BROADCOM®

ENABLING AI Infrastructure

OPEN // SCALABLE // POWER EFFICIENT

Safe Harbor Statement

This presentation contains forward-looking statements (including within the meaning of Section 21E of the United States Securities Exchange Act of 1934, as amended, and Section 27A of the United States Securities Act of 1933, as amended) that are based on current expectations and beliefs of Broadcom's management, current information available to Broadcom's management, and current market trends and market conditions and involve risks and uncertainties, that may cause actual results to differ materially from those contained in forward-looking statements. Accordingly, undue reliance should not be placed on such statements.

Our filings with the SEC, which are available without charge at the SEC's website at https://www.sec.gov, discuss some of the important risk factors that may affect our business, results of operations and financial condition. Actual results may vary from the estimates provided. We undertake no intent or obligation to publicly update or revise any of the estimates and other forward-looking statements made in this presentation, whether as a result of new information, future events or otherwise, except as required by law.

Agenda

- **Enabling AI Infrastructure:** Charlie Kawwas
- **Scalable AI Networks:** Ram Velaga
- **AI Server Interconnect:** Jas Tremblay
- **Optical Interconnects:** Near Margalit
- **Foundational Technology for AI Interconnect:** Vijay Janapaty
- **Custom AI Accelerators:** Frank Ostojic
- **Closing Remarks:** Charlie Kawwas

