



A Foundation for Next-Gen Cloud Networks: Introducing Industry's First Cloud-Optimized 51.2 Tbps Networking Platform

March 2, 2023

Our focus: data infrastructure

Data center



Every major
cloud

Carrier
infrastructure



4 of 5 top
wireless OEMs

Automotive



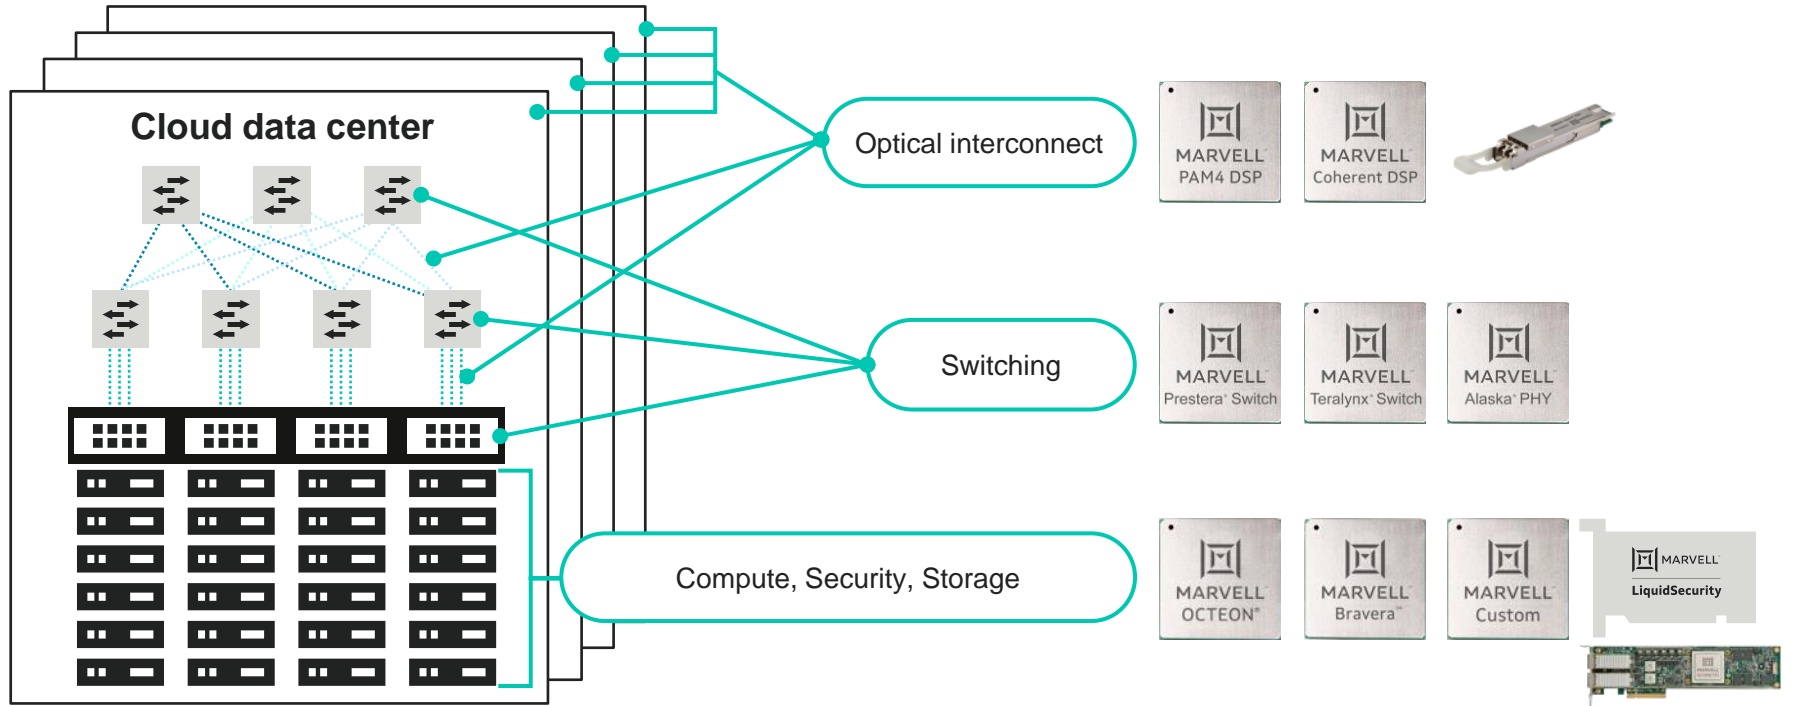
8 of 10 top
auto OEMs

Enterprise
networking



17 of 20 top
networking OEMs

Cloud-optimized data center portfolio



