Evolving data center optics to higher speeds

100G and 400G considerations

Cisco Knowledge Network, October 13, 2022

Paymon Mogharabi, Product Manager, Cisco Transceiver Module Group
Rick Gomez, Cisco Global Optics Strategy Leader
Yoav Schreiber, Customer Solutions Marketing Manager

Agenda

Engage with slido

Smartphone

Web
Slido.com
Event Code
3571970

- Data center connectivity growth drivers
- Decision factors
- Switch to server connectivity
- Switch to switch connectivity
- 100G single wavelength
- How to migrate to 400G
- Next steps

Increasing demands on data centers

300%

more applications will run in the data center and edge locations

500 million

cloud-native applications and services will be developed

40%

of new enterprise IT infrastructure will be deployed at the edge

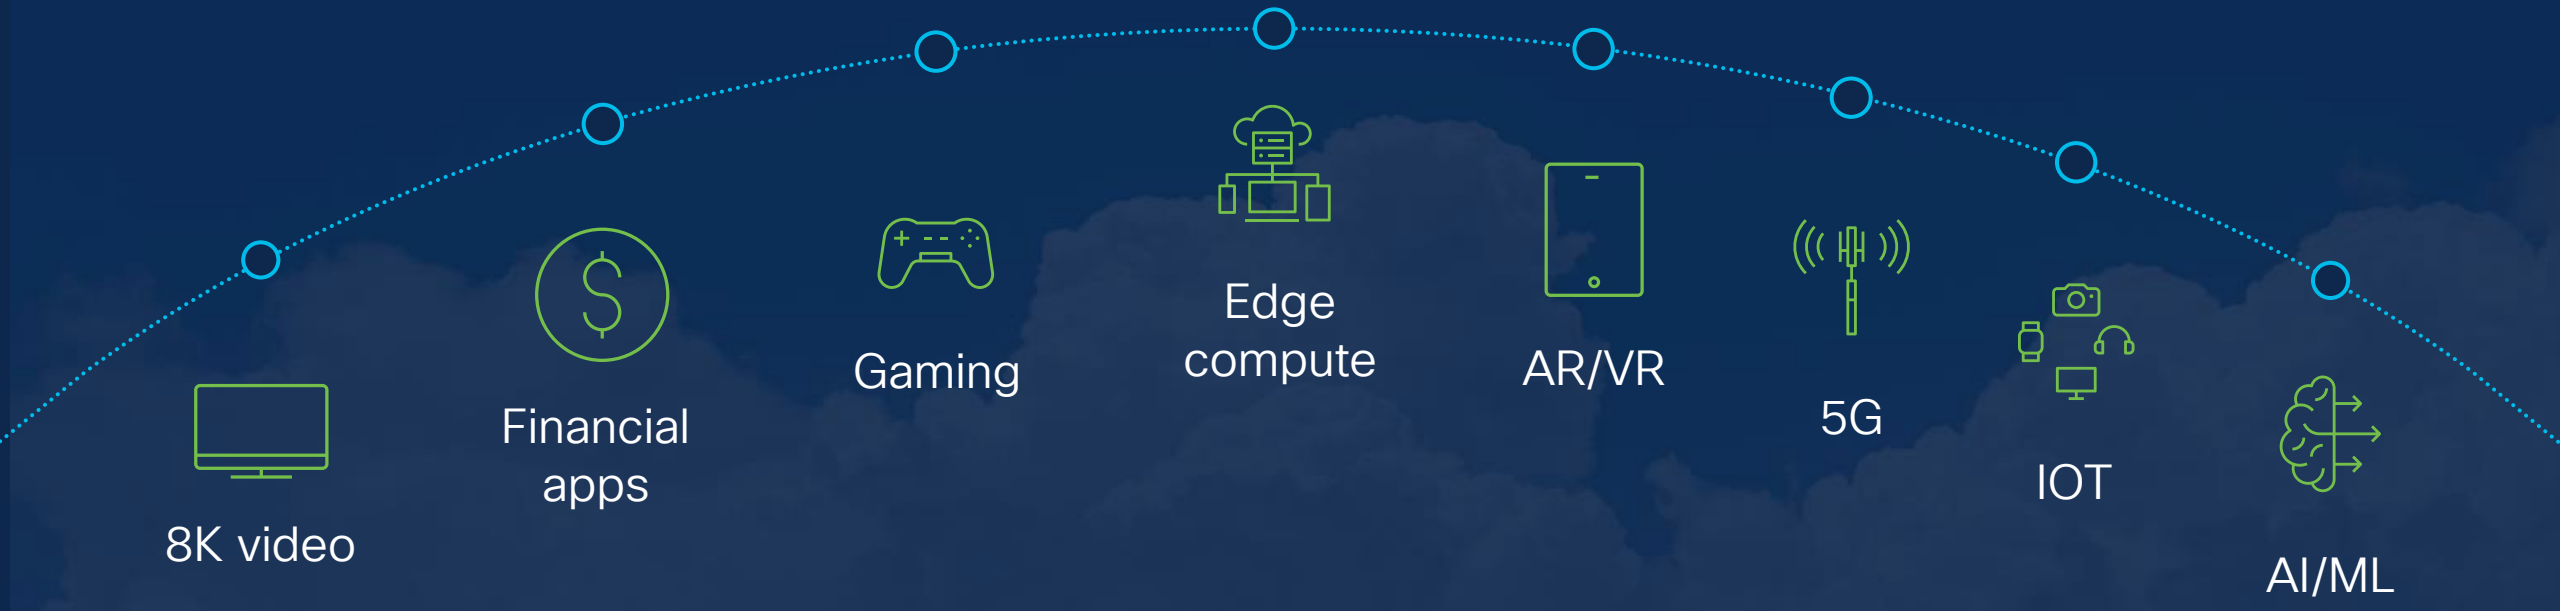

Data center operator top of mind

Increasing data center capacity and sustainability

Preserve investments in existing optics infrastructure and cabling

Simplify operations and management of optical links

Preparing for 400G capacity expansion

Requirements for denser and faster data center infrastructure

HPC/HFT

- Compute Intensive Data Processing at high speeds
- Financial services, government platforms, university R&D, oil and gas, and utilities

Edge Computing

- Edge computing - bringing compute, storage and networking, closer to the source driving a need for higher speed at the edge
- 5G, IOT, high-speed mobile network, gaming, streaming services

Infrastructure Virtualization

- Virtualized infrastructure drives VM proliferation, placing great stress on server connectivity
- Network, storage, data, CPU, and GPU pooling for optimal use of resources

Machine Learning and AI

- Mission-critical workloads such as Artificial Intelligence and Machine Learning are driving increasing amounts of computation traffic in the data center
- Retail, Manufacturing, Image recognition, healthcare, smart cars and financial services

