400G and 800G in the Data Center

Presented by: Hardev Singh

Director – Data Center and Campus Platform Product Management

Presented under NDA

ARISTA

Ethernet Switch – Data Center Shipments

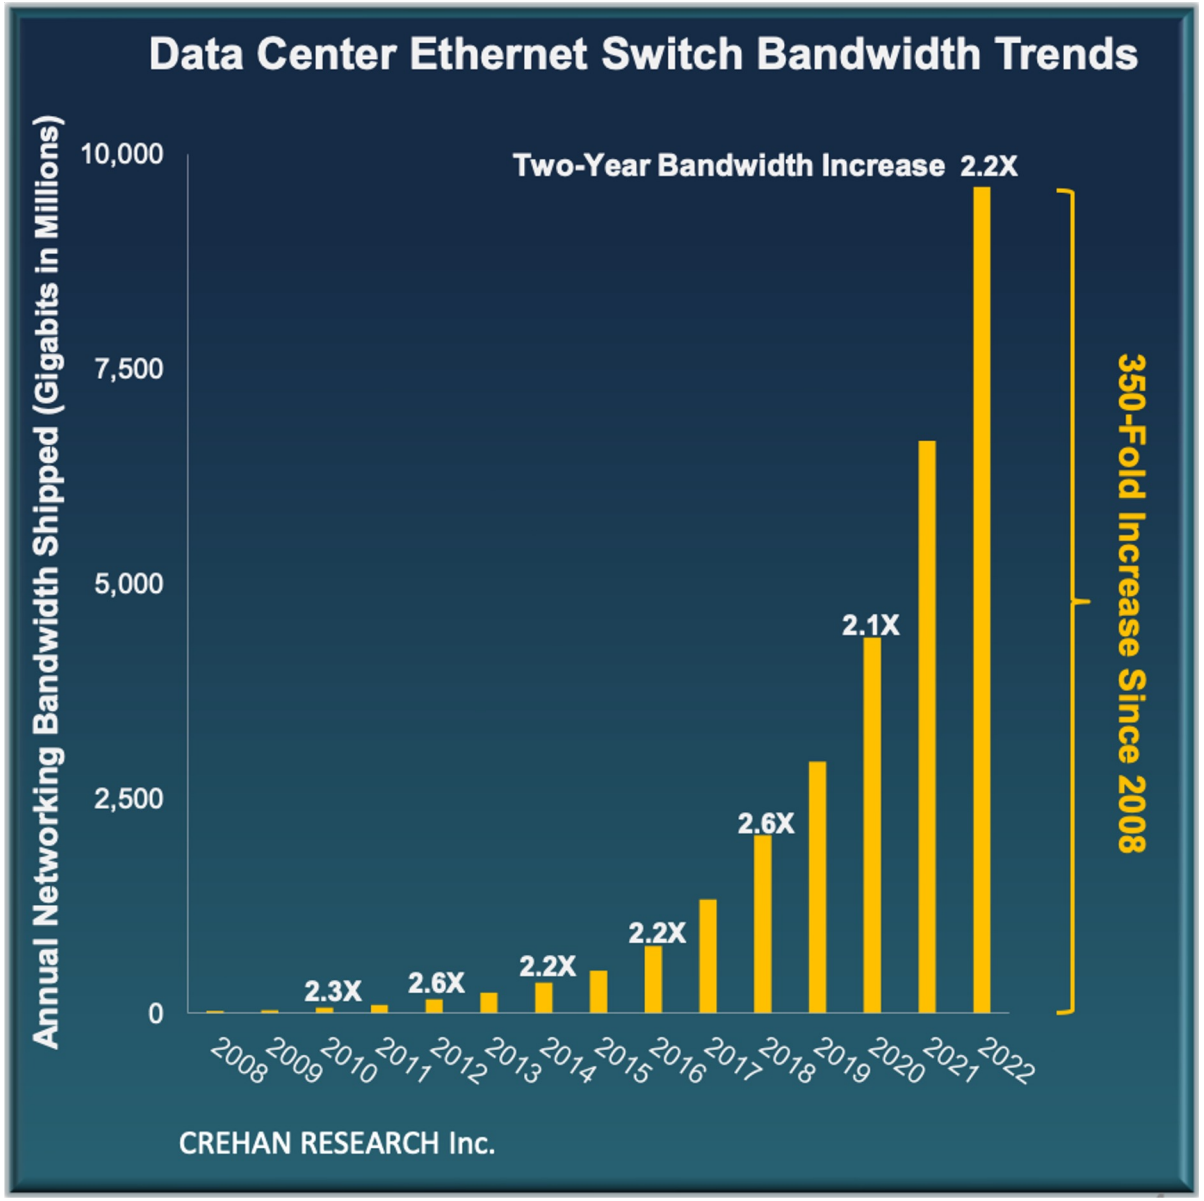

Source: Crehan Research – 2023

Networking bandwidth doubles every two years

Modern Network DC Trends and Key Drivers

Hyperscale Trends

Large scale AI/ML clusters

- Largest radix, low latency networks
- 200/400G Compute with 800G interconnect
- Backend networks for AI compute

Increase intercloud bandwidth

- 100G -> 400G -> 800G and beyond
- MACsec and IPSec at high speeds

Reduce Infrastructure: Compute

- Large scale, high diameter 2-tier networks
- Lowest power and cost per Gbps
- High Efficiency with 100G SerDes

Enterprise Trends

Migrate compute from 10/25G to 50/100G

- Performance NICs with PCIe Gen4
- Backwards compatibility
- Investment Protection

Native 400G links at Fabric layers

- Flat topologies for East-West traffic
- Systems with mix of 25G to 400G speeds

HPC clustered applications

- Requiring 100G for All Flash Array Connectivity, NVMe-oF
- Increased use of AI / ML, Cloud native and Containerized

What drives this growth?

