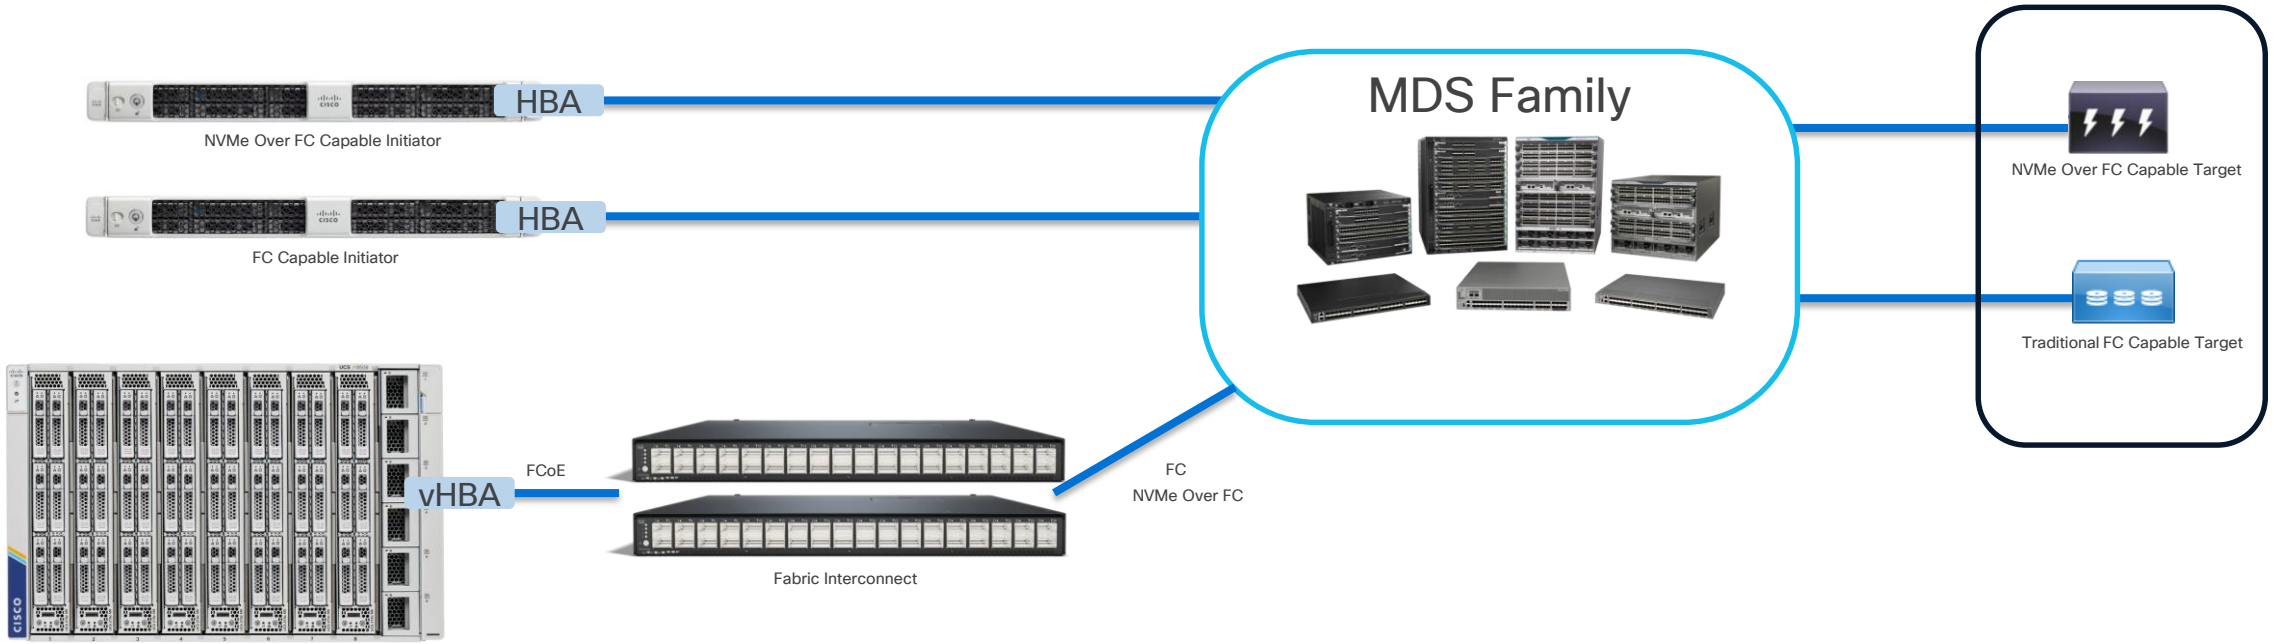

NVMe Over FC Intersight Configuration

Best Practice

- Add a vHBA to the server profile, the vHBA Type should be fc-nvme-initiator.
- Create a new Fiber Channel Adapter Policy and click on **Select Cisco Provided Configuration**.
- Select **FCNVMeInitiator**.
- Change the **SCSI I/O Queues** to **16** if the value isn't set in the Fiber Channel Adapter Policy.