BROADCOM

ENABLING AI Infrastructure

Charlie Kawwas, Ph. D.
President, Broadcom

OPEN // SCALABLE // POWER EFFICIENT

Heritage of **Technology Innovation**

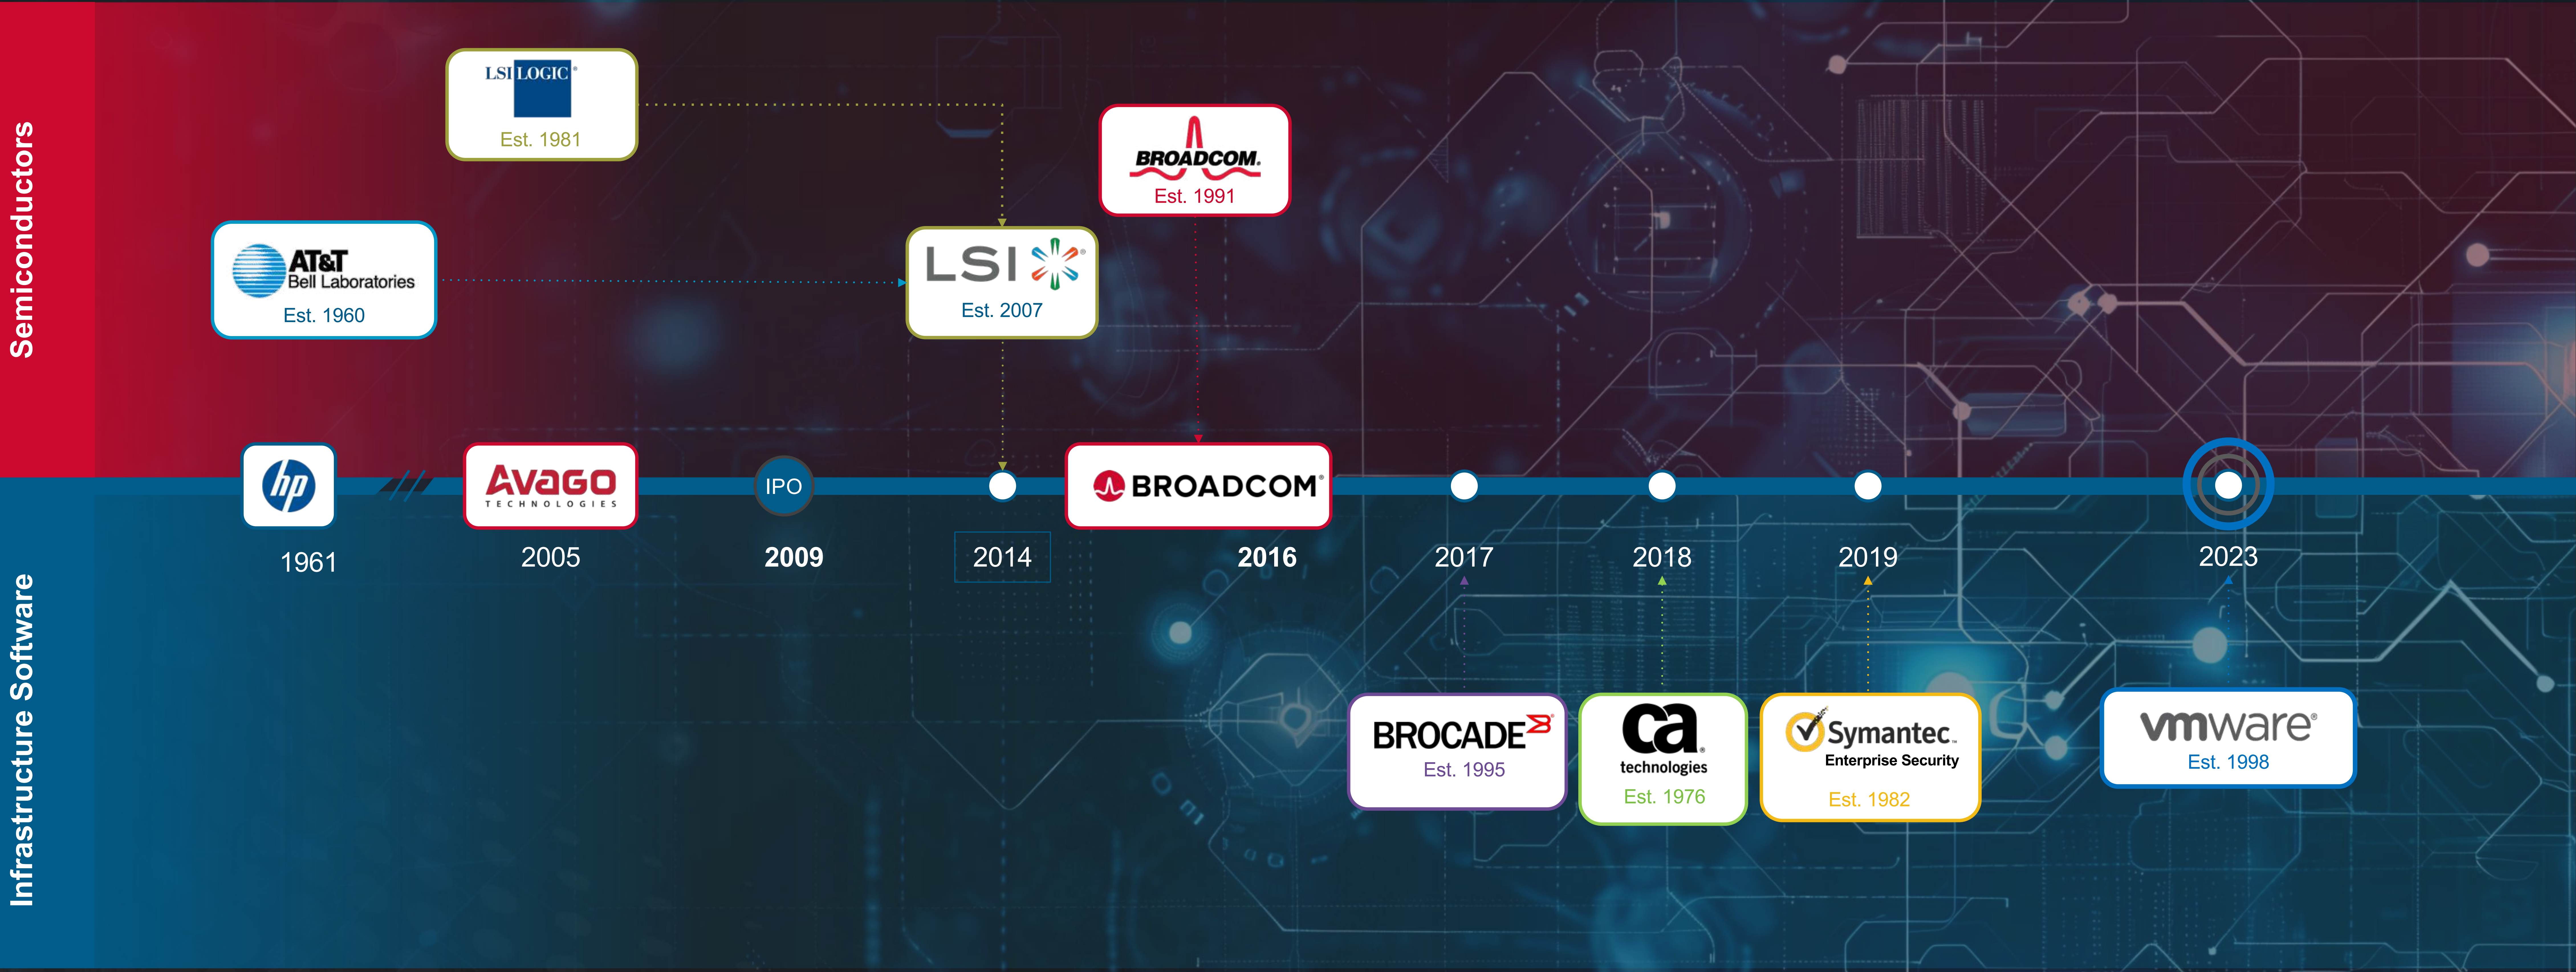

Semiconductors

AT&T Bell Laboratories
Est. 1960

LSI LOGIC
Est. 1981

LSI
Est. 2007

BROADCOM.
Est. 1991

hp
1961

AVAGO TECHNOLOGIES
2005

IPO
2009

2014

BROADCOM
2016

2017

2018

2019

2023

Infrastructure Software

BROCADE
Est. 1995

ca technologies
Est. 1976

Symantec. Enterprise Security
Est. 1982

vmware
Est. 1998

OPEN // SCALABLE // POWER EFFICIENT

Semiconductors

FY23 Revenue **$28B**

2019-2023 CAGR **13%**

Annual R&D spend **$3B+**

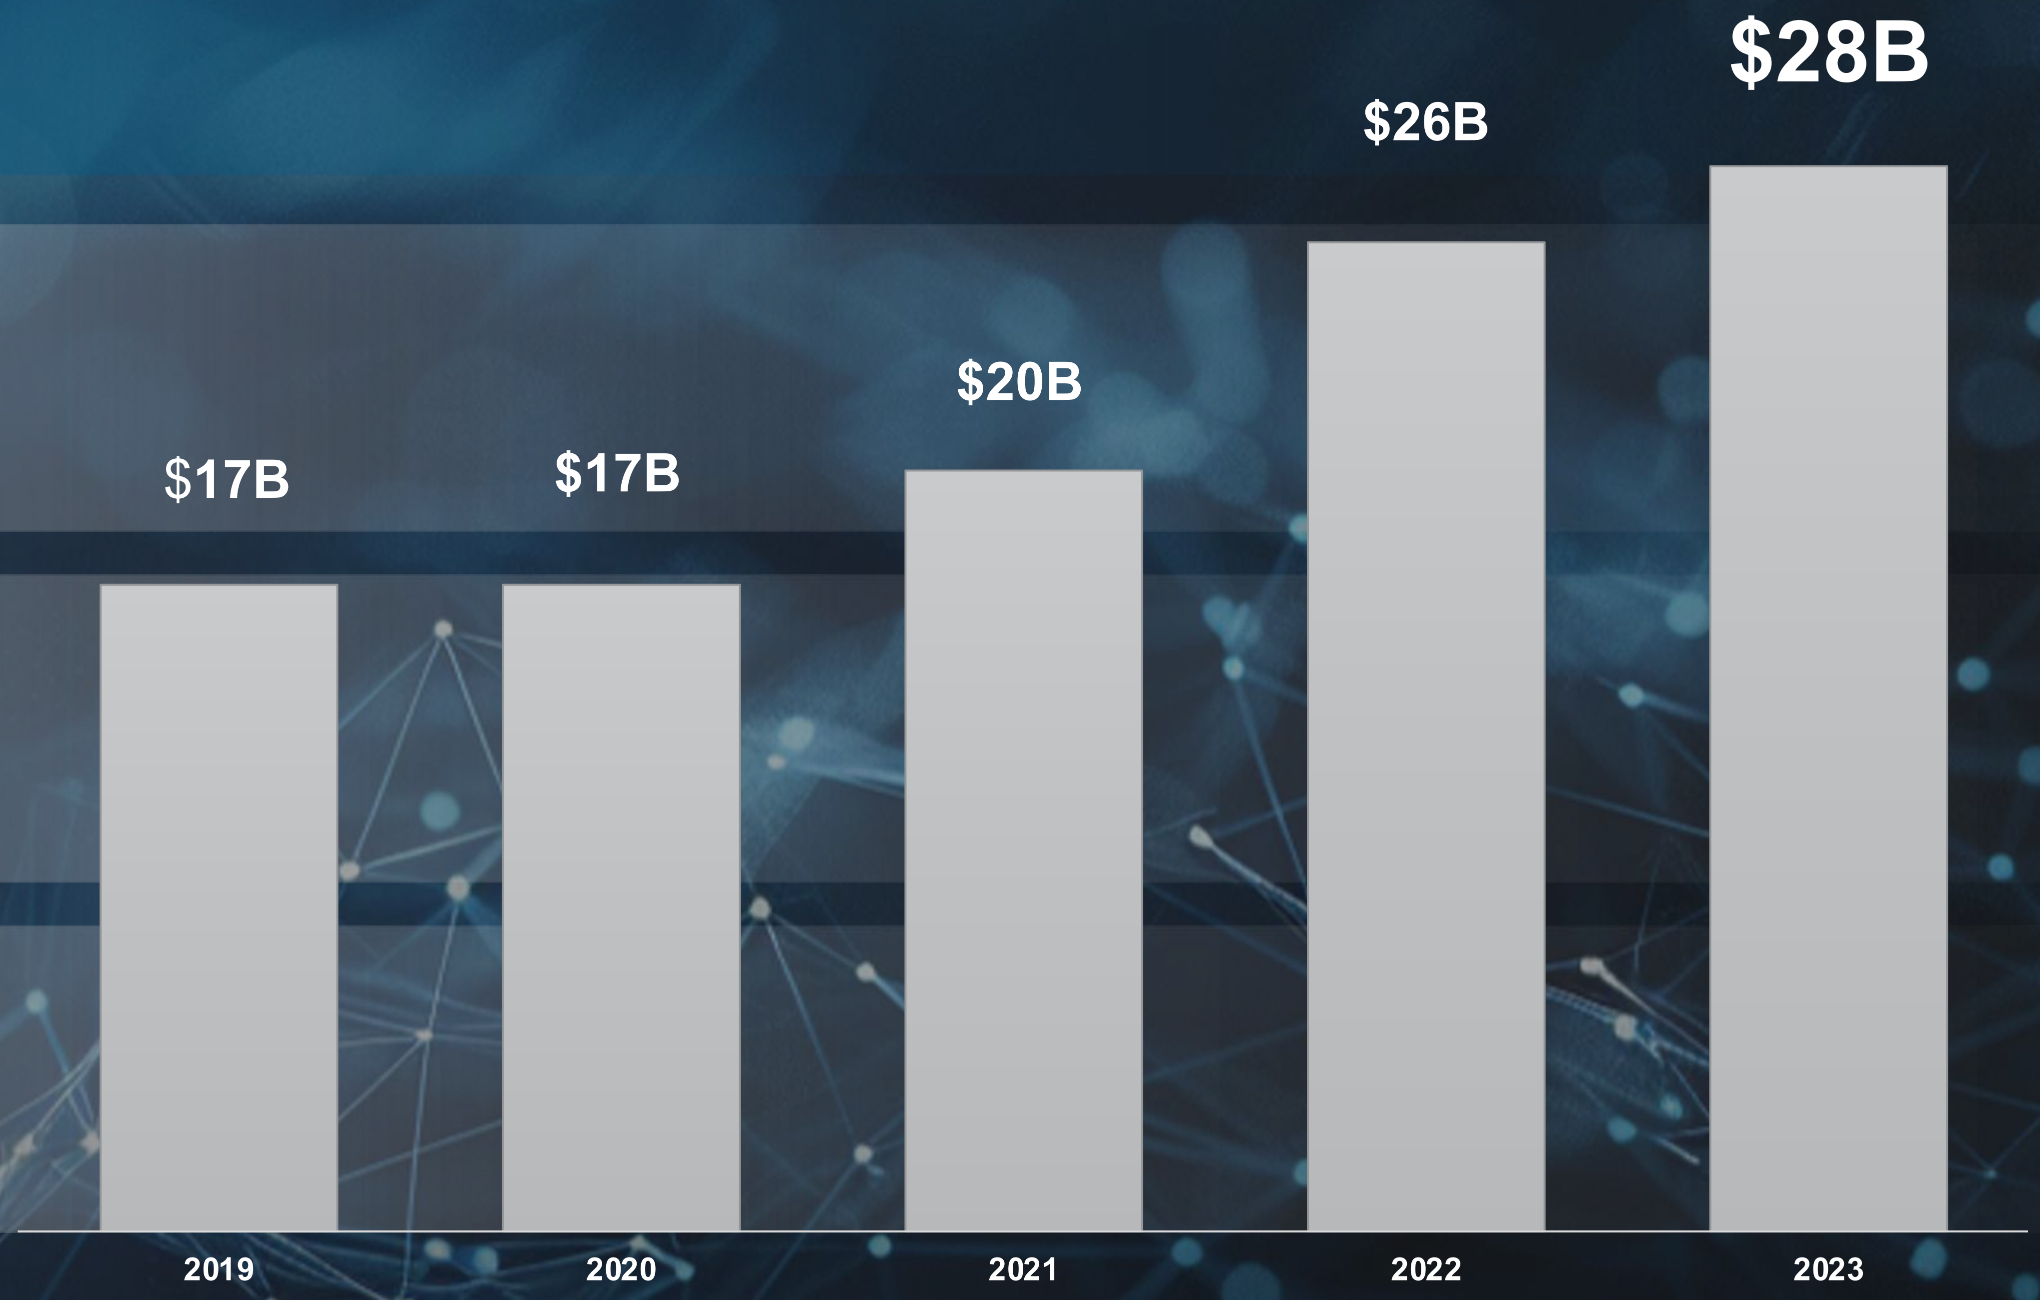

Technology Leadership → Organic Growth

OPEN // SCALABLE // POWER EFFICIENT

Strategy

Market
10 Year Horizon

Technology
Leadership

Execution
Excellence

Sustainable Franchise

Strategy

Market
10 Year Horizon

Technology
Leadership

Execution
Excellence

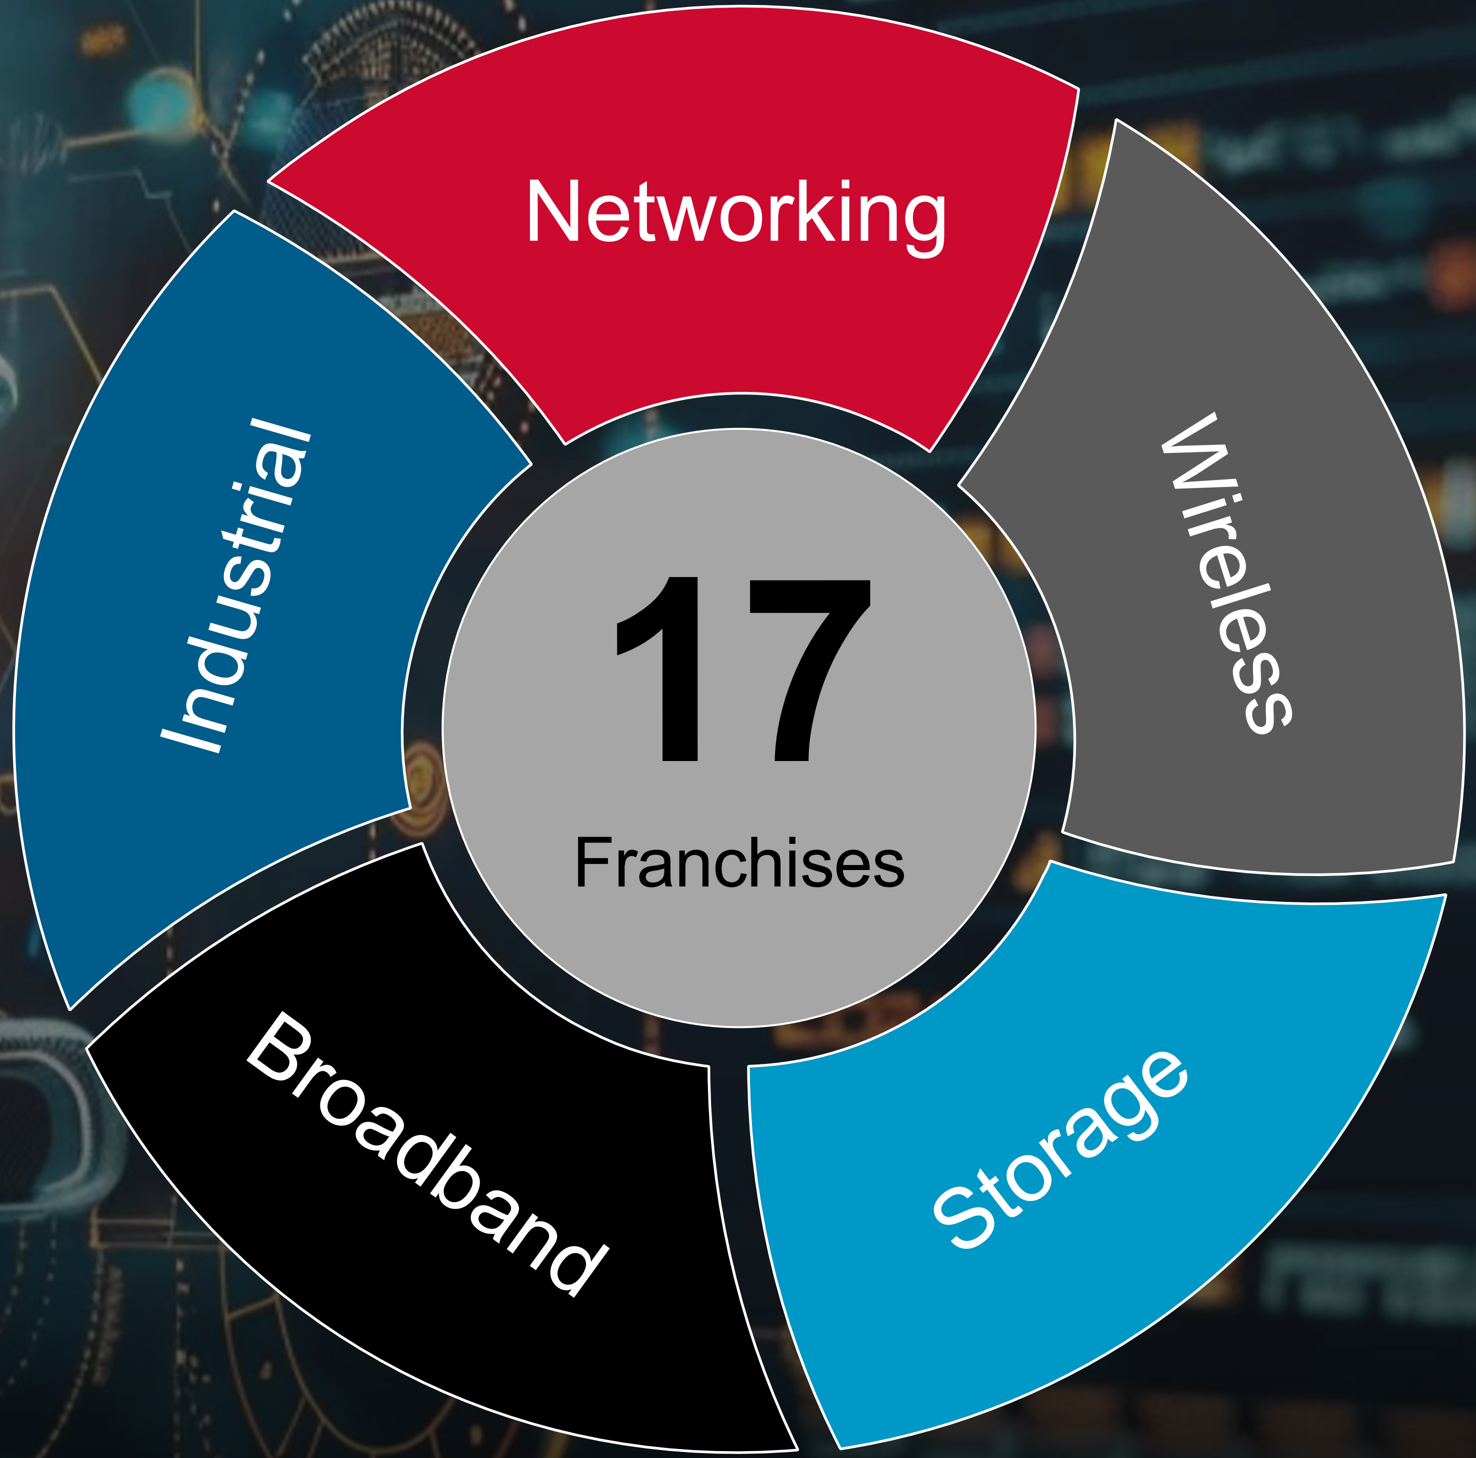

OPEN // SCALABLE // POWER EFFICIENT

17

Franchises

Networking

Wireless

Storage

Broadband

Industrial

Enabling AI infrastructure

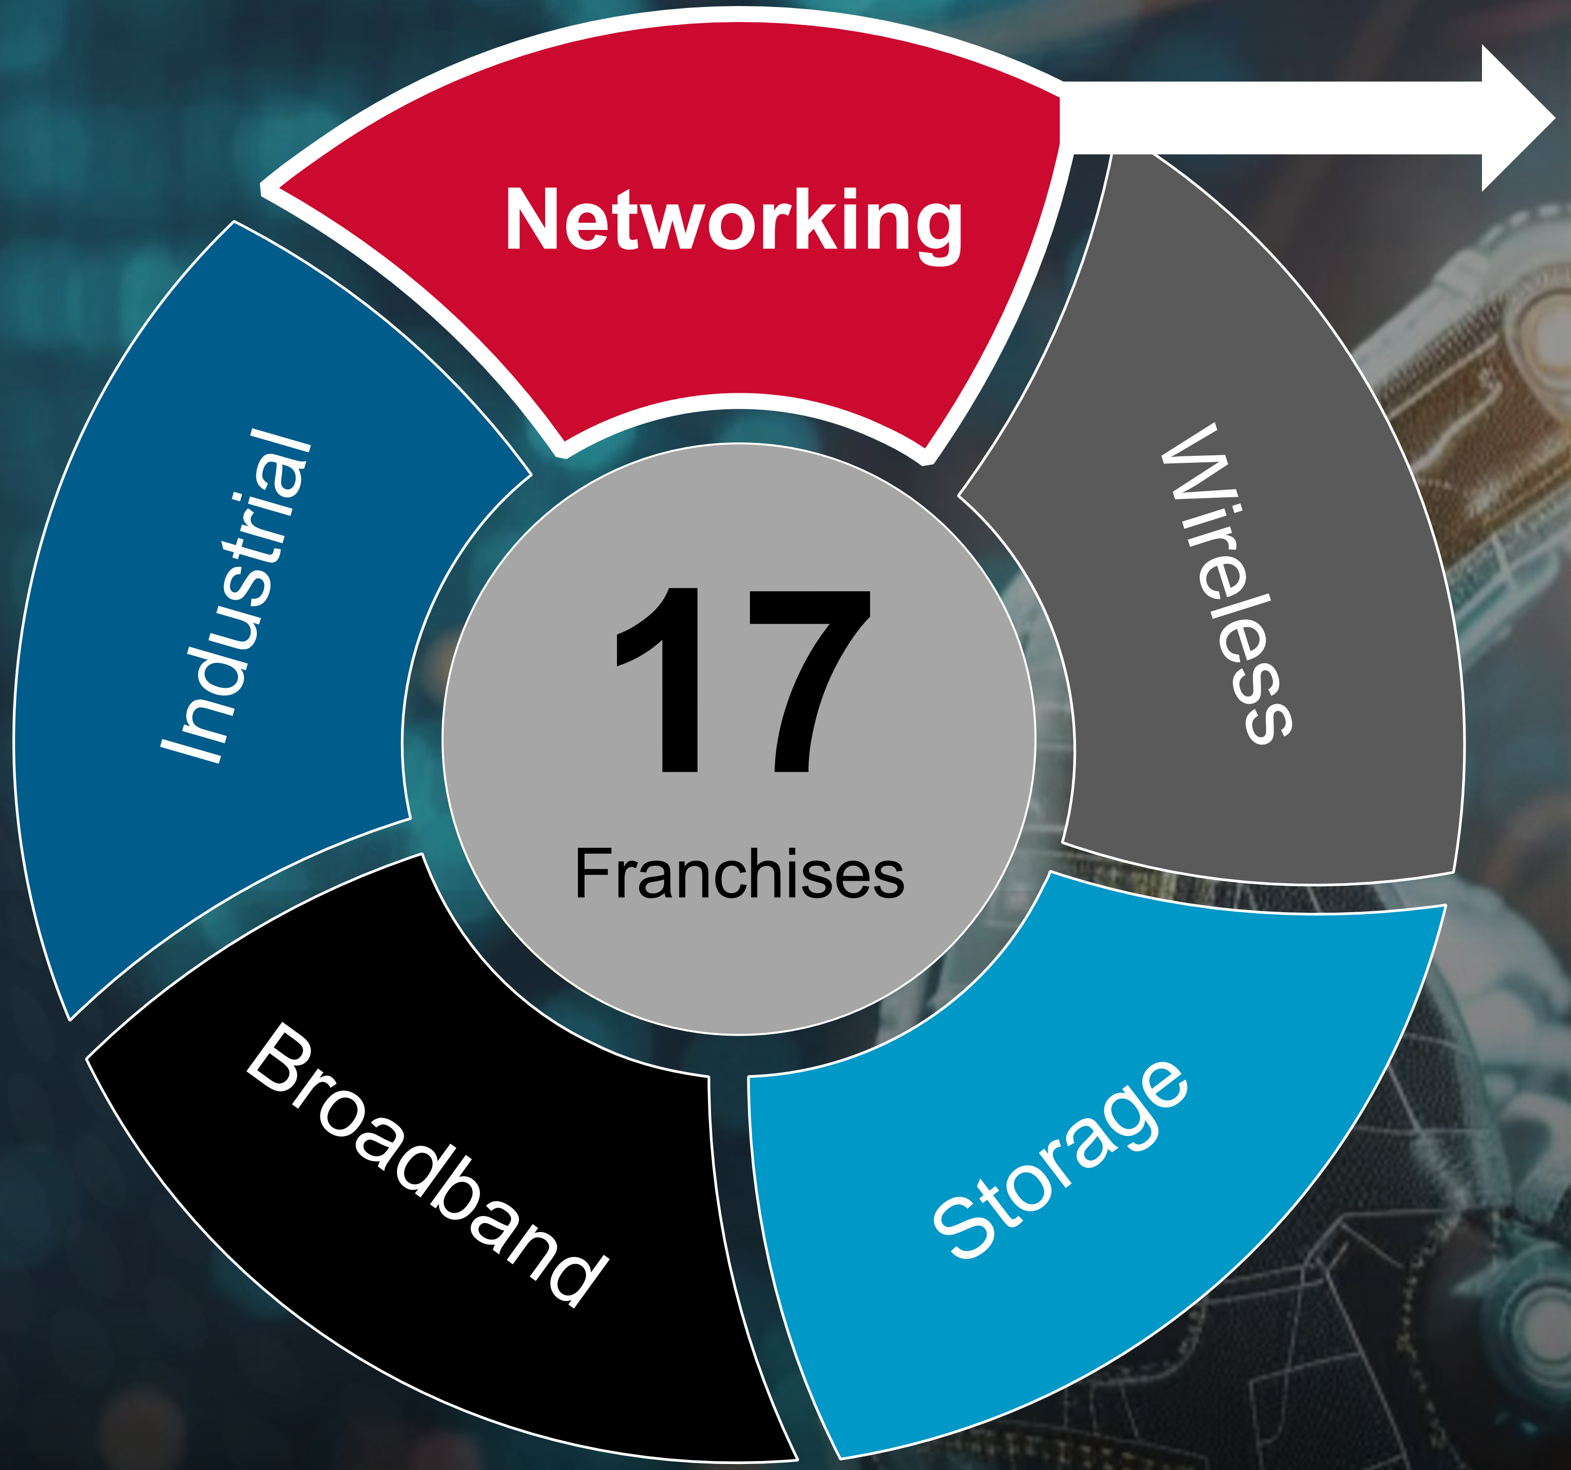

OPEN // SCALABLE // POWER EFFICIENT

Two AI Markets

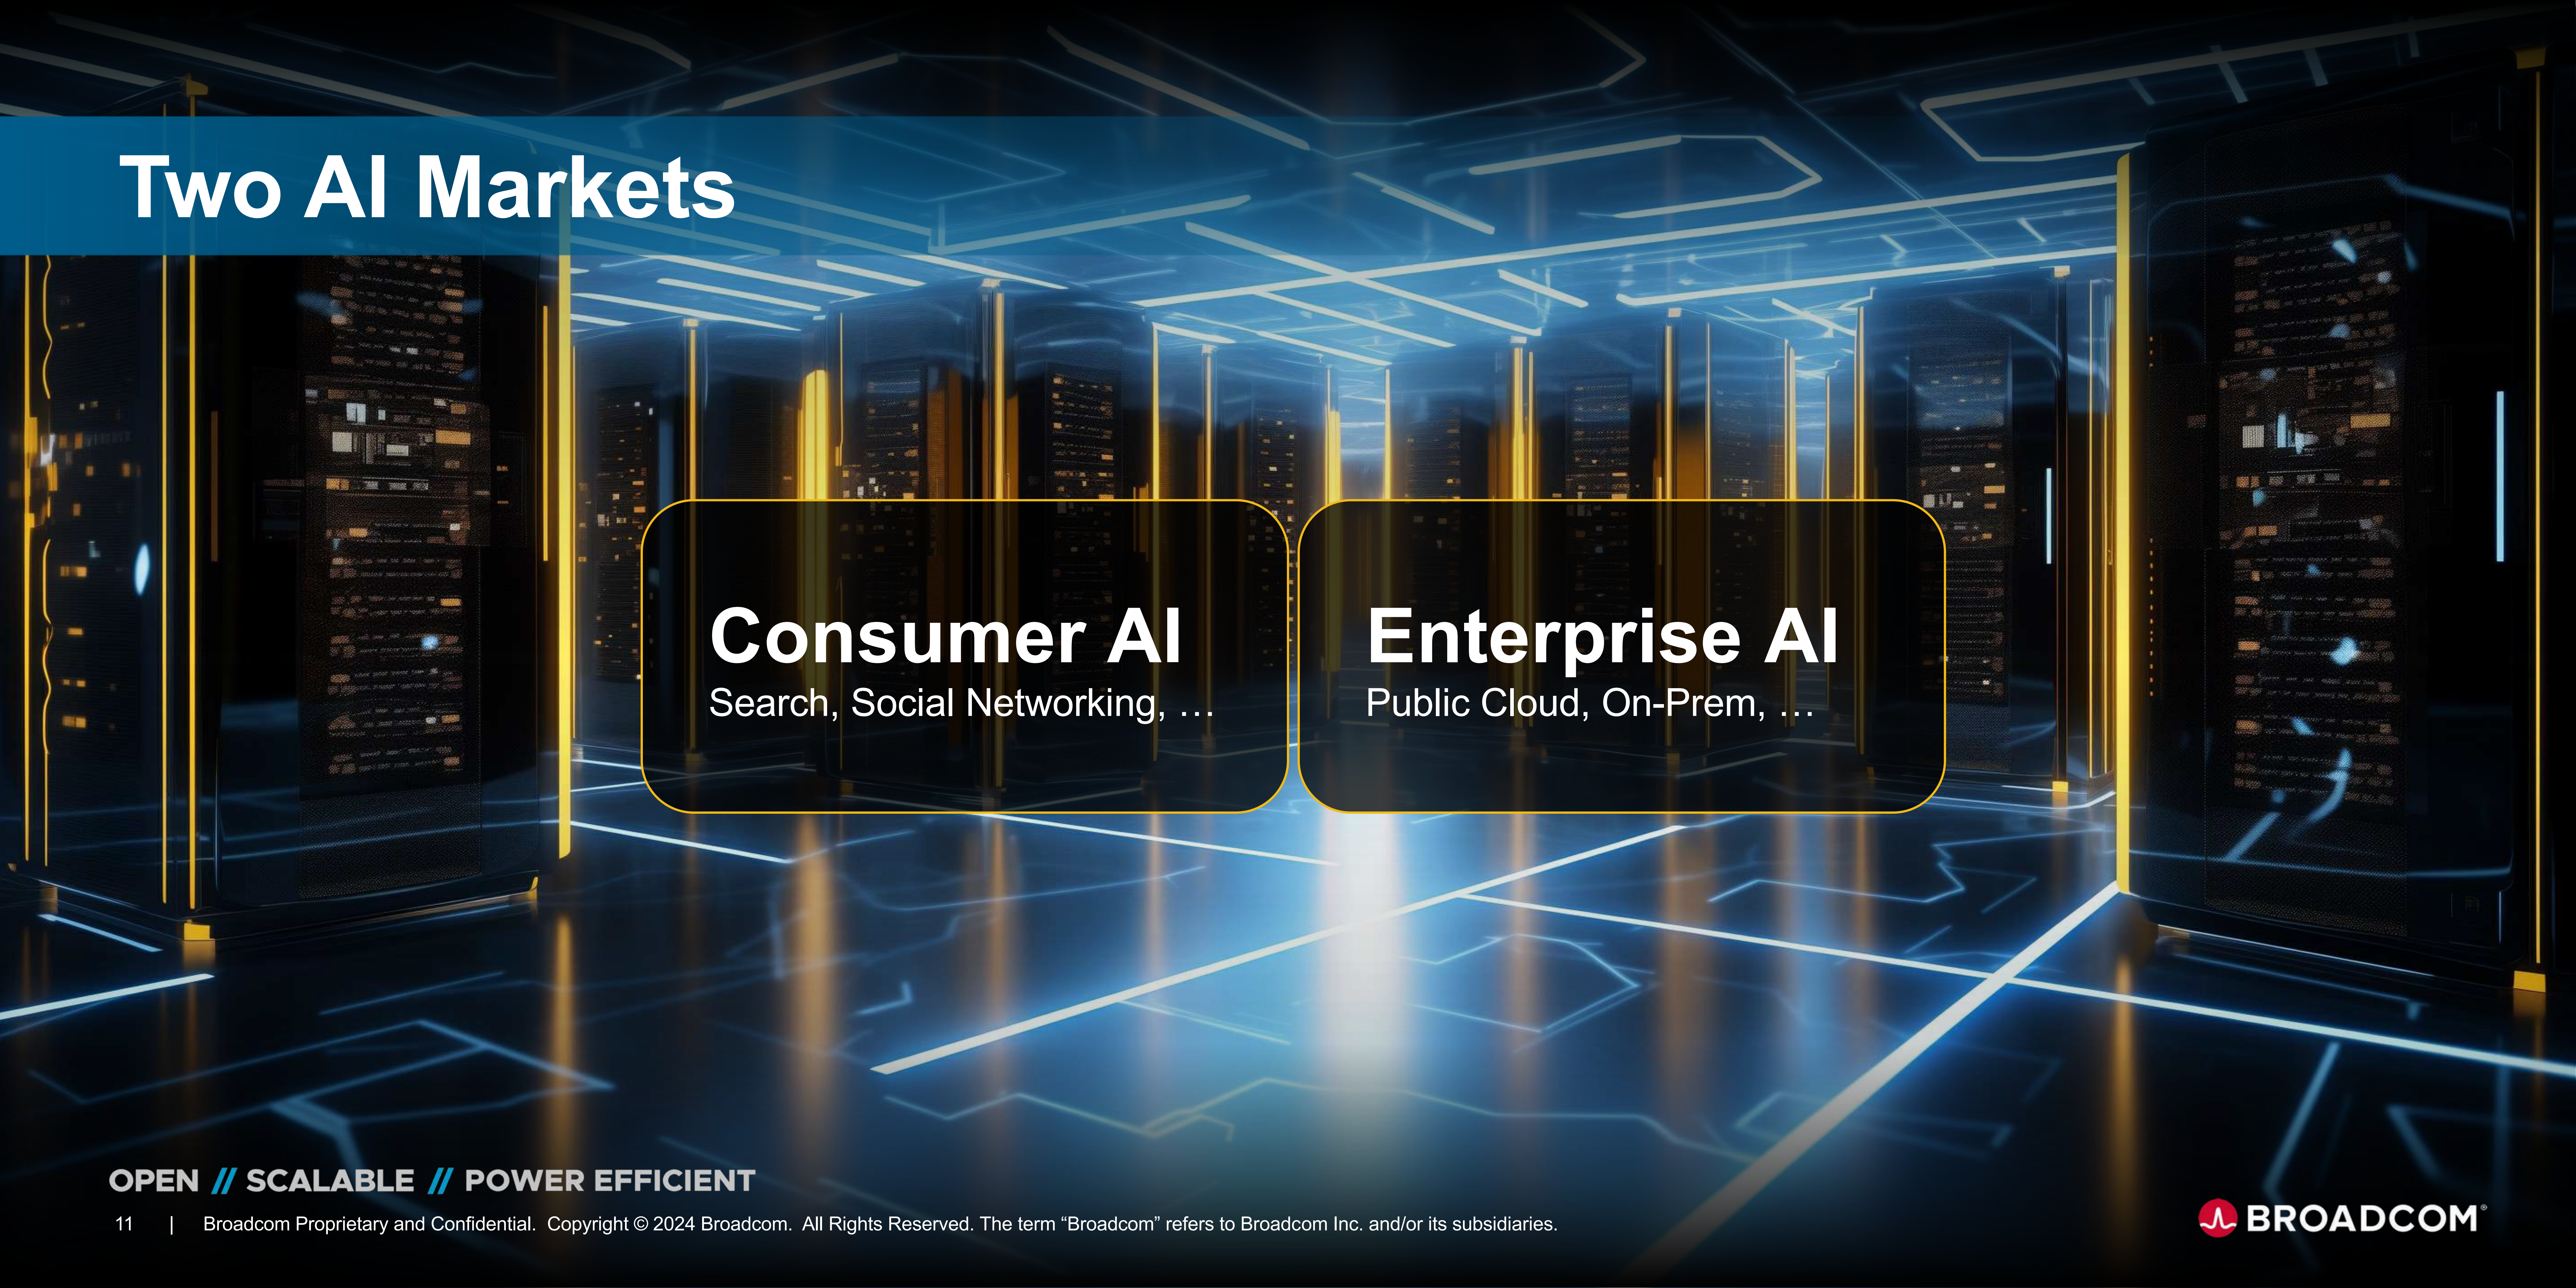

Consumer AI

Search, Social Networking, …

Enterprise AI

Public Cloud, On-Prem, …

OPEN // SCALABLE // POWER EFFICIENT

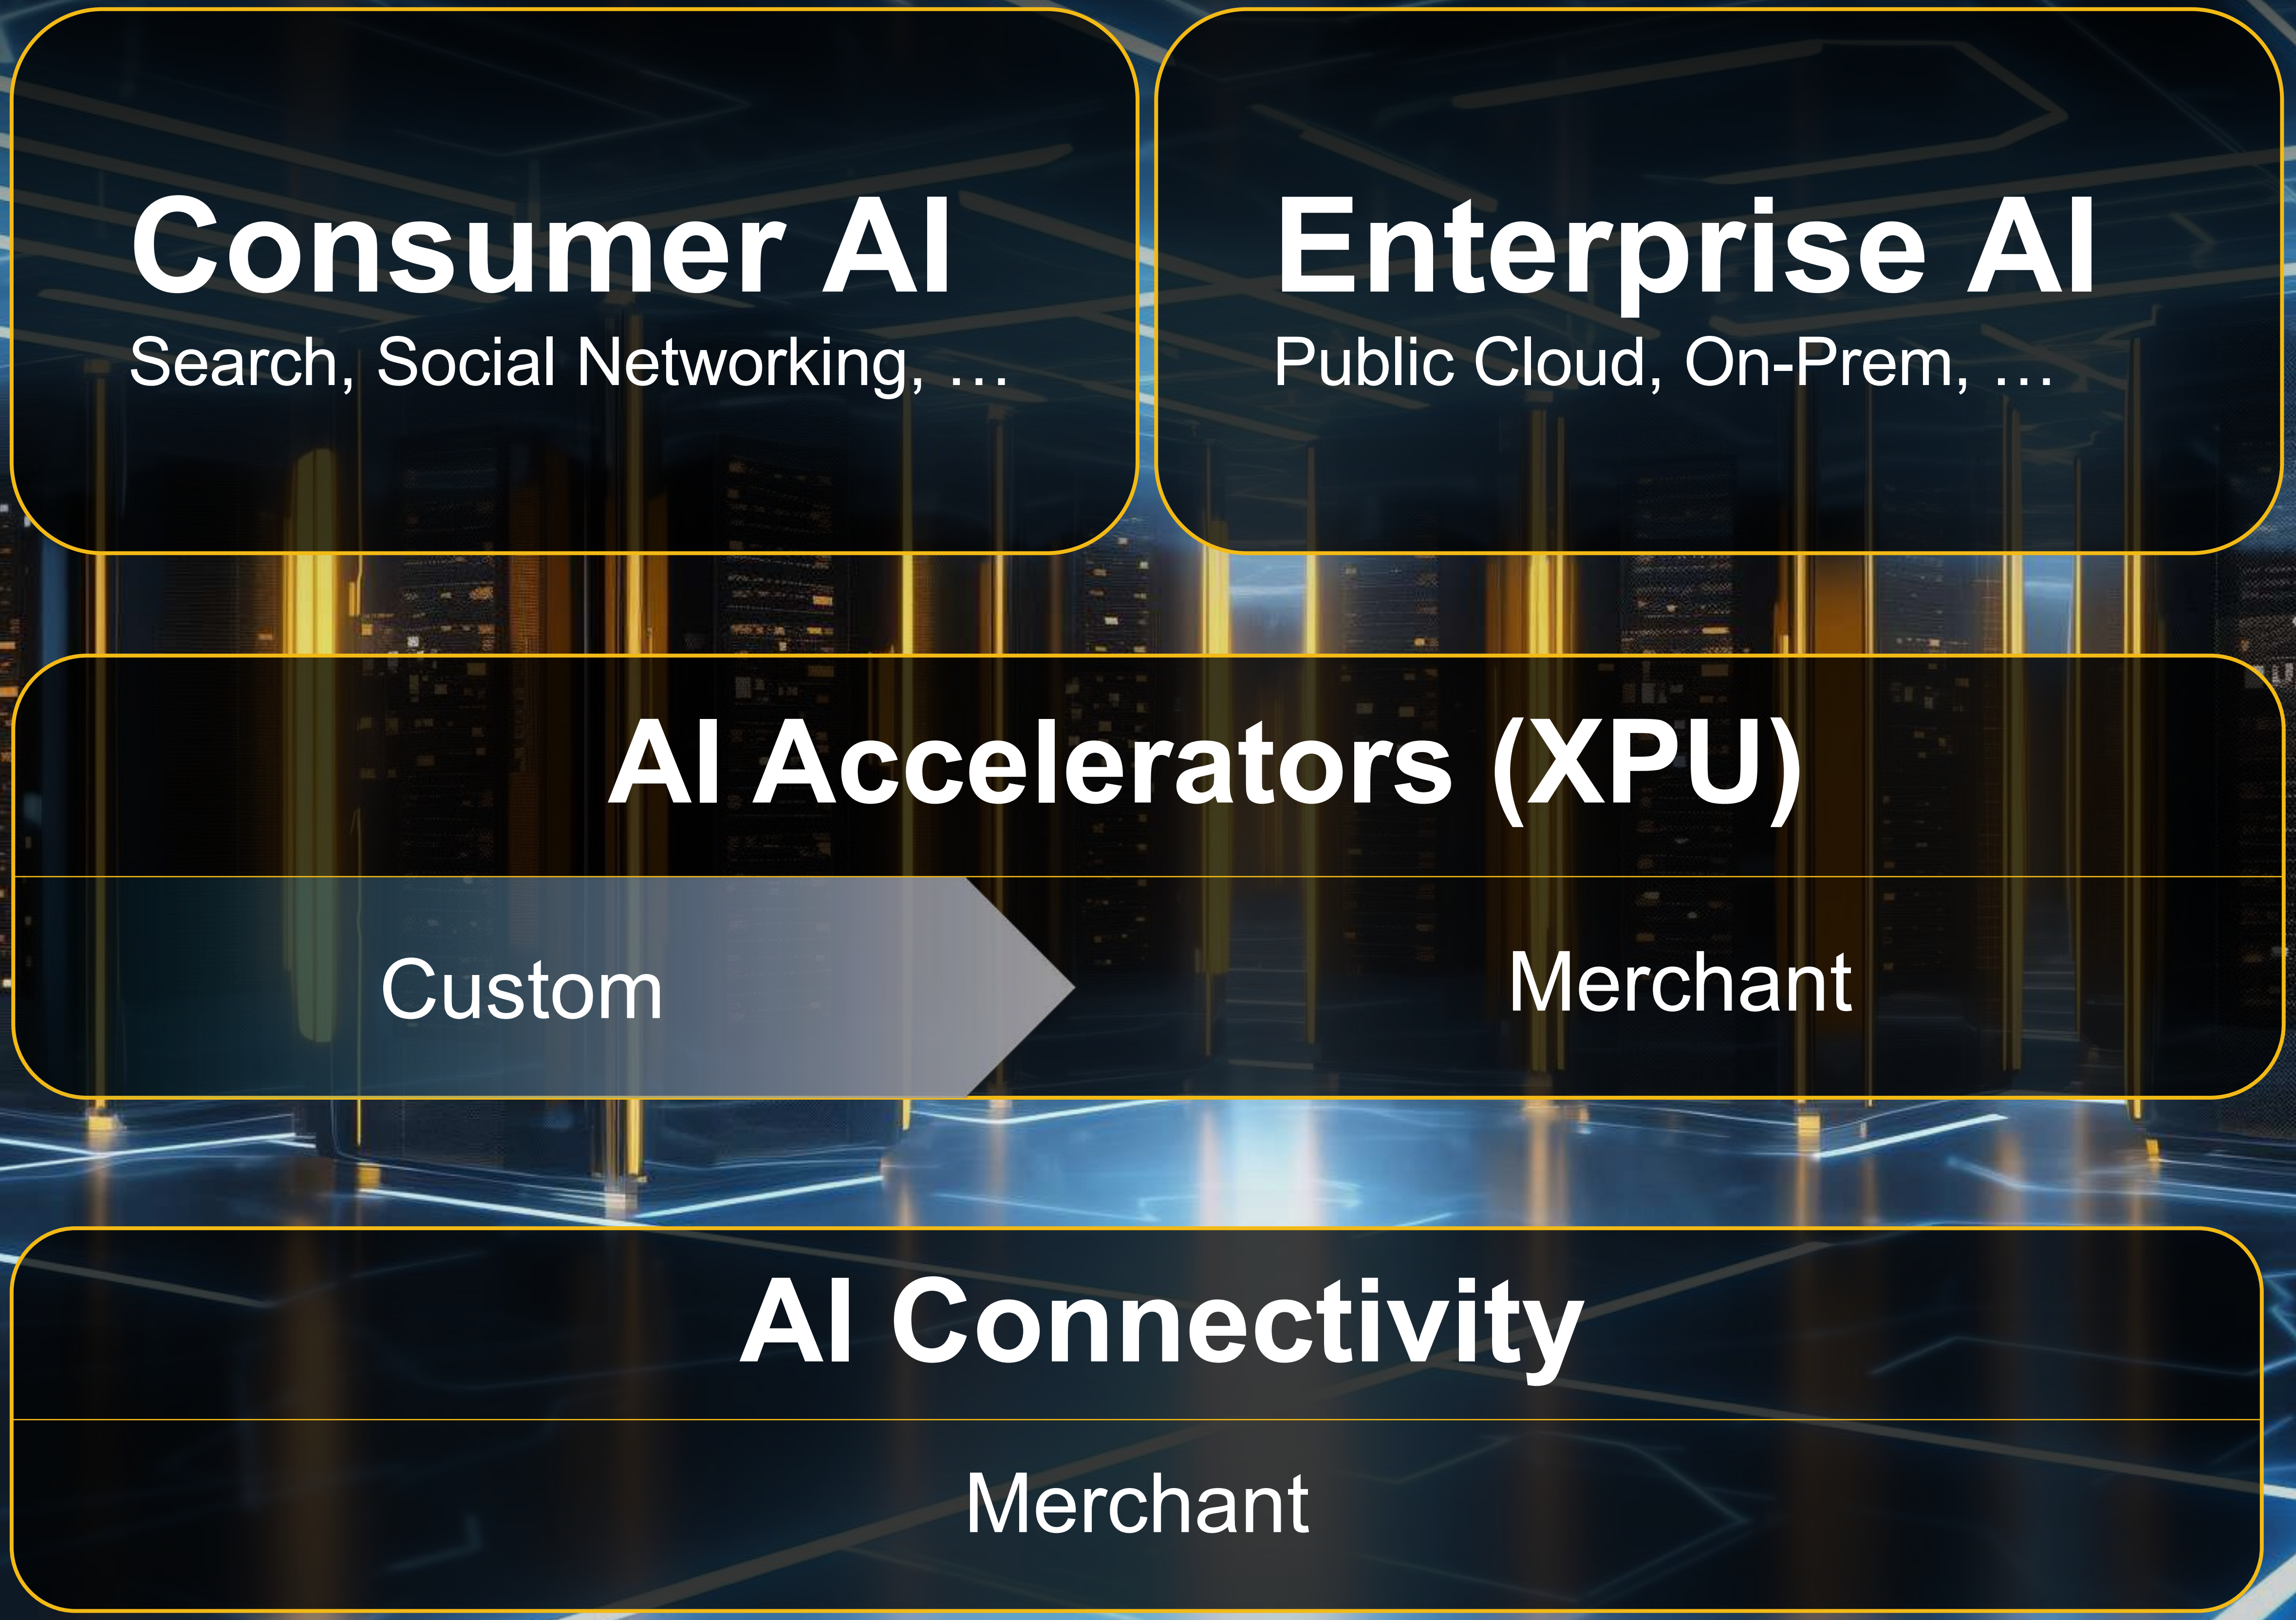

Consumer AI
Search, Social Networking, …

Enterprise AI
Public Cloud, On-Prem, …

AI Accelerators (XPU)

Custom → Merchant

AI Connectivity

Merchant

Consumer AI

Search, Social Networking, …

Enterprise AI

Public Cloud, On-Prem, …

AI Accelerators (XPU)

Custom

Merchant

AI Connectivity

Merchant

Broadcom's AI Journey…

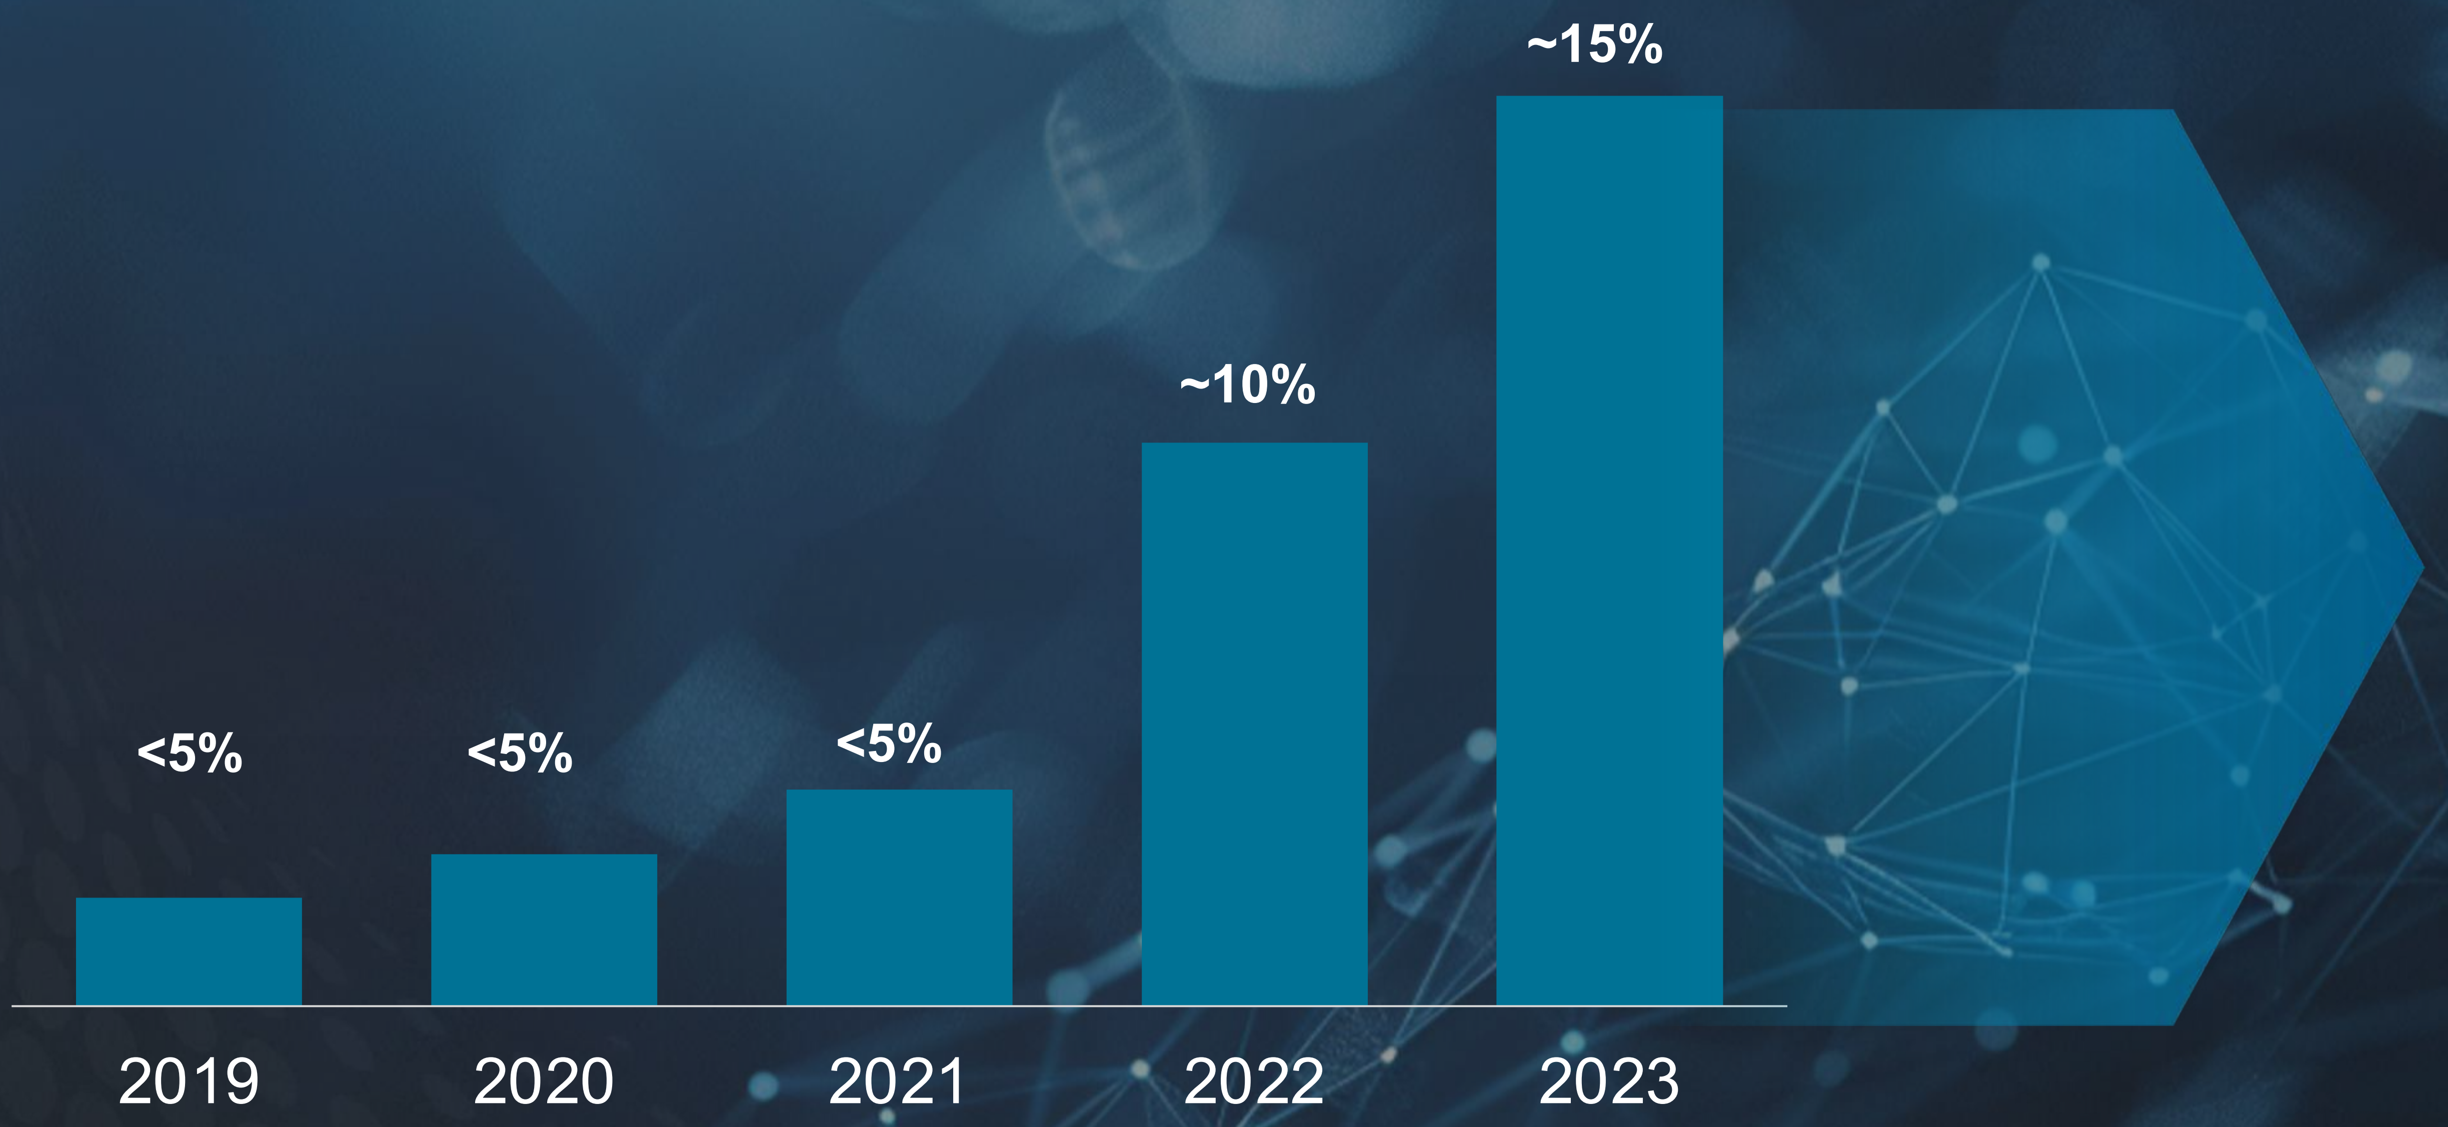

AI Revenue as % of Semiconductor Solutions Revenue

25%

FY2024 Target

OPEN // SCALABLE // POWER EFFICIENT

… is Accelerating

AI Revenue as % of Semiconductor Solutions Revenue

$10B+

Revised FY2024 Target

OPEN // SCALABLE // POWER EFFICIENT

XPUs → Multi-Generational strategic engagements

	2024	2025	2026
Customer #1	Production	✓	✓
Customer #2	Production	✓	✓
Customer #3 *New*	Ramp → Production	✓	✓