Addresses unique needs of the largest data center operators

Key market dynamics

1

New AI applications

2

Network bottlenecks

3

Bandwidth growth

New AI drives more bandwidth

More AI

 **OpenAI** ChatGPT

 **Google** Bard

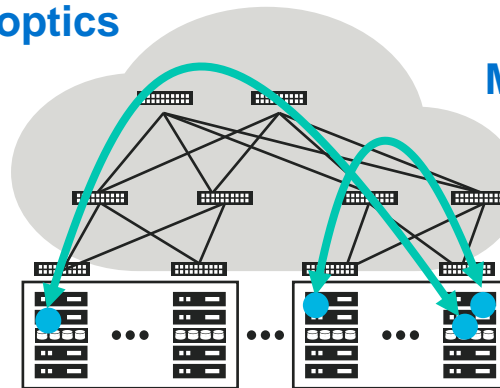
 **Microsoft** Bing AI

 **Baidu** Ernie Bot

More optics

More bandwidth

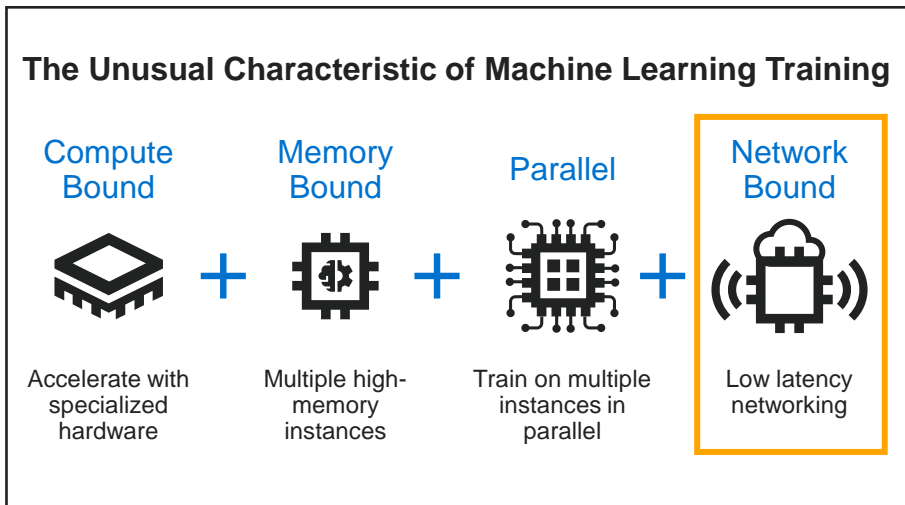
More switches



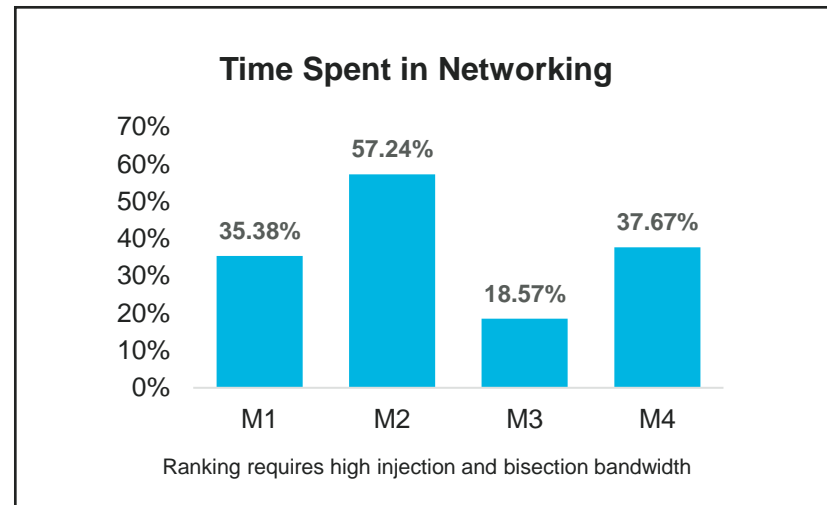
Cloud operators need networking to do more