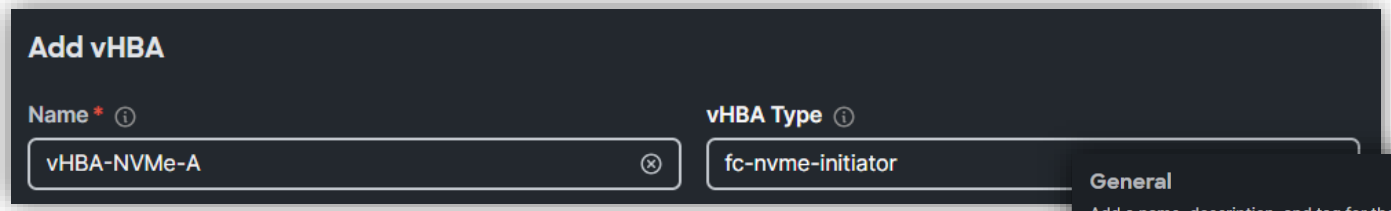

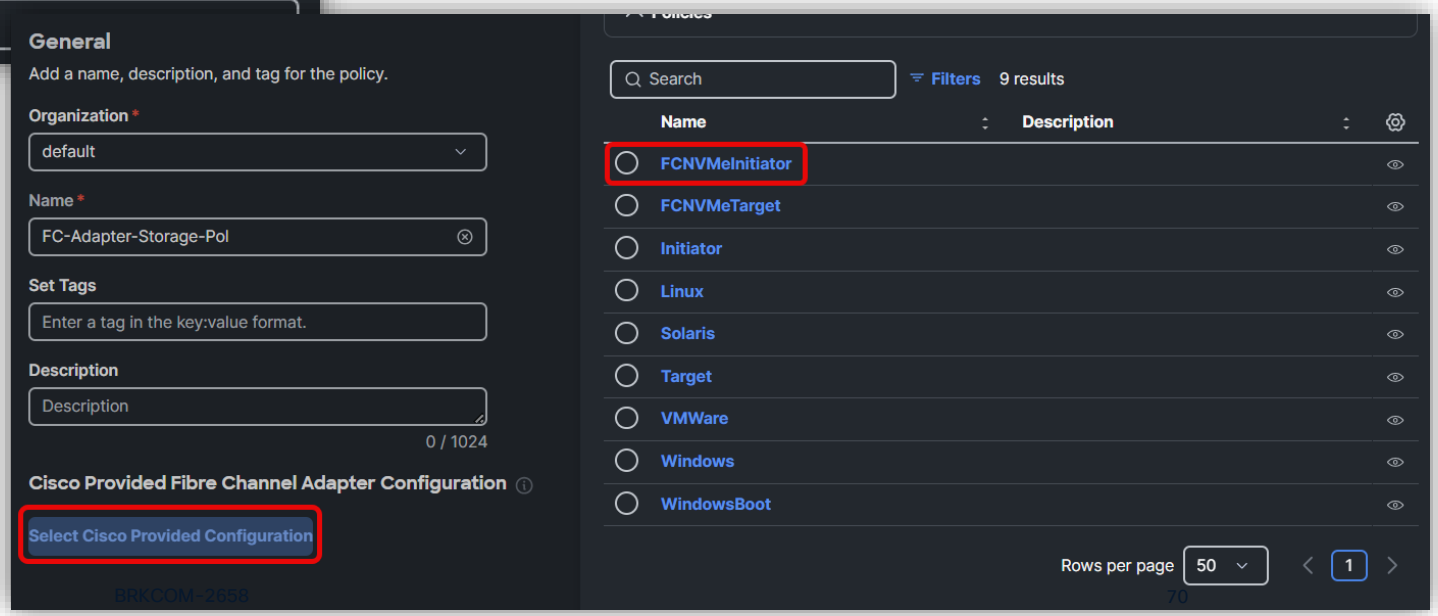

NVMe Over Fabrics Comparison

NVMe Over Fabric Types

NVMe Over TCP

General Applications
Easy to manage
Scalable

NVMe Over RoCEv2

HPC/AI/Big Data
Server Cluster
Performance
Low Latency
Access Large Storage Pool

NVMe Over FC

Enterprise Applications
Business Critical
High Availability
Performance
Reliability
Security