The Build-Out is On…

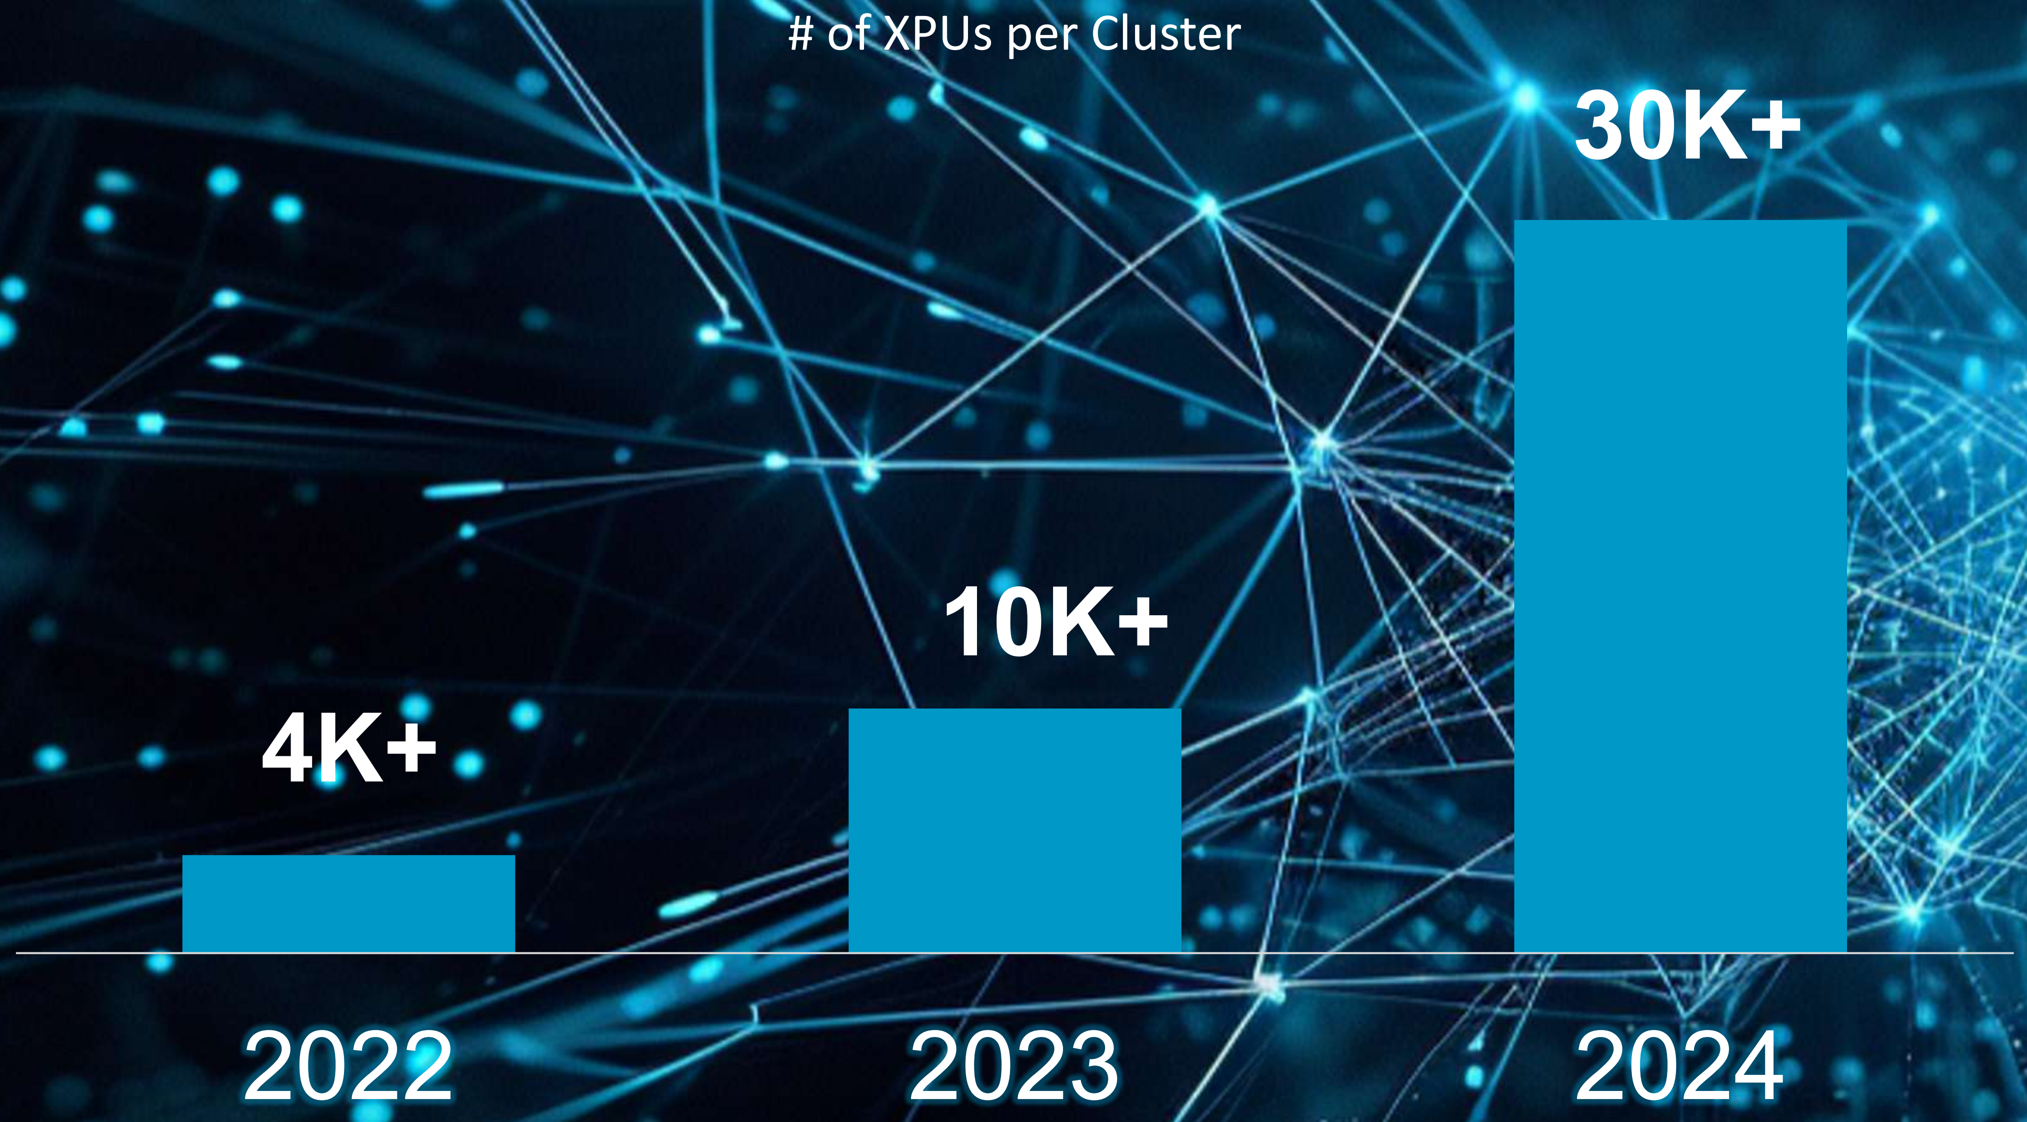

1M?