Data center server and switch ports moving to faster speeds

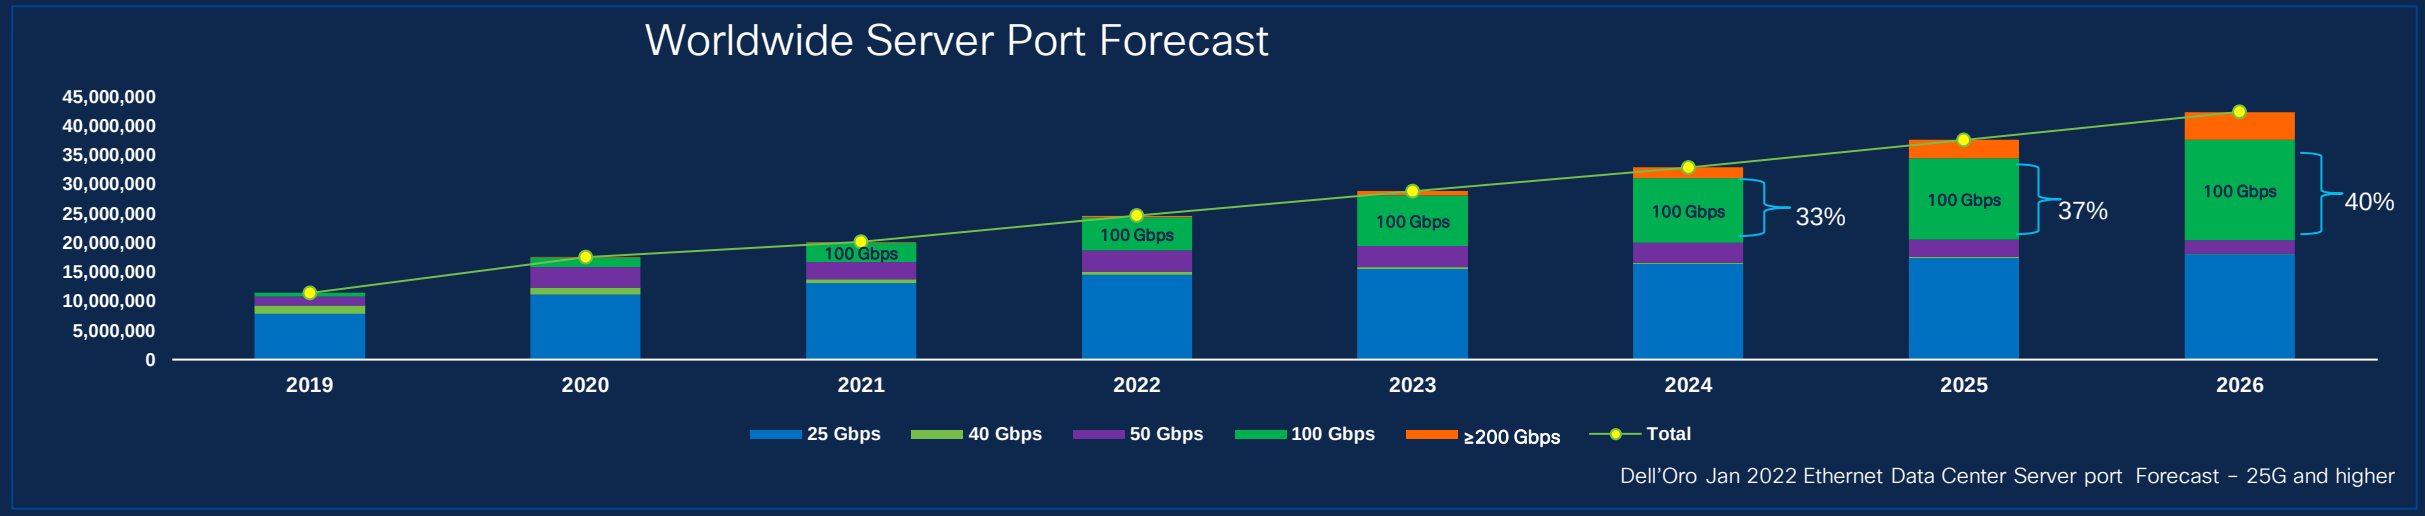

Dell'Oro Jan 2022 Ethernet Data Center Server port Forecast – 25G and higher

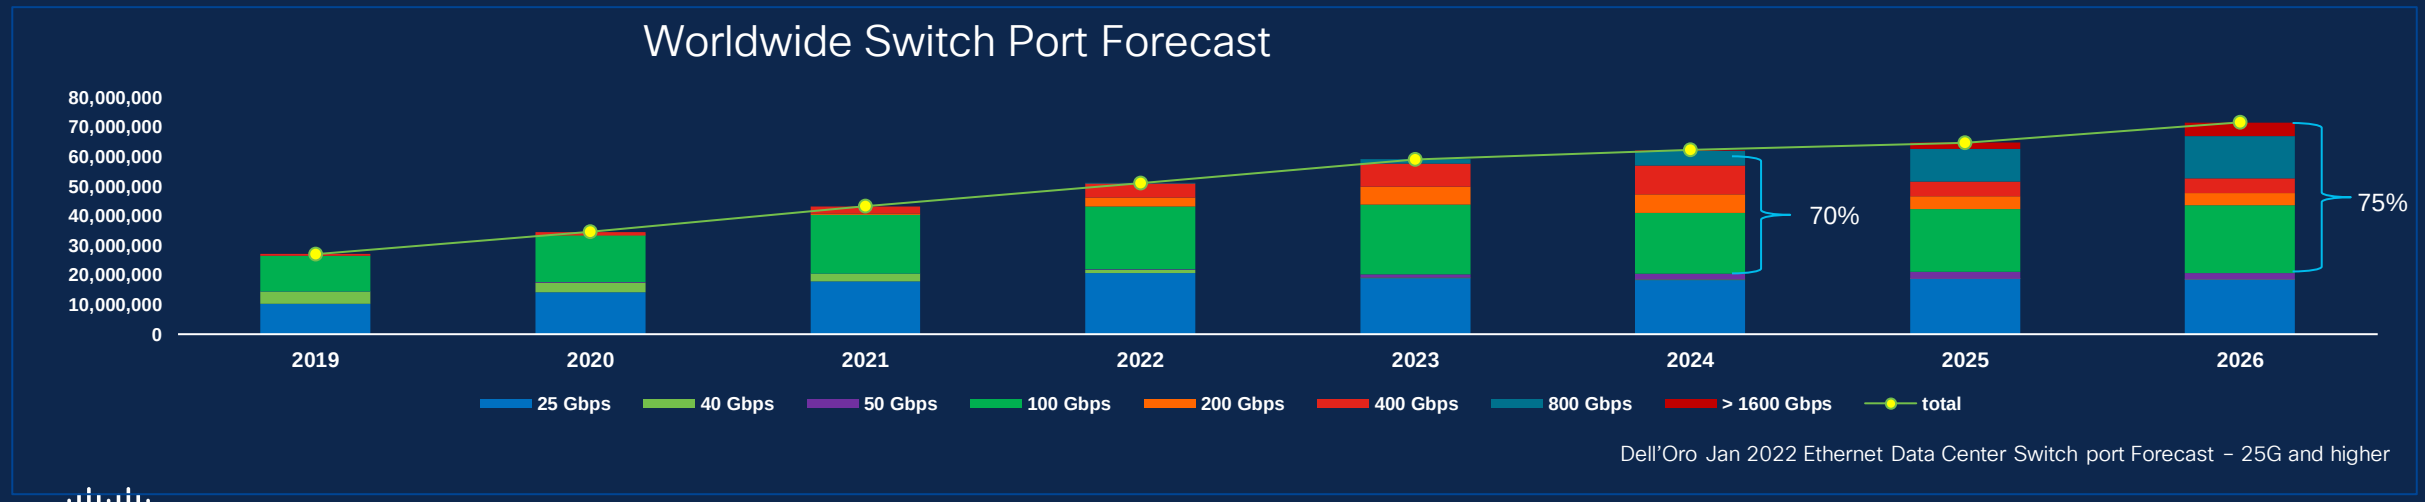

Dell'Oro Jan 2022 Ethernet Data Center Switch port Forecast – 25G and higher

Speed evolution in the data center

Switch fabric

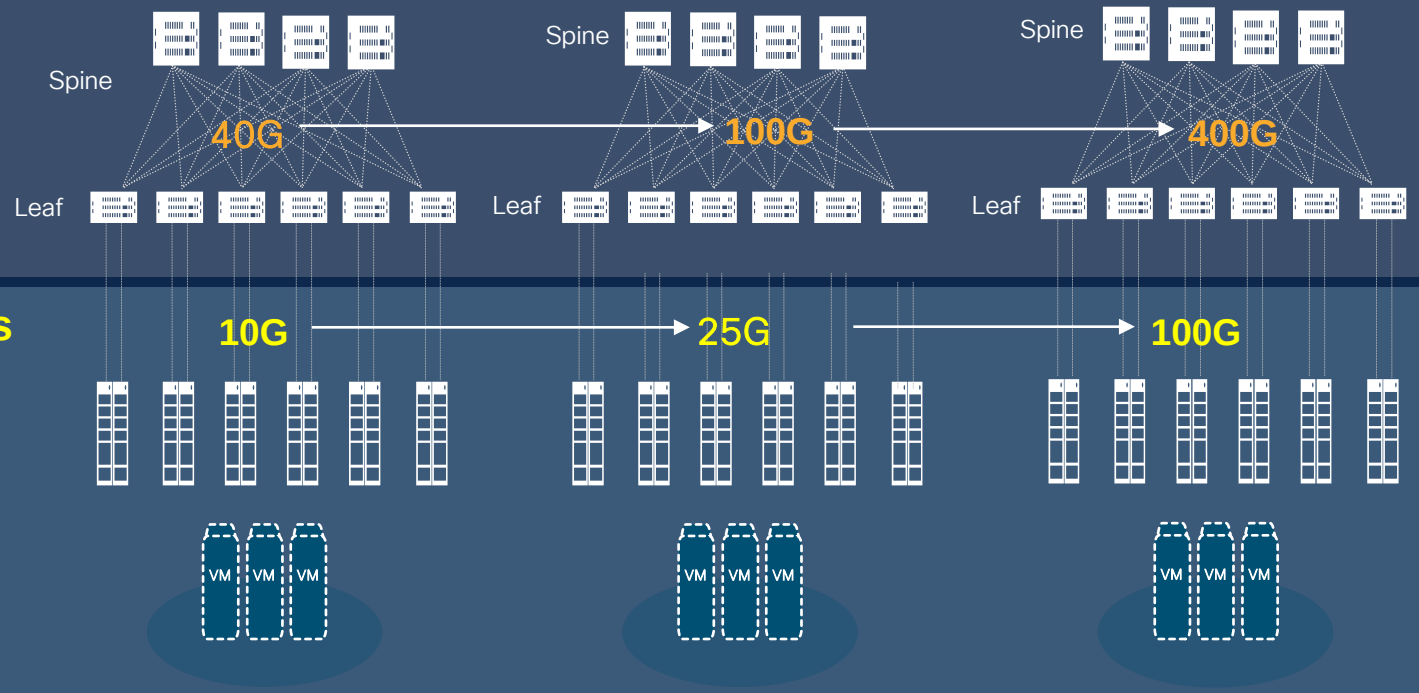

- Switch silicon bandwidth growing due to higher Radix and faster Serdes speeds
- Switch ASIC throughput growing: 6.4 Tbps to 12.8 Tbps to 25.6 Tbps to 51.2 Tbps (future)
- Optics increasing from 40Gbps to 100G Gbps to 400Gbps to 800Gbps

Servers

- Server network connectivity evolves with server processor upgrade cycles as data center traffic grows
- Server port speed is transitioning from 1/10 Gbps to 25 Gbps to 100 Gbps

Decision factors

Decision Factors

Industry Evolution	Technology Options	Return on Investment	Capacity Expansion
Maturity and falling cost of 100GbE networking equipment	Preassembled vs. pluggable	Maintain existing cable infrastructure	Maximize port bandwidth