Low latency critical for demanding applications

AWS keynote at re:Invent 2022



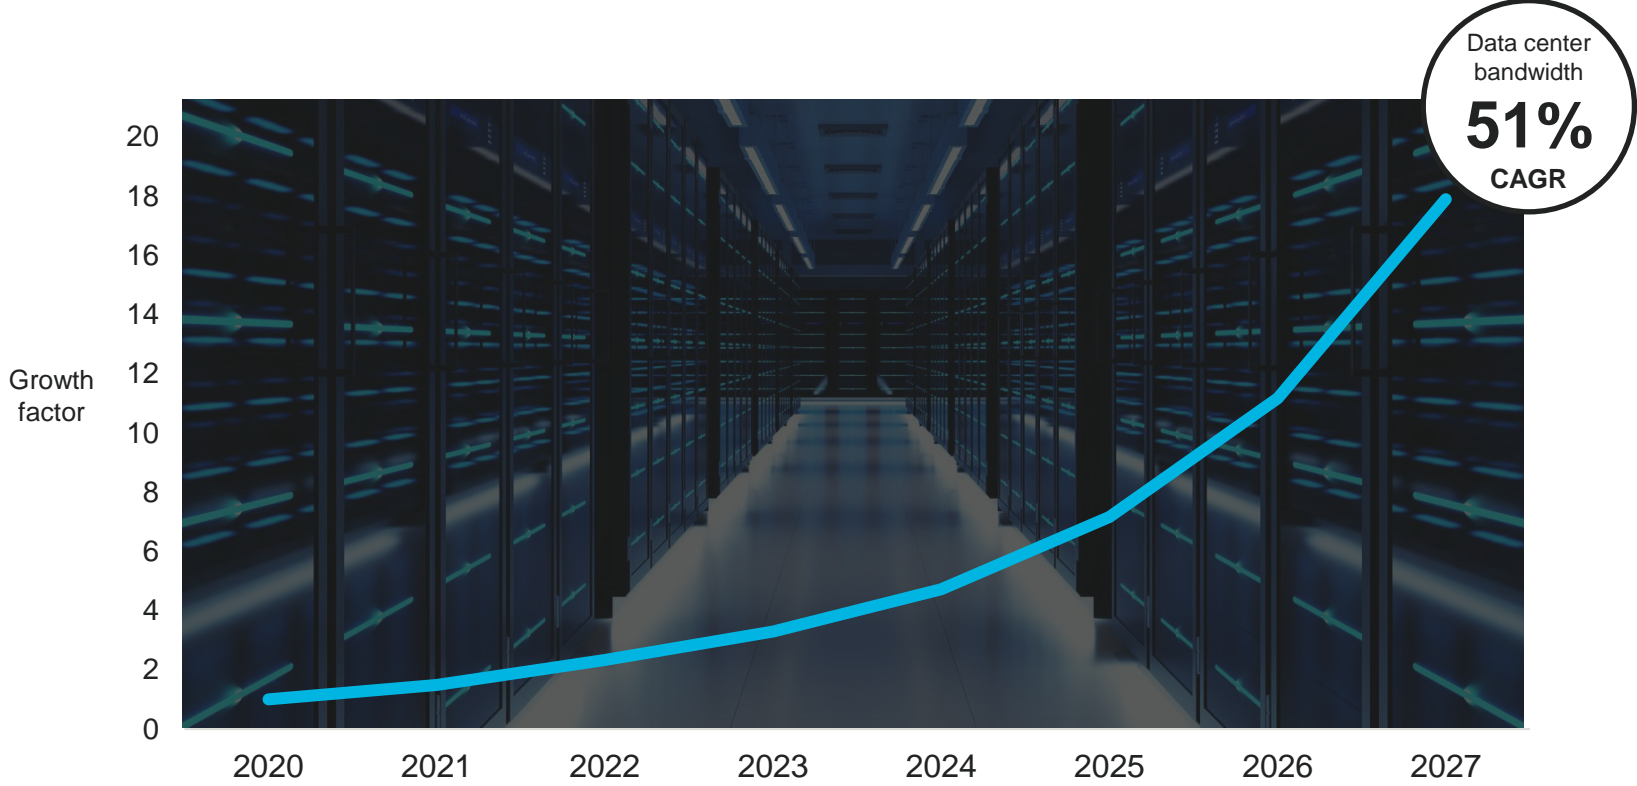
Meta keynote at OCP Summit 2022



Mx = ML training models

Network bottlenecks limit performance, hurt revenue

Data center bandwidth growth accelerating



Source: Marvell estimates based on industry analyst forecasts

A foundation for next-gen data center networks