2027+

OPEN // SCALABLE // POWER EFFICIENT

XPU

Scale Up

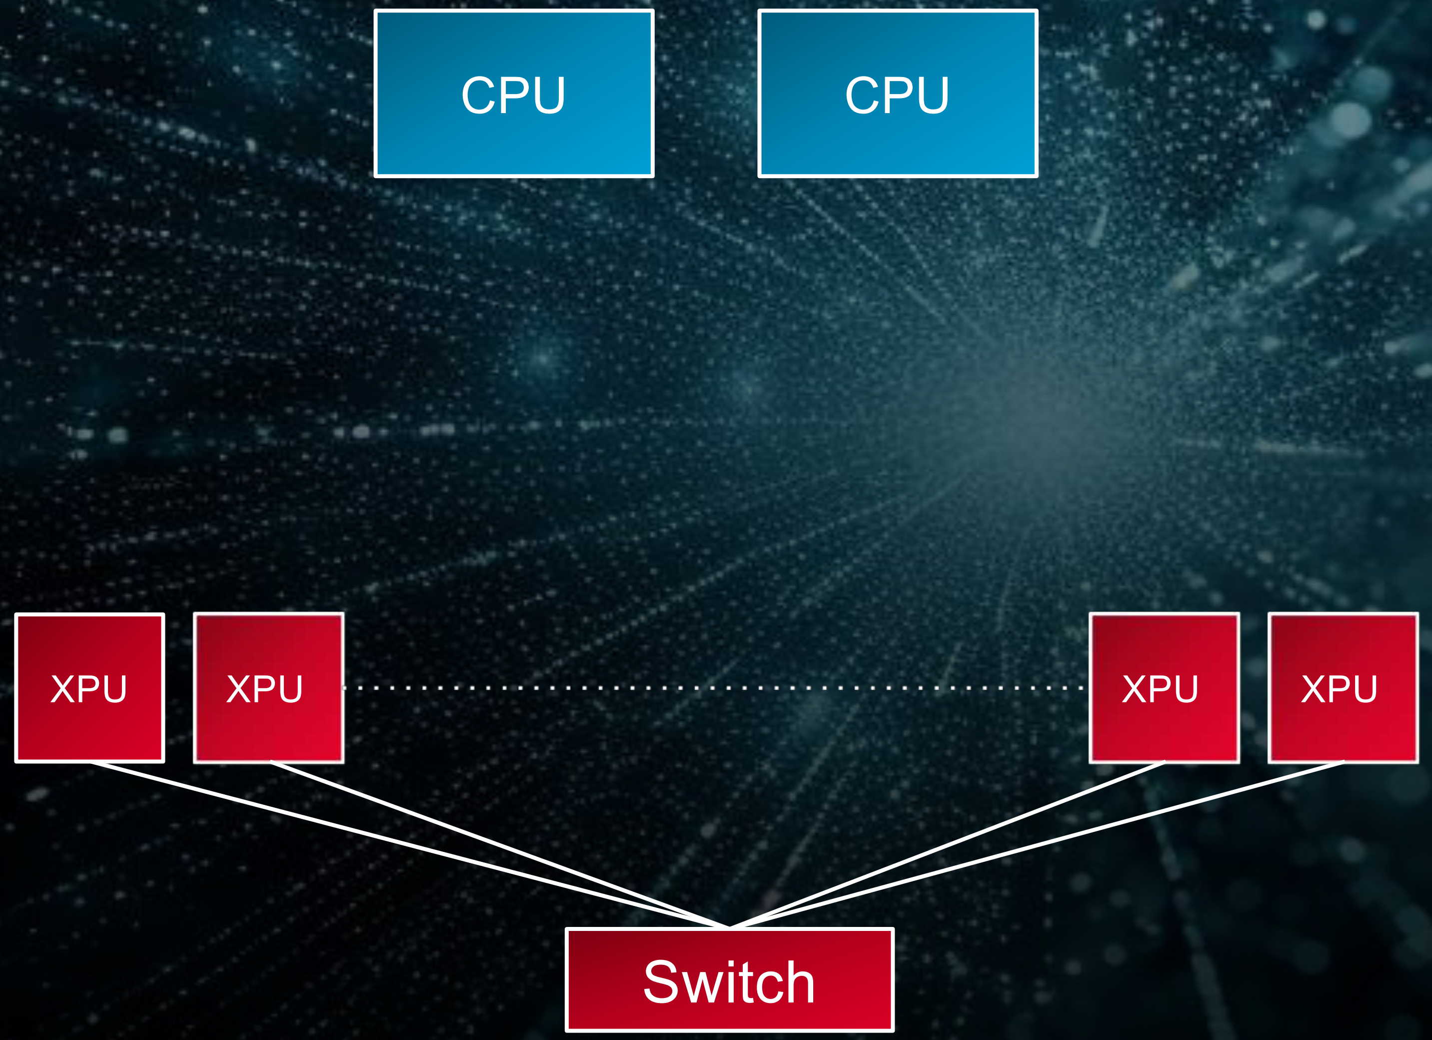

Scale Up

Interconnect

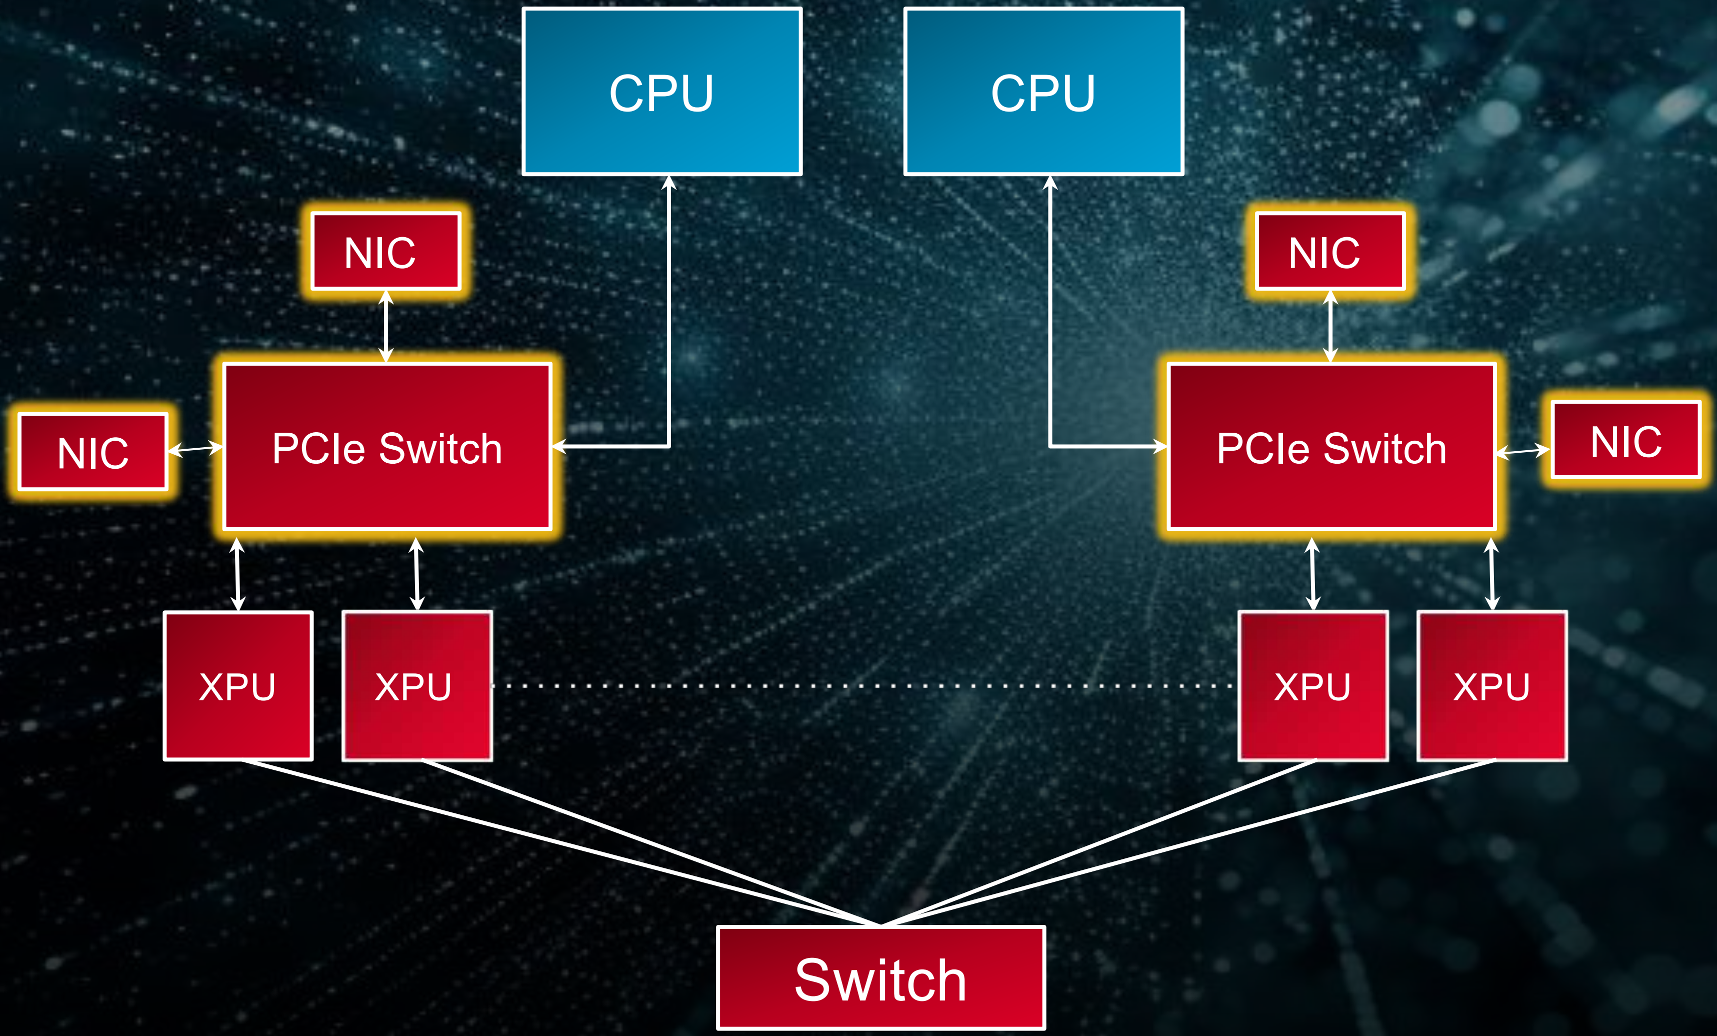

Scale Up

OPEN // SCALABLE // POWER EFFICIENT

Interconnect

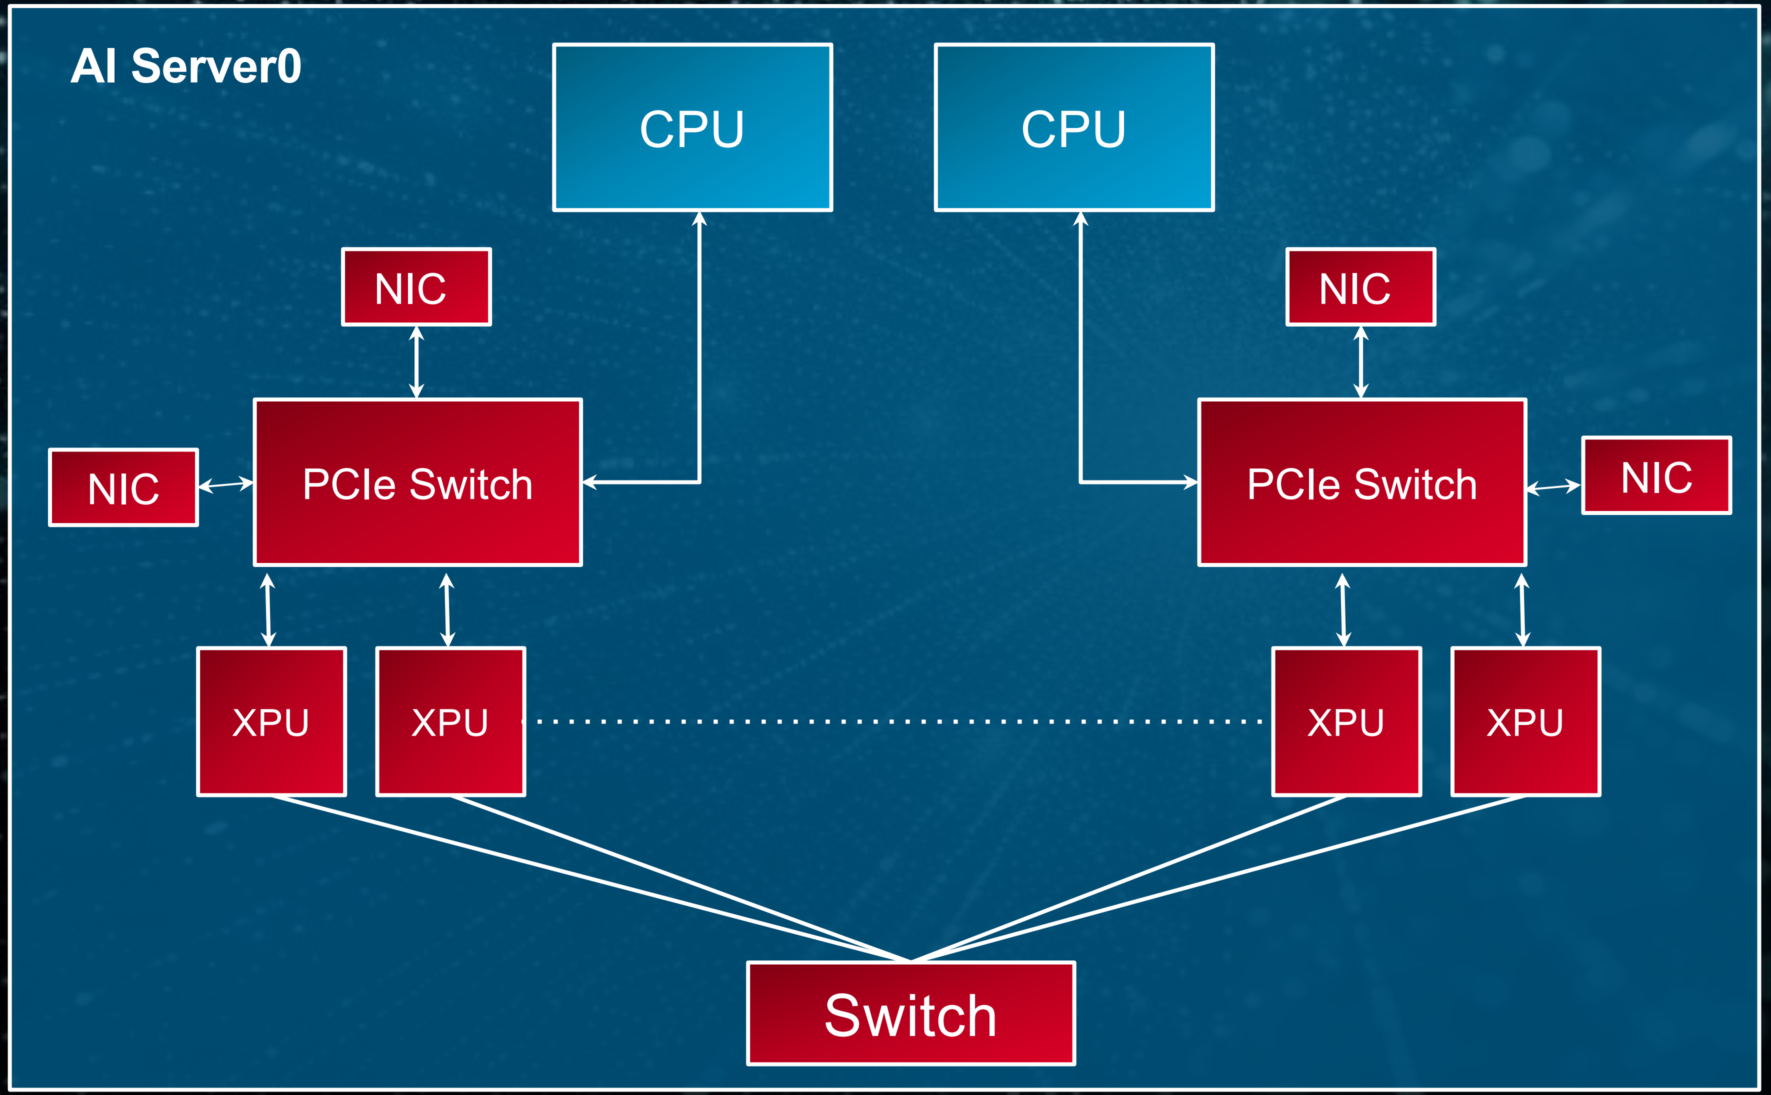

→Broadcom SAM

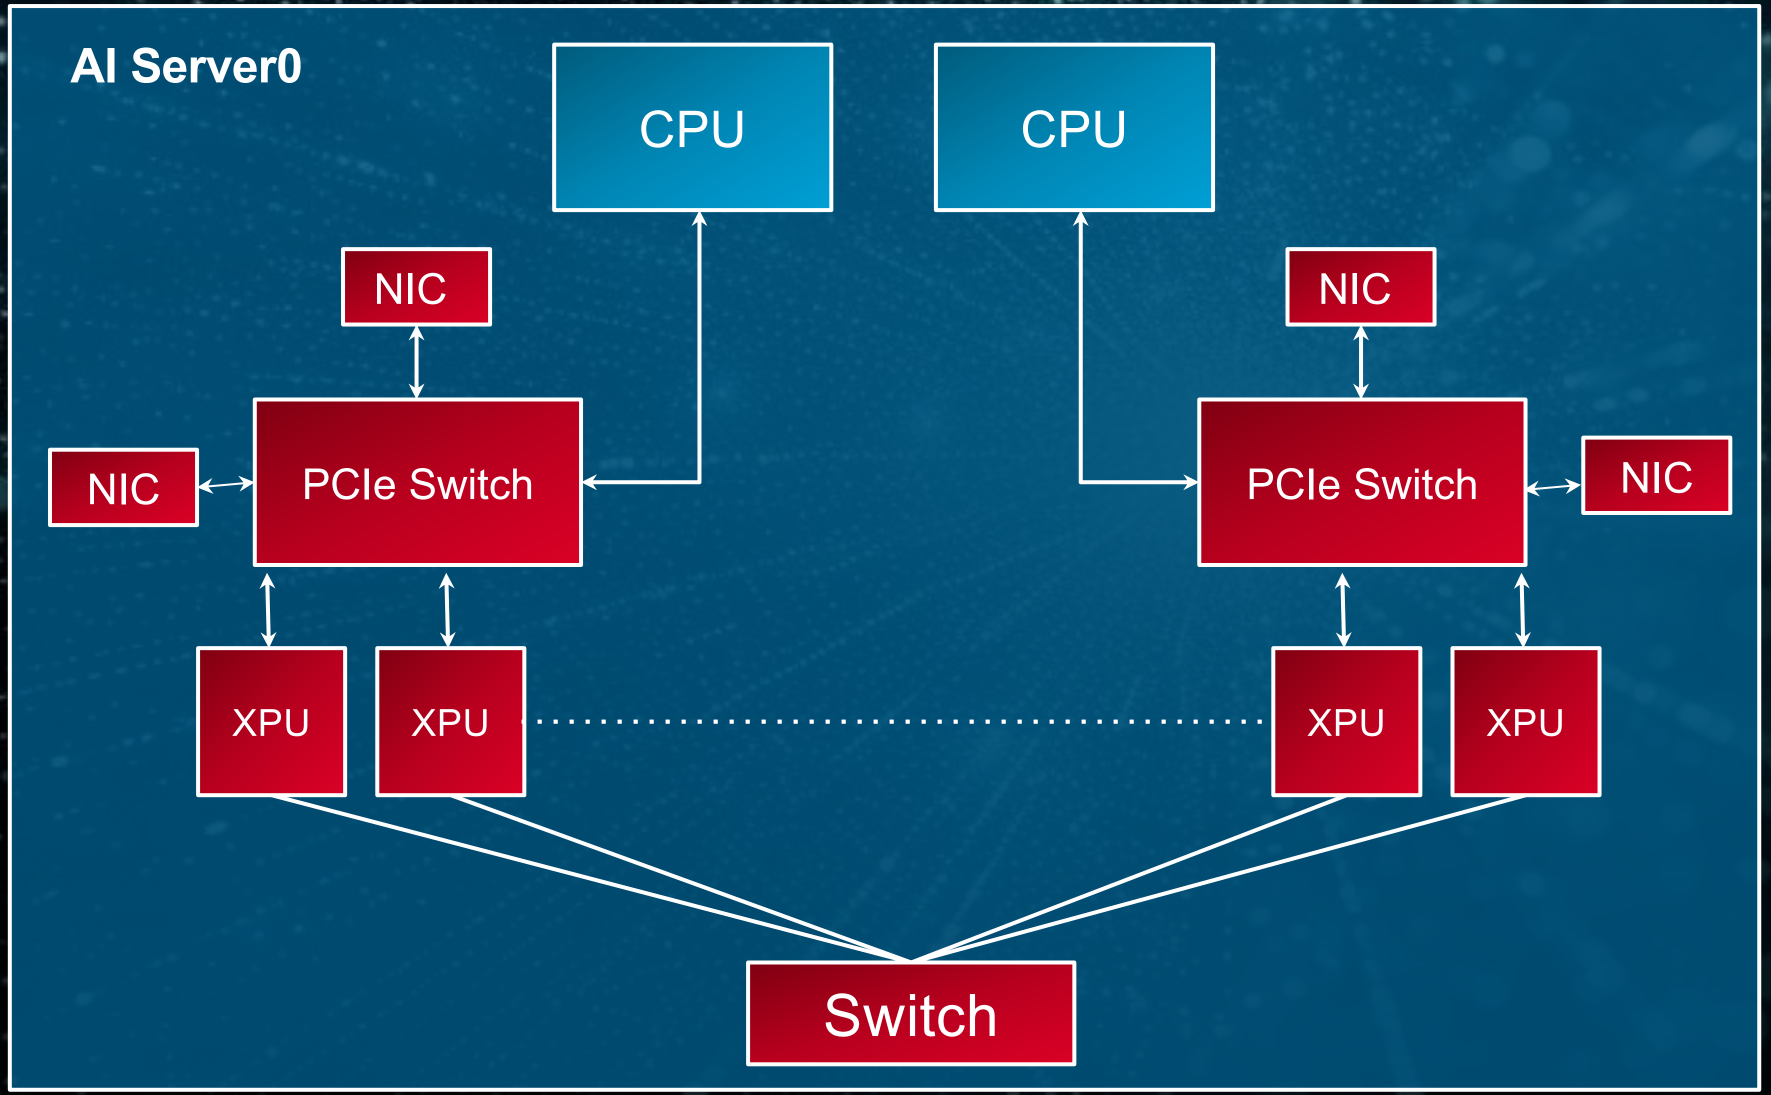

OPEN // SCALABLE // POWER EFFICIENT

Scale Out


AI Server
CPU
CPU
NIC
NIC
NIC
PCIe Switch
PCIe Switch
NIC
XPU
XPU
XPU
XPU
Switch

Scale Out

OPEN // SCALABLE // POWER EFFICIENT

Scale Up + Scale Out = Back End

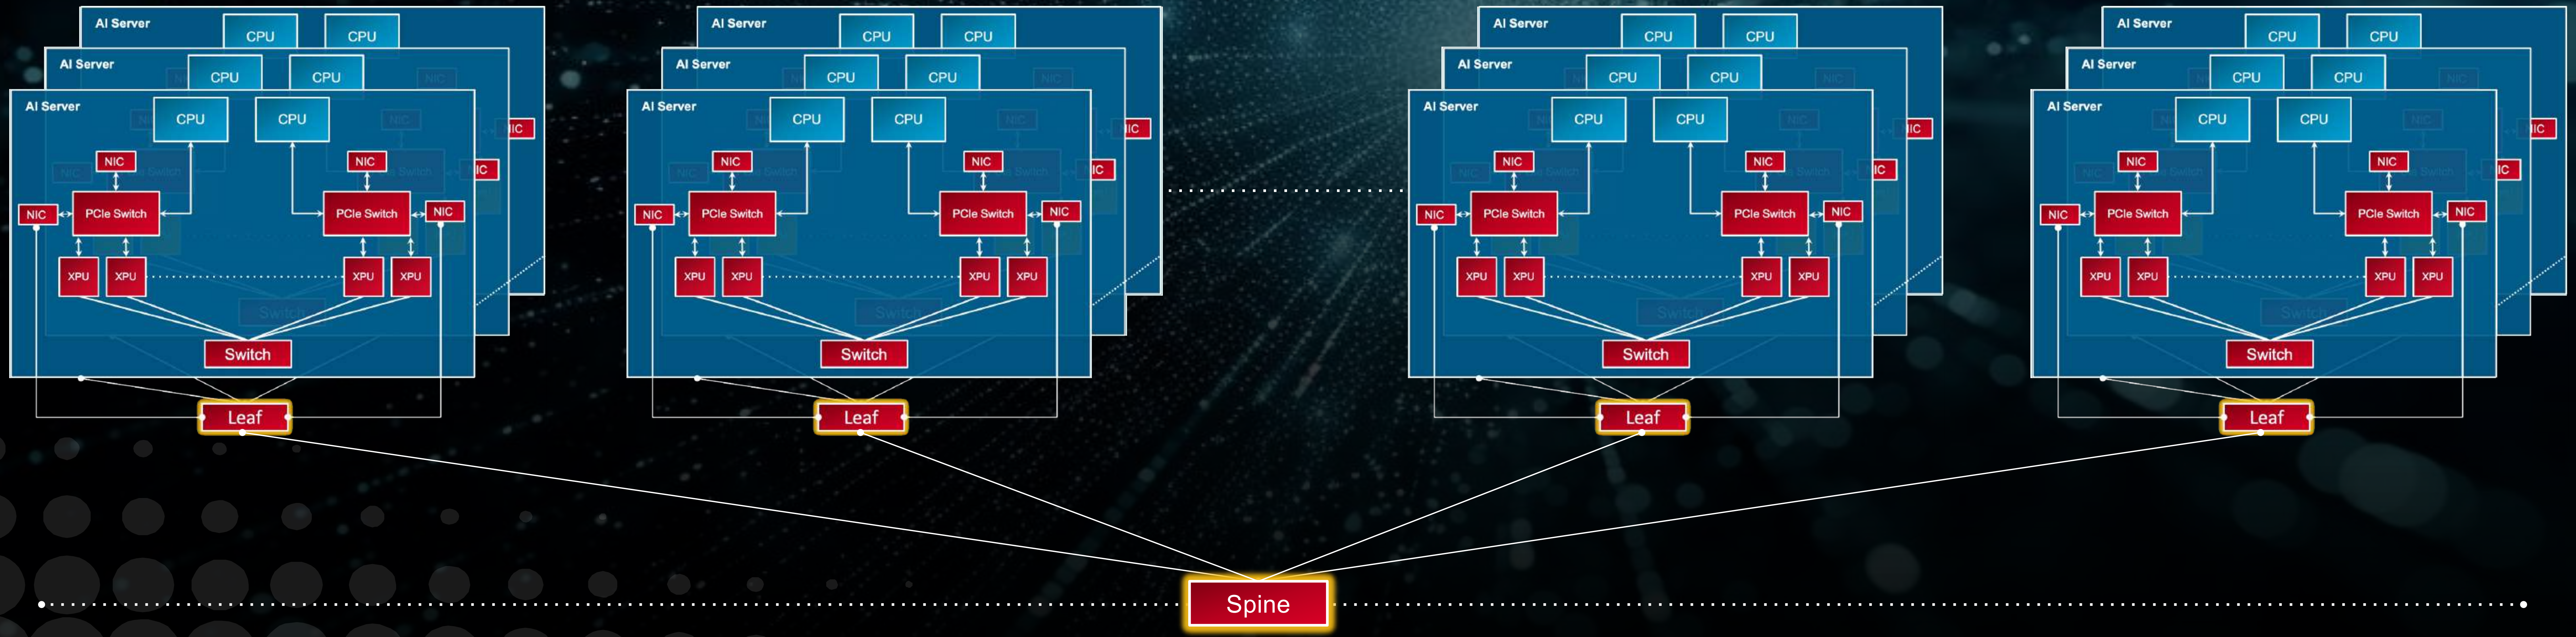

OPEN // SCALABLE // POWER EFFICIENT

Front End

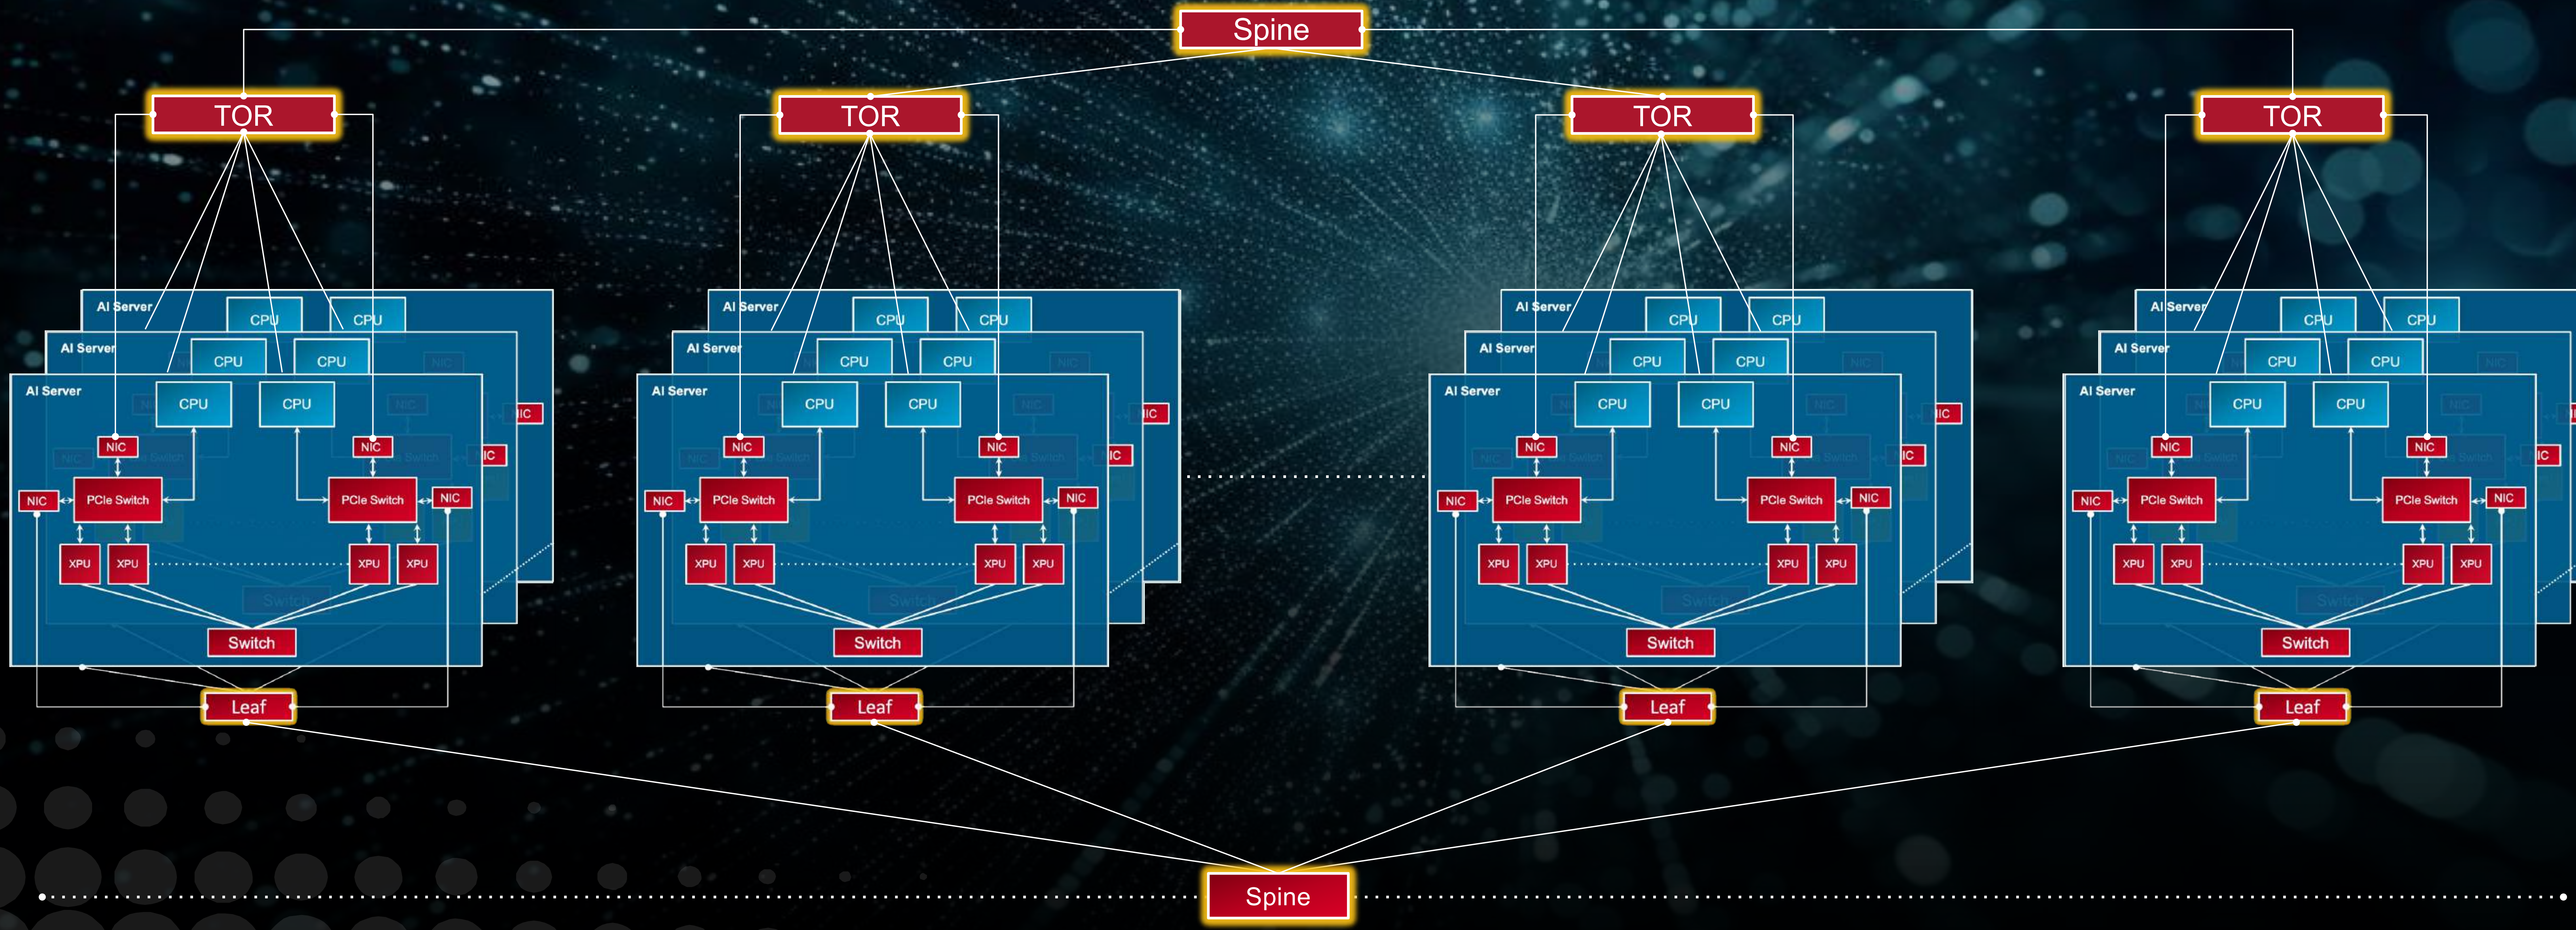

OPEN // SCALABLE // POWER EFFICIENT

Interconnect

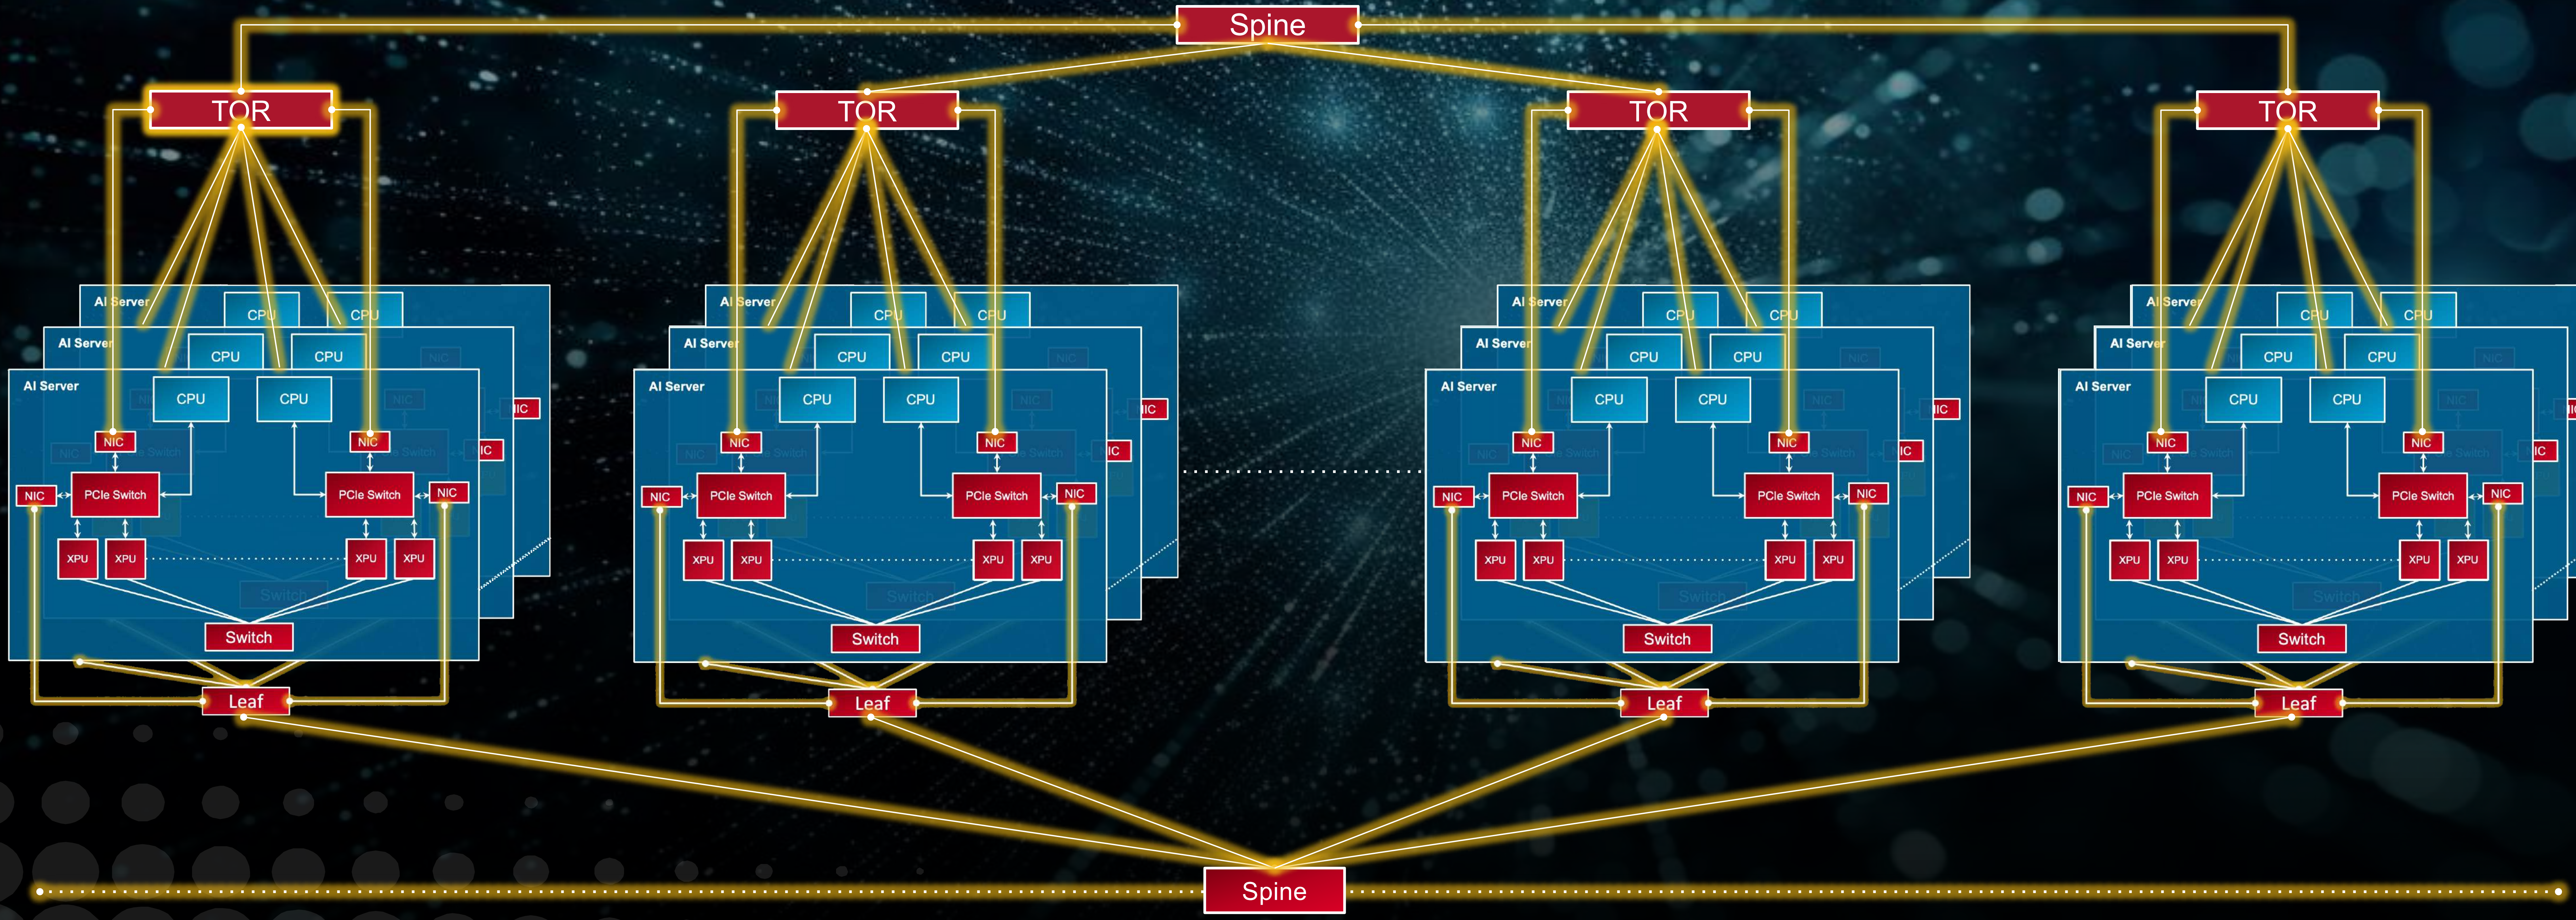

OPEN // SCALABLE // POWER EFFICIENT

→Broadcom SAM

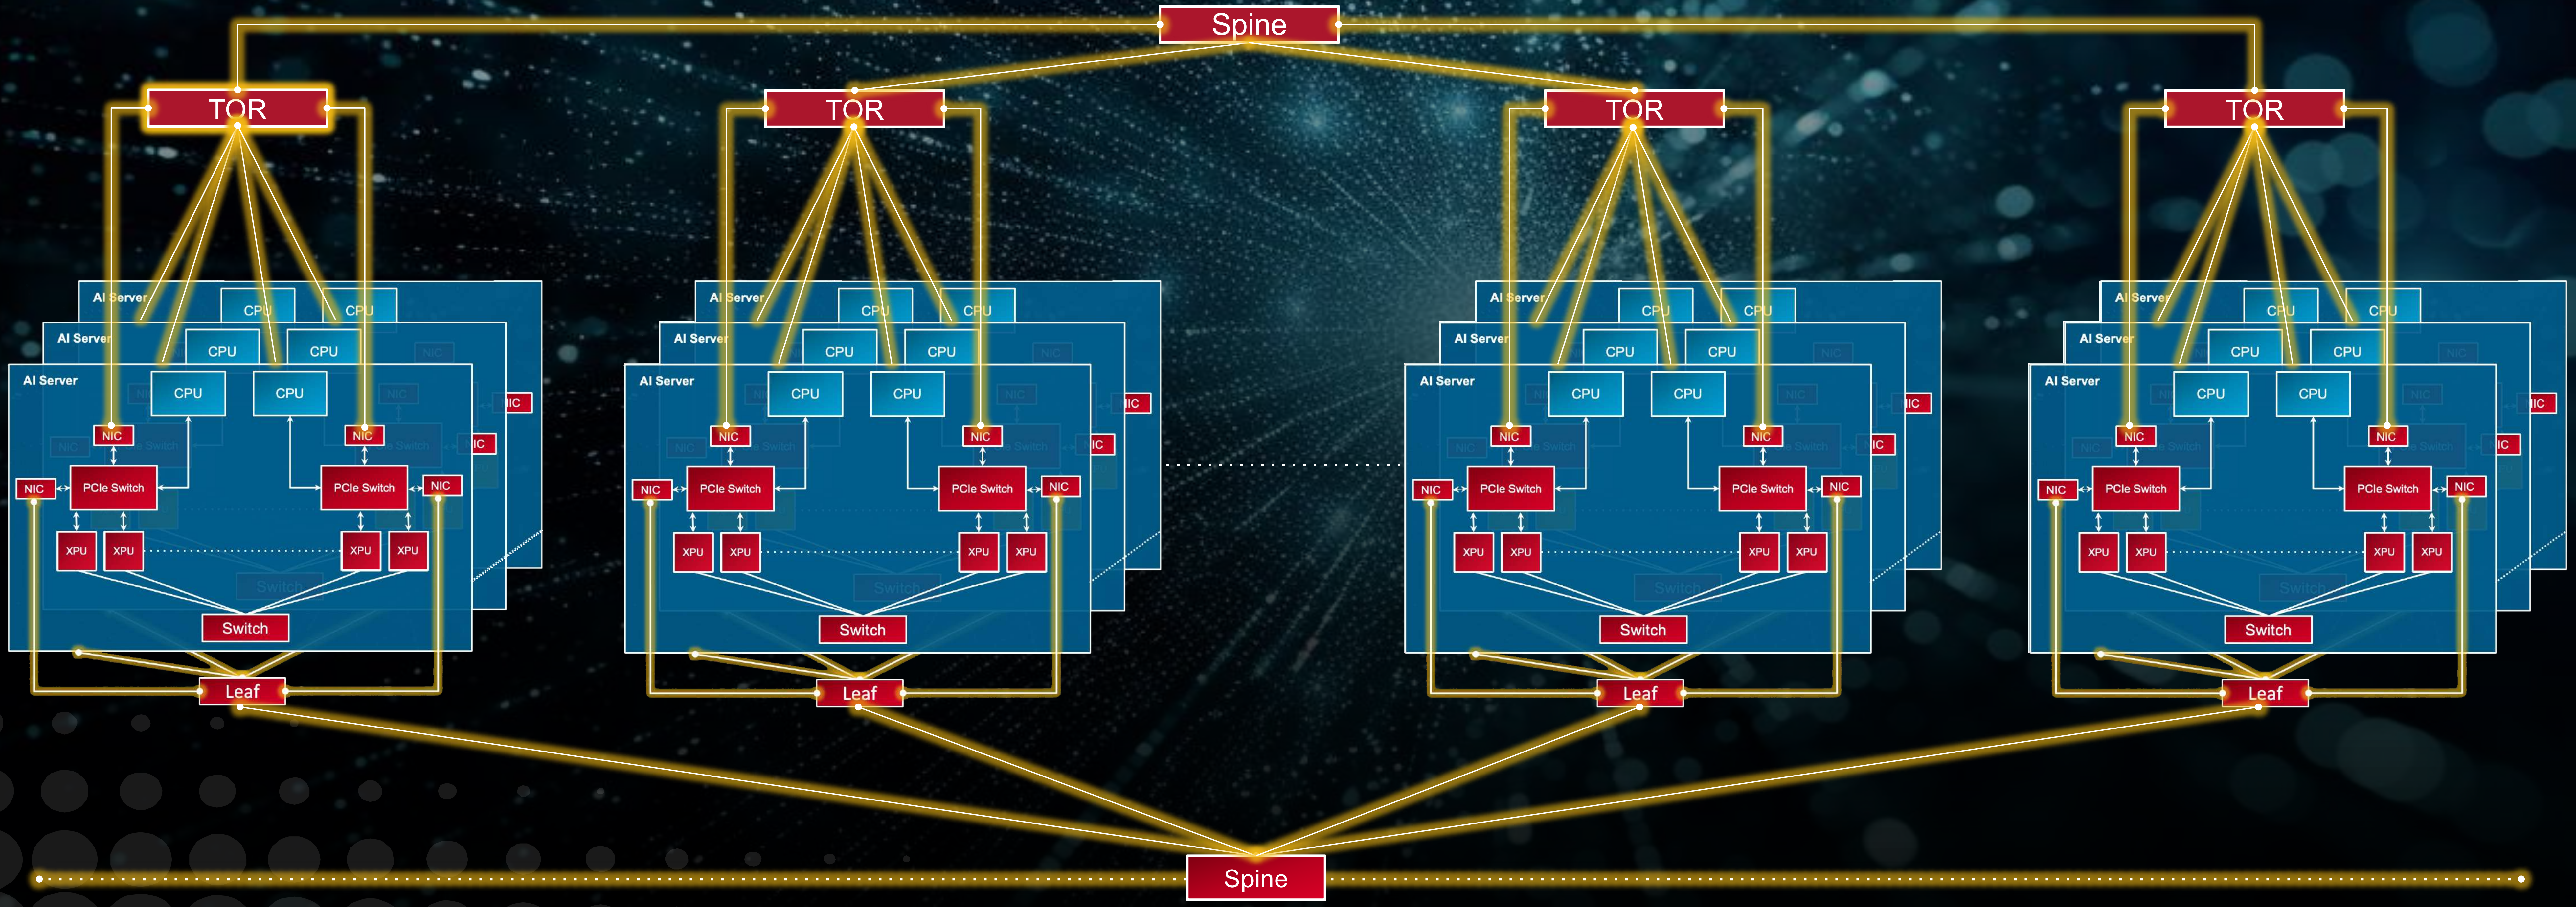

OPEN // SCALABLE // POWER EFFICIENT

1. Power 2. Complexity 3. Ecosystem

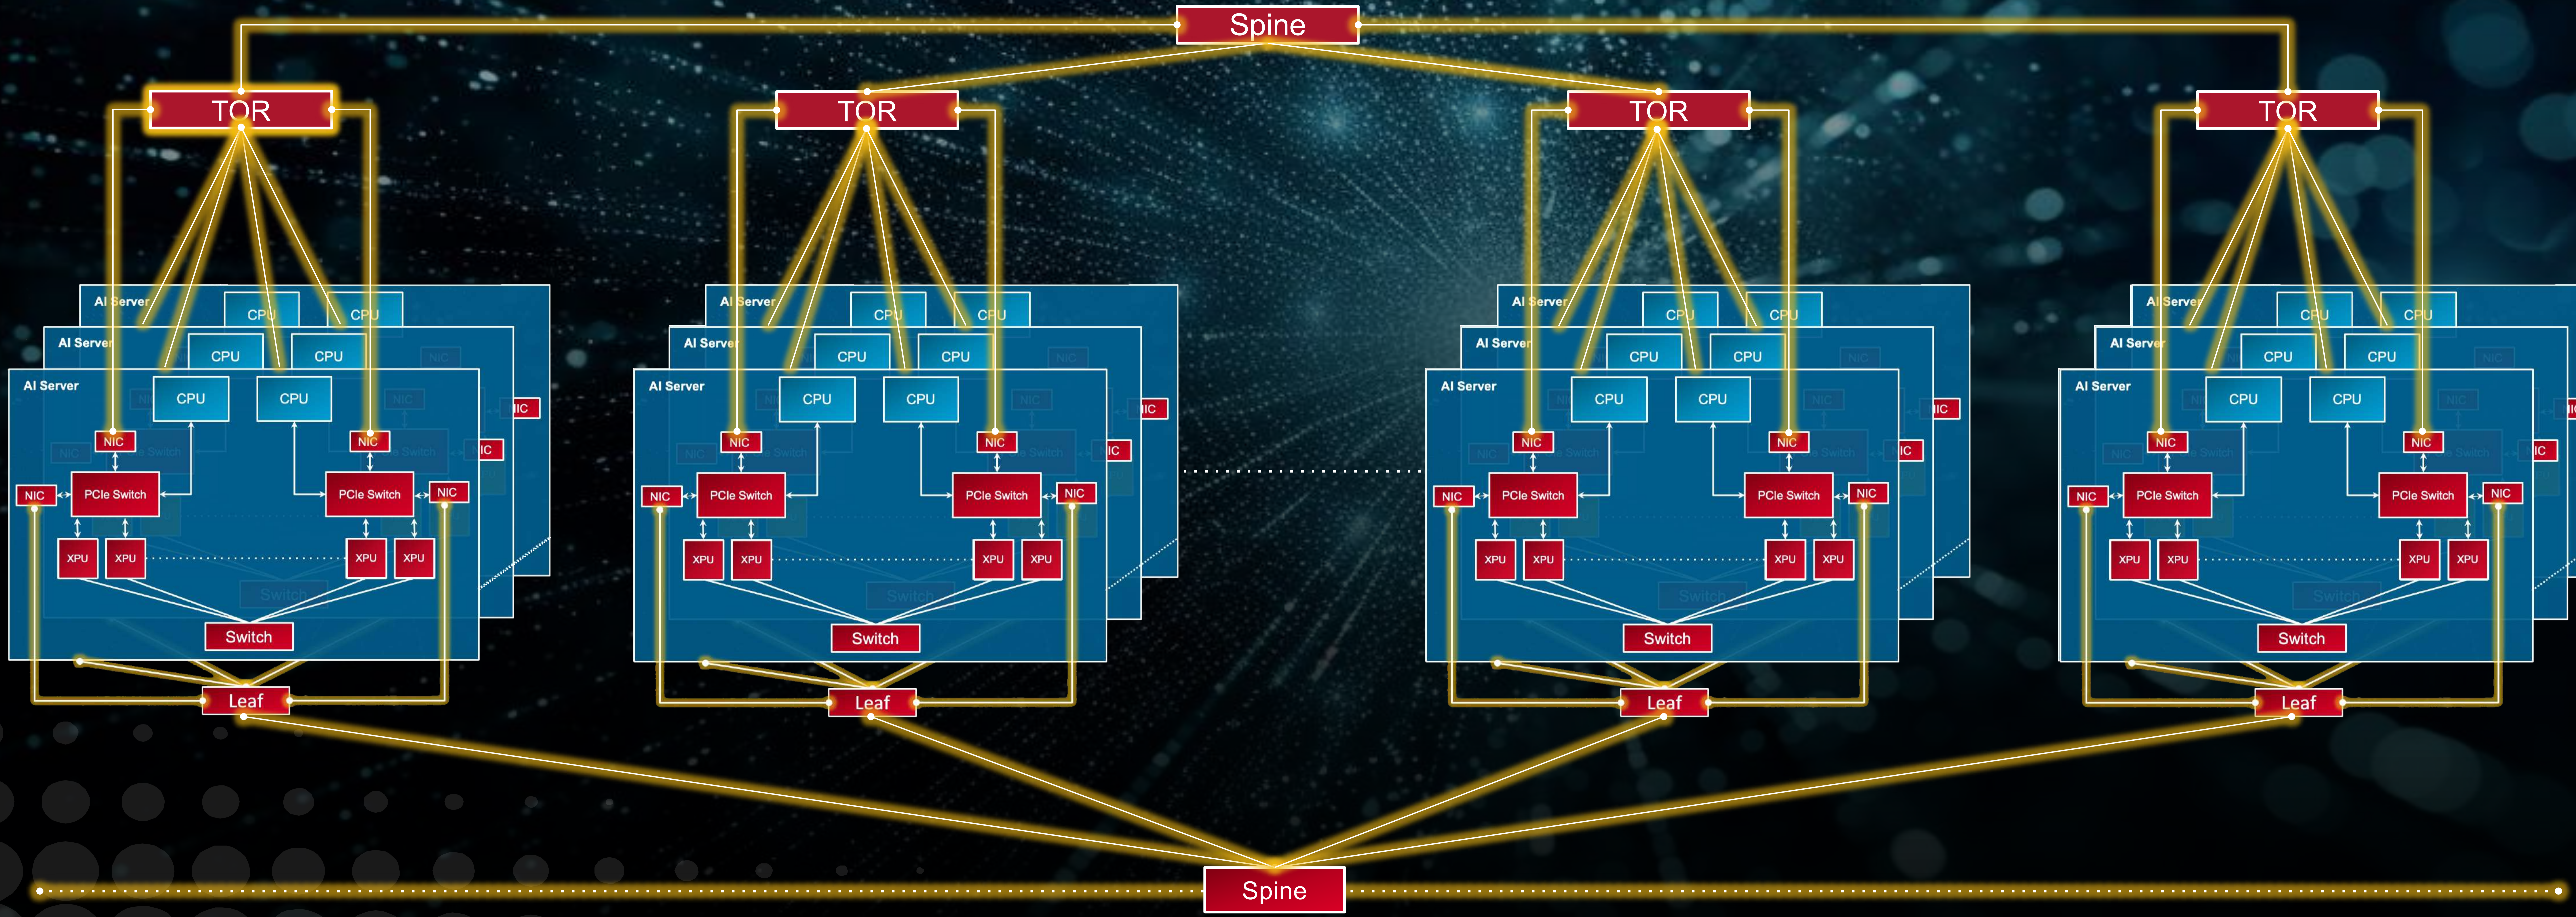

OPEN // SCALABLE // POWER EFFICIENT

Open Scalable Power Efficient

Sustainable Technology Leadership

Ethernet PCIe Optics SerDes DSP Custom