Increasing use of 100G NICs for 4x server uplink expansion	Copper vs. fiber	Optics backwards compatibility	Prepare for 400G

Evolution in NIC and server performance

PCIe bandwidth expansion driving higher Ethernet port speeds in the NIC

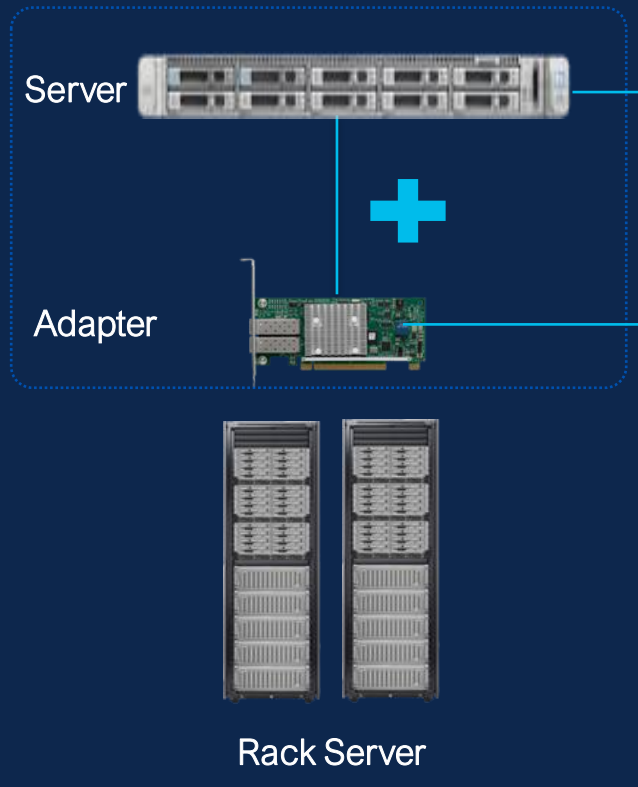

- Faster CPU
- More and faster memory speeds
- Higher Power

- Increase PCIe Lane performance (GEN 3.0/4.0 and future Gen5.0)
- Increase PCIe Lanes from x8 to x16 and x32
- Increase uplink interface speeds from 10G to 25/100G/200G and 400G

Ports	Modulation	BW per Lane (Single direction)	x8	x16	x32
PCIe Gen2	NRZ	5 Gbps	40 Gbps	80 Gbps	160 Gbps
PCIe Gen3	NRZ	8 Gbps	64 Gbps	128 Gbps	256 Gbps
PCIe Gen4	NRZ	16 Gbps	128 Gbps	256 Gbps	512 Gbps
PCIe Gen5	NRZ	32 Gbps	256 Gbps	512 Gbps	1,024 Gbps
PCIe Gen6	PAM-4	64 Gbps	512 Gbps	1,024Gbps	2,048 Gbps

Increasing data rates drive shift from copper to fiber cables

Inter Building (2km)

Within Building (500m)

Within Floor(100m)

Inter Rack MOR/EOR (15m-30m or more)

Inter Rack TOR (5m-10m)

Inter Rack (3m-5m)

Within Rack (2m-3m)

Within Rack (1m-2m)

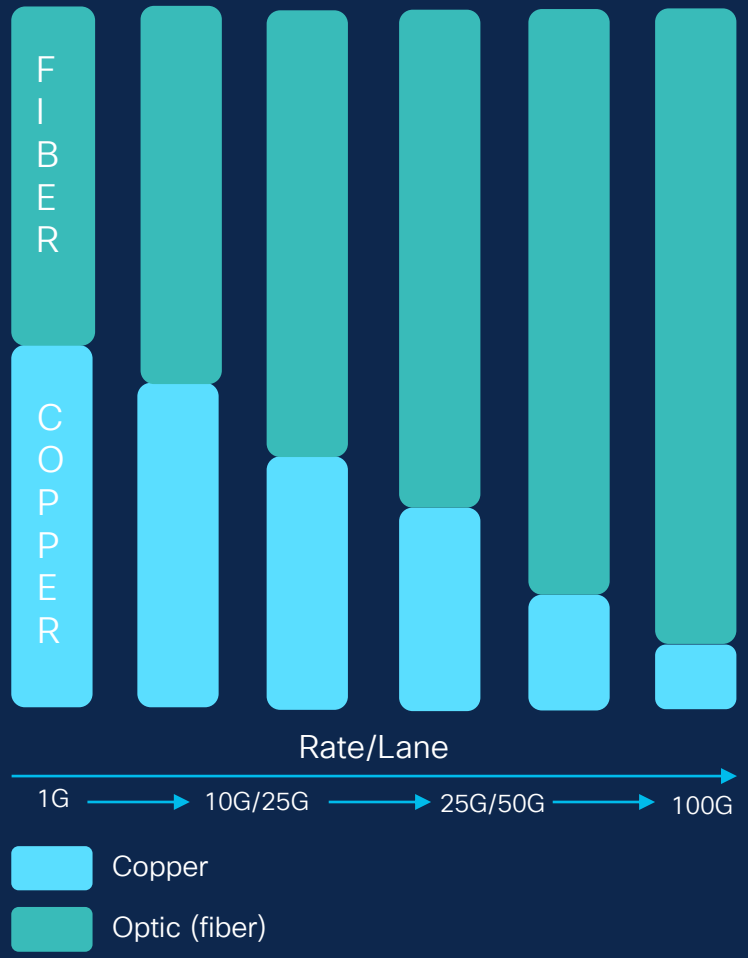

Why the move to fiber and pluggable?