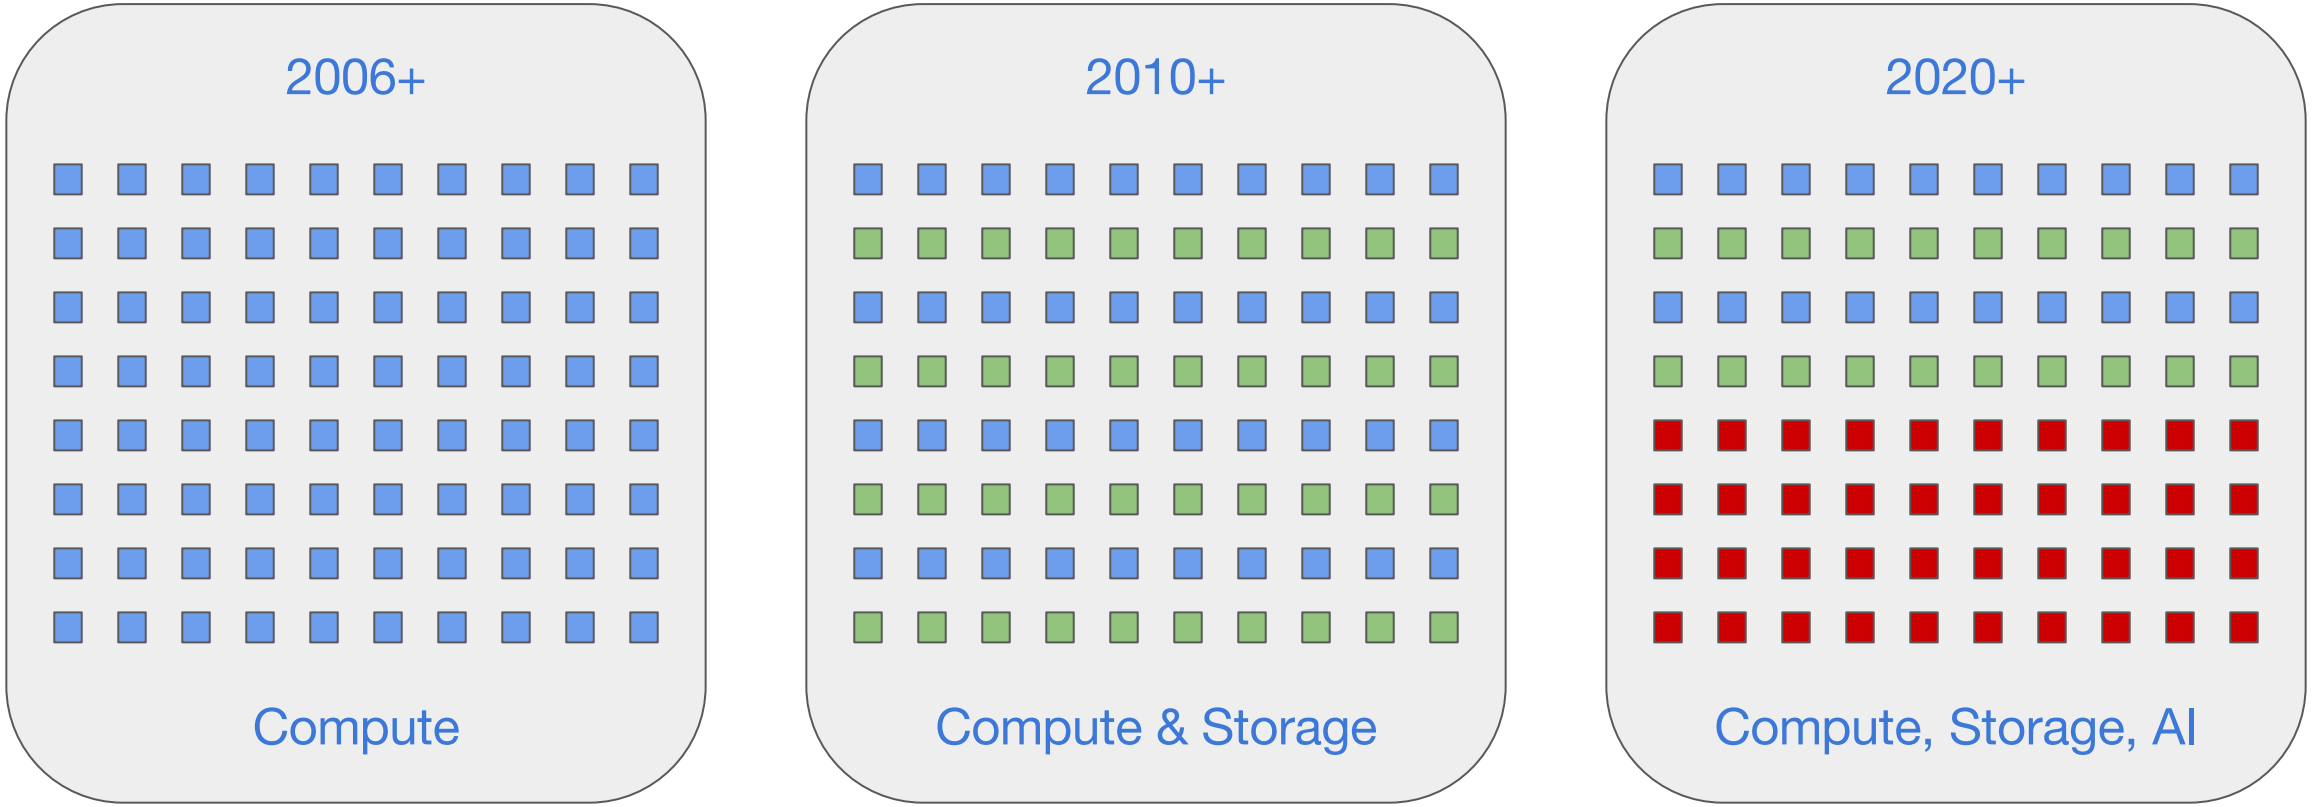

Ethernet Switch – Data Center Port Shipments

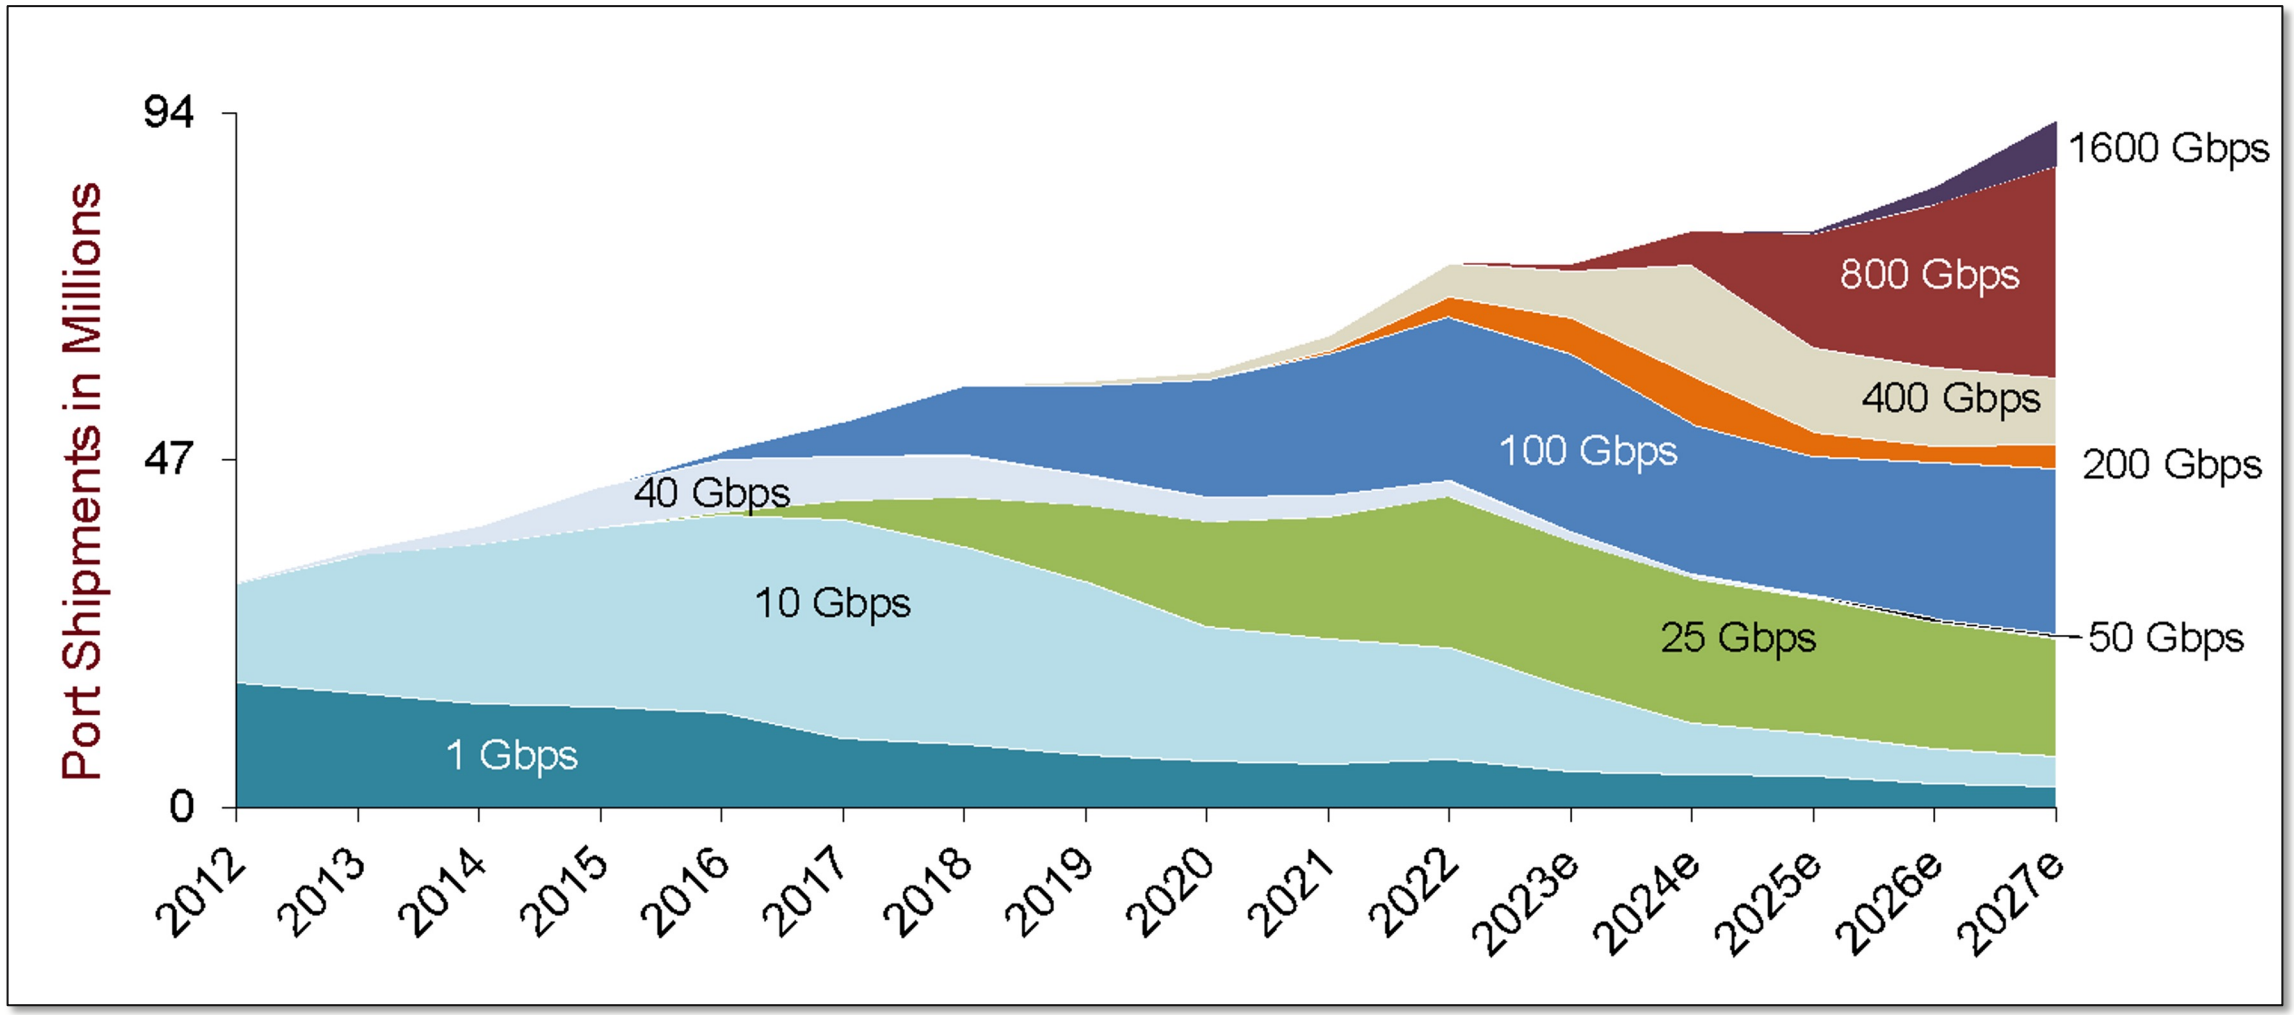

Courtesy of Dell'Oro – 2023

Presented under NDA