1.6 Tbps PAM4 DSP



200G per lambda

51.2 Tbps Ethernet switch



Ultra-low latency, programmable

Cloud-optimized platform enables cloud networks to scale

Nova: Industry's first 1.6T PAM4 DSP



Sampling now

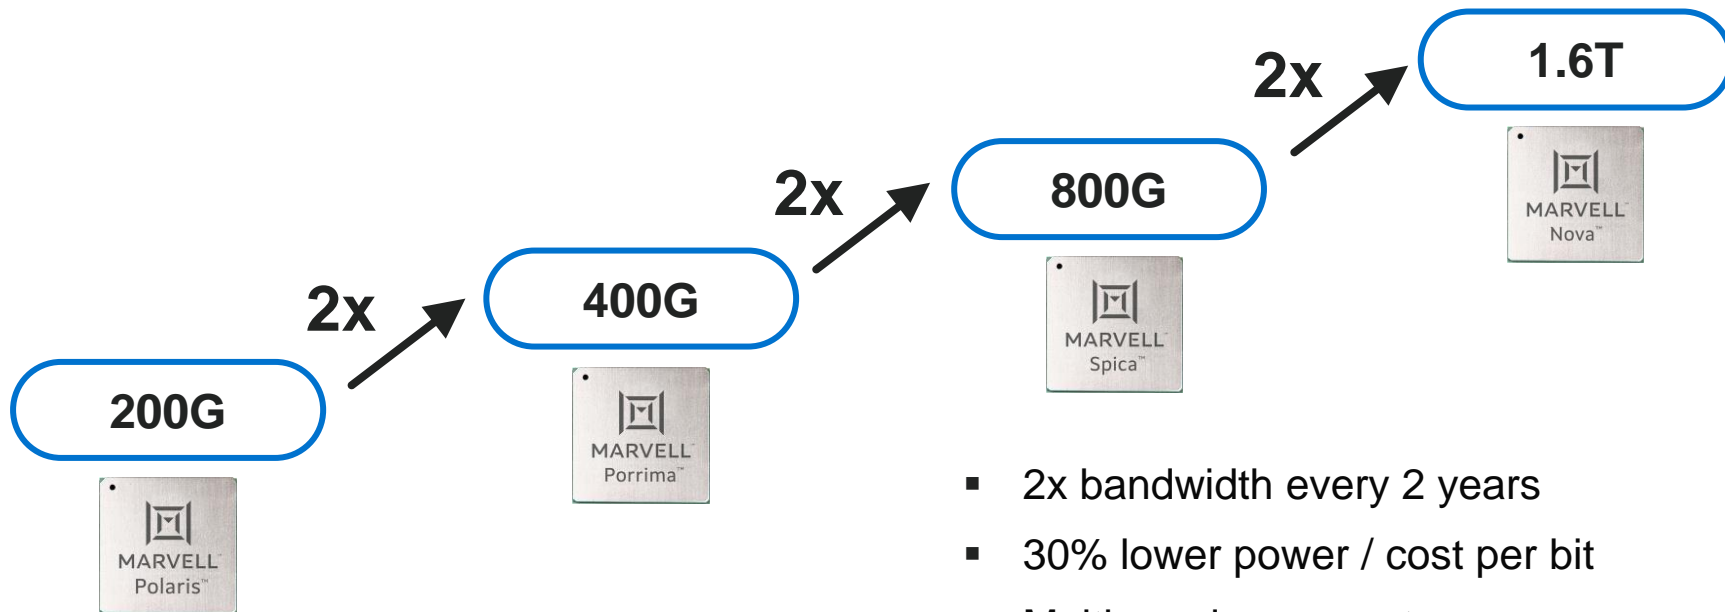
- **8 x 200G / wavelength (λ)**
- 30% lower cost/bit*
- 30% lower power/bit*
- Half module count
- 2x more reliable optics**
- Multi-vendor