- Copper Reach limitation at higher speeds
  - Reach over passive copper cable may be less than 2m for 100Gb/s lanes
- Copper weight and thickness limitation

NIC Ports	SerDes Data Rate	Modulation	Passive Cu Cable Reach (IEEE)
25G	1 x 25G	NRZ	5m
100G	4 x 25G	NRZ	5m
200G	4 x 50G	PAM-4	3m
400G	4 x 100G	PAM-4	2m

Switch to server connectivity

Switch to server connectivity options

Low Cost

Longer Reach
Space Optimized

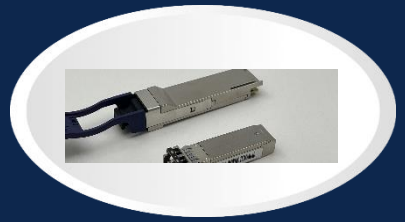

Longer Reach,
Space Optimized
Superior Manageability

Longer Reach,
Space Optimized,
Investment protection

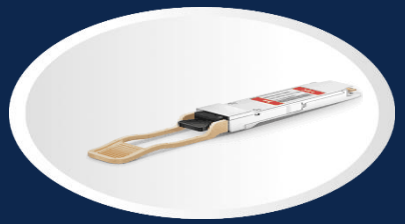

Longer Reach,
Space Optimized,
Breakout to 25G

	DAC	AOC	Short Link (SL)	SR1.2 BIDI	Short Reach (SR4)
Product	QSFP-100G-CU	QSFP-100G-AOC	QSFP-100G-SL4	QSFP-100G-SR1.2	QSFP-100G-SR4-S
Connector Type	cable	cable	MPO	LC	MPO
Speed	100G Only	100G Only	100G Only	100G Only	100G Only
Placement	Inter Rack	Inter and rack to rack	Inter and rack to rack	Rack to rack	Rack to rack
Weight/space	Heavy/Bulky	Light	Light	Light	Light
Field Replaceable	Entire Cable	Entire Cable	End-points only	End-points only	End-points only
Reach	0-5m	0-30m	0-30m	0-100m	0-100m
Cost/distance	Lowest Cost /Distance	Low cost	Low Cost	More Expensive	Most Expensive
Investment Protection	No	No	Yes	Yes	Yes
Connectivity to 400G	No	No	No	Yes	No

Why Short Link for 30m switch to server connectivity

Lower capex

An upgrade to higher speeds can be done while maintaining existing infrastructure

Reduce opex

Improved serviceability with field-upgradable end-points

Further reach and less bulky than copper cabling

Compatibility

Qualified on Cisco host platforms.

100G Short Link SL4 and 25G SL

Q-S: 100G to 4 x 25G

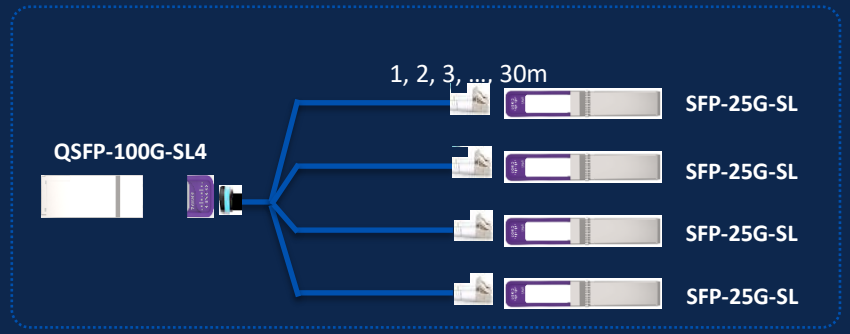

Q-Q: 100G to 100G

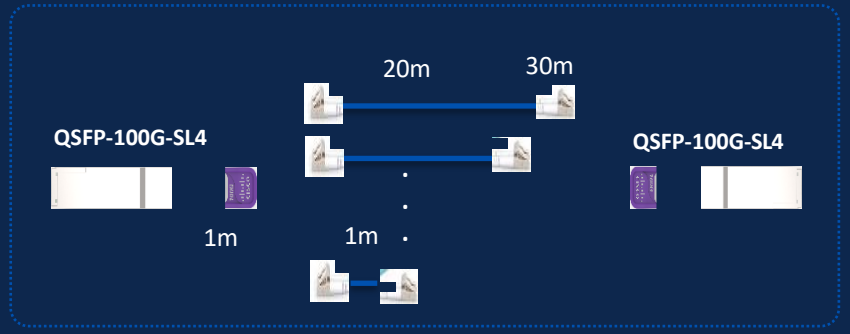

S-S: 25G to 25G

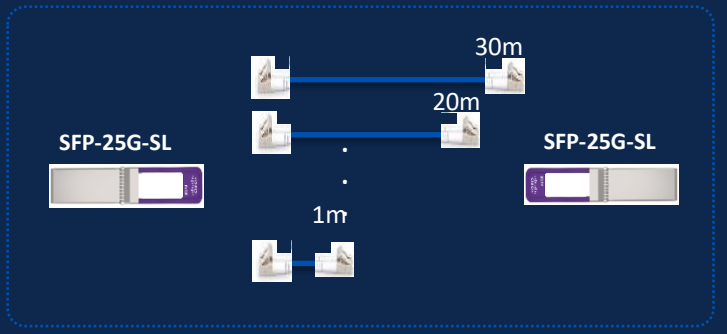

- Set of low-cost transceivers offered for data center server connectivity for lengths of up to 30m
- Option to purchase patch cords from Cisco

Connectivity options:

- Q-S: QSFP-100G-SL4 to 4 x SFP-25G-SL (Breakout)
- Q-Q: QSFP-100G-SL4 to QSFP-100G-SL4
- S-S: SFP-25G-SL to SFP-25G-SL