ARISTA

Arista's Market Leadership in 100G/200G/400G

Data Center Ethernet Switch Analysis and Forecast

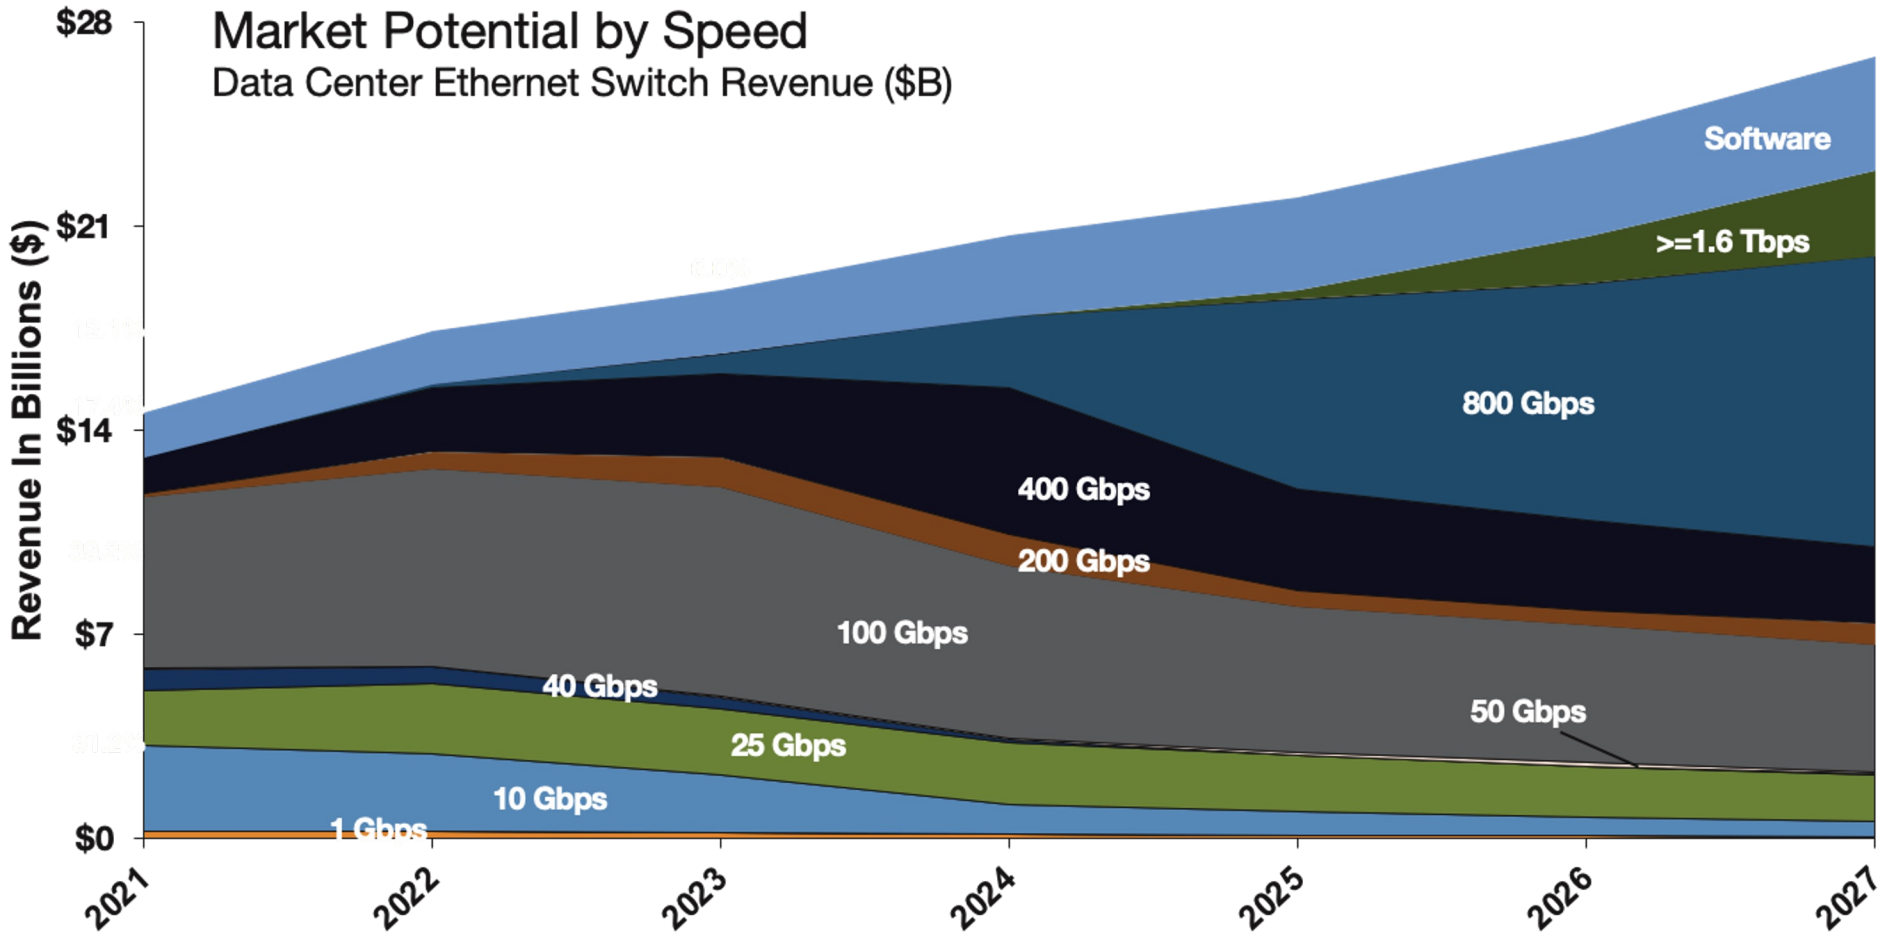

Source: Dell'Oro January 2023 - Long Term Ethernet Switch Forecast

Note - Initial 800 Gbps shipments will not be using 800 G Ethernet MAC and will be configured mostly as 2x400 Gbps or as 8x100 Gbps