*As compared to optical modules based on Marvell's previous PAM4 DSP generation

**Expected reliability improvement compared to the previous Marvell PAM4 DSP generation.

Doubles data center bandwidth for new AI/ML applications

Why PAM4 DSP?



- 2x bandwidth every 2 years
- 30% lower power / cost per bit
- Multi-vendor ecosystem
- High volume, high reliability
- Backward/forward compatibility


Next-gen 1.6T disrupts data interconnects

100G/λ 800G



2RU

64 modules



200G/λ 1.6T



1RU

32 modules



200G/λ enables 1RU design, improving bandwidth density


Next-gen 1.6T disrupts data interconnects




30%
lower
cost/bit*



30%
lower
power/bit*



50%
fewer
modules



2x
more reliable
optics**

Nova essential for cloud-optimized 51.2T-based networks

Teralynx 10: ultra-low latency 51.2T switch



Sampling: 2Q23

- **Ultra-low latency**
- Advanced telemetry
- Permutable flex-forwarding
- 80% power savings


Scales the cloud and addresses network bottlenecks

Ultra-low latency switch architecture




*Based on RFC 2544 FIFO latency with Teralynx 7 for 400G port

Optimized for AI/ML and data center fabrics



- **Ultra-low latency**
Reduces job completion time
- **Congestion-aware routing**
Minimizes congestion
- **Advanced telemetry**
Auto-tunes network in real-time
- **Permutable flex-forwarding**
Programs packet-forwarding as networks evolve

Industry-leading 112G SerDes




Best-in-class long reach

- Lowest bit error rate (BER)
- Flexible data rate
- Optimized designs without retimers
- Eliminates need for flyover cables

Enables lowest cost and power system design

The benefits of an all-Marvell solution...



...predictable deployment and faster time-to-market

Key takeaways

1 Industry's first cloud-optimized 51.2T networking platform

2 Nova, industry's first 200G/λ 1.6T PAM4 DSP

3 Teralynx 10, ultra-low latency 51.2T switch

4 Enables AI/ML and data center network scaling, lowers power/cost per bit

5 Accelerates time to market



Essential technology, done right™