Data Center Topology for 100G server connectivity

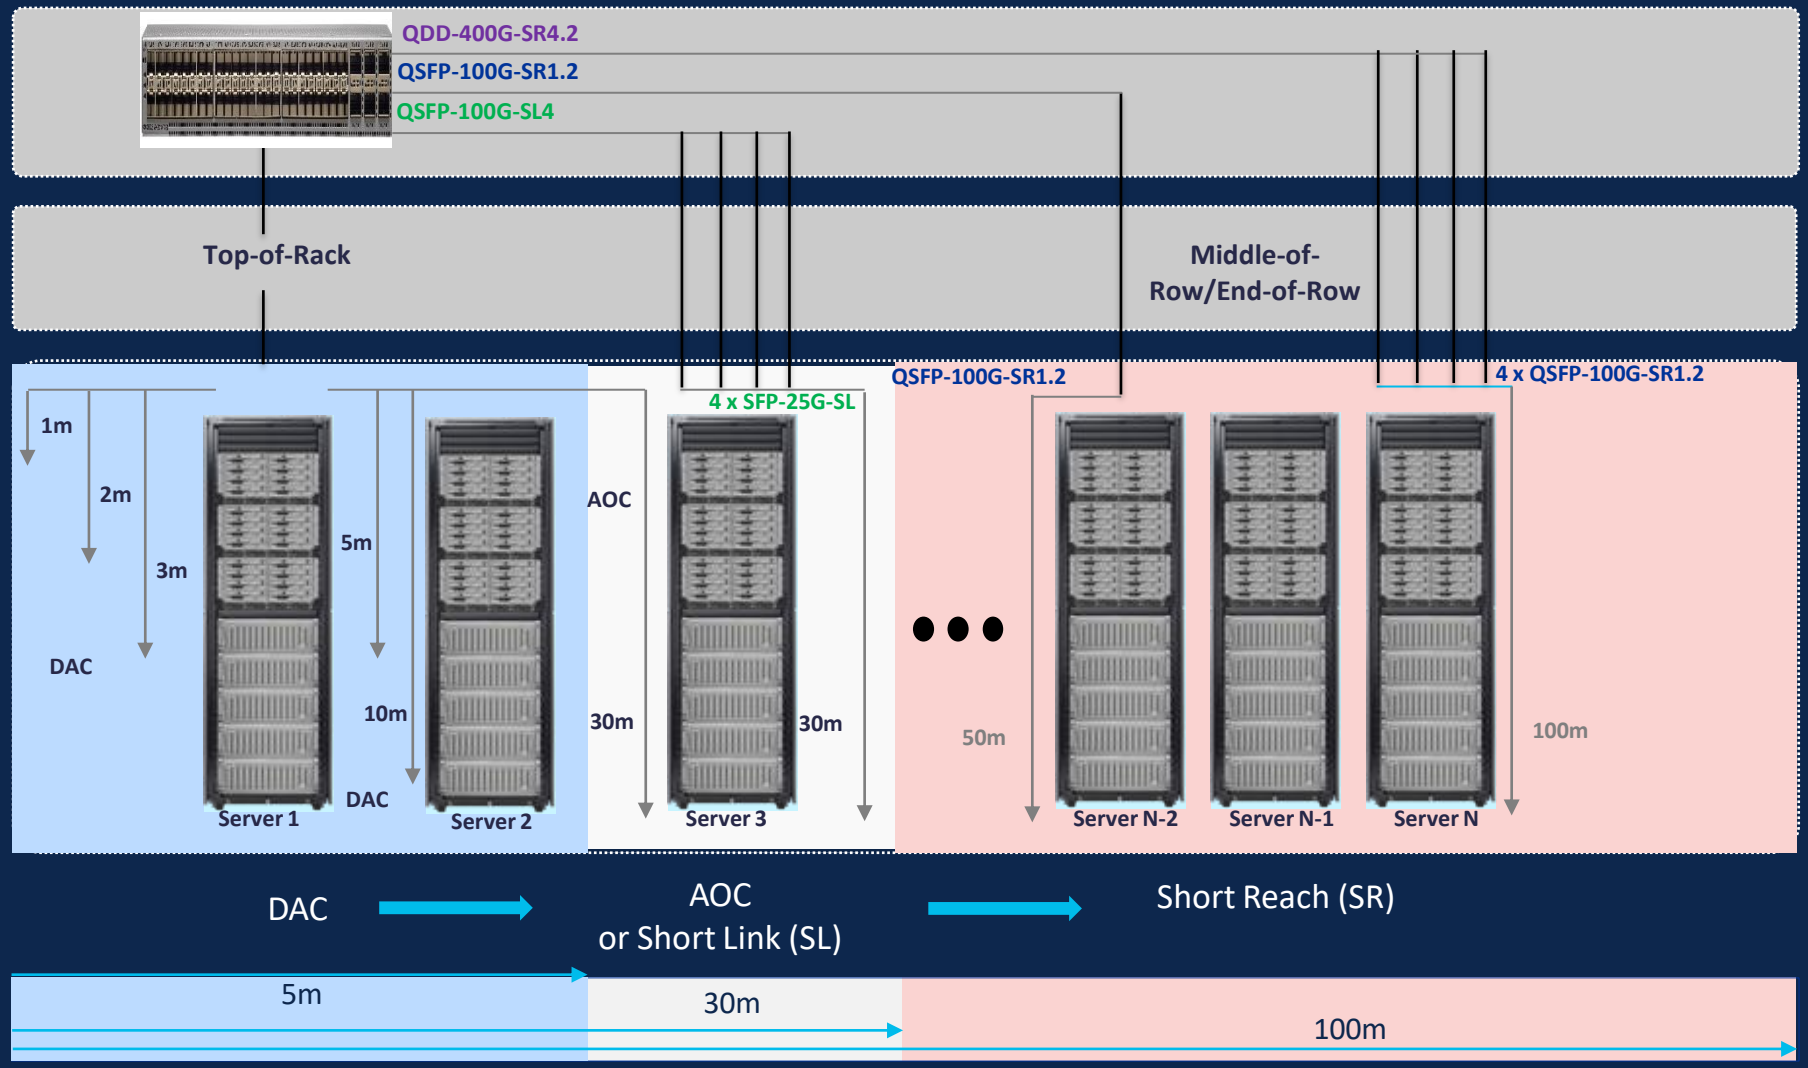

Switch to switch connectivity

Switch to switch connectivity options

40G BIDI

40/100G BIDI

100G SR1.2

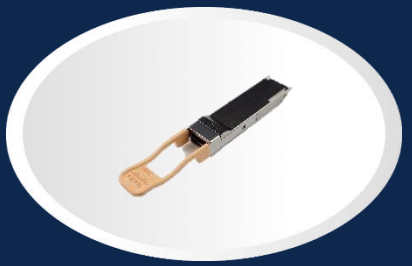

100G SR4

Product	QSFP-40G-SR-BD	QSFP-40/100-SRBD	QSFP-100G-SR1.2	QSFP-100G-SR4-S
Reach 40G OM3/OM4 100G OM3/OM4	100m/150m N/A	100m/150m 70m/100m	N/A 70m/100m	N/A 70m/100m
Connector	Duplex LC	Duplex LC	Duplex LC	MPO
Breakout	No	No	No	Yes
40G Rate	Yes	Yes (40G mode)	No	No
100G rate	No	Yes	Yes	Yes
Connectivity to 400G	No	No	Yes – (4x100G Breakout)	No

Why 100G SR1.2 BiDi

Lower capex

- Reuse existing 10/40 or 100G duplex LC multimode fiber (MMF) infrastructure
- Support link lengths of 70 and 100 meters on OM3, OM4 and OM5 multimode fibers

Reduce opex

- Seamless upgrade to 100Gbps rate from 10Gbps SR or 40Gbps BiDi

Future ready

- Provides connectivity to 400 Gbps data rate in a 4 x 100G breakout mode.