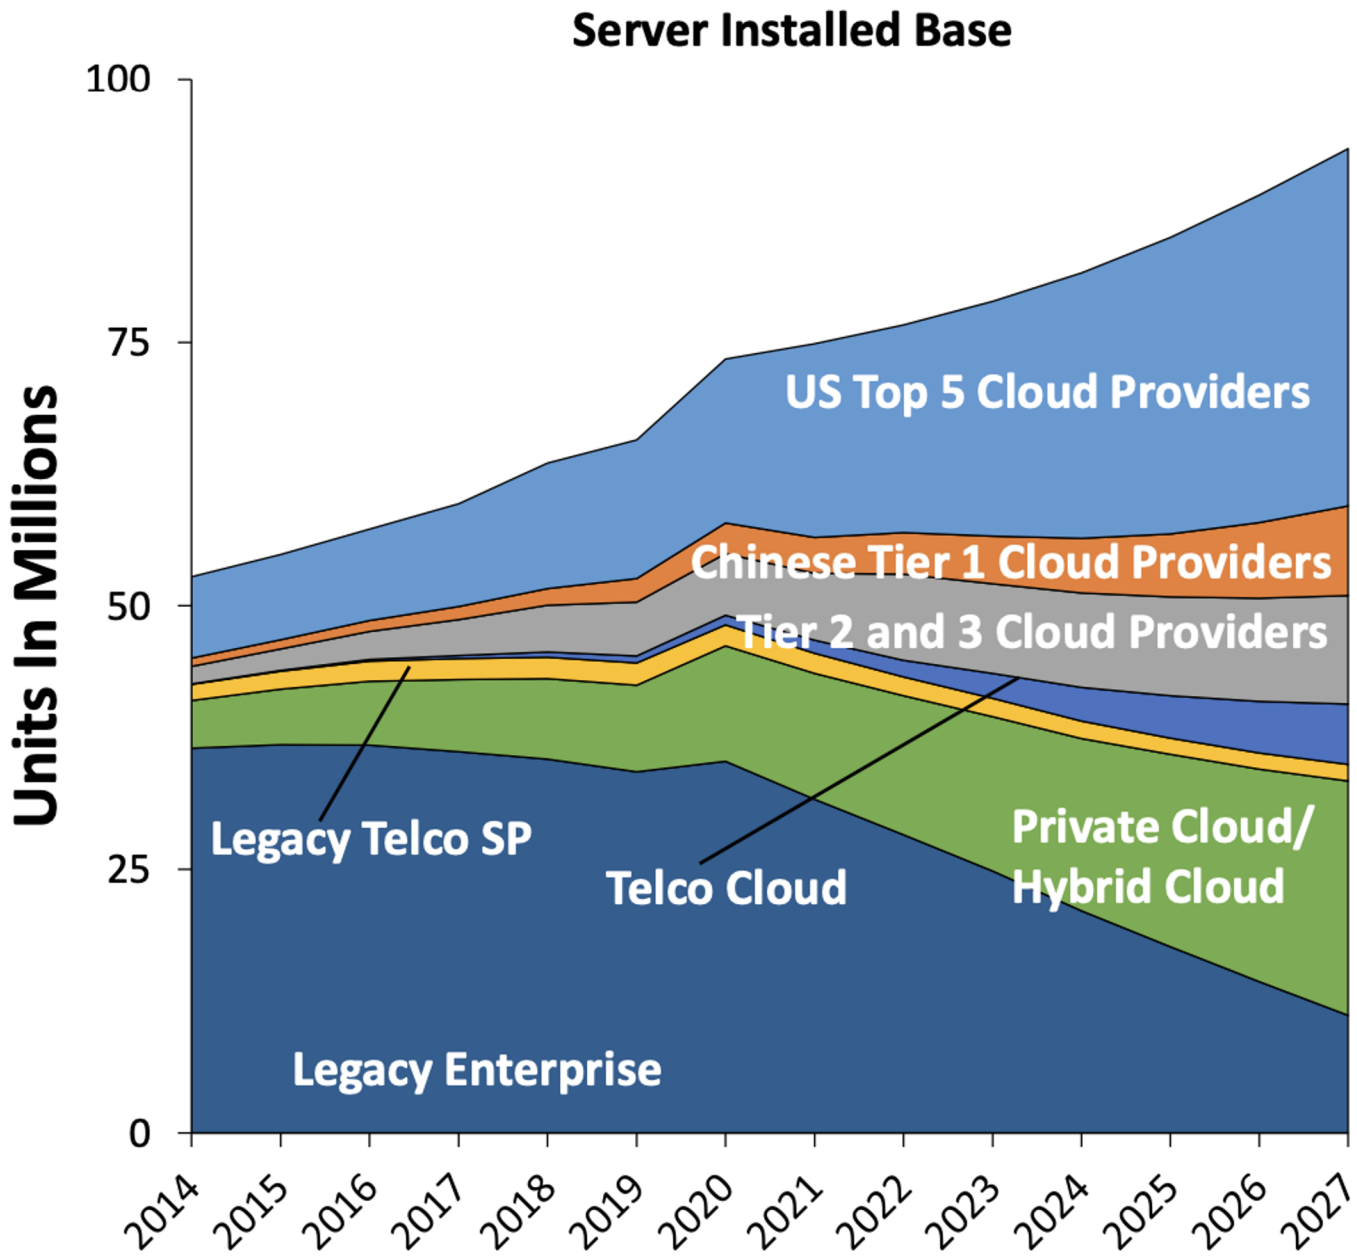

Source: 650 Group 2023 Switching Report

Is everything moving to the Cloud?

Server and NIC Speed Migrations

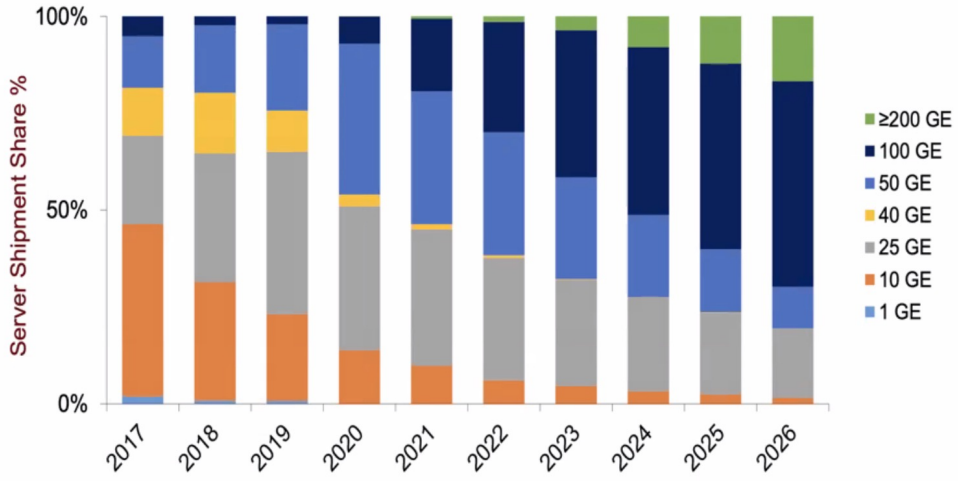

Cloud Link-Speed Forecast

Cloud: All about 100G+

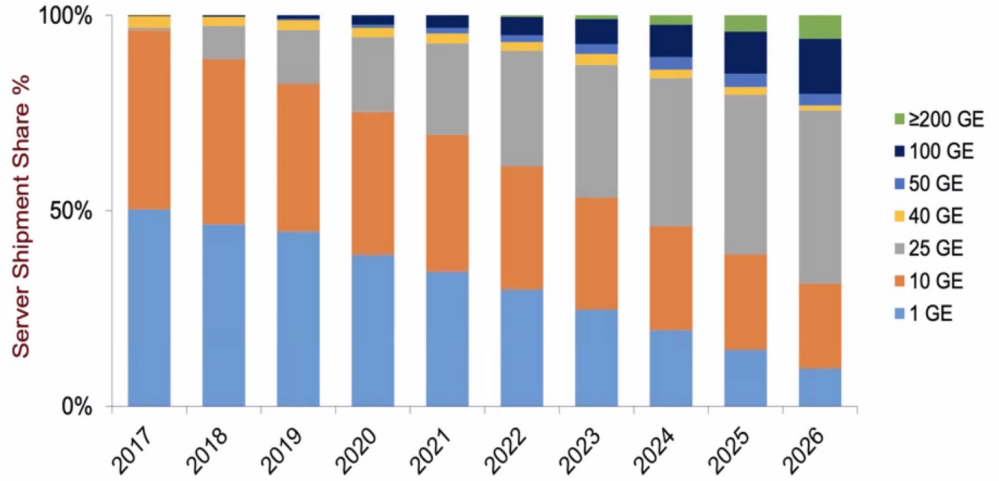

Enterprise Link-Speed Forecast

Enterprise: Mix of 10G, 25G, 100G

Both Cloud and Enterprise speed transitions underway

Network Switch demand based on Compute, Storage & HPC refresh and growth

New generation of 7050X4 and 7060X5 systems lower Opex and Capex

Source: Dell'Oro Group Inc. – Cloud Data Center Capex Servers and Storage Systems Controllers and Adapters

NIC Evolution to TOR Switch mapping

Widely Deployed

10G / 1G NRZ

10G SFP+ & 1G SFP

7050SX3

Deployed & Growing

25G NRZ

25G SFP

7050SX3

100G & 50G-2 QSFP (QSFP28)

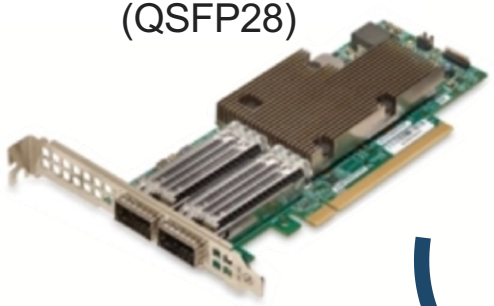

7050CX3

New Releases

50G PAM-4

50G SFP

200G & 100G-2 QSFP (QSFP56)

100G-2 SFP-DD or DSFP

7050CX4 and 7060X5 NEW

New 7050X4 and 7060X5 Series Provide Backward Compatibility and Investment Protection

Presented under NDA

ARISTA

2022/23: Shipping Silicon for Customer Networks

Feature Rich

Trident4

4X Higher Performance

Up to 12.8Tbps & 132MB Buffer

Programmable Pipeline

Highest Bandwidth

Tomahawk4

2X Higher Performance

Up to 25.6Tbps & 114MB Buffer

Scale Out and High Radix

Cloud Grade Scale

Jericho2C+

50% Higher Performance

7.2Tbps and 2.7Bpps

Deep Buffers and Extensible

Consistent High Performance and Extensible EOS

Next Generation 800G Silicon for AI Applications

Highest Bandwidth

Tomahawk5

2x Higher Performance

Up to 51.2Tbps & 165MB Buffer

Supports Linear Drive Optics

Scale Out and High Radix

Cloud Grade Scale

Jericho3

2x Higher Performance

14.4Tbps and 5.4Bpps

Supports Linear Drive Optics

Deep Buffers and Extensible

Consistent High Performance and Extensible EOS

ARISTA

Comprehensive 400G Platforms

Universal Leaf and Spine

7800R3 Series
Up to 460Tbps
576 x 400G ports
36 x 400G, 48 x 100G LC
Wirespeed, 2.5M Routes