BROADCOM

ENABLING AI Infrastructure

Scalable AI Networks

Ram Velaga

General Manager
Core Switching Group
Broadcom

OPEN // SCALABLE // POWER EFFICIENT

New AI Models → 100X+ Scale ➔ Distributed Computing

1 Million

AI Accelerators

BLOOM

cohere

DeepMind Gopher

Gemini

GPT-4

LLaMA

MISTRAL AI_

PaLM 2

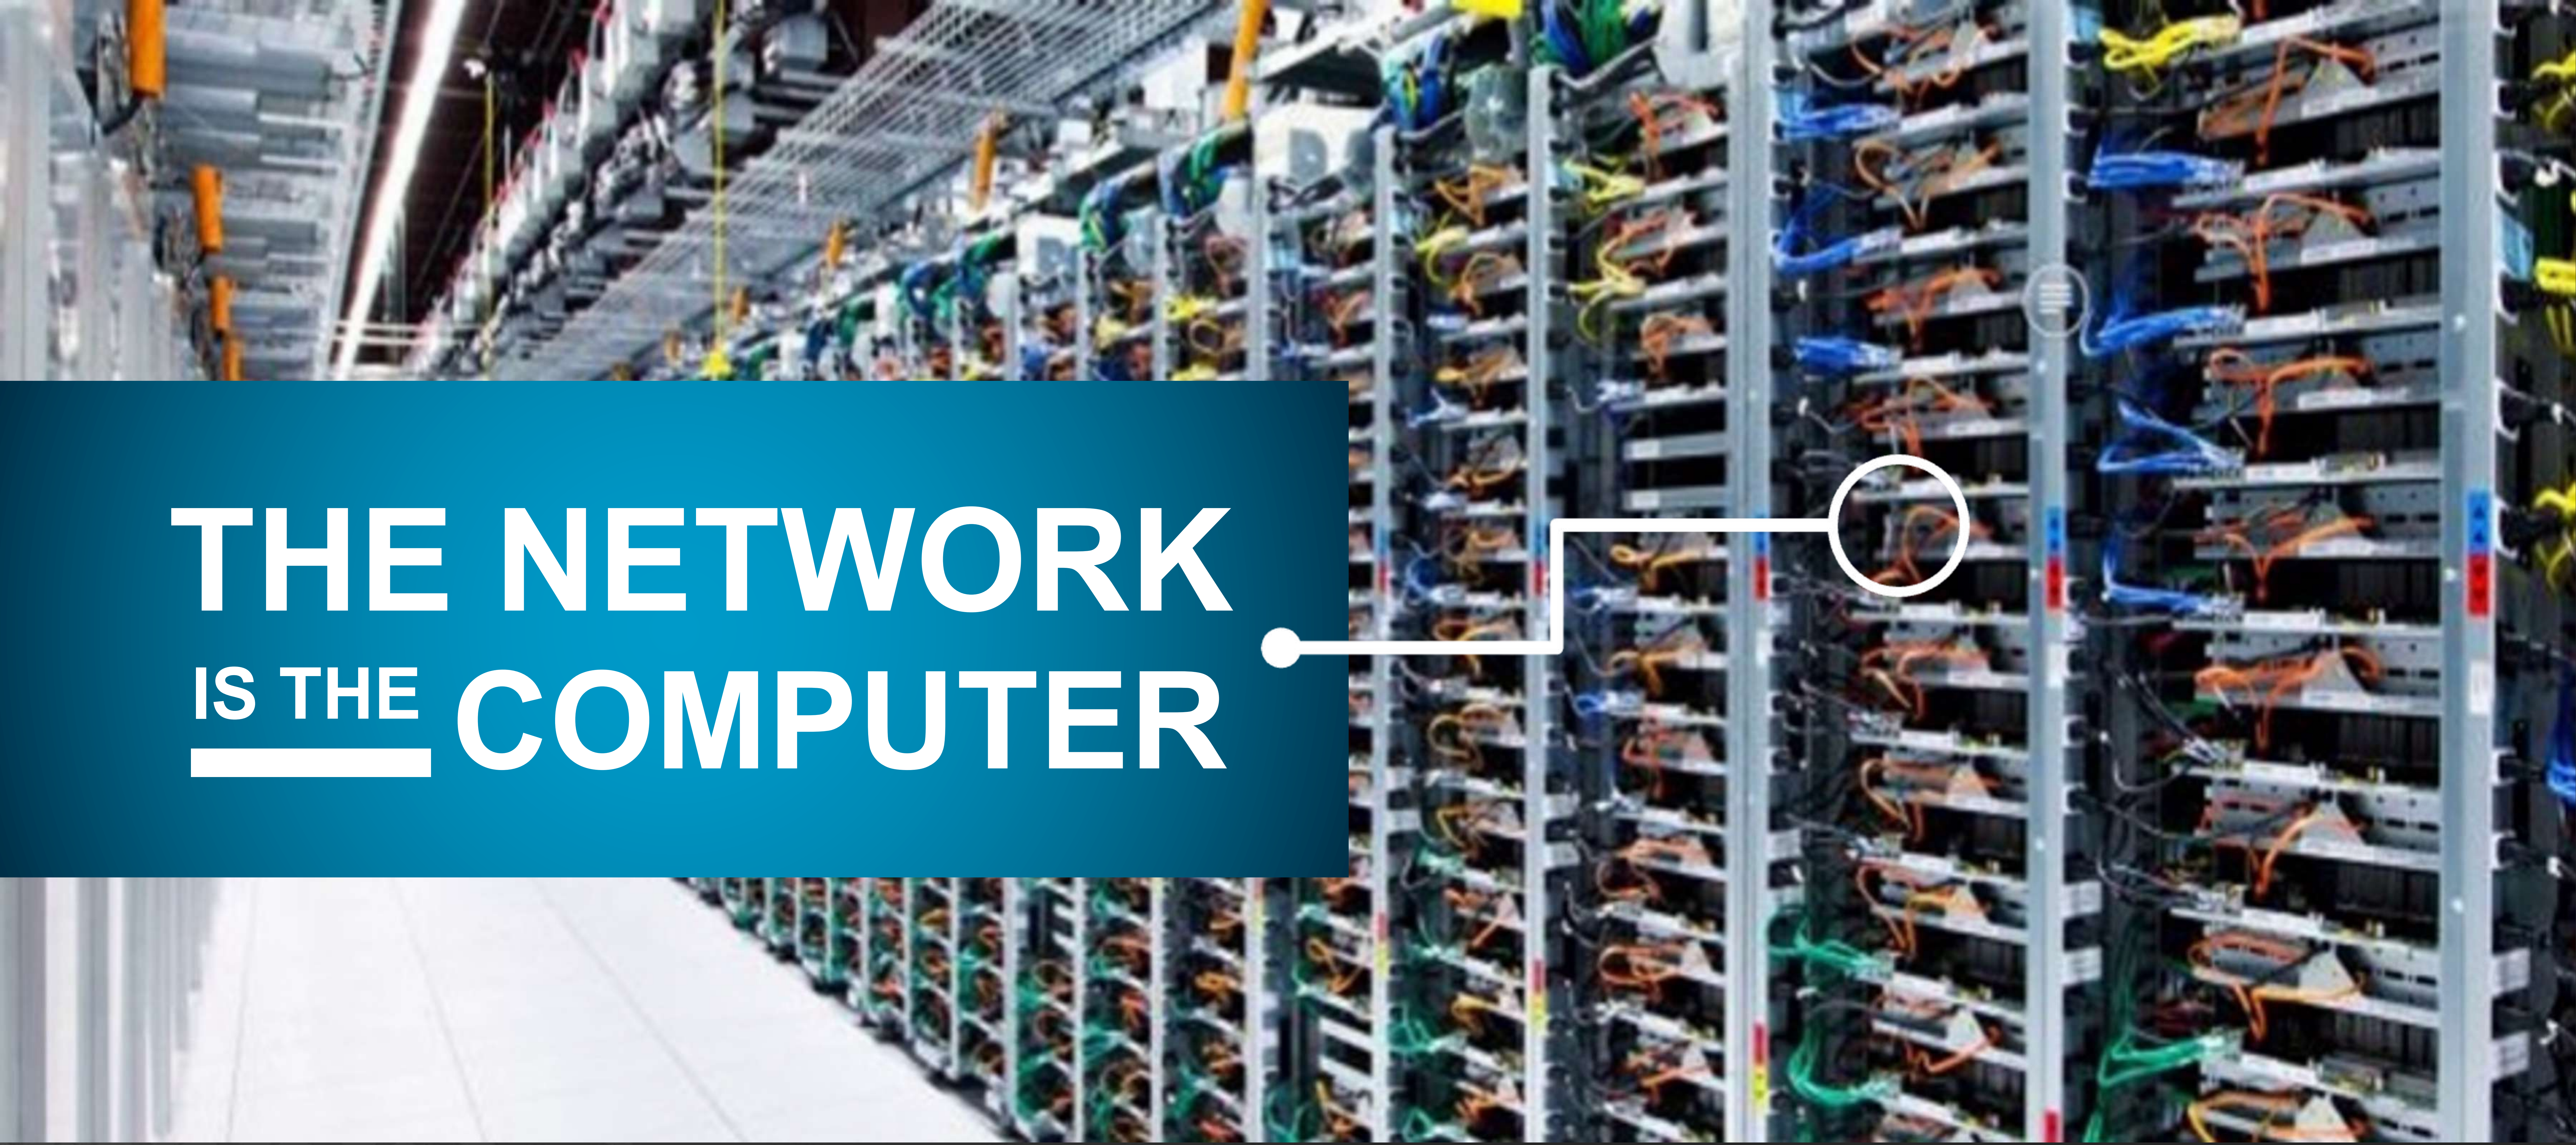

THE NETWORK IS THE COMPUTER

OPEN // SCALABLE // POWER EFFICIENT

> “Network I/O is Key for Recommendation Workloads.”

Meta, OCP 2022 Global Summit

TIME SPENT IN NETWORKING

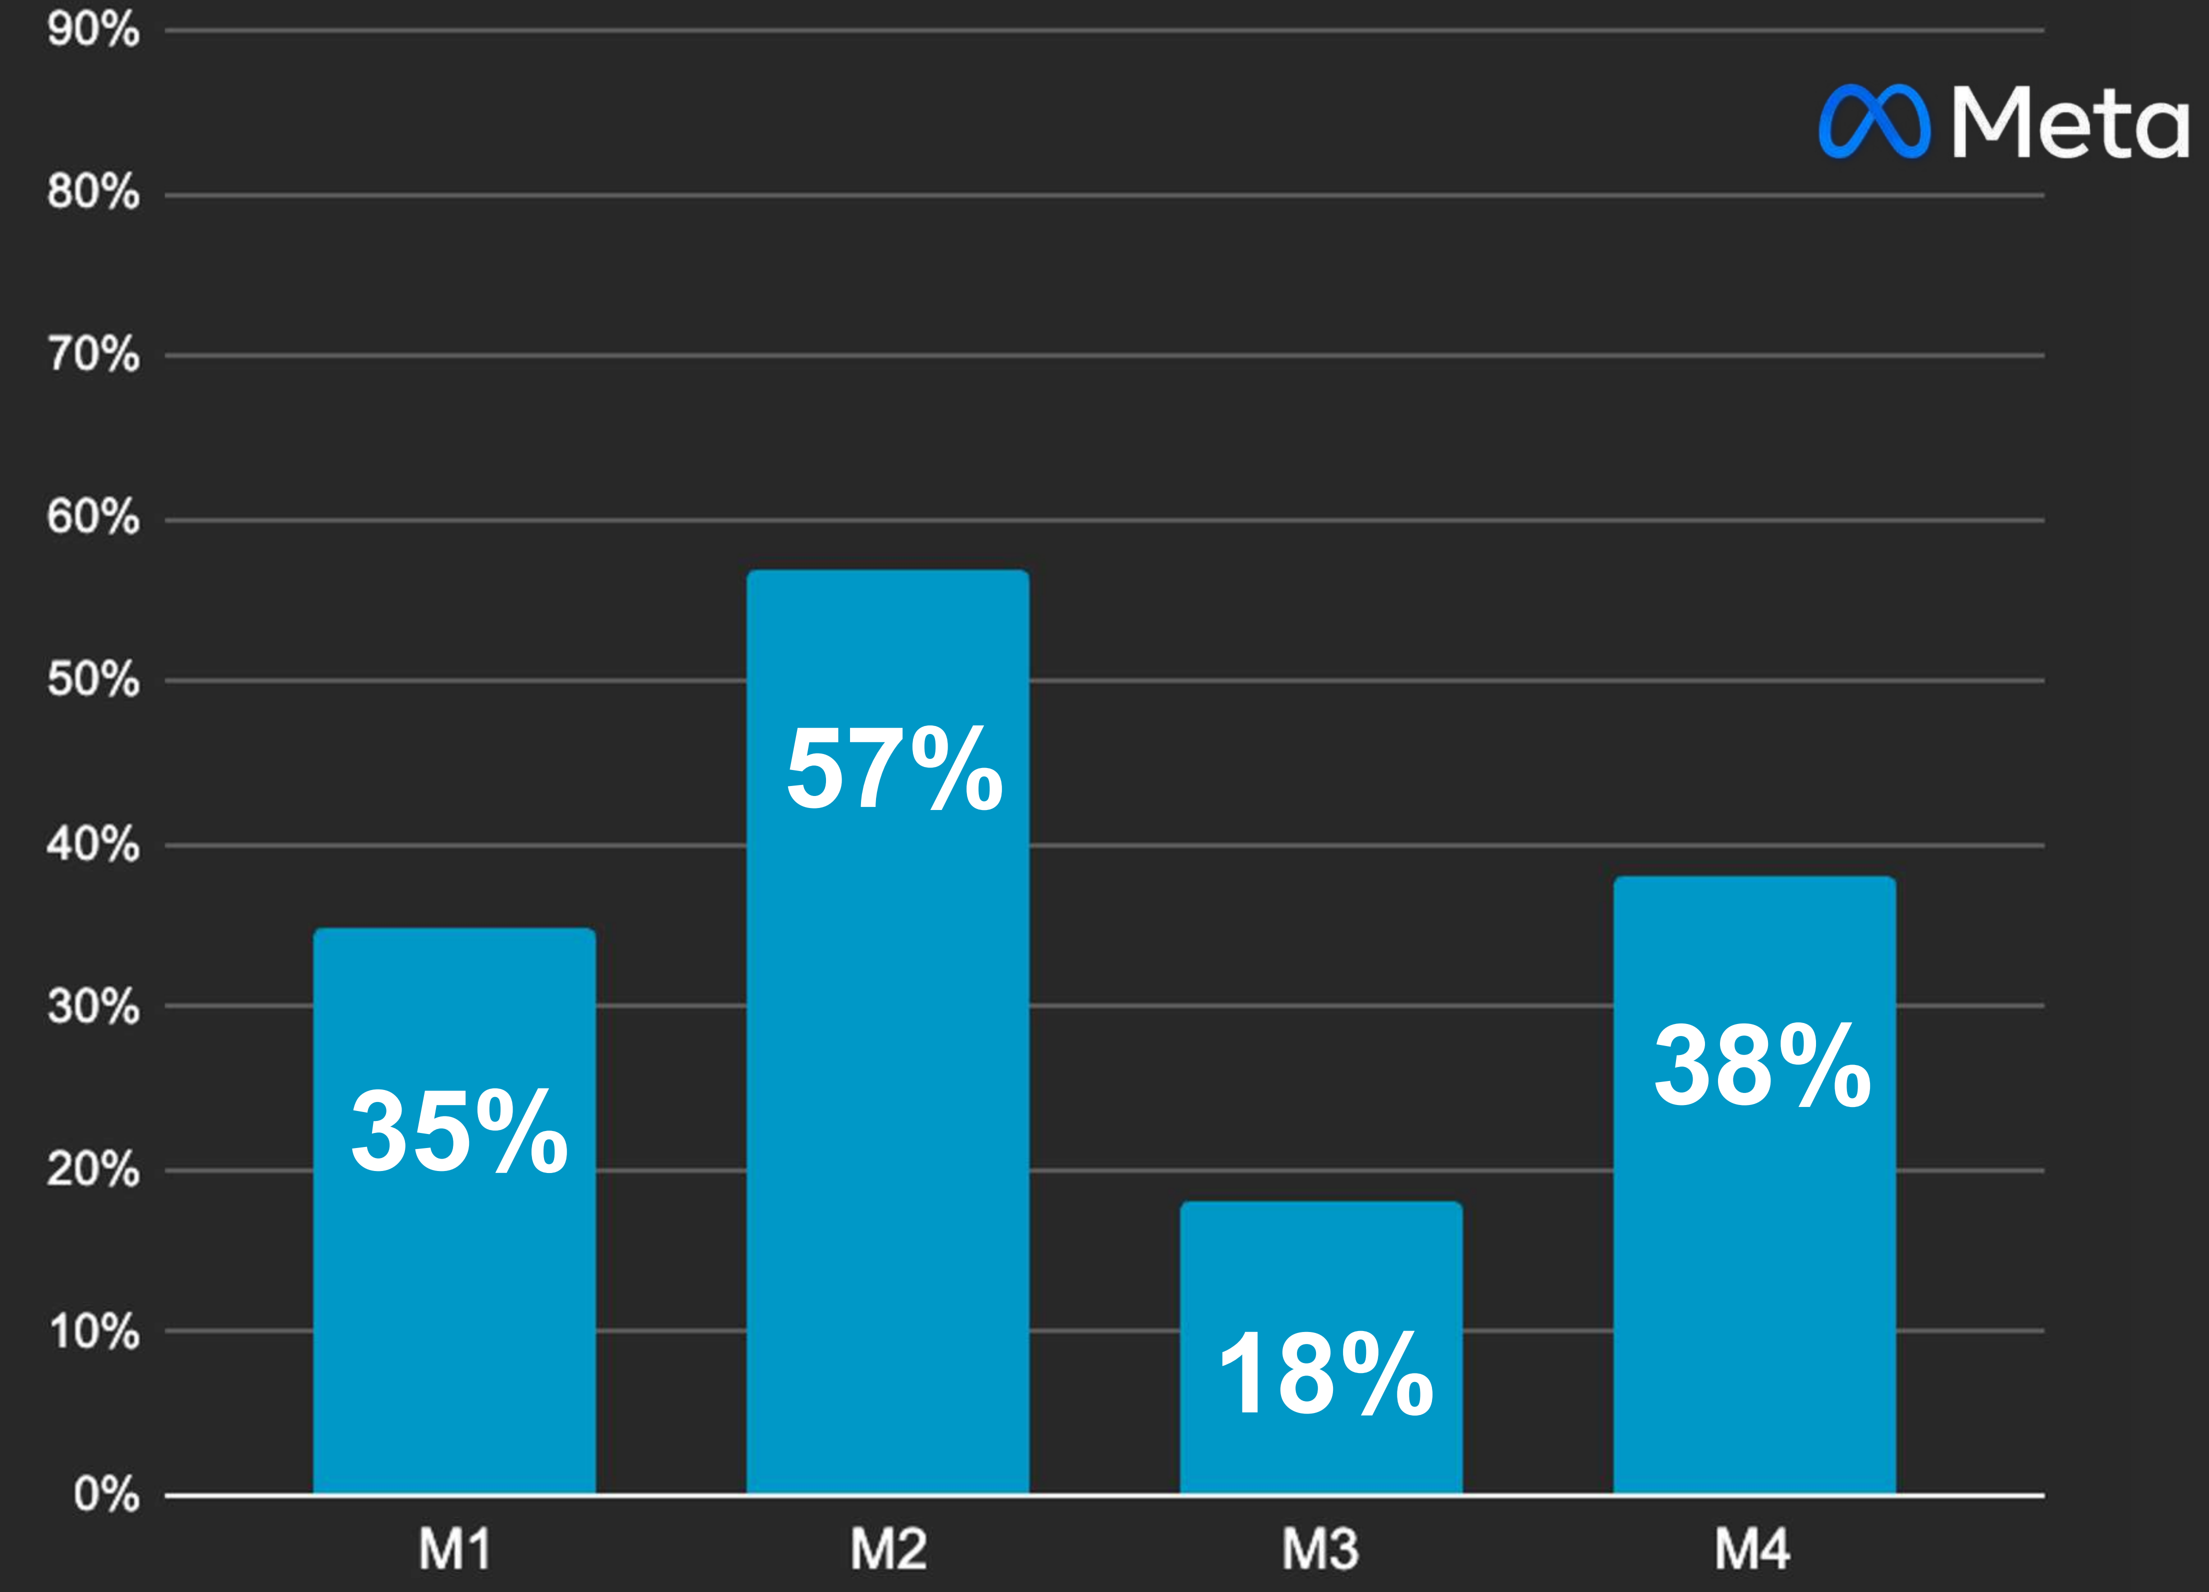

M# = ML model #

Ranking requires high injection & bisection bandwidth

OPEN // SCALABLE // POWER EFFICIENT

What Makes AI Networks Unique?

- Very high bandwidth
- RDMA traffic ➔ Bulk data transfers
- Intermittent data surges
- Straggler data significantly impacts job completion time
- Training jobs run over long durations (hours, days)

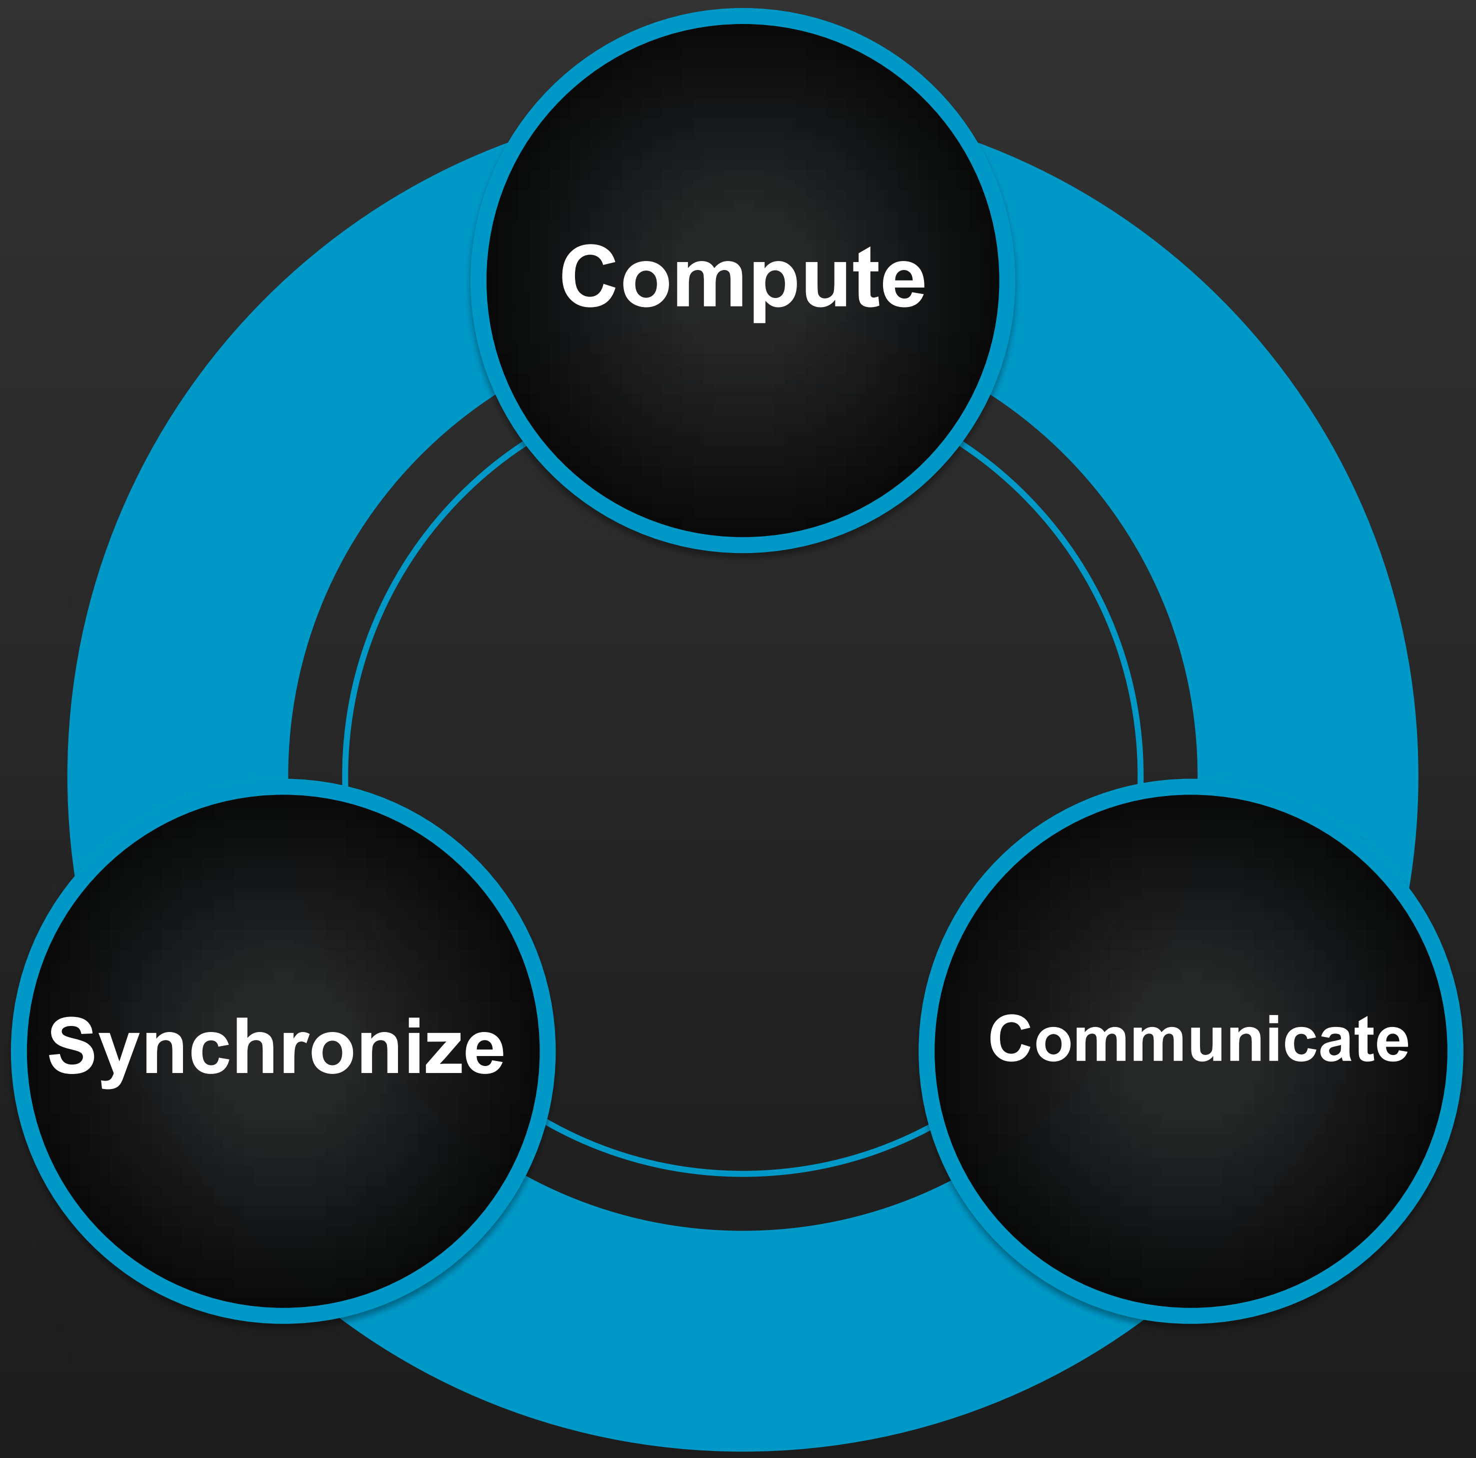

Broadcom’s AI Network Solutions

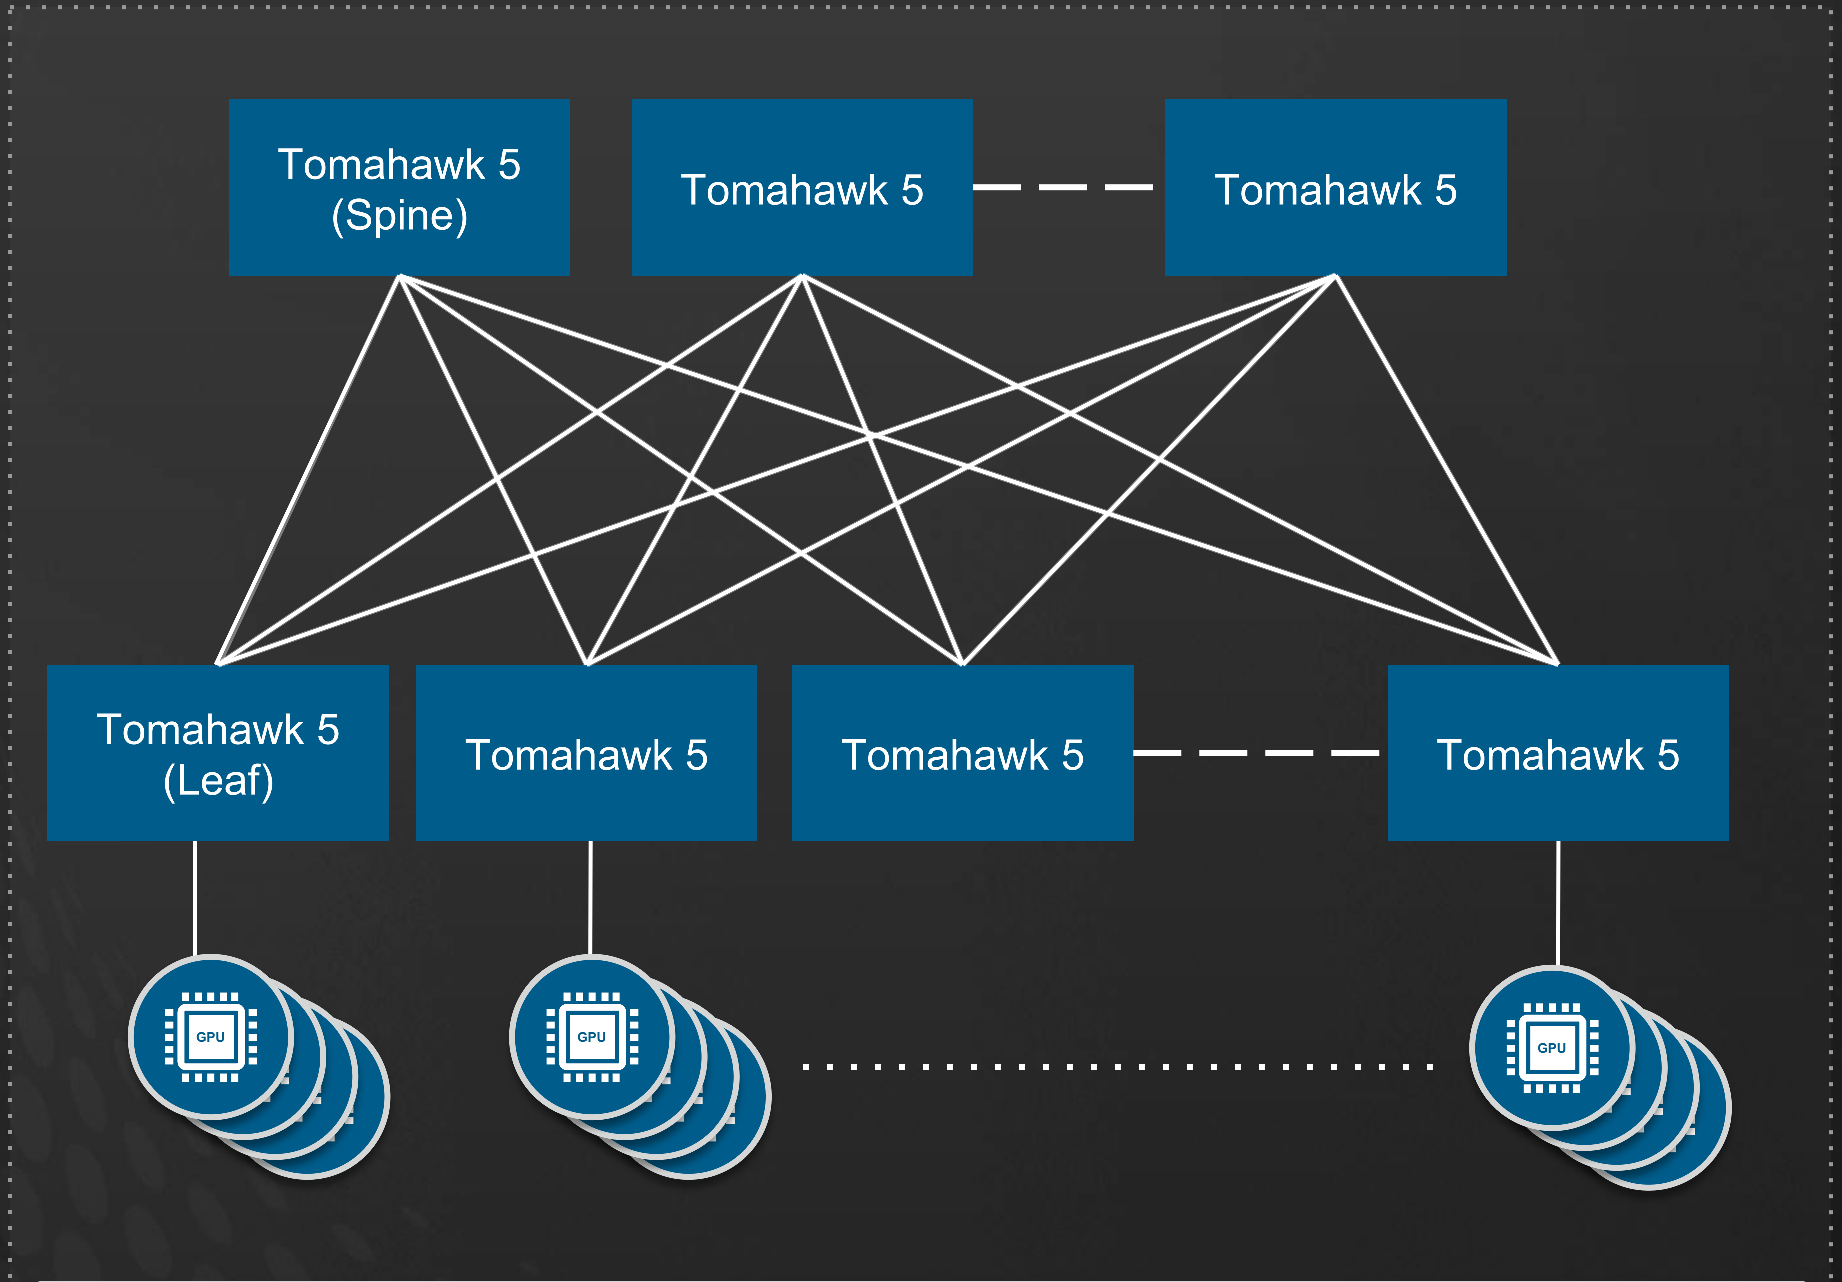

Tomahawk 5 (Spine)

Tomahawk 5

Tomahawk 5

Tomahawk 5 (Leaf)