400G migration path with SR1.2

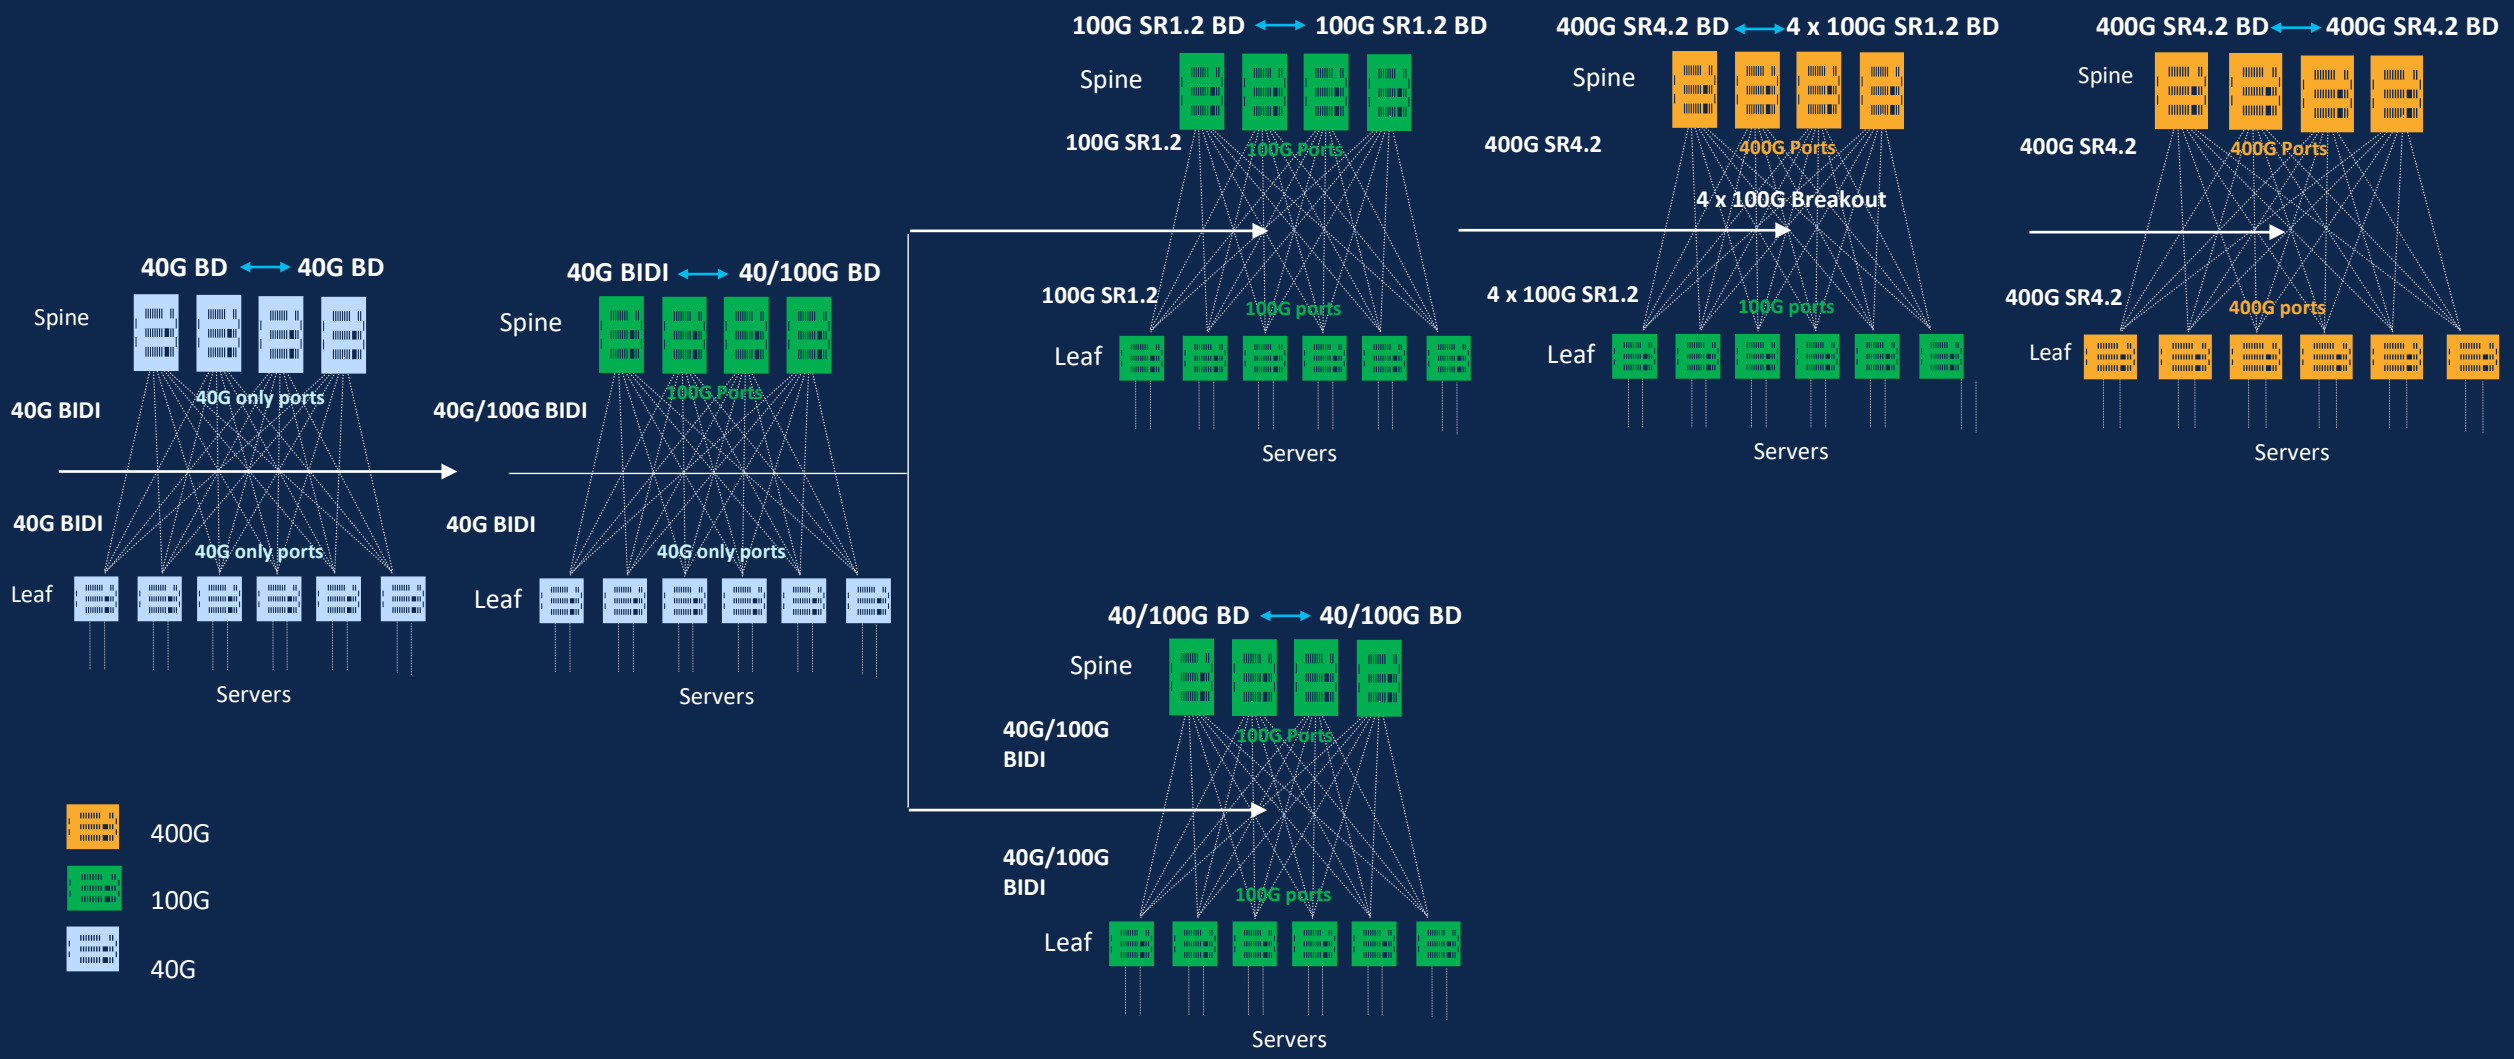

Optics interoperability and backwards compatibility

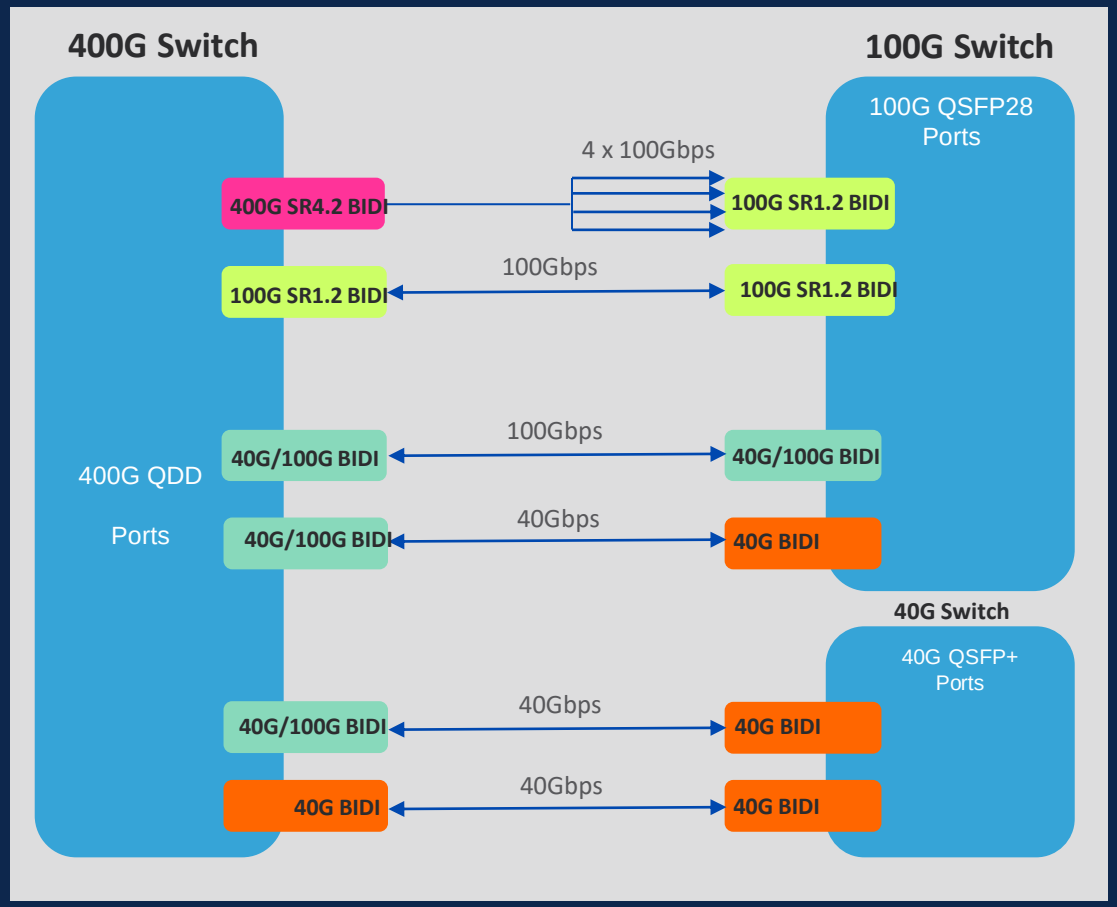

- 400G interoperability
  - Connect QSFP-400G-SR4.2 to QSFP-100G-SR1.2 using 4 x 100G breakout
- QSFP-DD port backwards compatibility
  - SR4.2, SR1,2, 40G/100G BiDi, 40G BiDi
- QSFP28 backwards compatibility
  - SR1.2, 40G/100G BiDi, 40G BiDi

100G single wavelength

Why 100G single wavelength optics

- Interoperability with 400G QSFP-DD DR4 and others
- Fully utilize 400G port bandwidth via breakout

- Single-wavelength reduces component count
- Single lane data path aligned with future single-lane 100G ports (SFP112)

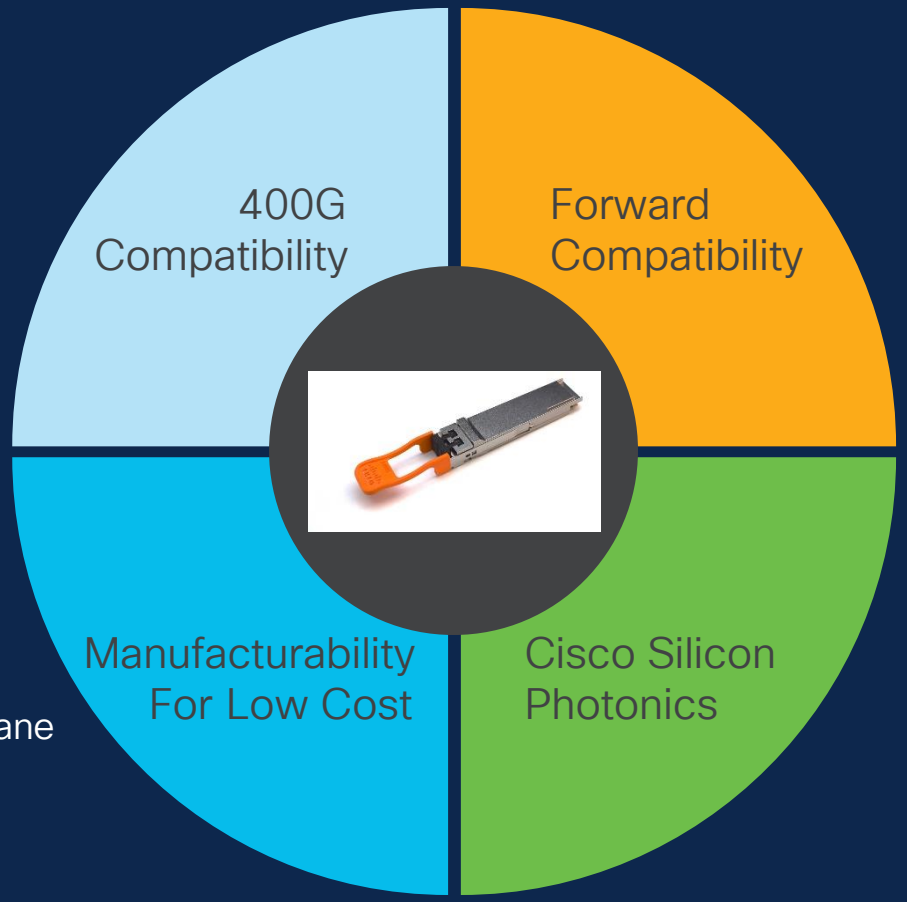

- Optically interoperable with future 100G form factors
- Future proofs investment in 100G network

- Leverages Cisco investment in optics technology
- ER-Lite 25km reach is Cisco value-add

Complete portfolio for 100G single wavelength optics

DR (500m)

QSFP-100G-DR-S

- Leaf switch connectivity to high density spine switch
- Nexus bundles

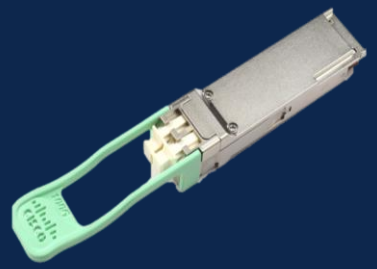

FR (2km)

QSFP-100G-FR-S

- Alternative to CWDM4
- 400G compatibility for incremental upgrade

LR (10km)

QSFP-100G-LR-S

- Alternative to LR4
- 400G compatibility for full utilization of port bandwidth

ER-Lite (25km)

QSFP-100G-ERL-S

- Low-cost alternative to ER4-Lite, CFP2-DCO
- Cisco Silicon Photonics performance advantage