7500R3 Series
Up to 230Tbps
288 x 400G ports
24 x 400G, 36 x 100G LC
Wirespeed, 2.5M Routes

7280R3 Series
Up to 24 x 400G
Wirespeed, 2.5M Routes
Carrier Grade Routing

Modular Spine

7388X5 Series
Up to 128 x 200G
64 x 400G OSFP, QSFP-DD
MACsec ready

7368X4 Series
Up to 128 x 100G
32 x 400G OSFP, QSFP-DD

7358X4 Series
Up to 128 x 100G
32 x 400G OSFP or QSFP-DD
Feature Rich High Scale

Fixed Leaf & Spine

7060X5 Series
64 x 400G / 128 x 200G
Hyperscale Performance
I/O Intensive environments

7060X4 Series
32 x 400G / 128 x 100G
Hyperscale Performance

7050X4 Series
32 x 400G / 128 x 100G
Feature Rich High Scale

ARISTA

400G-ZR and ZR+ for DCI

400G-ZR and Arista's Pluggable Line system (Shipping)

Simple and cost-effective DCI for <120km

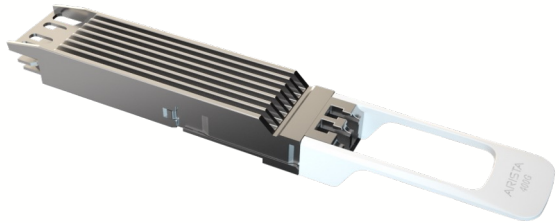

400G-ZR Transceiver

Fiber Mux / Demux

OSFP or QSFP Line System

400G-ZR+ optics for long-haul DWDM (3Q 2023)

Interoperates with 3rd party line systems for 500 to 1000km+ reach

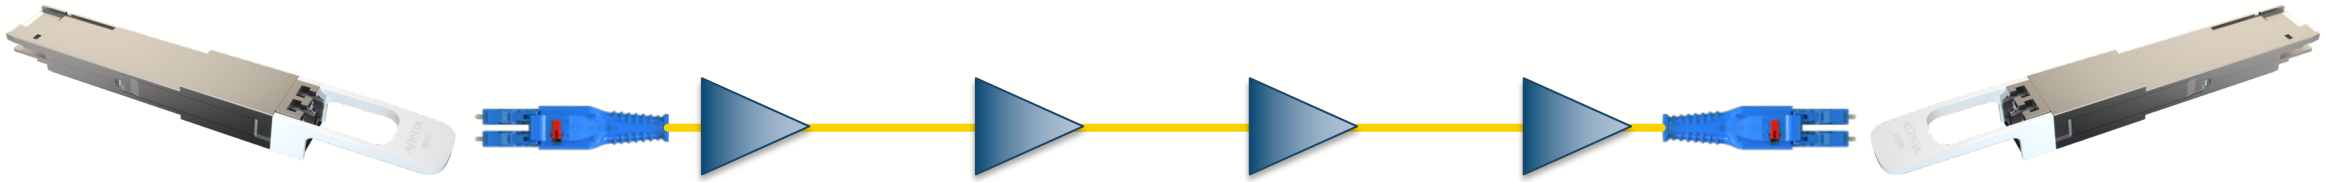

500km to 1000km+ depending on line system and mode of operation

High Speed Optics Evolution

Form Factor	Elec. Interface		Optical Interface Speeds
100G QSFP28	4x 25G Serdes	QSFP 	4x 25GE or 1x 100GE
400G OSFP or 400G QSFP-DD	8x 56G Serdes	OSFP  QSFP-DD 	8x 50GE, 4x 100GE or 1x 400GE
800G OSFP or 800G QSFP-DD	8x 112G Serdes	OSFP800  QSFP-DD800 	8x 100GE, 2x 400GE or 1x 800GE

ARISTA

What's after 400G?

800G / port using 100G serdes

First 800G / port products = 2x 400GbE / port

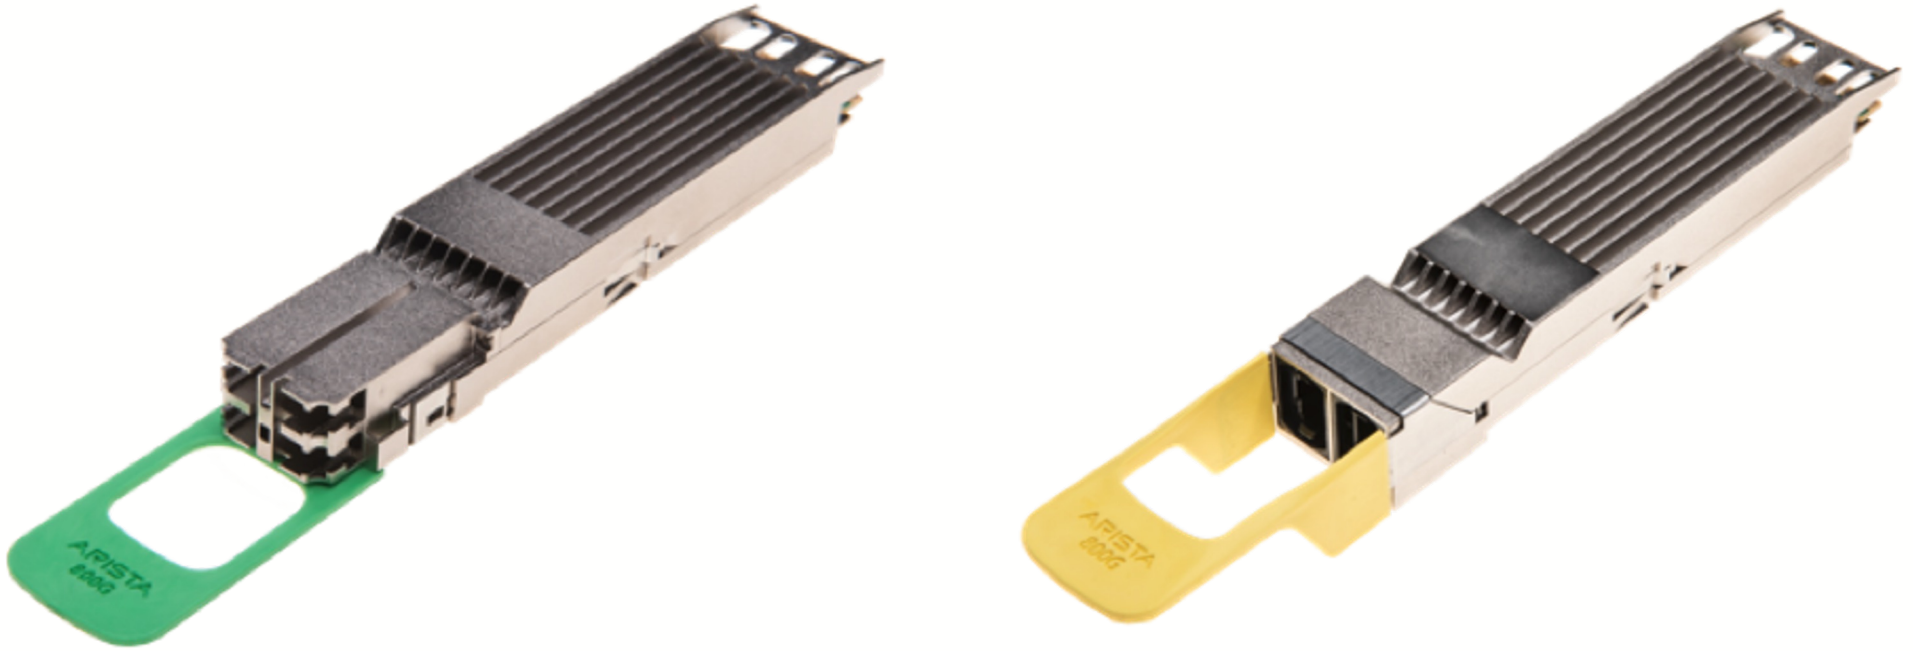

- 800G OSFPs use 2x "regular" LC and MPO12 connectors without any loss in faceplate density:
- **Deploy 400G OSFP today → easy upgrade to 800G in the future**