Tomahawk 5

Tomahawk 5

Tomahawk 5

Endpoint Scheduled

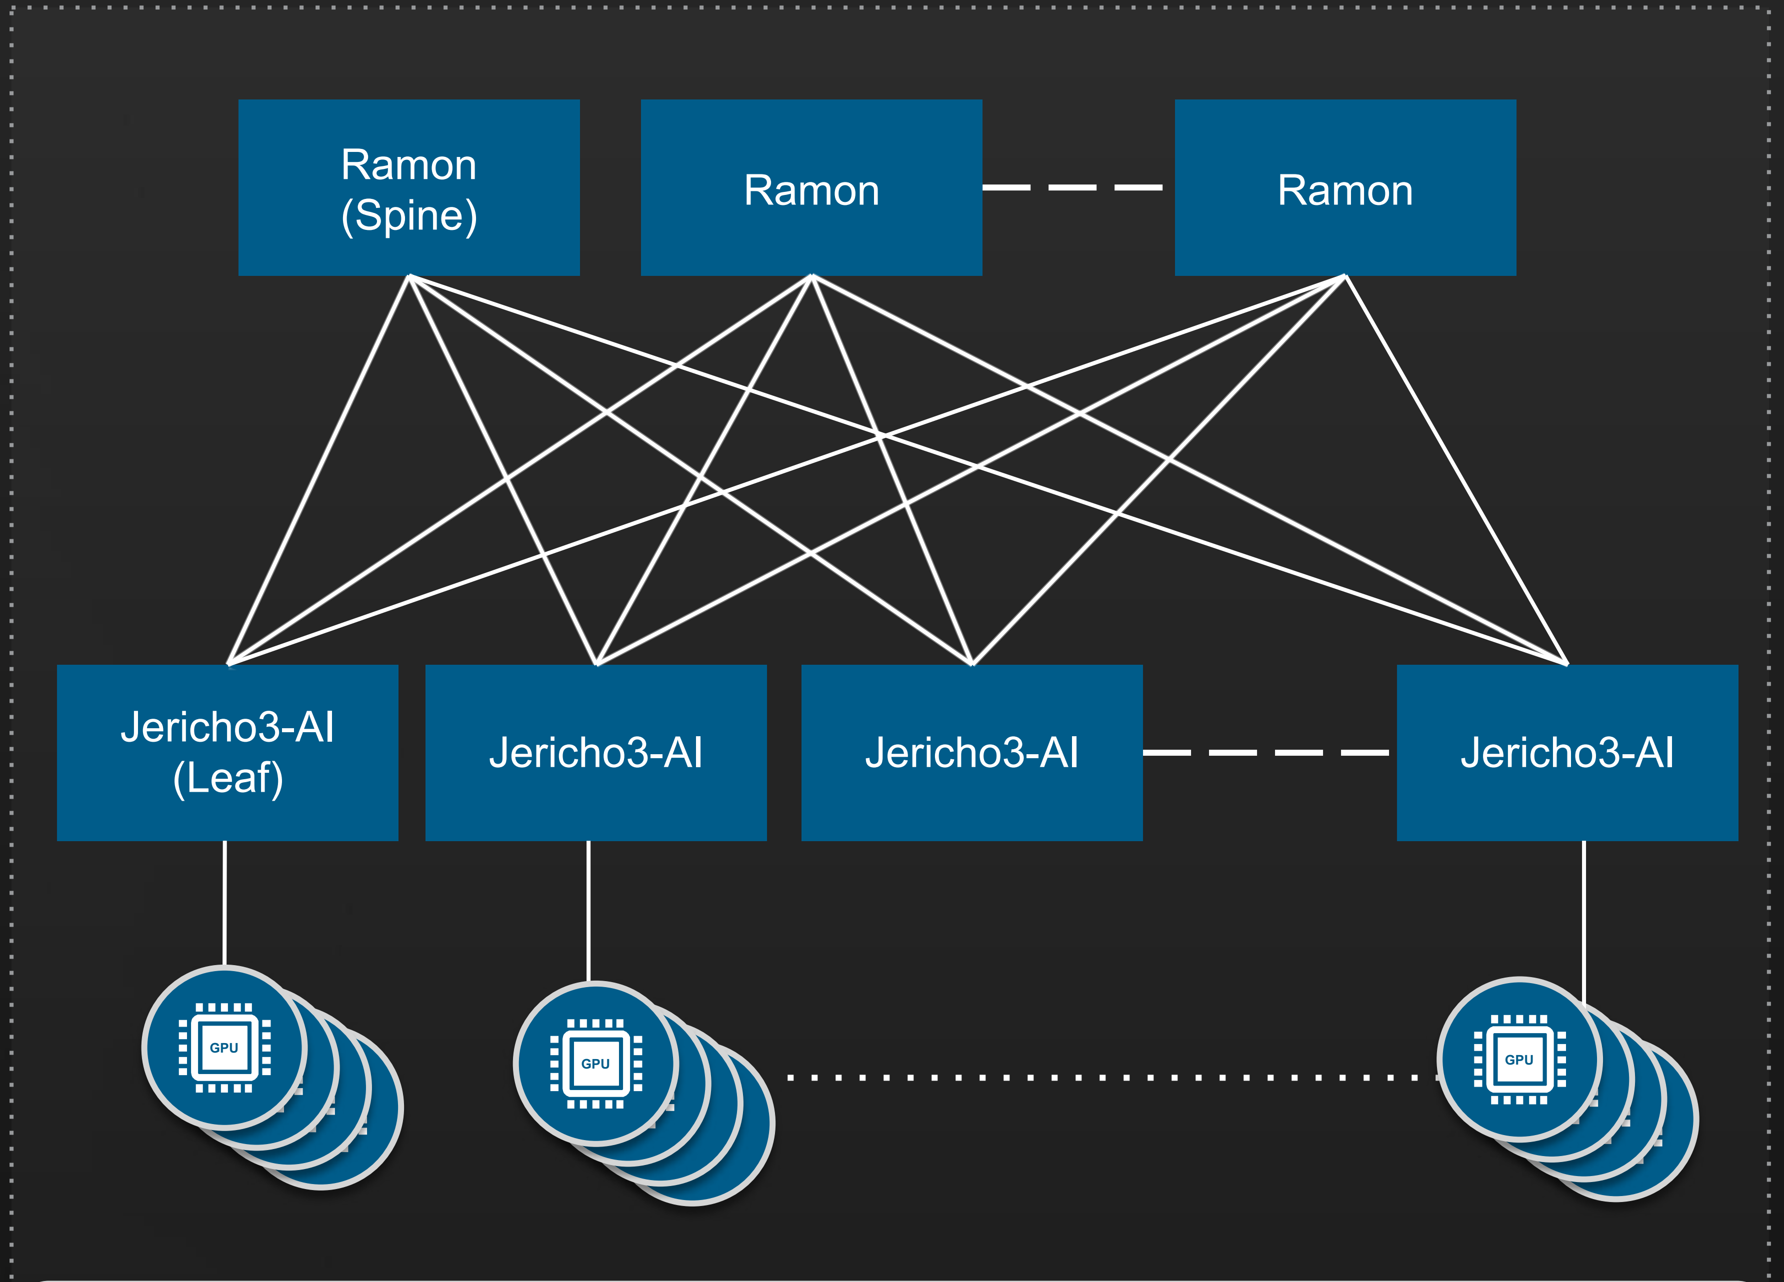

Ramon (Spine)

Ramon

Ramon

Jericho3-AI (Leaf)

Jericho3-AI

Jericho3-AI

Jericho3-AI

Switch Scheduled

Jericho3-AI Fabric: Switch Scheduled Ethernet Network

10% Performance improvement = network more than pays for itself

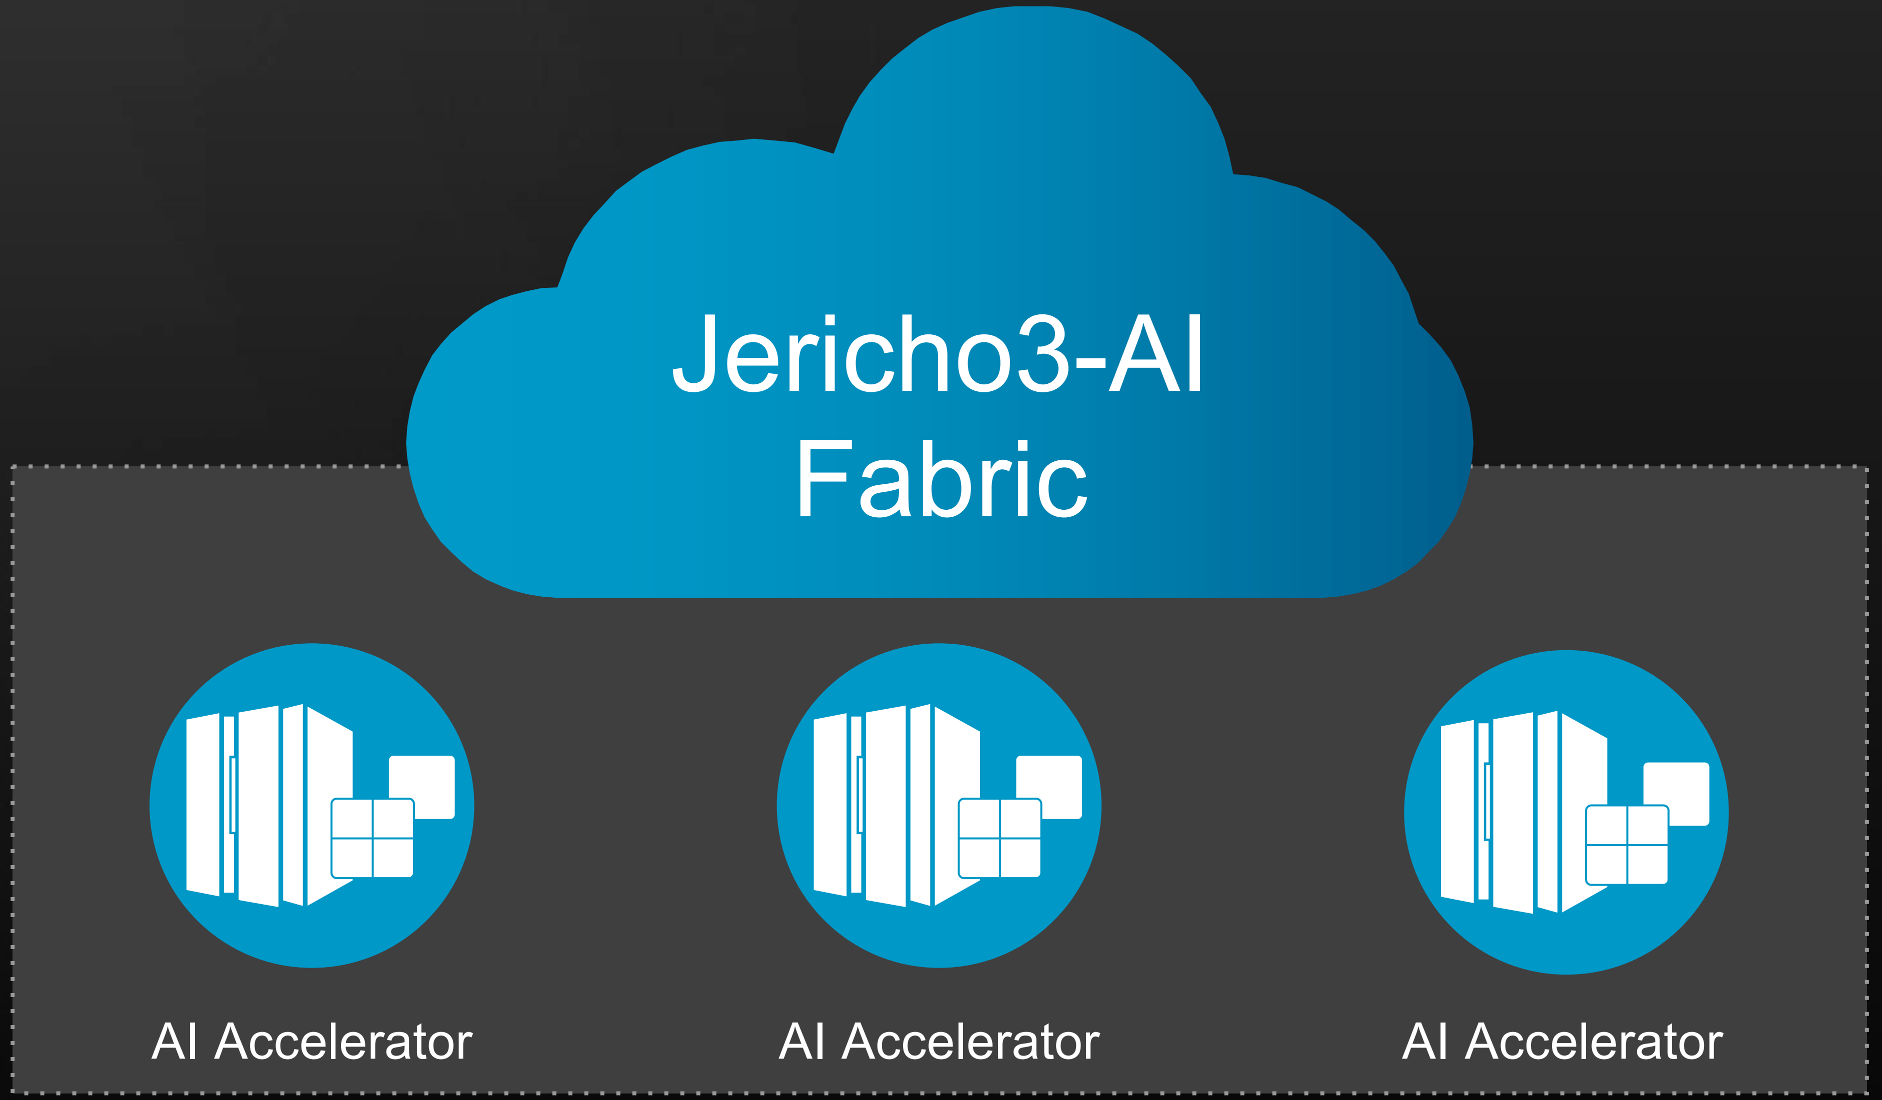

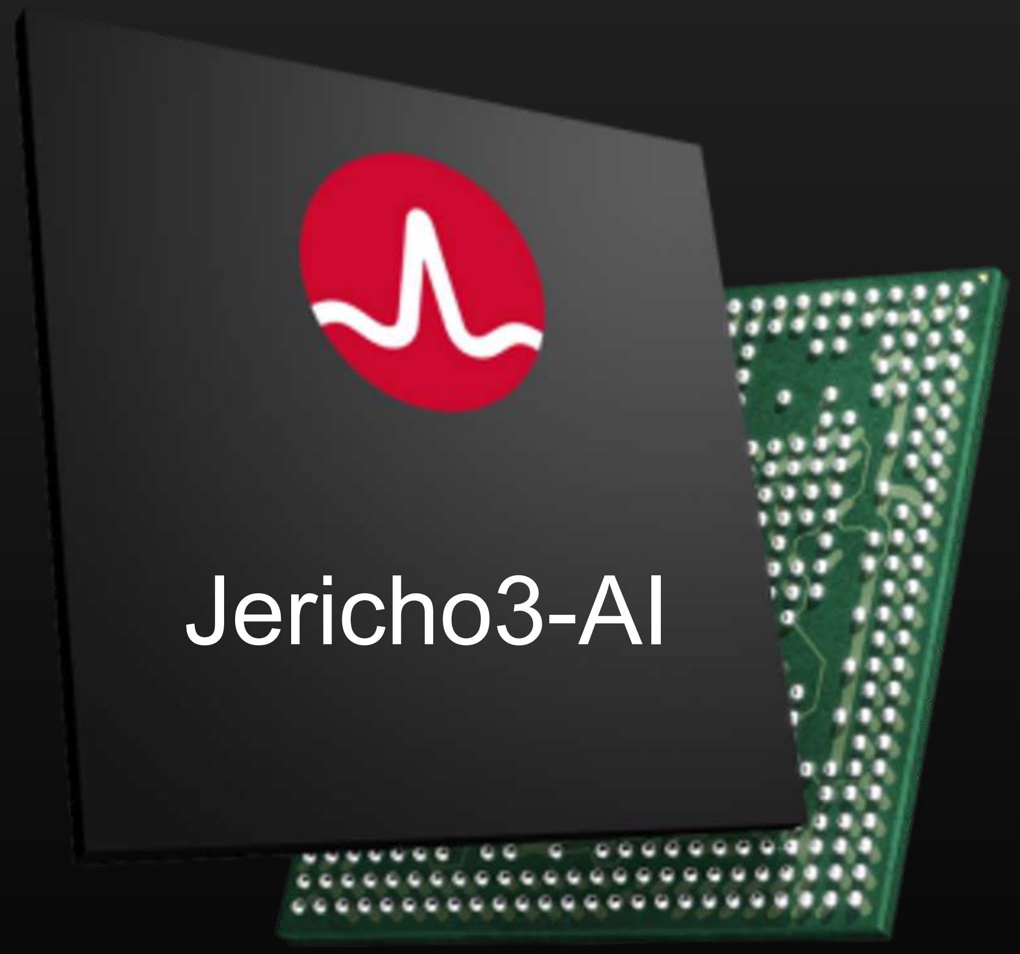

32,000 AI Accelerators** at **800Gbps each

Lowest time spent in networking

OPEN // SCALABLE // POWER EFFICIENT

Tomahawk: Endpoint Scheduled Ethernet Network

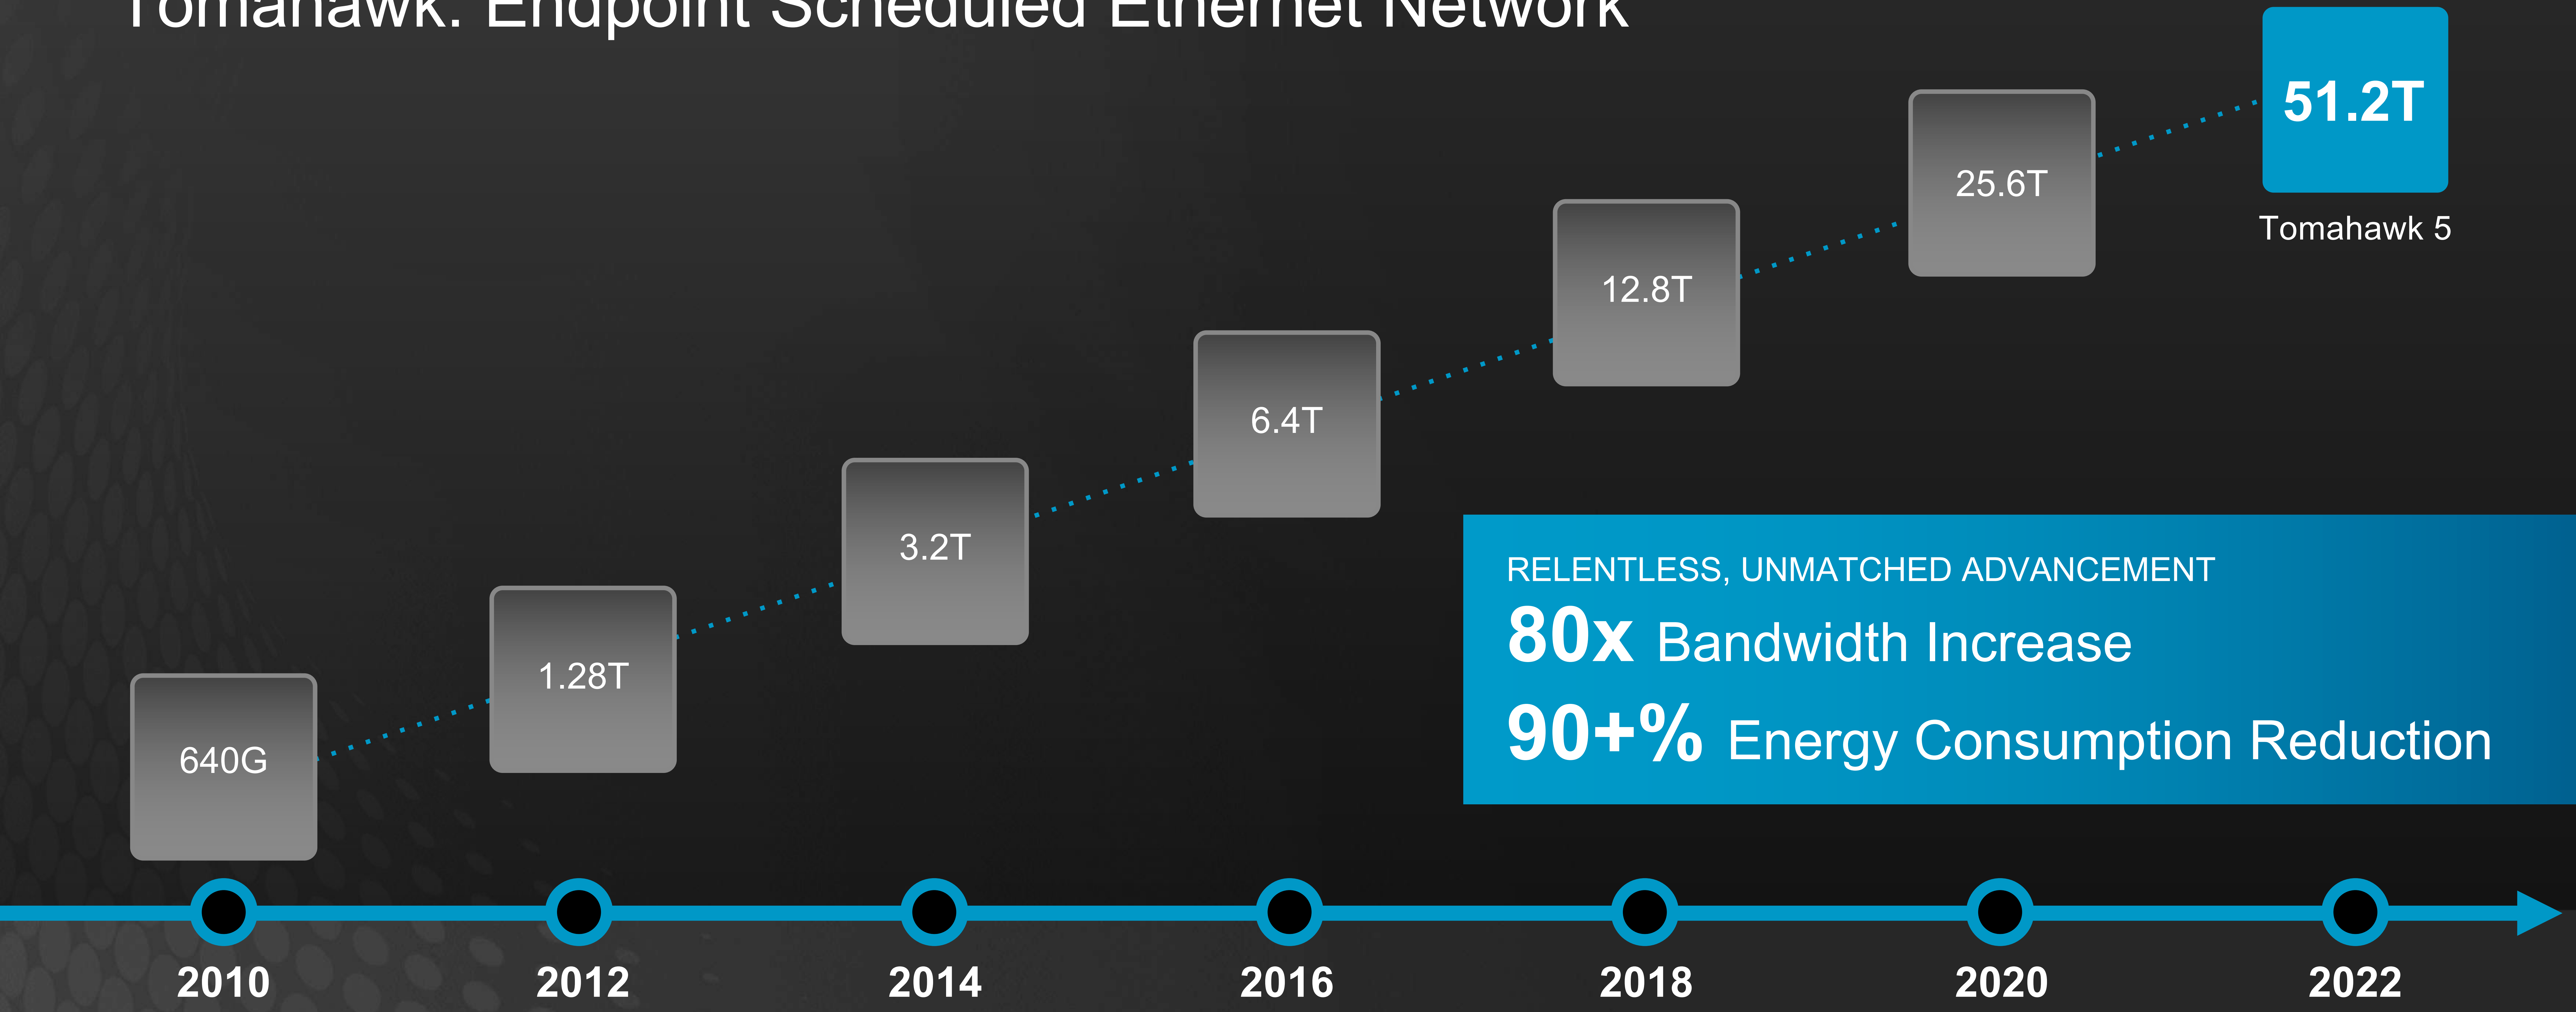

2010 — 640G

2012 — 1.28T

2014 — 3.2T

2016 — 6.4T

2018 — 12.8T

2020 — 25.6T

2022 — **51.2T** Tomahawk 5

RELENTLESS, UNMATCHED ADVANCEMENT

80x Bandwidth Increase

90+% Energy Consumption Reduction

* G=Gbps, T=Tbps

OPEN // SCALABLE // POWER EFFICIENT

Introducing Thor2: AI Optimized NIC

Thor2

- 400G high-performance NIC
- High-scale RDMA
- Industry’s lowest power
- Longest reach 100G Serdes

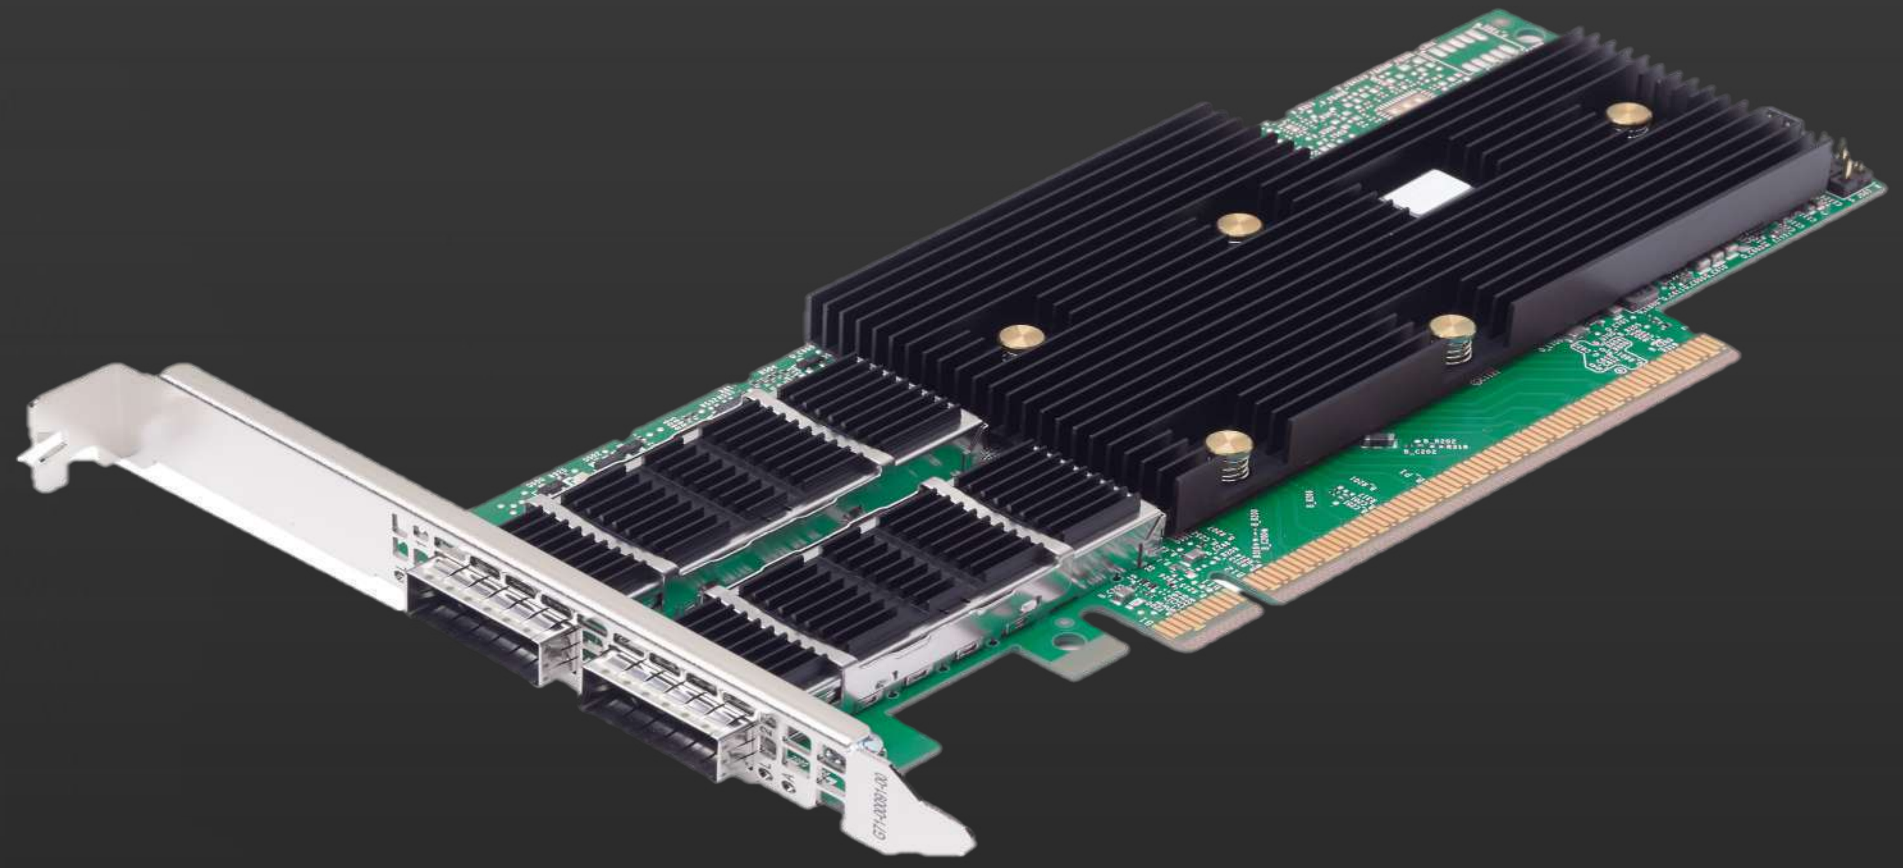

OPEN // SCALABLE // POWER EFFICIENT

Thor2 Consumption Models

Board

Chiplet

IP

End-to-End High Performance Ethernet AI Network

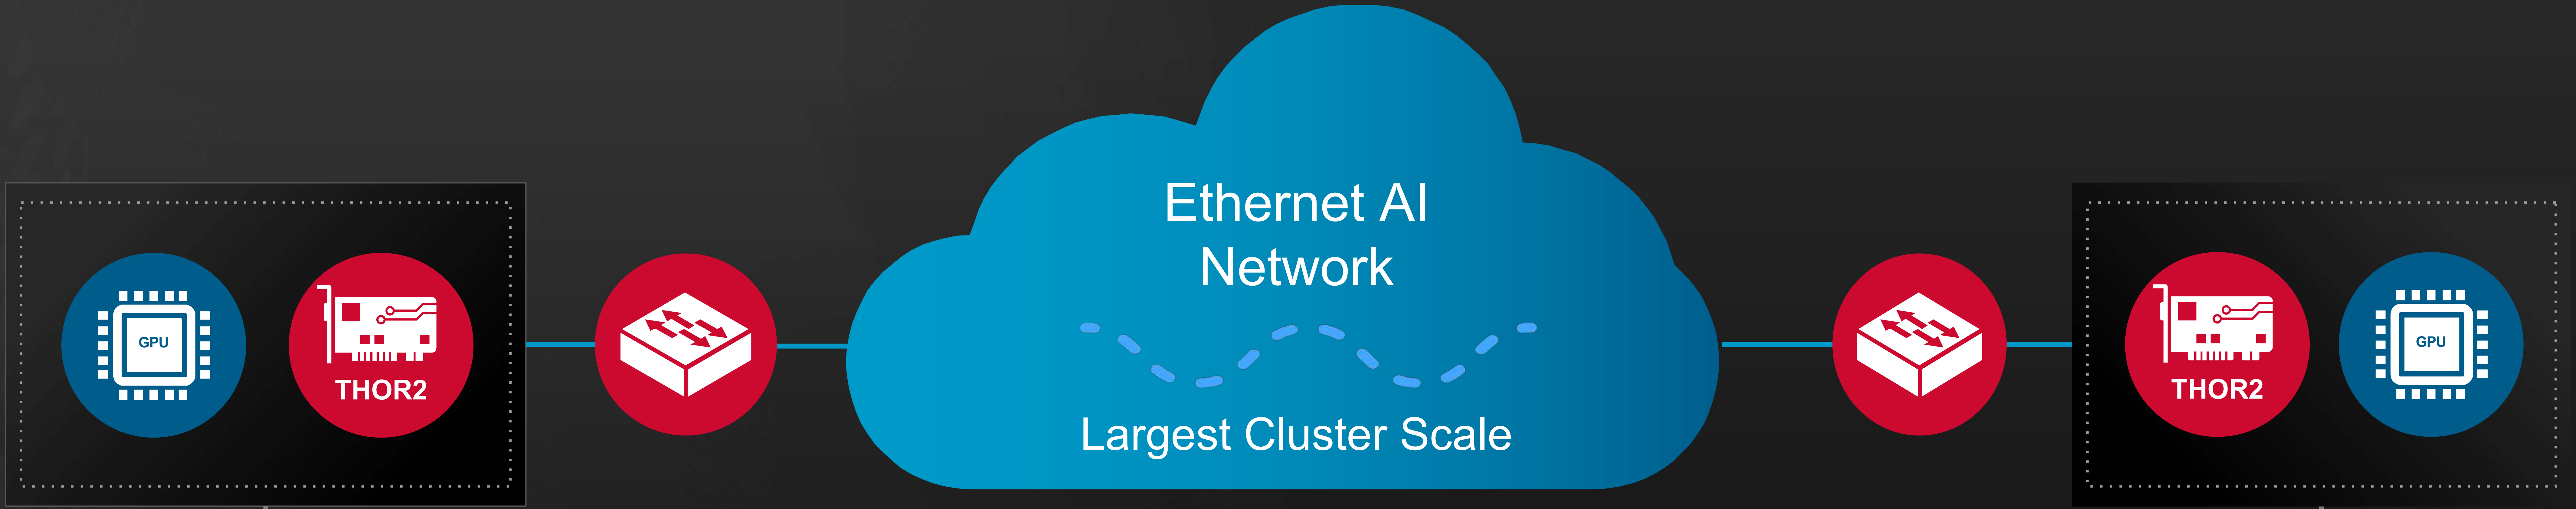

Perfectly load balanced fabric

Programmable E2E Congestion Control

Zero Impact Failover (ZIF)