Types of NVMe Over Fabric

FYI

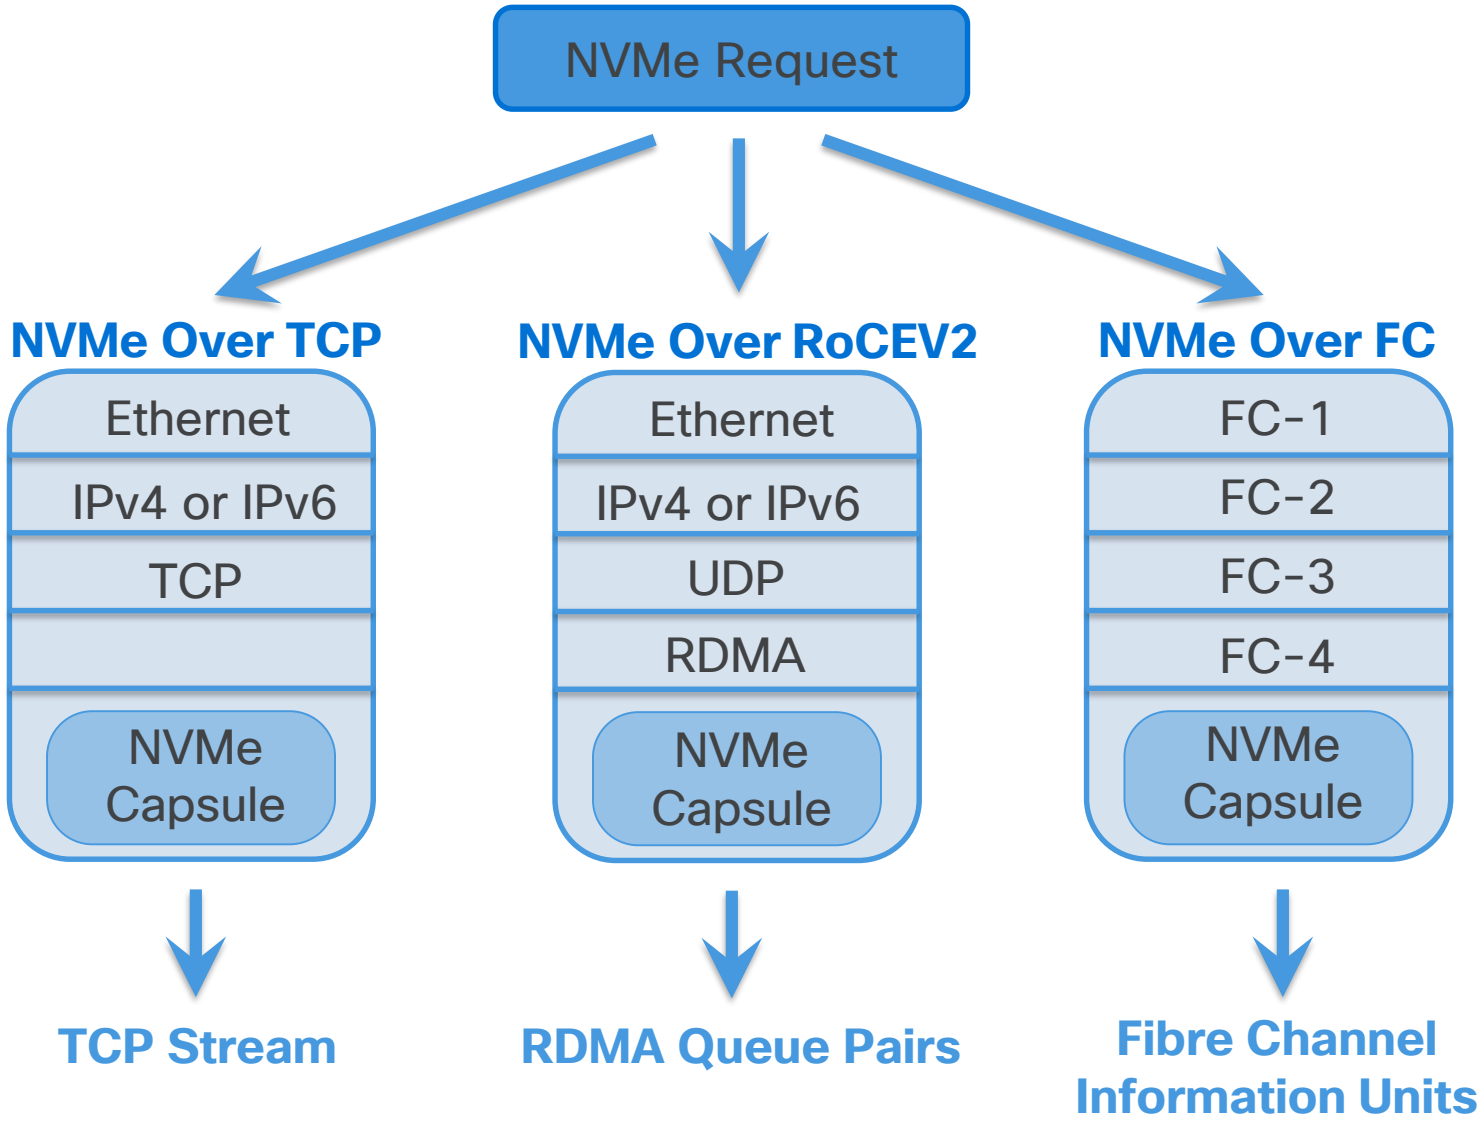

Key Takeaways

6th Gen FI:
64G FC End to End

6th Gen FI:
Replacement of FI 6454

Use NVMe for Storage

Complete Your Session Evaluations

Complete a minimum of 4 session surveys and the Overall Event Survey to be entered in a drawing to win 1 of 5 full conference passes to Cisco Live 2026.

Earn 100 points per survey completed and compete on the Cisco Live Challenge leaderboard.

Level up and earn exclusive prizes!

Complete your surveys in the Cisco Live mobile app.

Continue your education

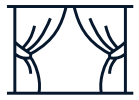

Visit the Cisco Showcase for related demos

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

Contact me at: jvdmade@cisco.com

Thank you

cisco Live !


CISCO