OSFP & QSFP-DD 800G (Dual 400G) Optics

Released and Orderable

800G-2FR4
800G-2LR4

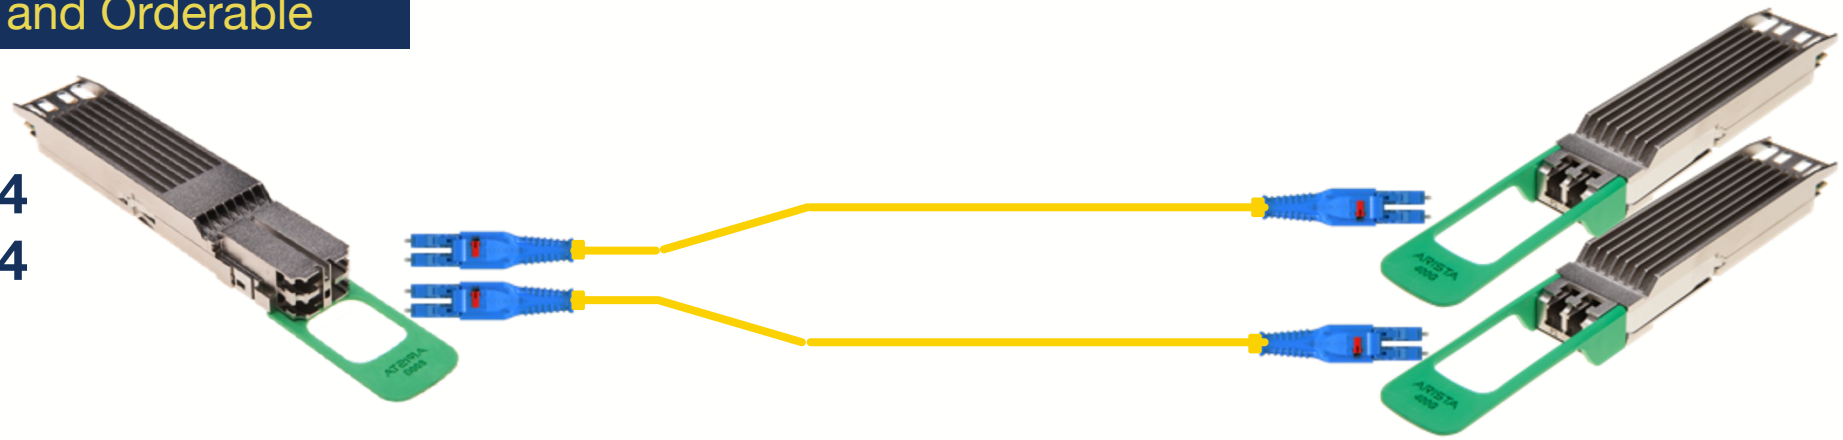

400G-FR4/LR4

400G-FR4/LR4

No change to customer fiber plant

800G-2XDR4
800G-2PLR4

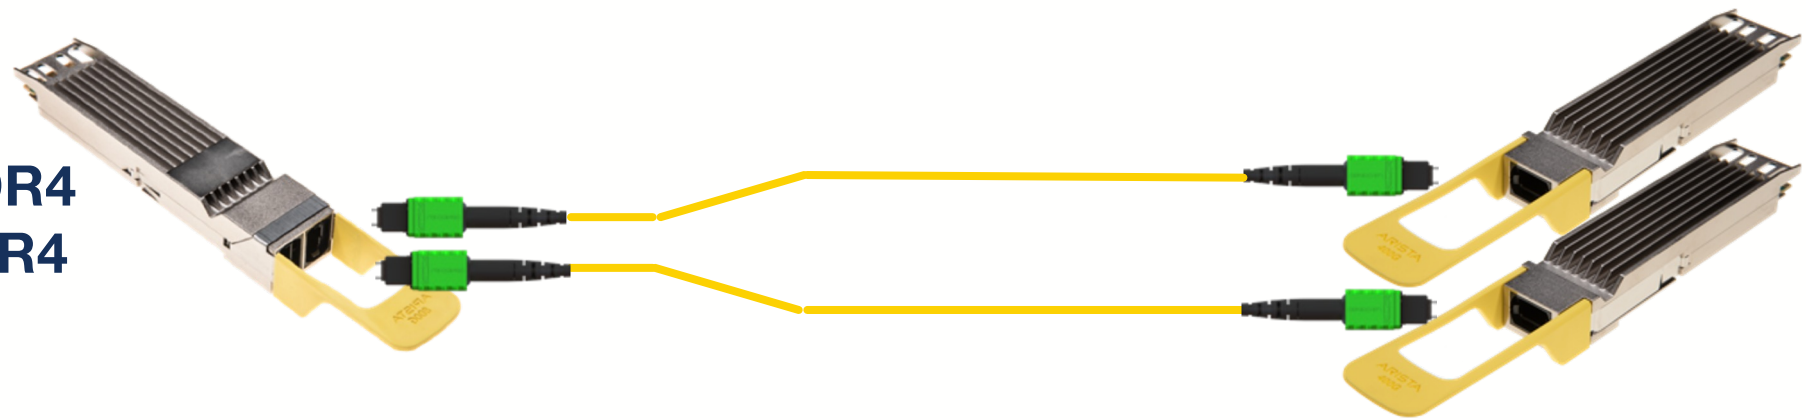

400G-DR4/XDR4/PLR4

400G-DR4/XDR4/PLR4

Presented under NDA
ARISTA

Linear Drive Pluggable Optics Modules

1. **Linear Drive means no DSP or CDR in Module**

   Just a linear driver to drive electro-optical modulator
2. **Requires high-performance switch SERDES**

   And very careful motherboard signal integrity design
3. **Achieves same power efficiency as direct drive CPO**

   While retaining the many advantages of pluggable optics modules

Compared to the latest DSP Optics, Linear Drive Optics reduce optics power by up to 50%, switch power by 25%