Secure & Multi-tenant fabric

Efficient Interconnects

OPEN // SCALABLE // POWER EFFICIENT

Ethernet Beats InfiniBand: 10+% Improvement in Job Completion Time

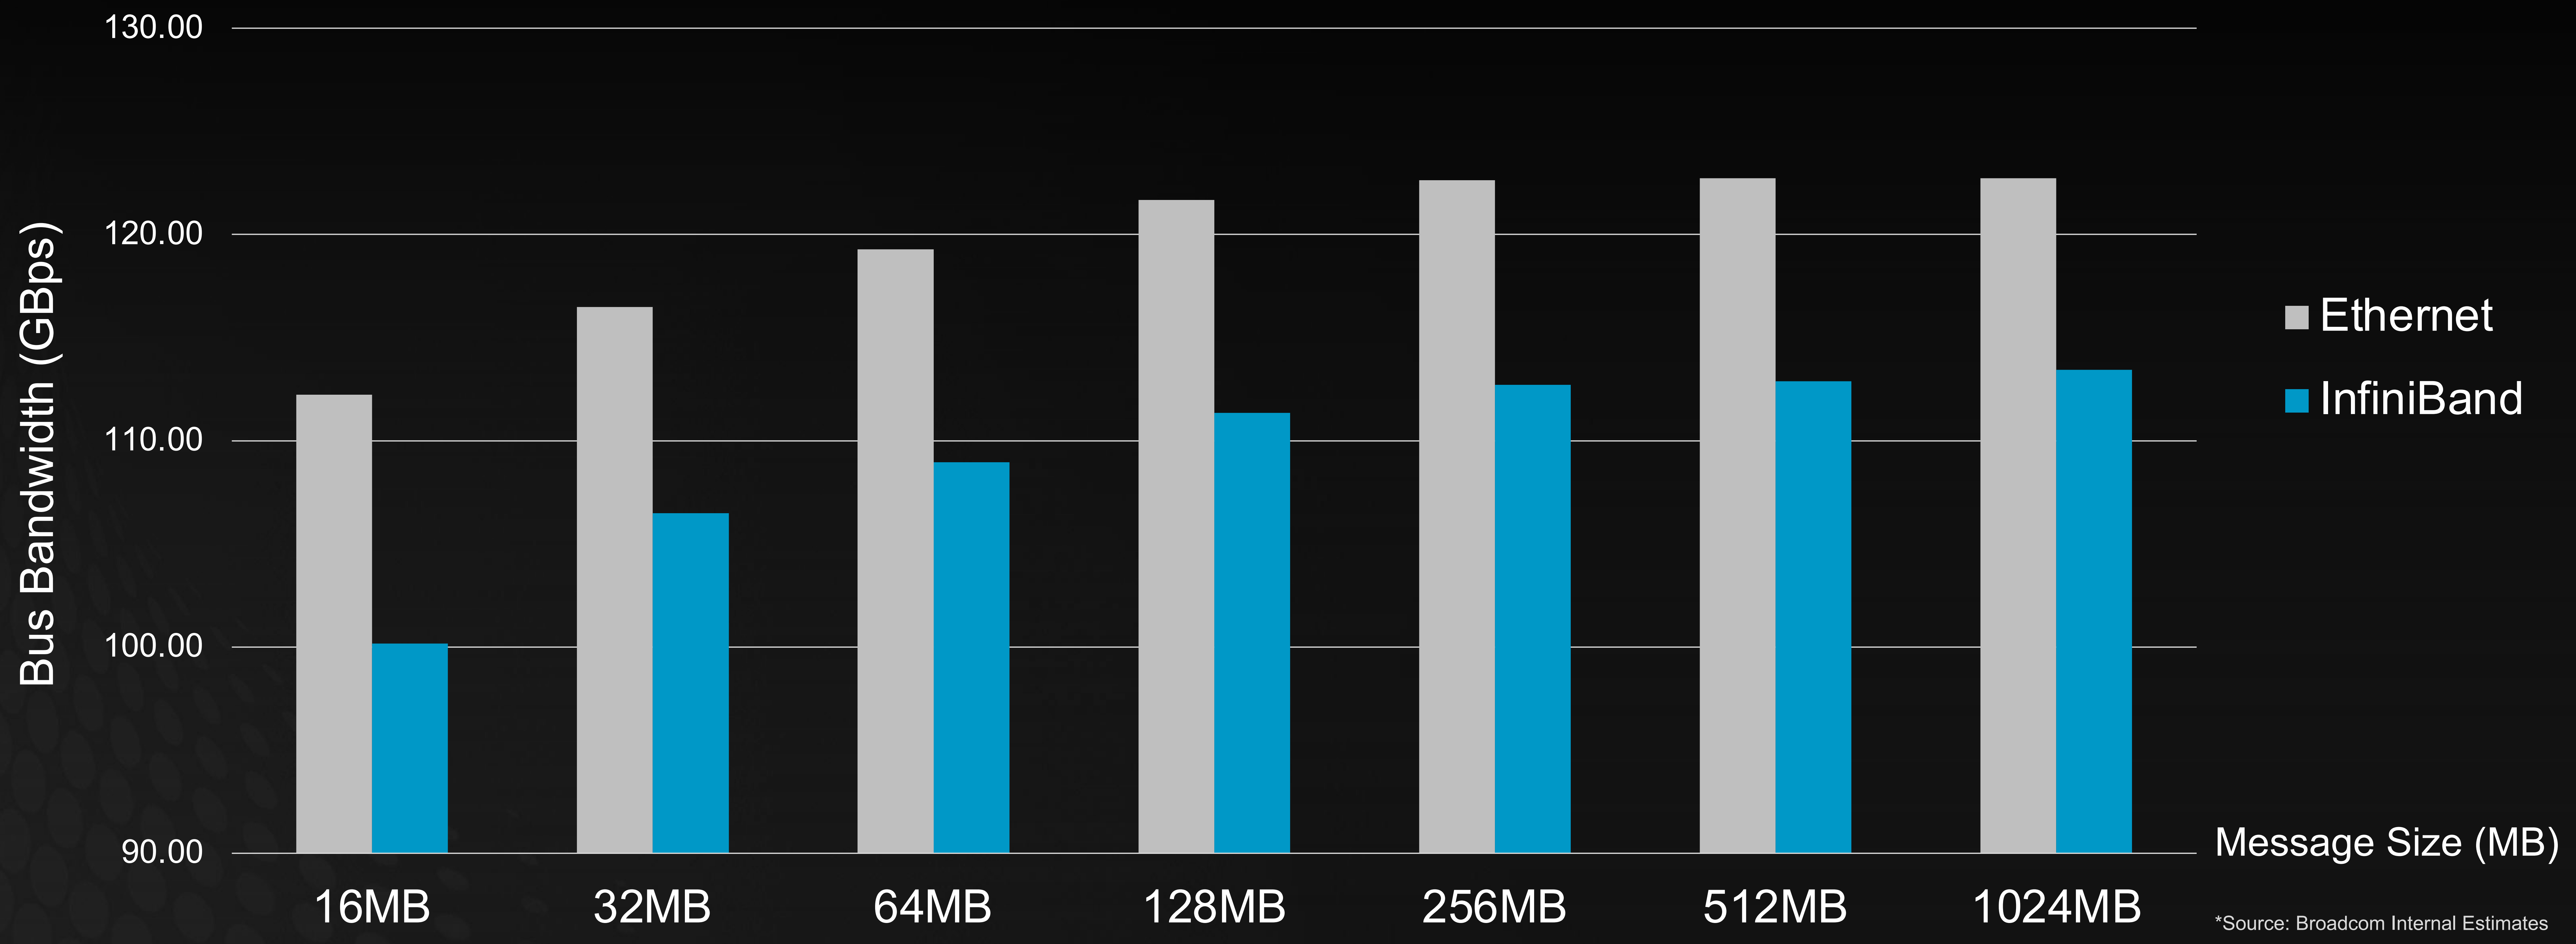

*Source: Broadcom Internal Estimates

OPEN // SCALABLE // POWER EFFICIENT

Ethernet Provides 30x Faster Failover than InfiniBand

2%

Optics Annual Failover Rate*

15

Failures per Month**

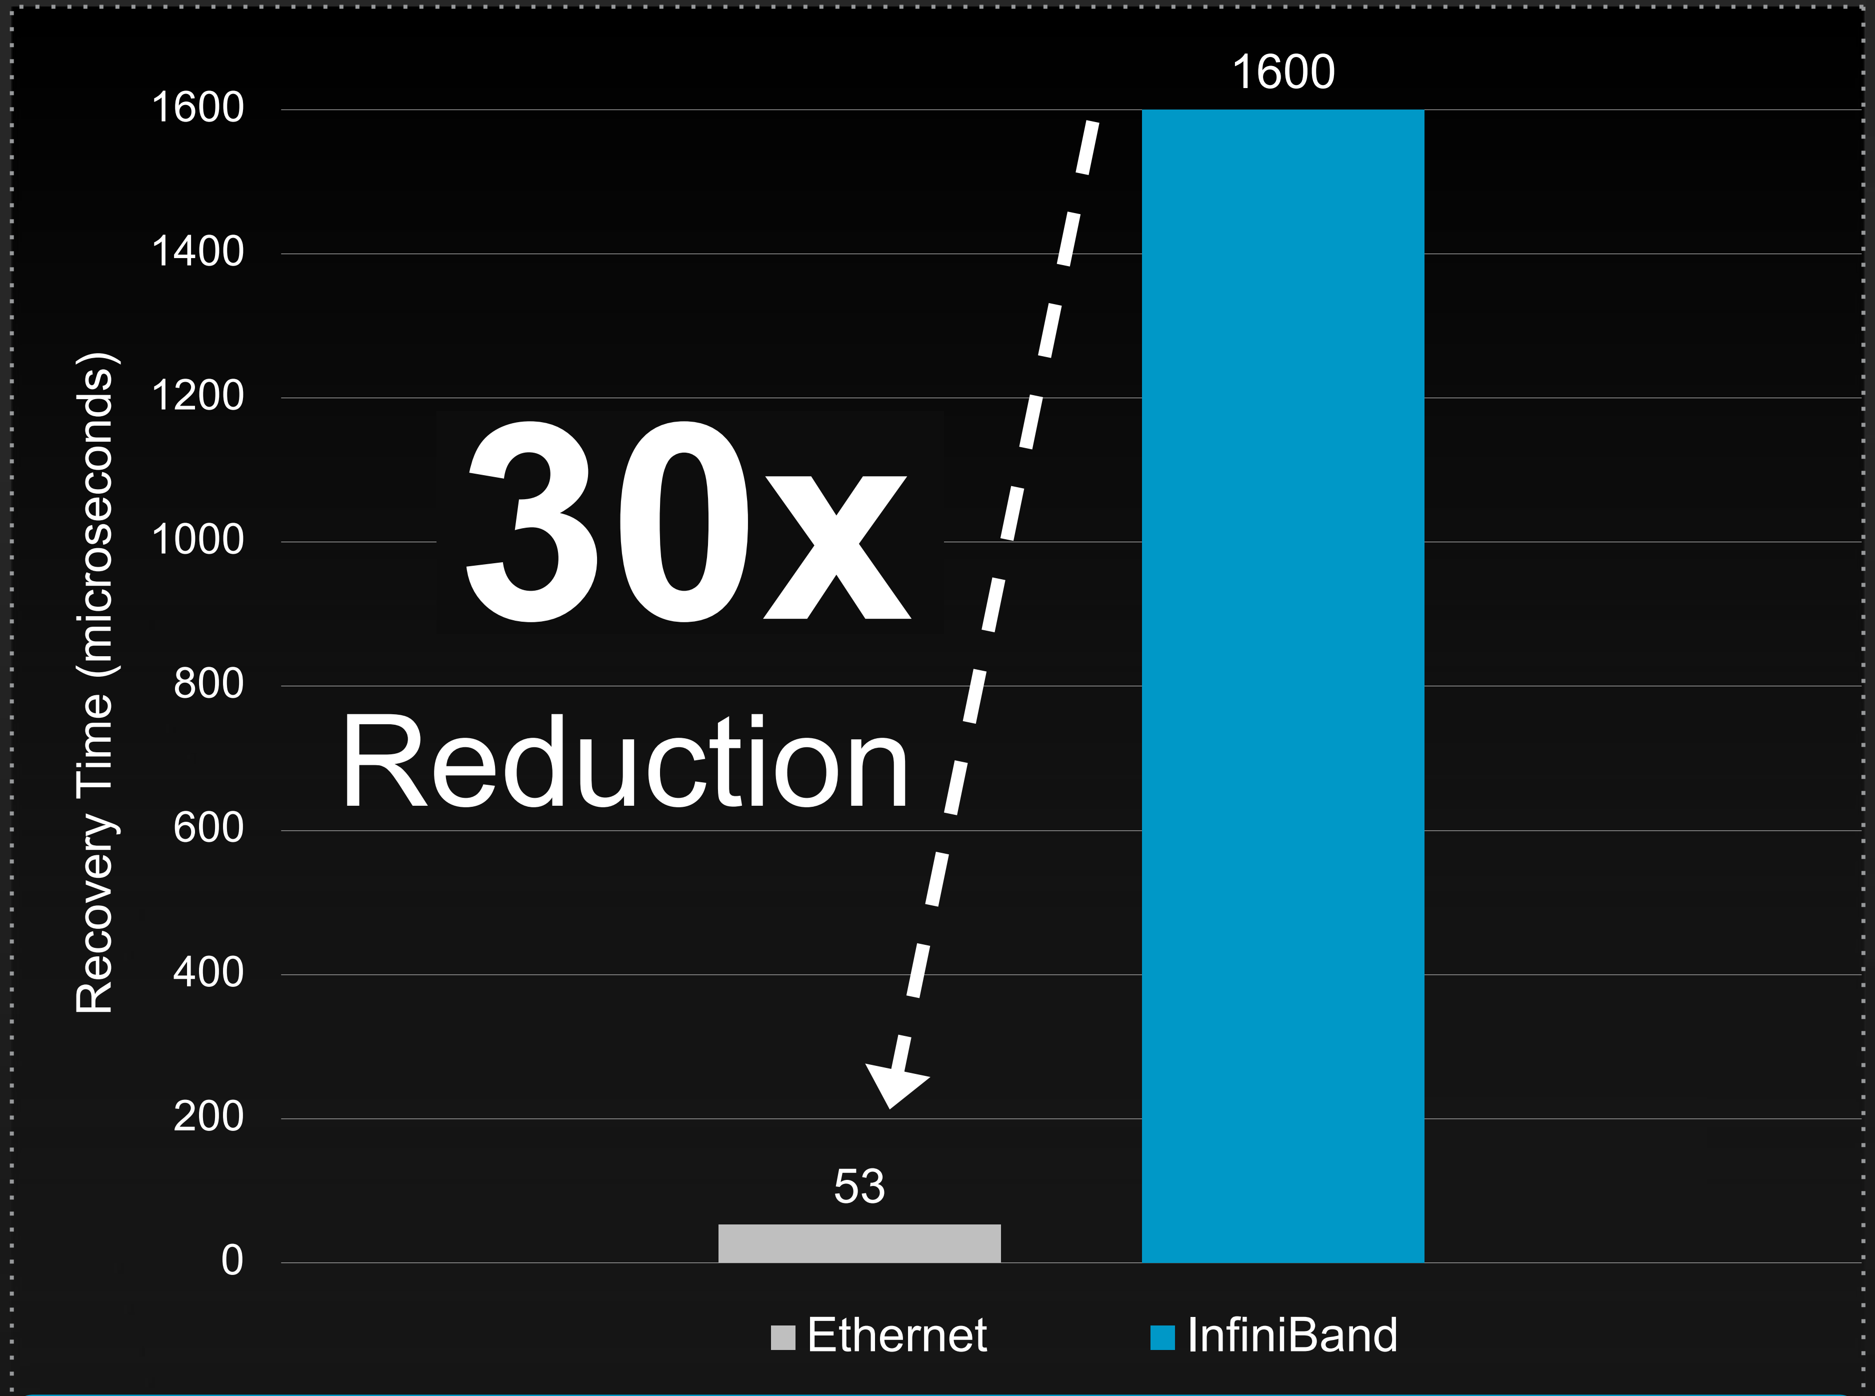

Fast recovery reduces job completion time

Source: Broadcom Internal Estimates

* Typical industry failure rate. ** Assuming 4K node cluster using 9.2K optic modules

OPEN // SCALABLE // POWER EFFICIENT

Reducing AI Interconnect Cost and Power

Extended Reach for Copper Cables

4m+ DAC (2x IEEE spec)

Linear Pluggable Optics

33% Lower Power

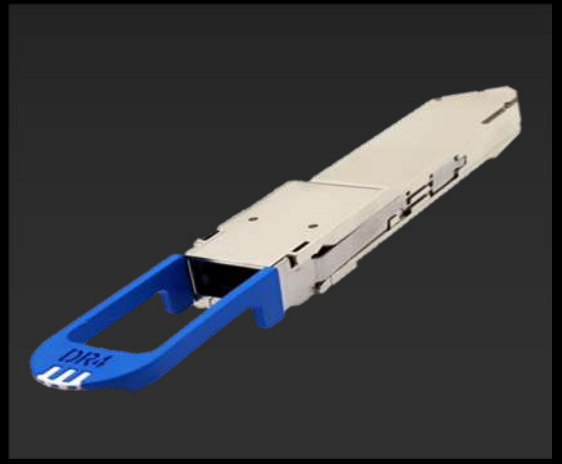

Co-Packaged Optics

Lowest Power & Cost

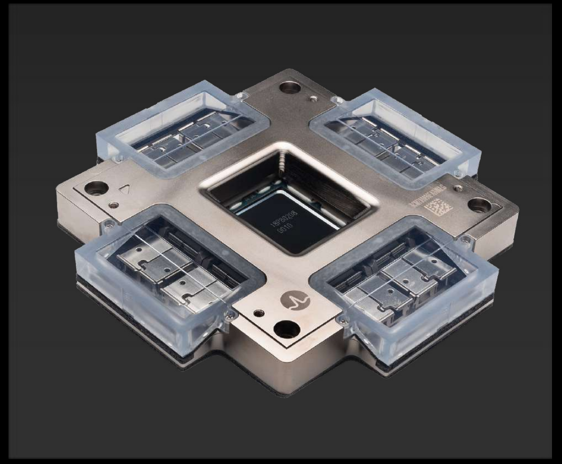

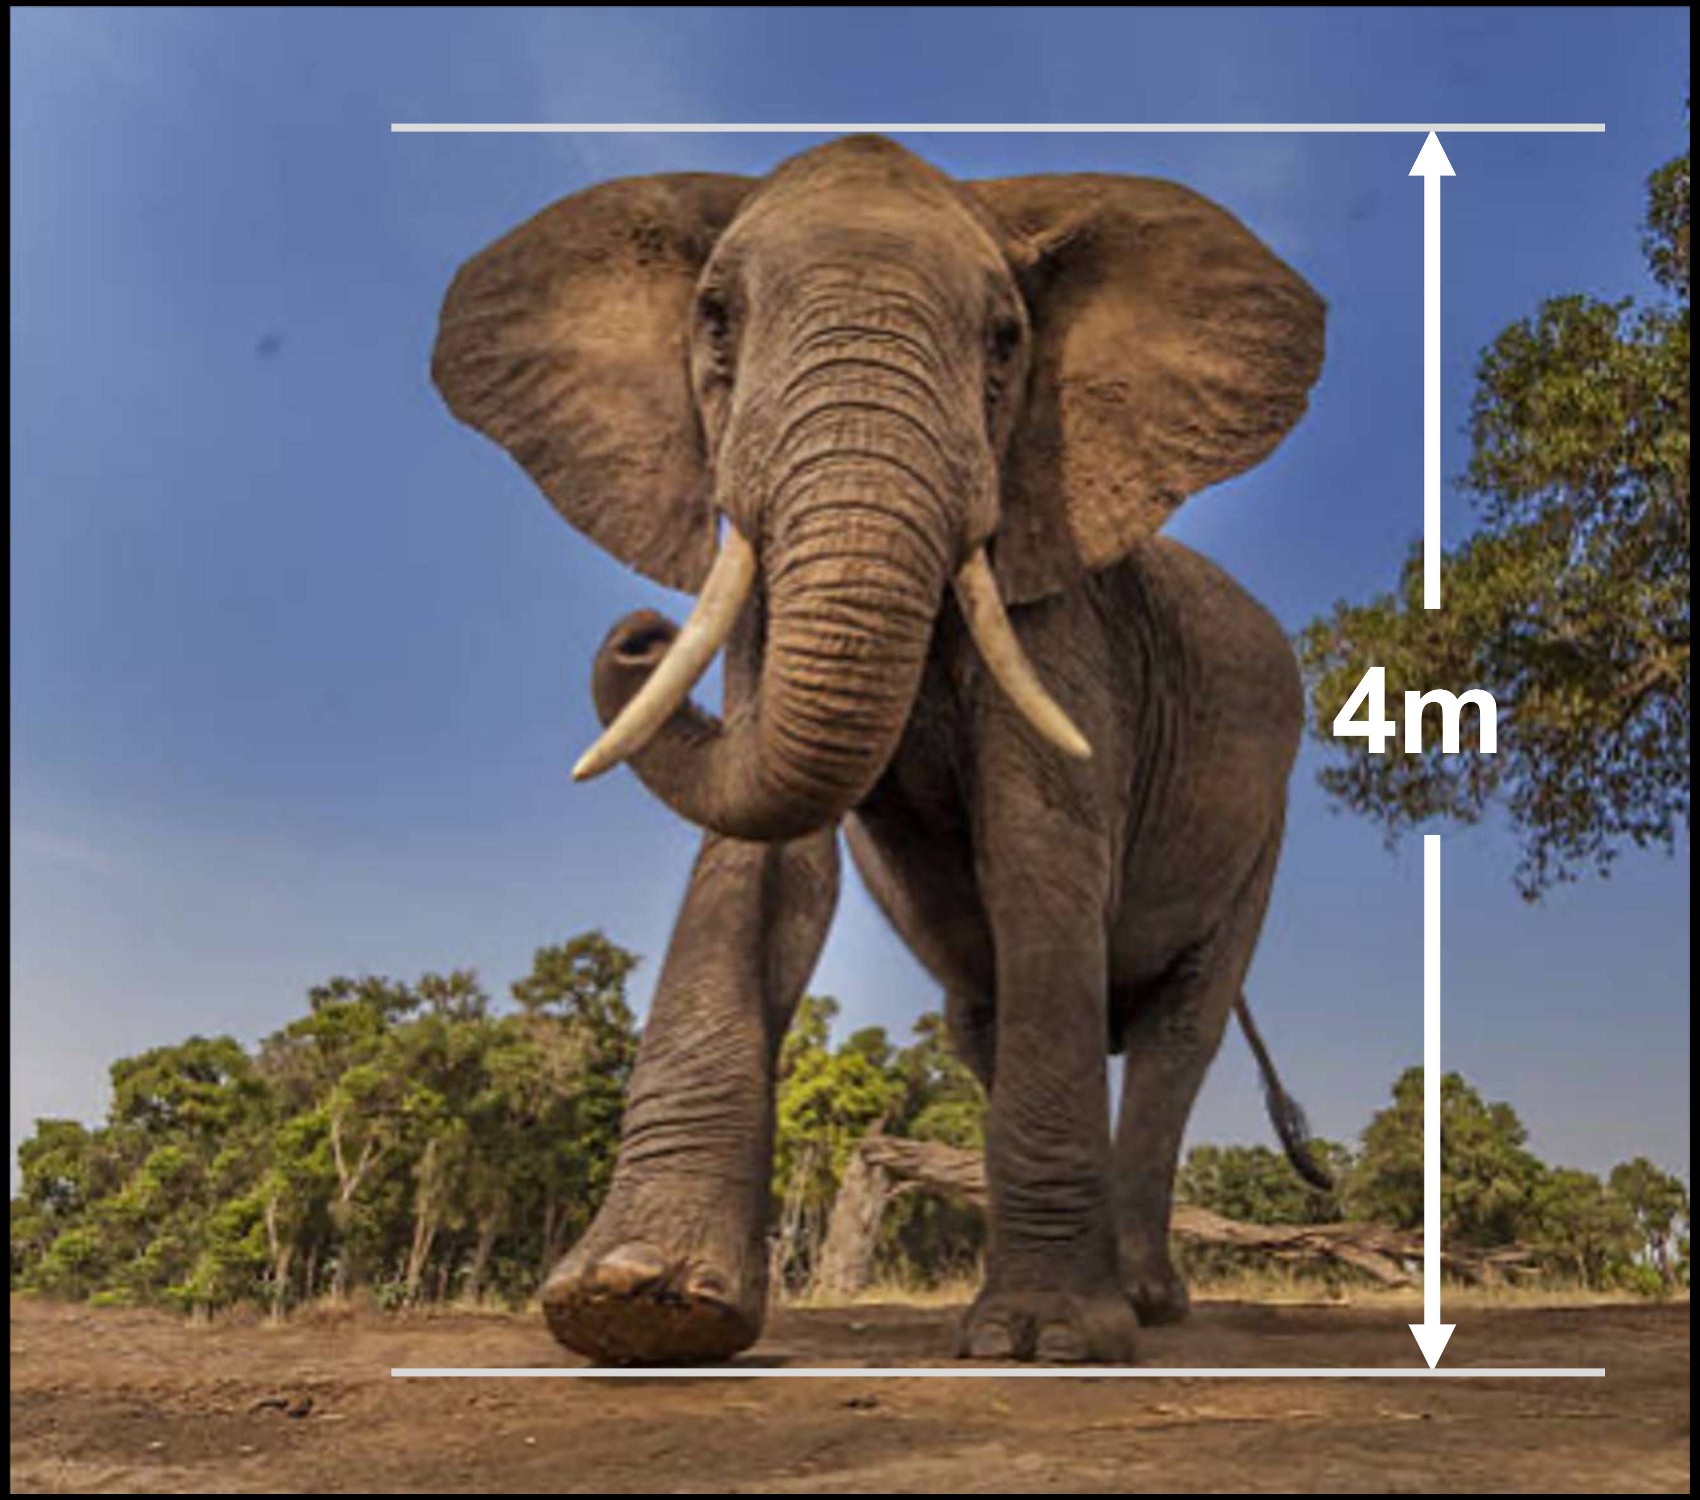

OPEN // SCALABLE // POWER EFFICIENT

Ethernet is the De-facto AI Network

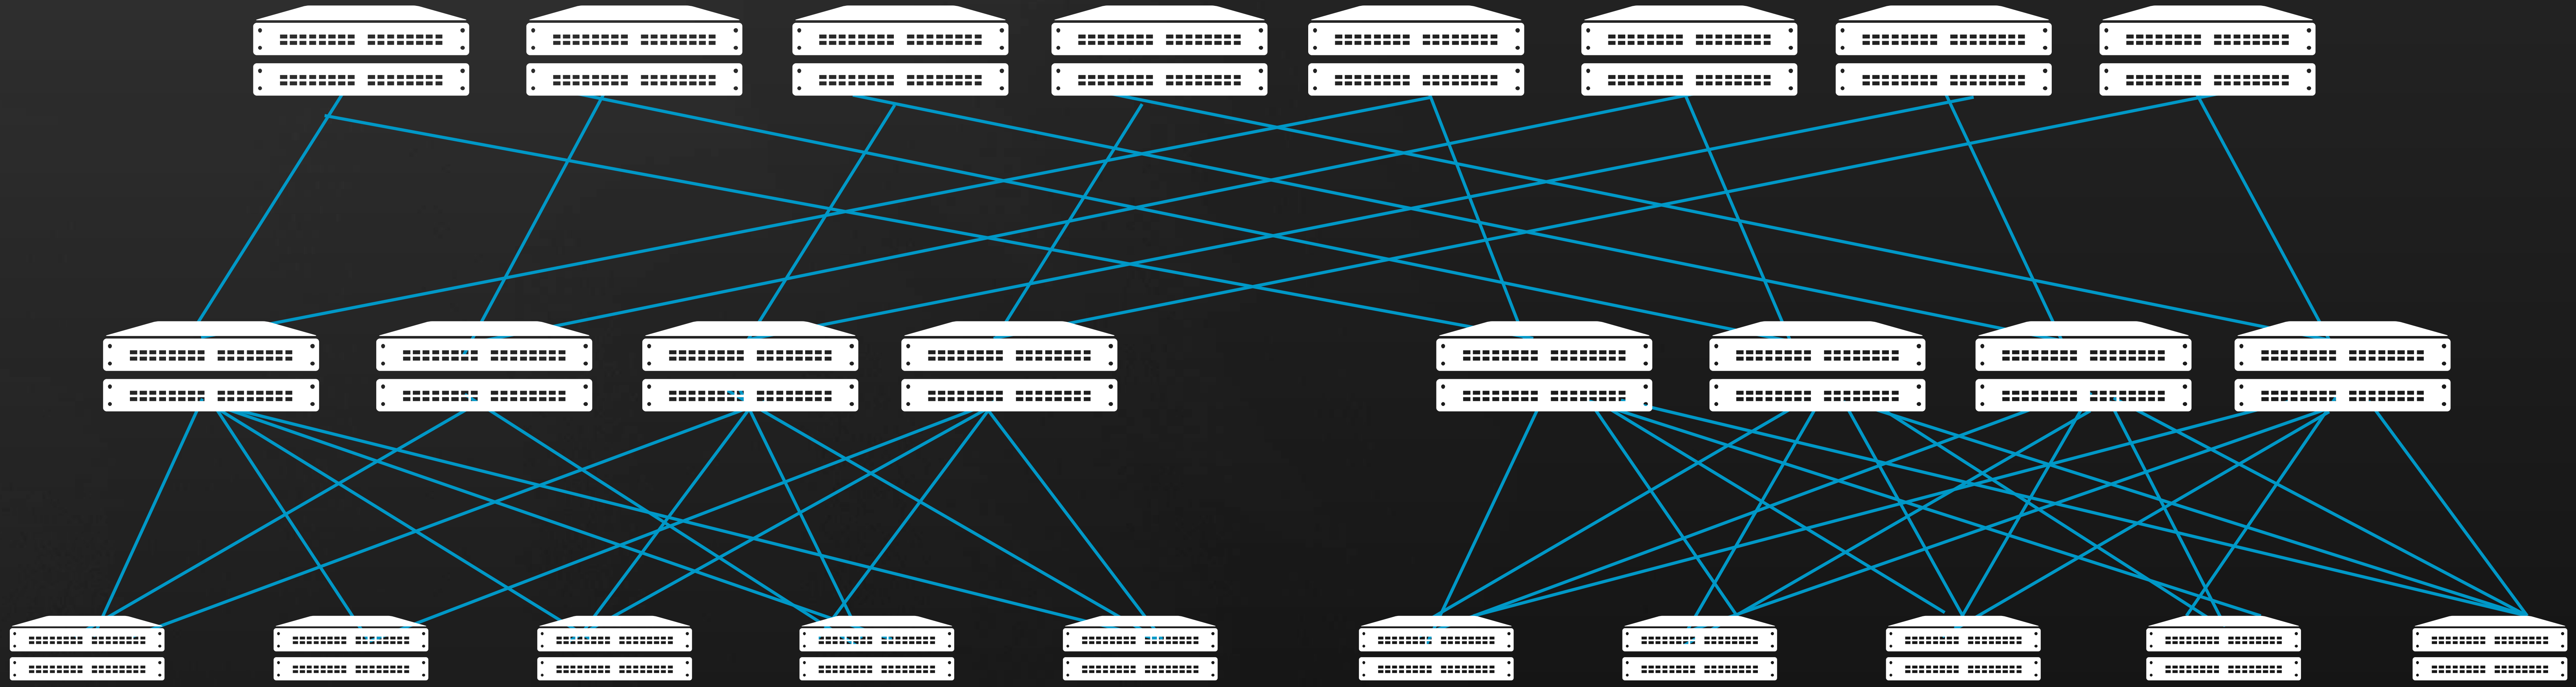

Alibaba Group 阿里巴巴集团

amazon

ByteDance

Google

Meta

ORACLE

Tencent 腾讯

OPEN // SCALABLE // POWER EFFICIENT

Large Ethernet AI Clusters

amazon	60,000+
ORACLE	30,000+
Meta	20,000+
ByteDance	10,000+

Ultra Ethernet: AI at Scale

AMD
ARISTA
BROADCOM

CISCO
EVIDEN
an atos business
Hewlett Packard
Enterprise
intel

Meta
Microsoft
ORACLE

Incredibly Strong Industry Reception: 55+ companies

Modernizing RDMA…

Classic RDMA

- In-order packet delivery
- Go-back-n → inefficient
- No multipathing
- DCQCN → hard to tune