400G migration with single mode fiber

Step 1: Prepare for 400G journey with 100G Single Lambda optics

Existing 100G network

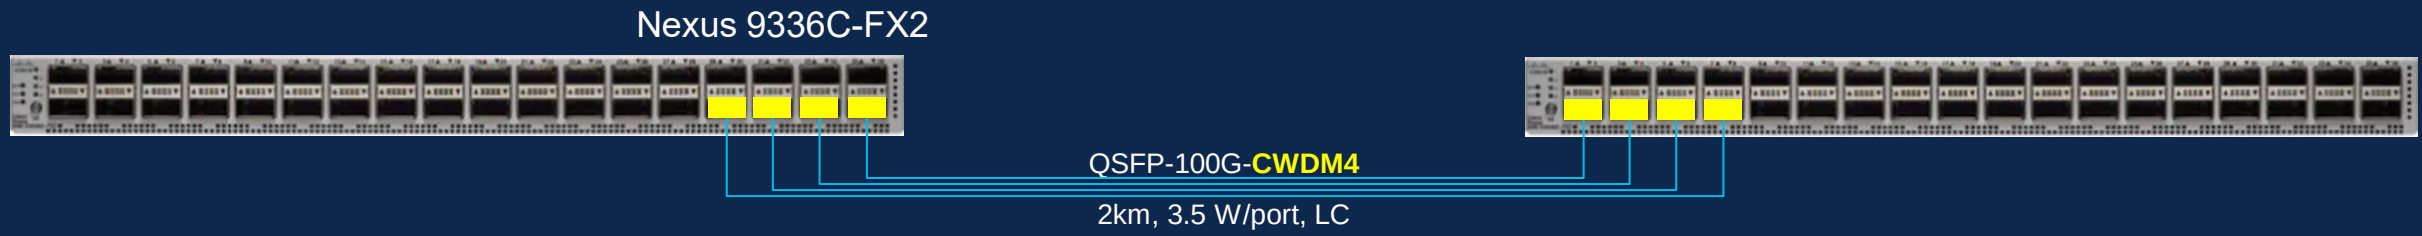

- SMF Cabling, LC Connector

Transition to 100G Single Lambda, 400G ready

- Transition to QSFP-100G-FR-S
- Reuse SMF cabling, LC connector

Step 2: Migrate one side to 400G Switch

✓ Improve efficiency: Use 100G Single Lambda

- Top 400G Switch: N9K-C93600CD-GX
- Reuse QSFP-100G-**FR**-S
- Reuse SMF cabling, LC connector
- Power consumption: 4x 4.3W = 17.2W

✗ Avoid low efficiency: don't use 100G Legacy Optics on 400G ports

- Top 400G Switch: N9K-C93600CD-GX
- CAUTION: don't use QSFP-100G-CWDM4-S on 400G port
- Legacy optics use 4x25G optics, not able to do 4x100G

Step 3: Maximize your port with 4x100G or 2x100G breakout

✓ Maximize Ports with 4x100G breakout

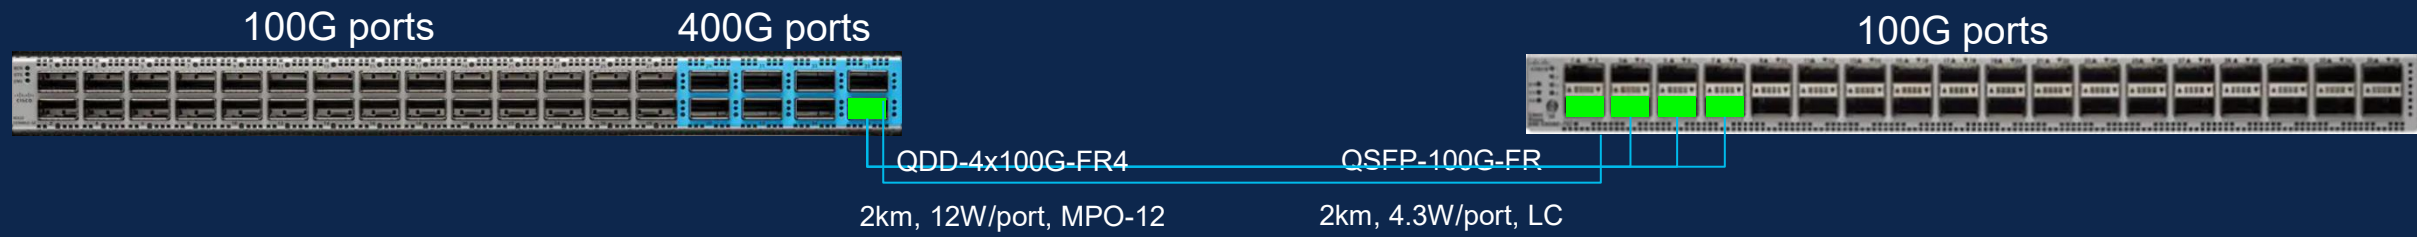

- Maximize efficiency: transition to 400G Optics,
- CAUTION: 400G port is MPO-12 connector, use Breakout cable for MPO to LC conversion
- Save 5W per 400G. 12W on QDD-4x100G-FR vs 17.2W for QSFP-100G-FR
- Save 3 ports

✓ If you must use legacy, maximize with 2x100G

- Use when remote end only supports legacy optics
- CAUTION: 2x100G optic is CS connector, use Breakout cable for CS to LC conversion
- Saves 2 ports

TMG Matrix to validate optics compatibility and interoperability

Go to: https://tmgmatrix.cisco.com

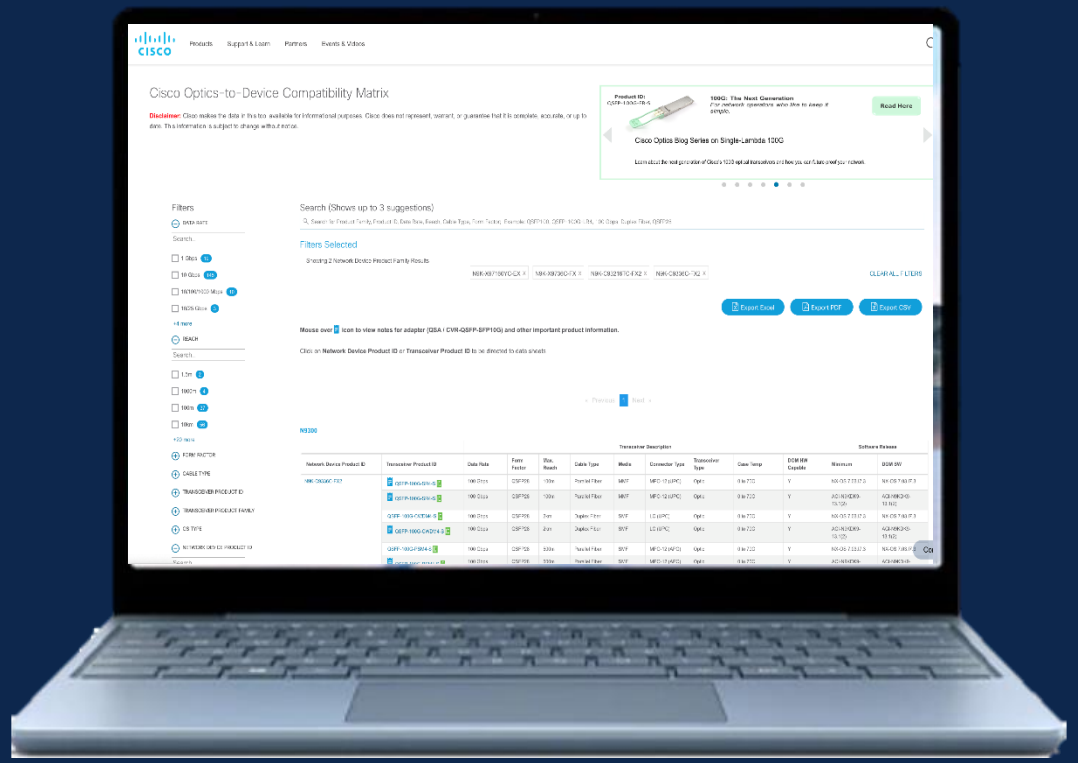

- ✓ Cisco hardware
- ✓ Cisco validated optics specifications
- ✓1-click to data sheet

Get Started Today

3 Get help with optical transceiver selection

Go to: cs.co/askoptics

2 Discover the optical transceivers to fit your data center network design needs

Visit: tmgmatrix.com.com

1 Learn more about the impact of higher speeds on data center cabling and optics

Go to: cisco.com/go/optics

The bridge to possible