Switch Power Savings with LD SiPh Optics

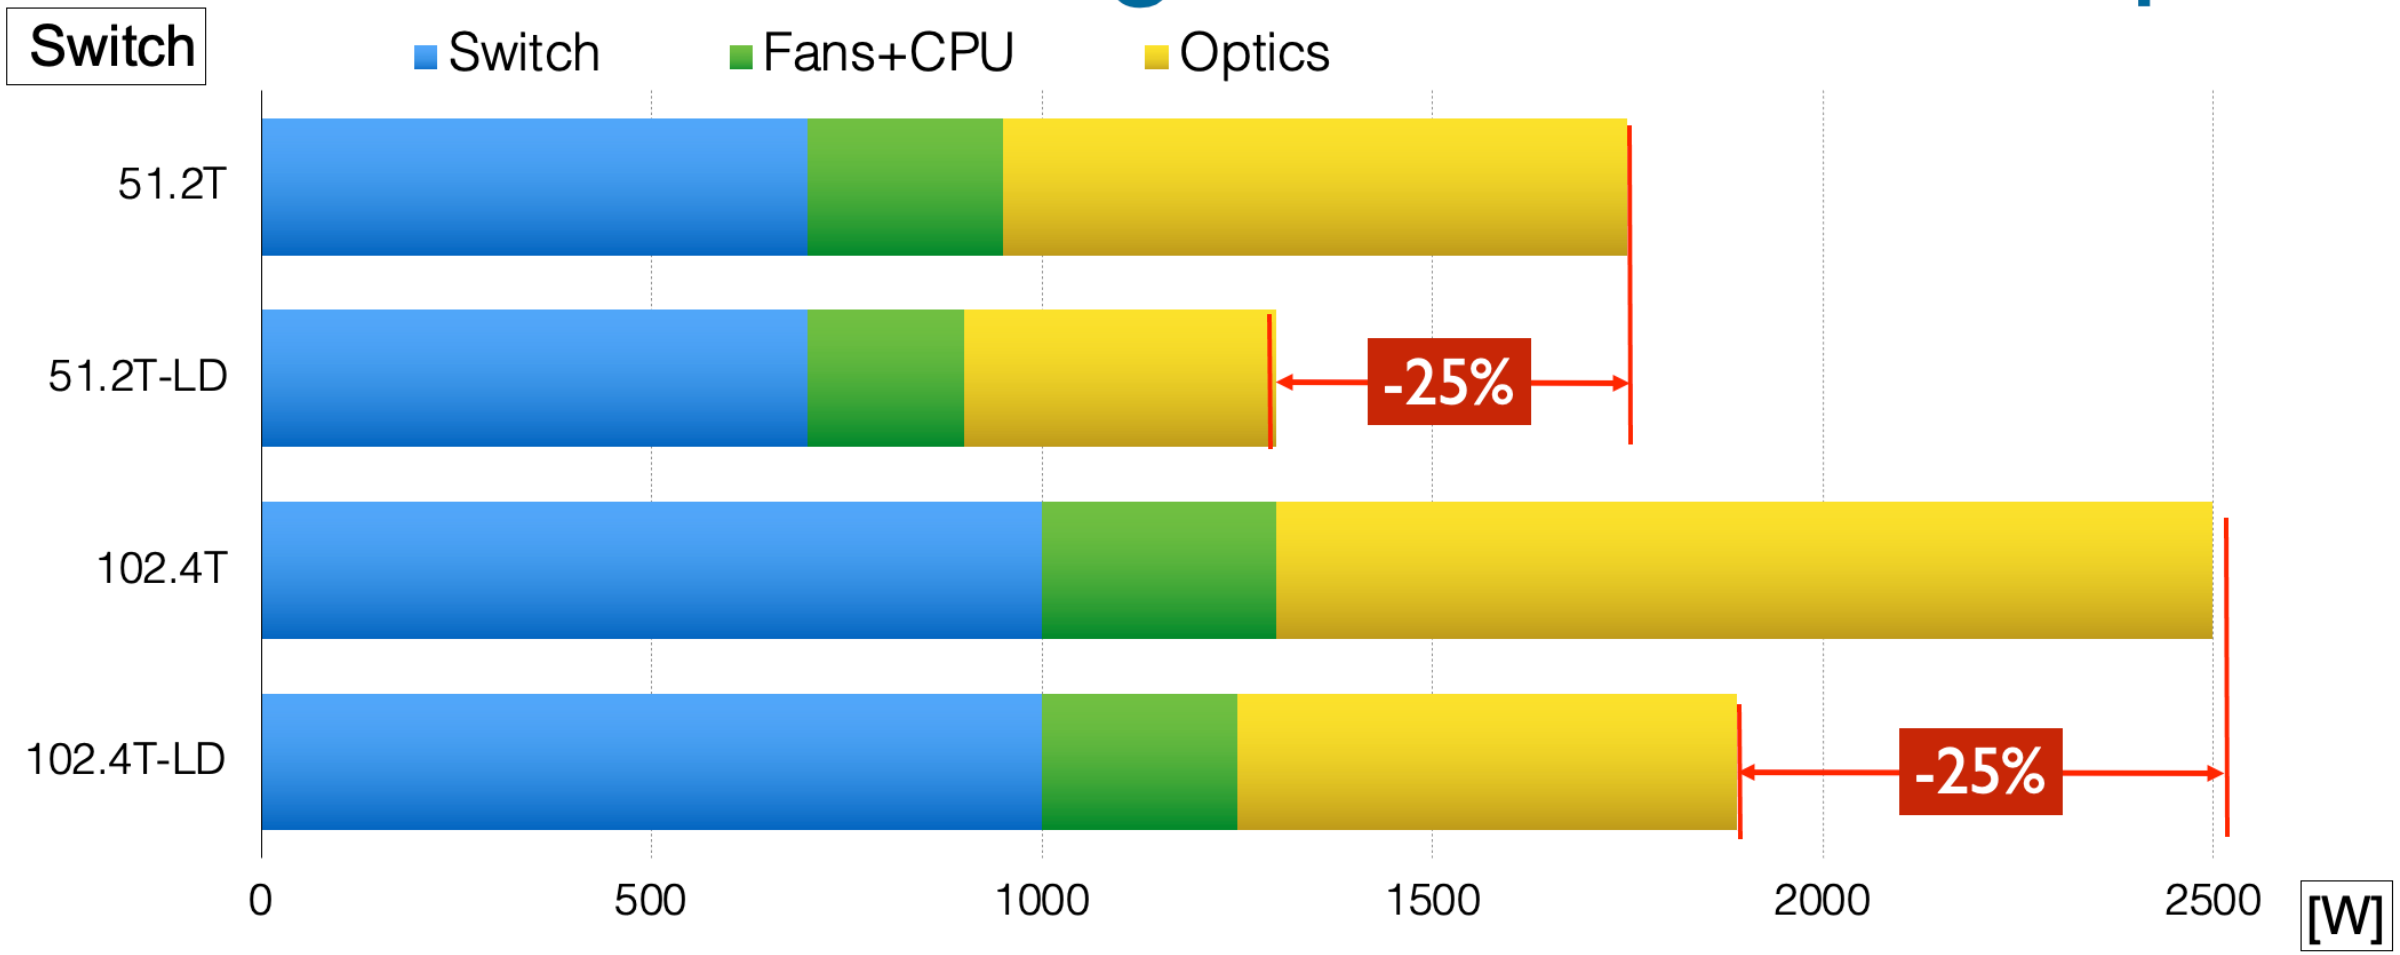

51.2T Power measured, 102.4T power projected

Linear Drive Pluggable Summary

1. **Same low power as CPO in a Pluggable Form Factor**

   While retaining all the advantages of Pluggable Modules
2. **Significant cost reduction compared to DSP Optics**

   DSP is the most expensive part in an optics module
3. **Requires careful system design to achieve low BER**

   SerDes performance and electrical signal integrity is key

LPO is the lowest risk way to reduce optics power and cost while retaining all the advantages of pluggable optics modules