UltraEthernet

- Out-of-order placement, in-order message completion
- Selective Ack and retransmit
- Packet-level multipathing
- Configuration-free congestion control

Ethernet Network for AI

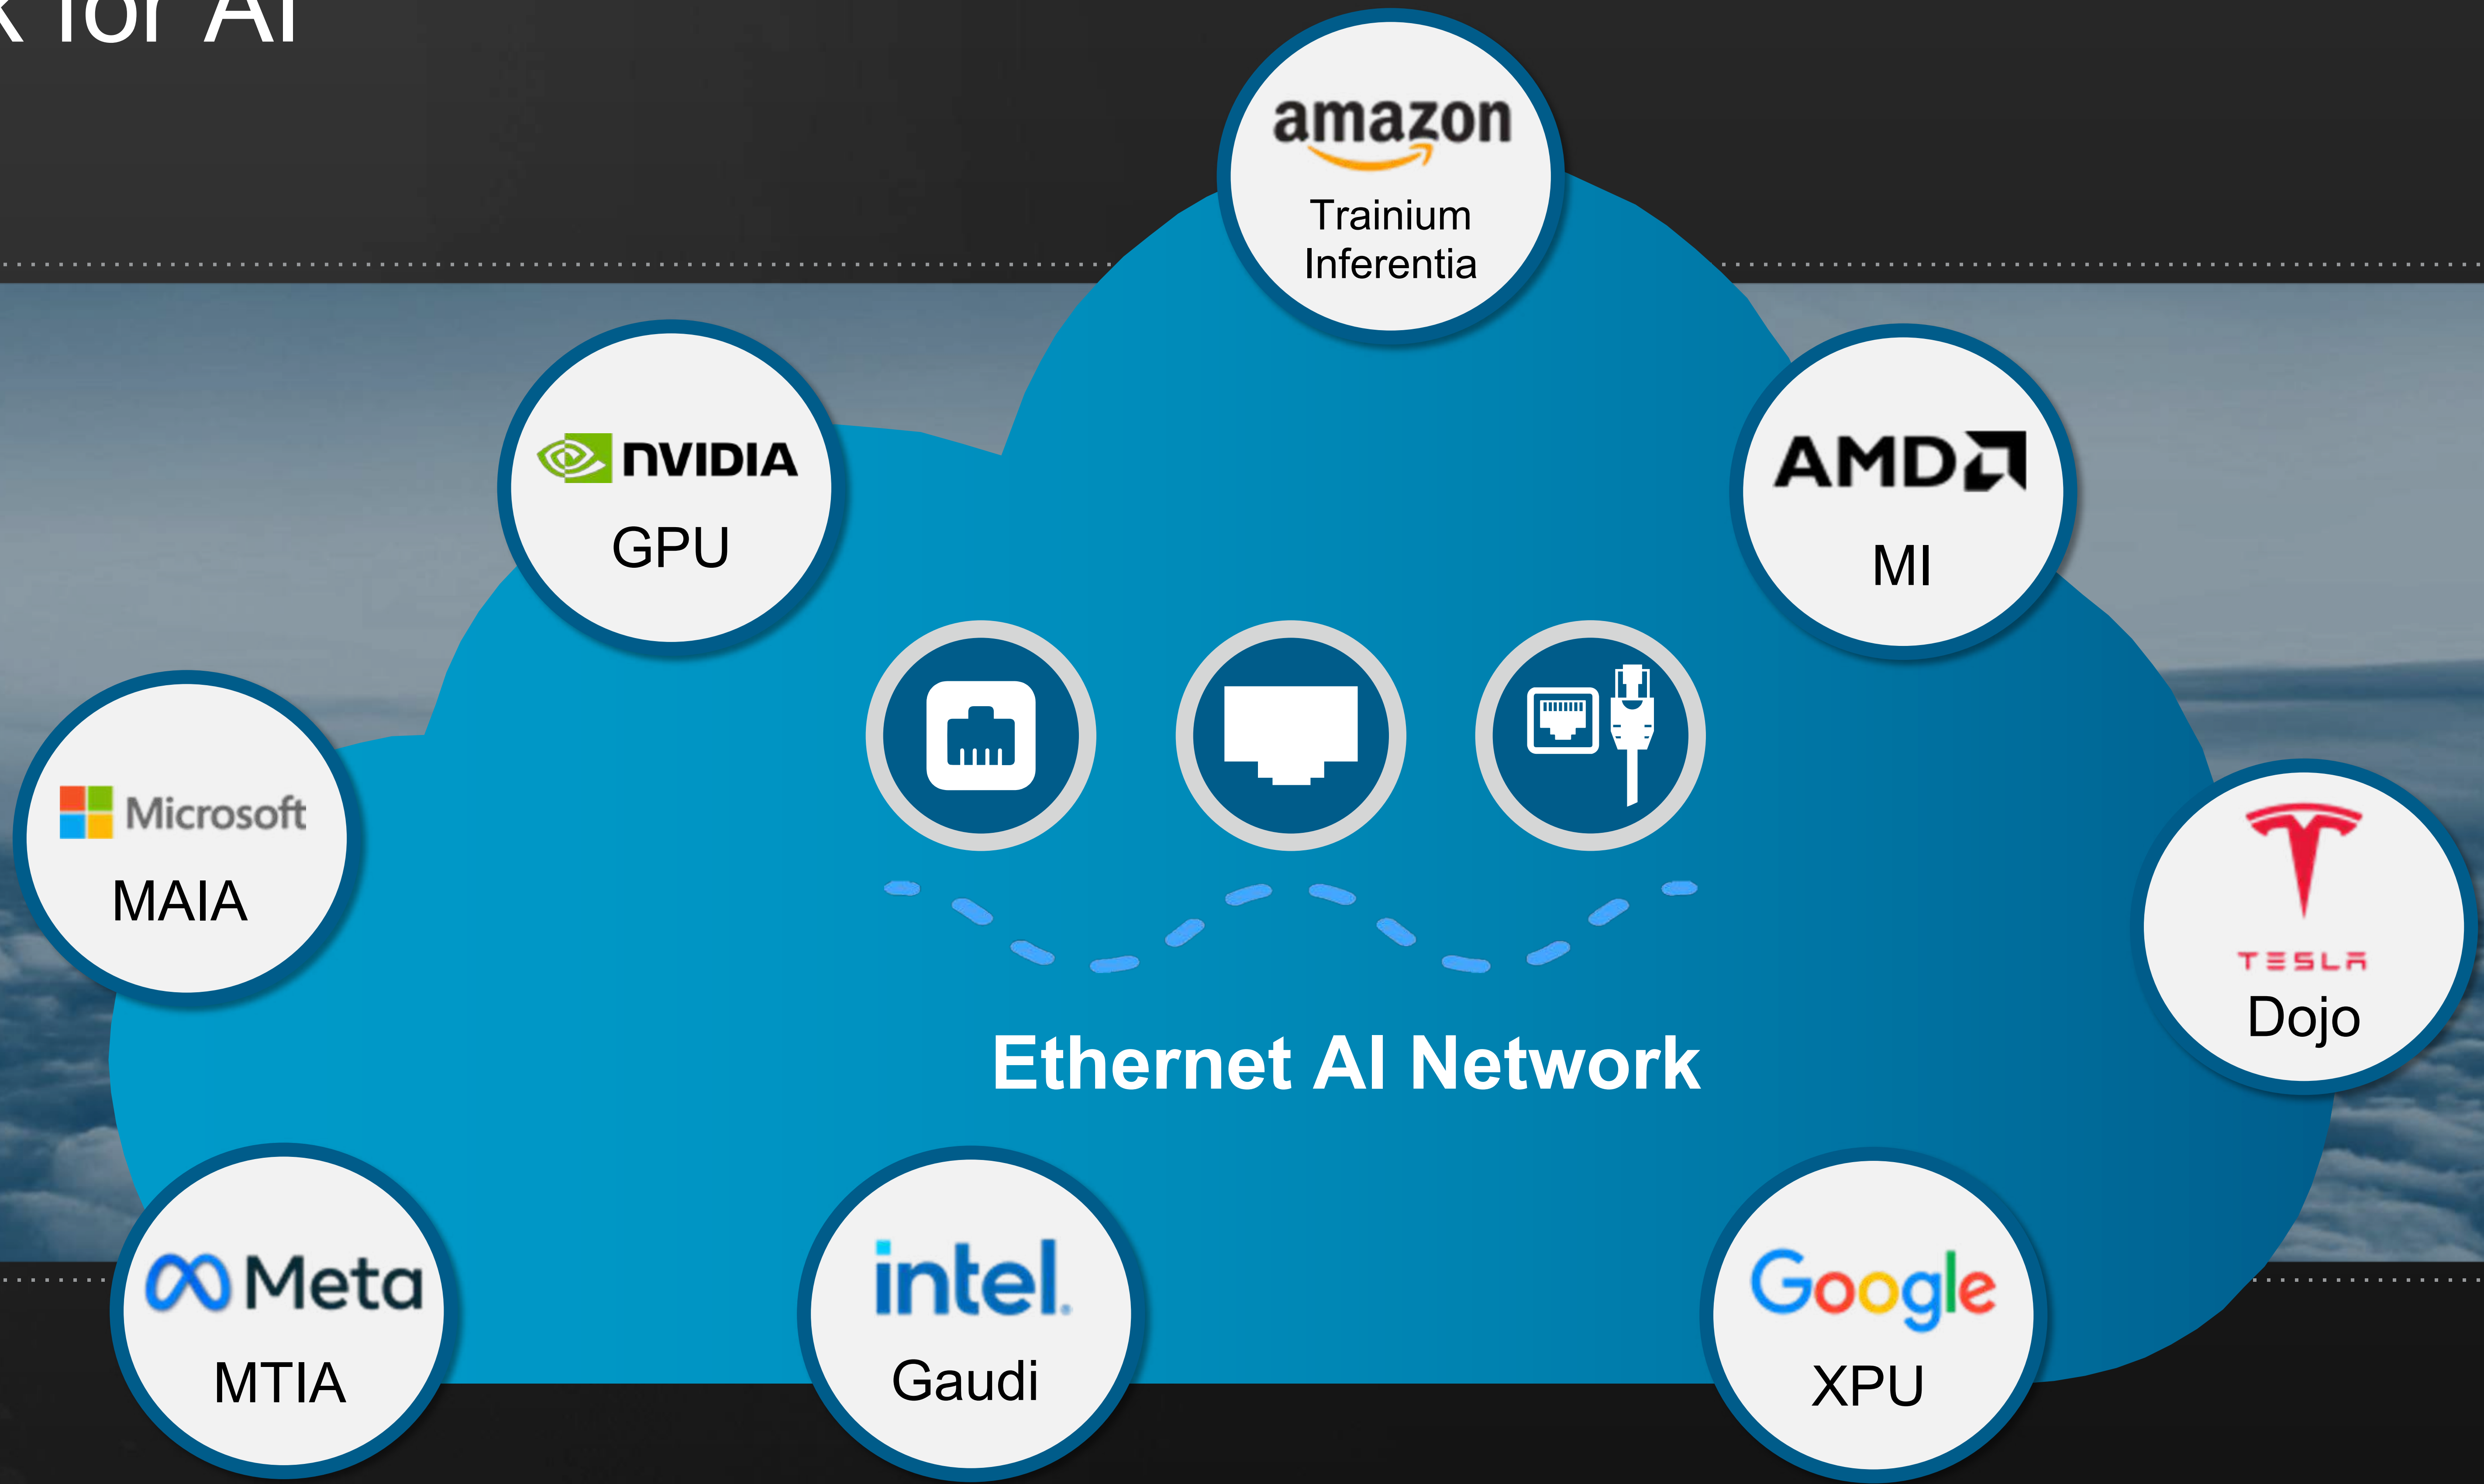

Ethernet AI Network

Comprehensive Ethernet Portfolio and Ecosystem

Network Control & Automation

- CHEF
- apstra
- ANSIBLE
- lagopus
- NEC
- nuagenetworks From Nokia
- OPEN DAYLIGHT
- puppet
- openstack
- vmware by Broadcom

Operation System

- ARISTA
- ARRCUS
- CISCO
- DELL
- Extreme networks
- H3C
- Hewlett Packard Enterprise
- Infoblox NEXT LEVEL NETWORKING
- ipinfusion
- JUNIPER NETWORKS
- Ruijie
- SONiC

Hardware Platform

- Accton Making Partnership Work
- ALPHA Networks
- ARISTA
- Celestica
- CISCO
- DELL
- DNI A Delta Group Company
- Hewlett Packard Enterprise
- H3C
- JUNIPER NETWORKS
- OPEN Compute Project
- QCT
- Ruijie

Tomahawk

Trident

Jericho

Thor

OPEN // SCALABLE // POWER EFFICIENT

BROADCOM®

ENABLING AI Infrastructure

AI Server Interconnect

Jas Tremblay
General Manager
Data Center Solutions Group
Broadcom

OPEN // SCALABLE // POWER EFFICIENT

Server Interconnect History

PCI Standards Group: Founded in 1992 with 900+ Companies

Server Evolution

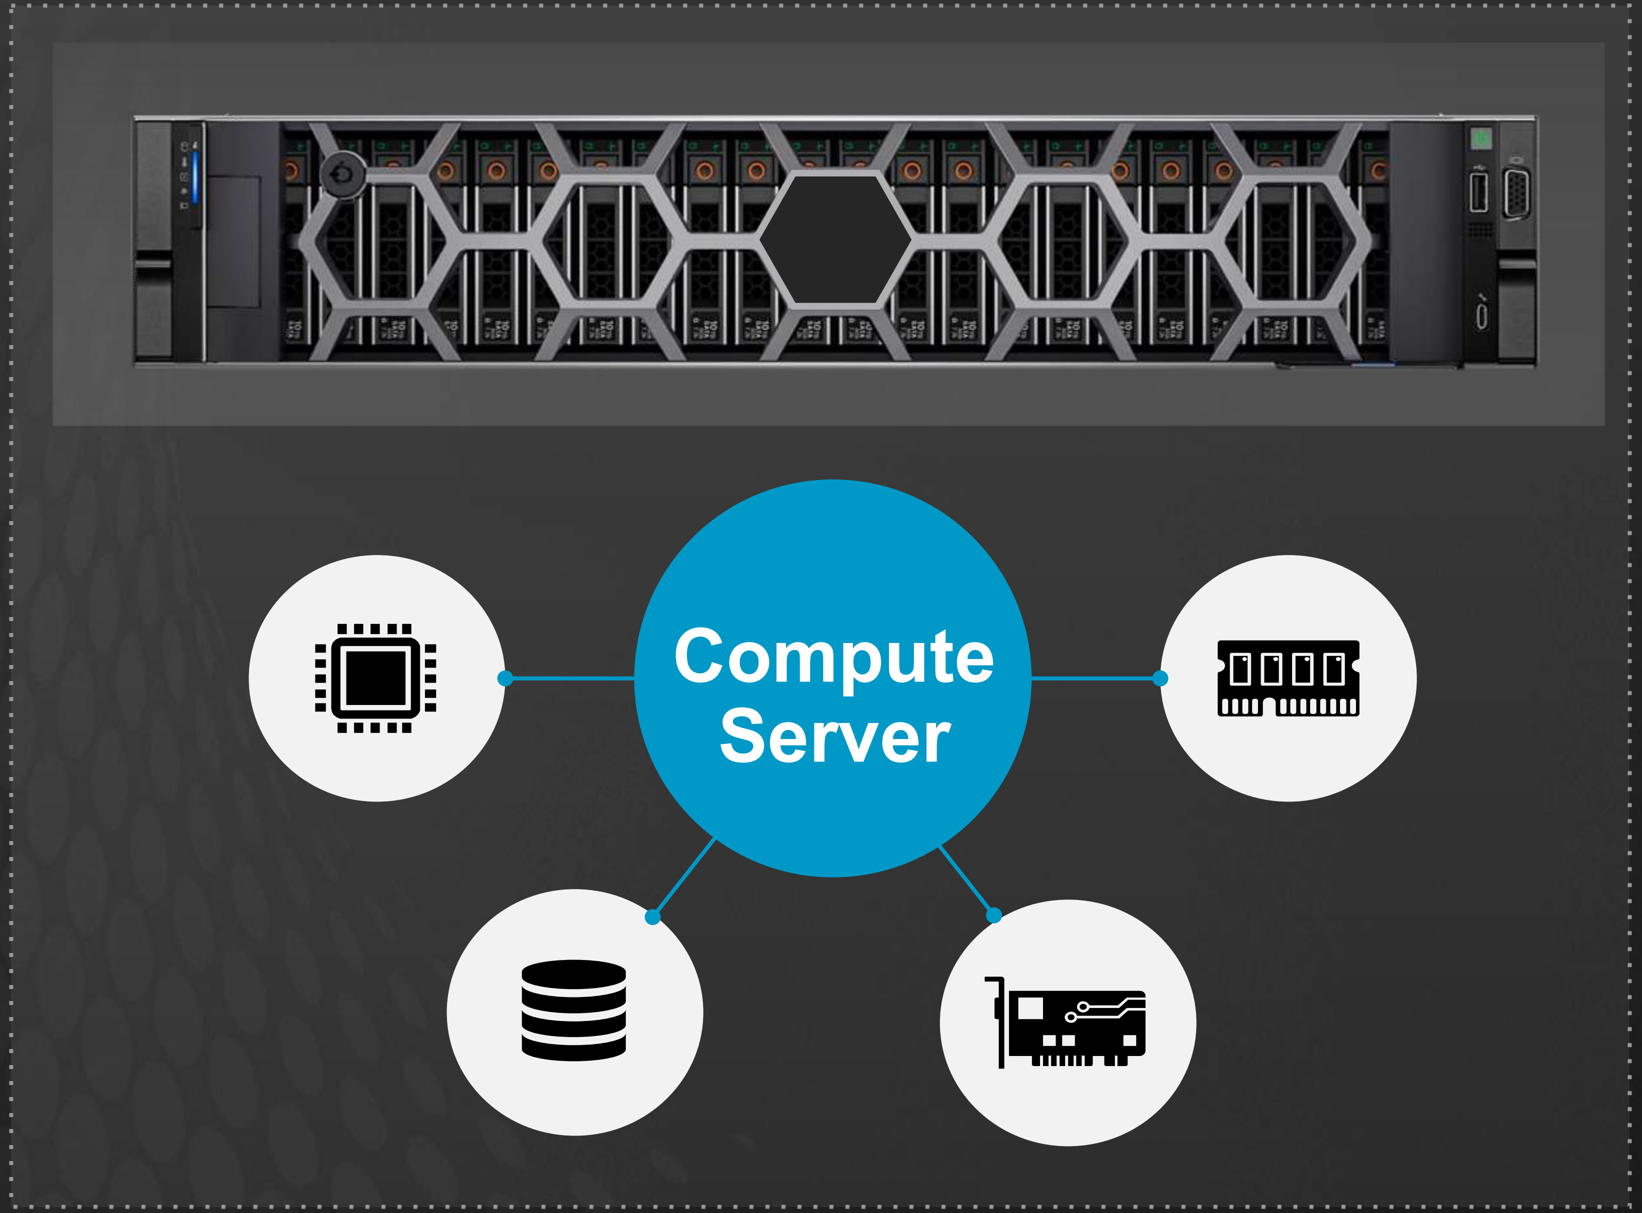

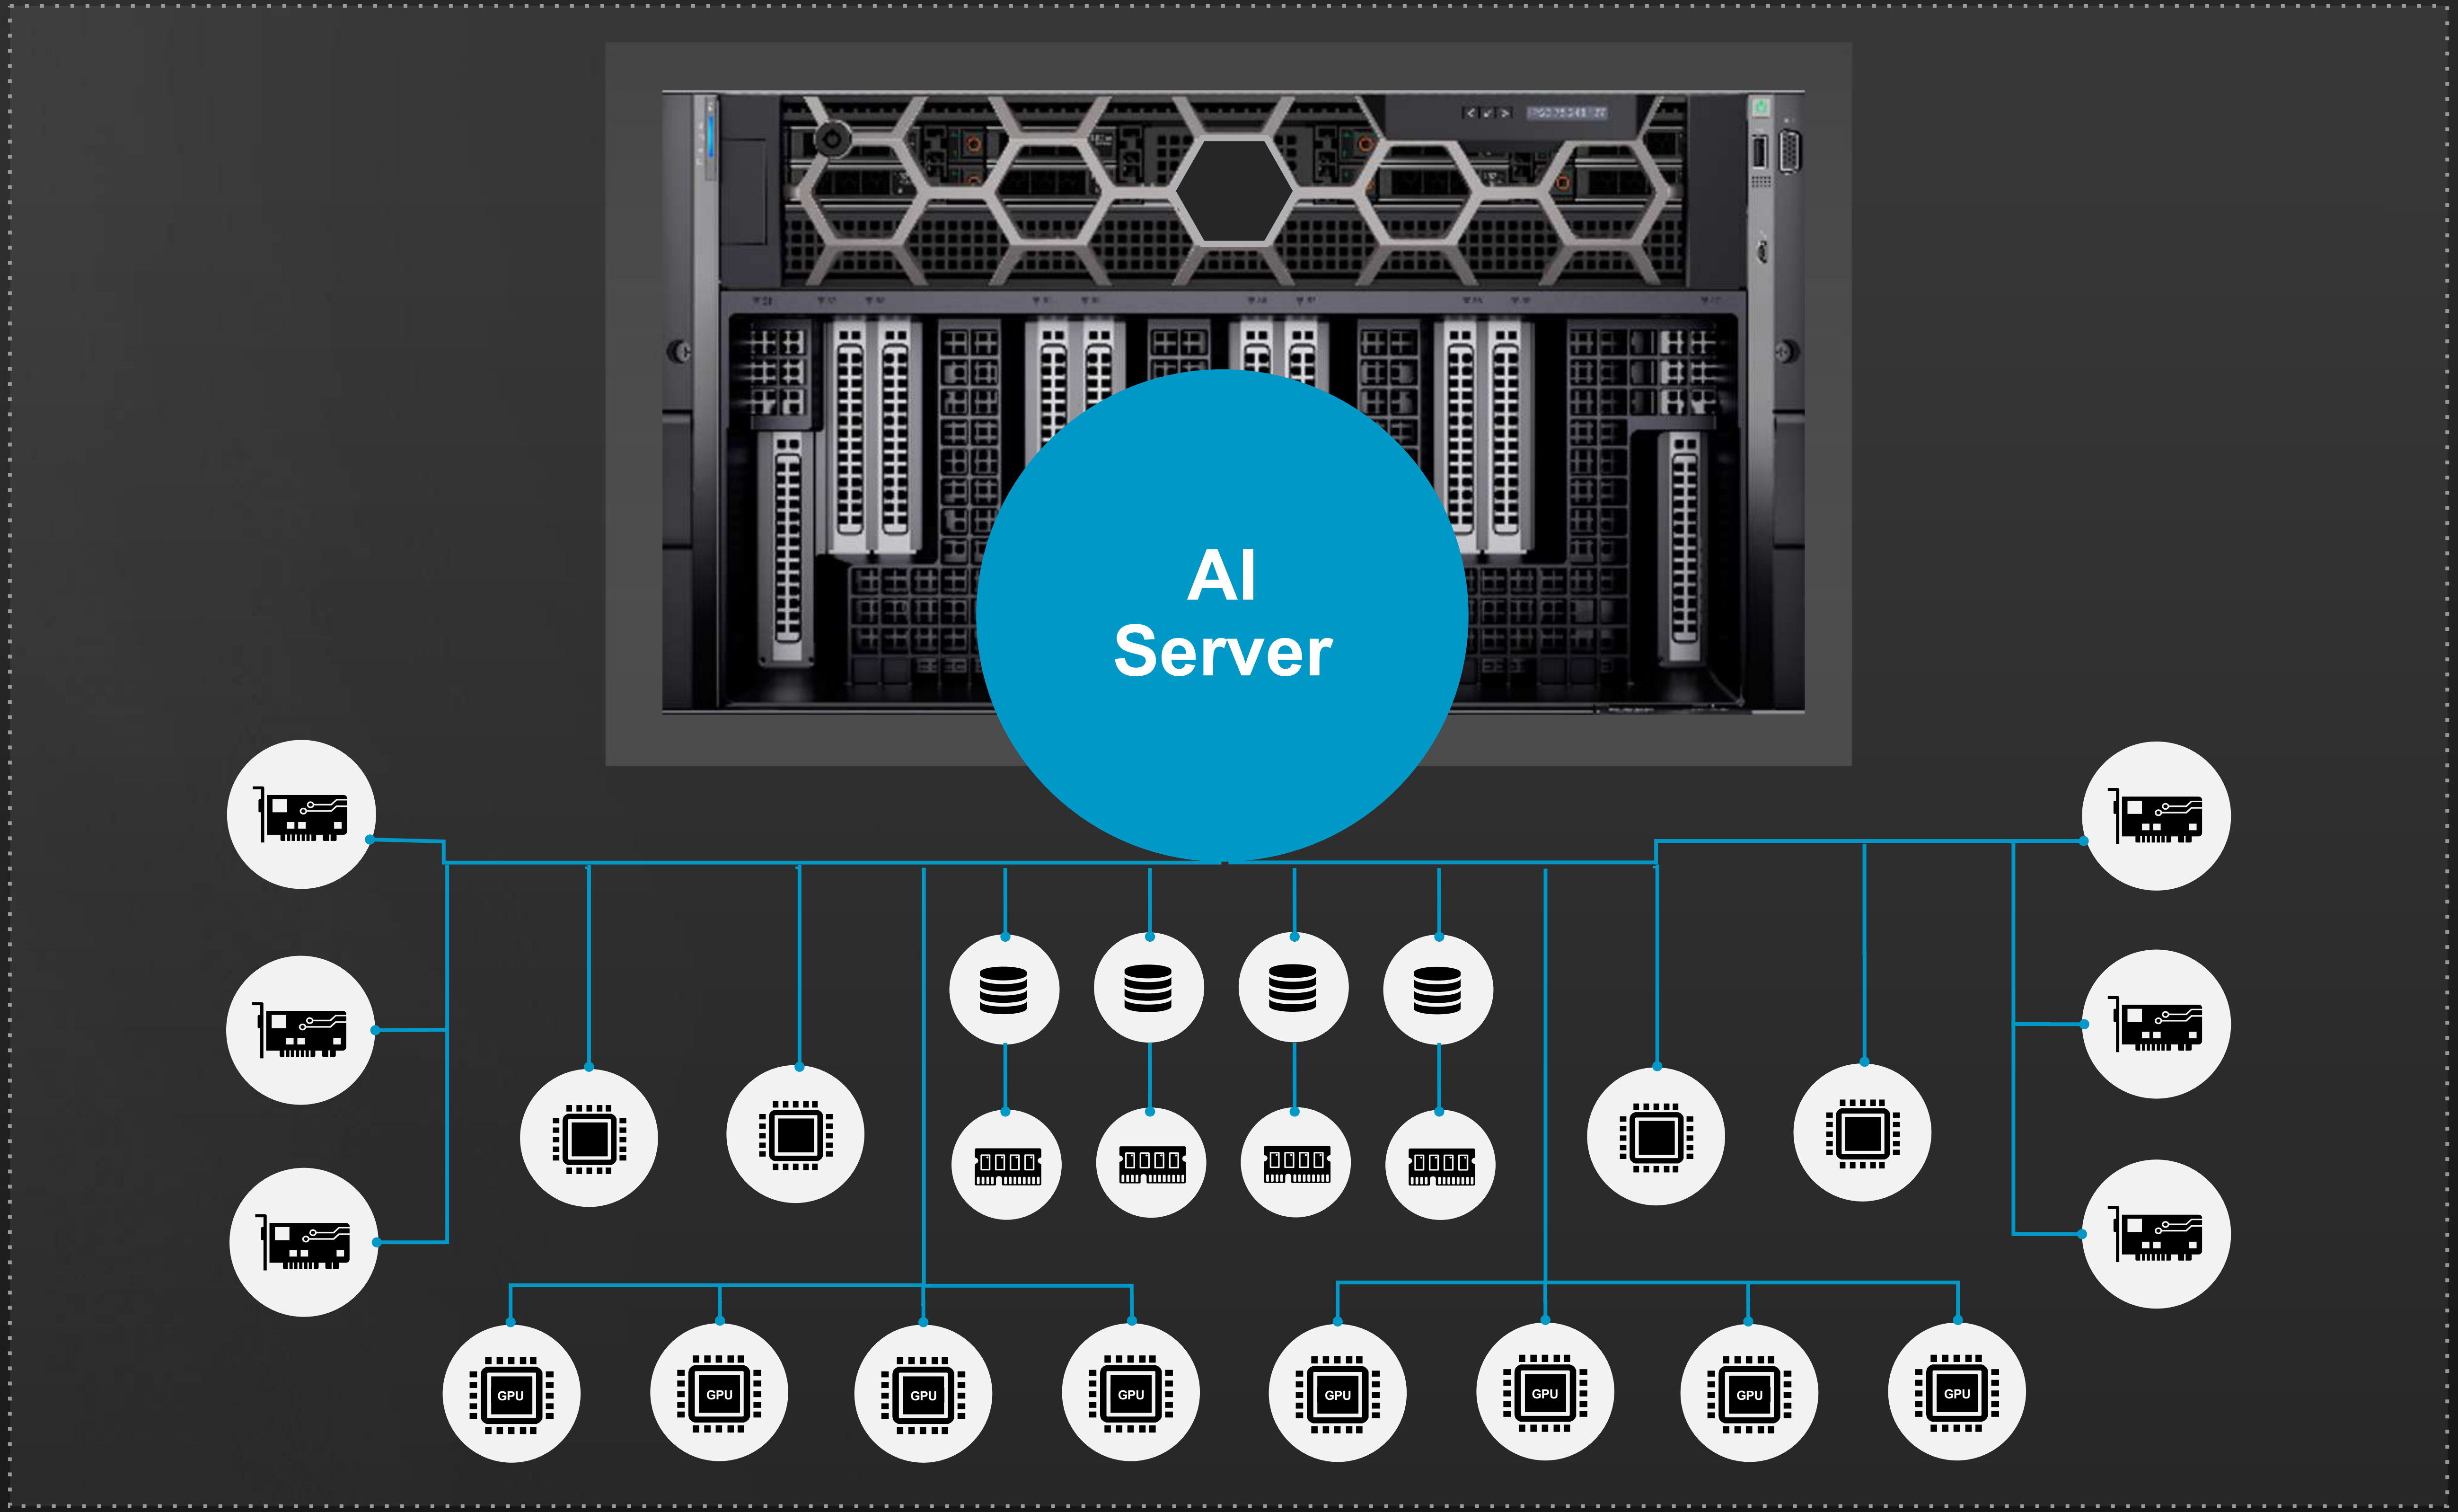

Compute Centric ➡ **High-Performance Interconnect**

OPEN // SCALABLE // POWER EFFICIENT

AI Server: Open → Options

Accelerators

a AMD BROADCOM
intel G Meta
Microsoft NVIDIA

Central Processing Unit

a AMD
intel

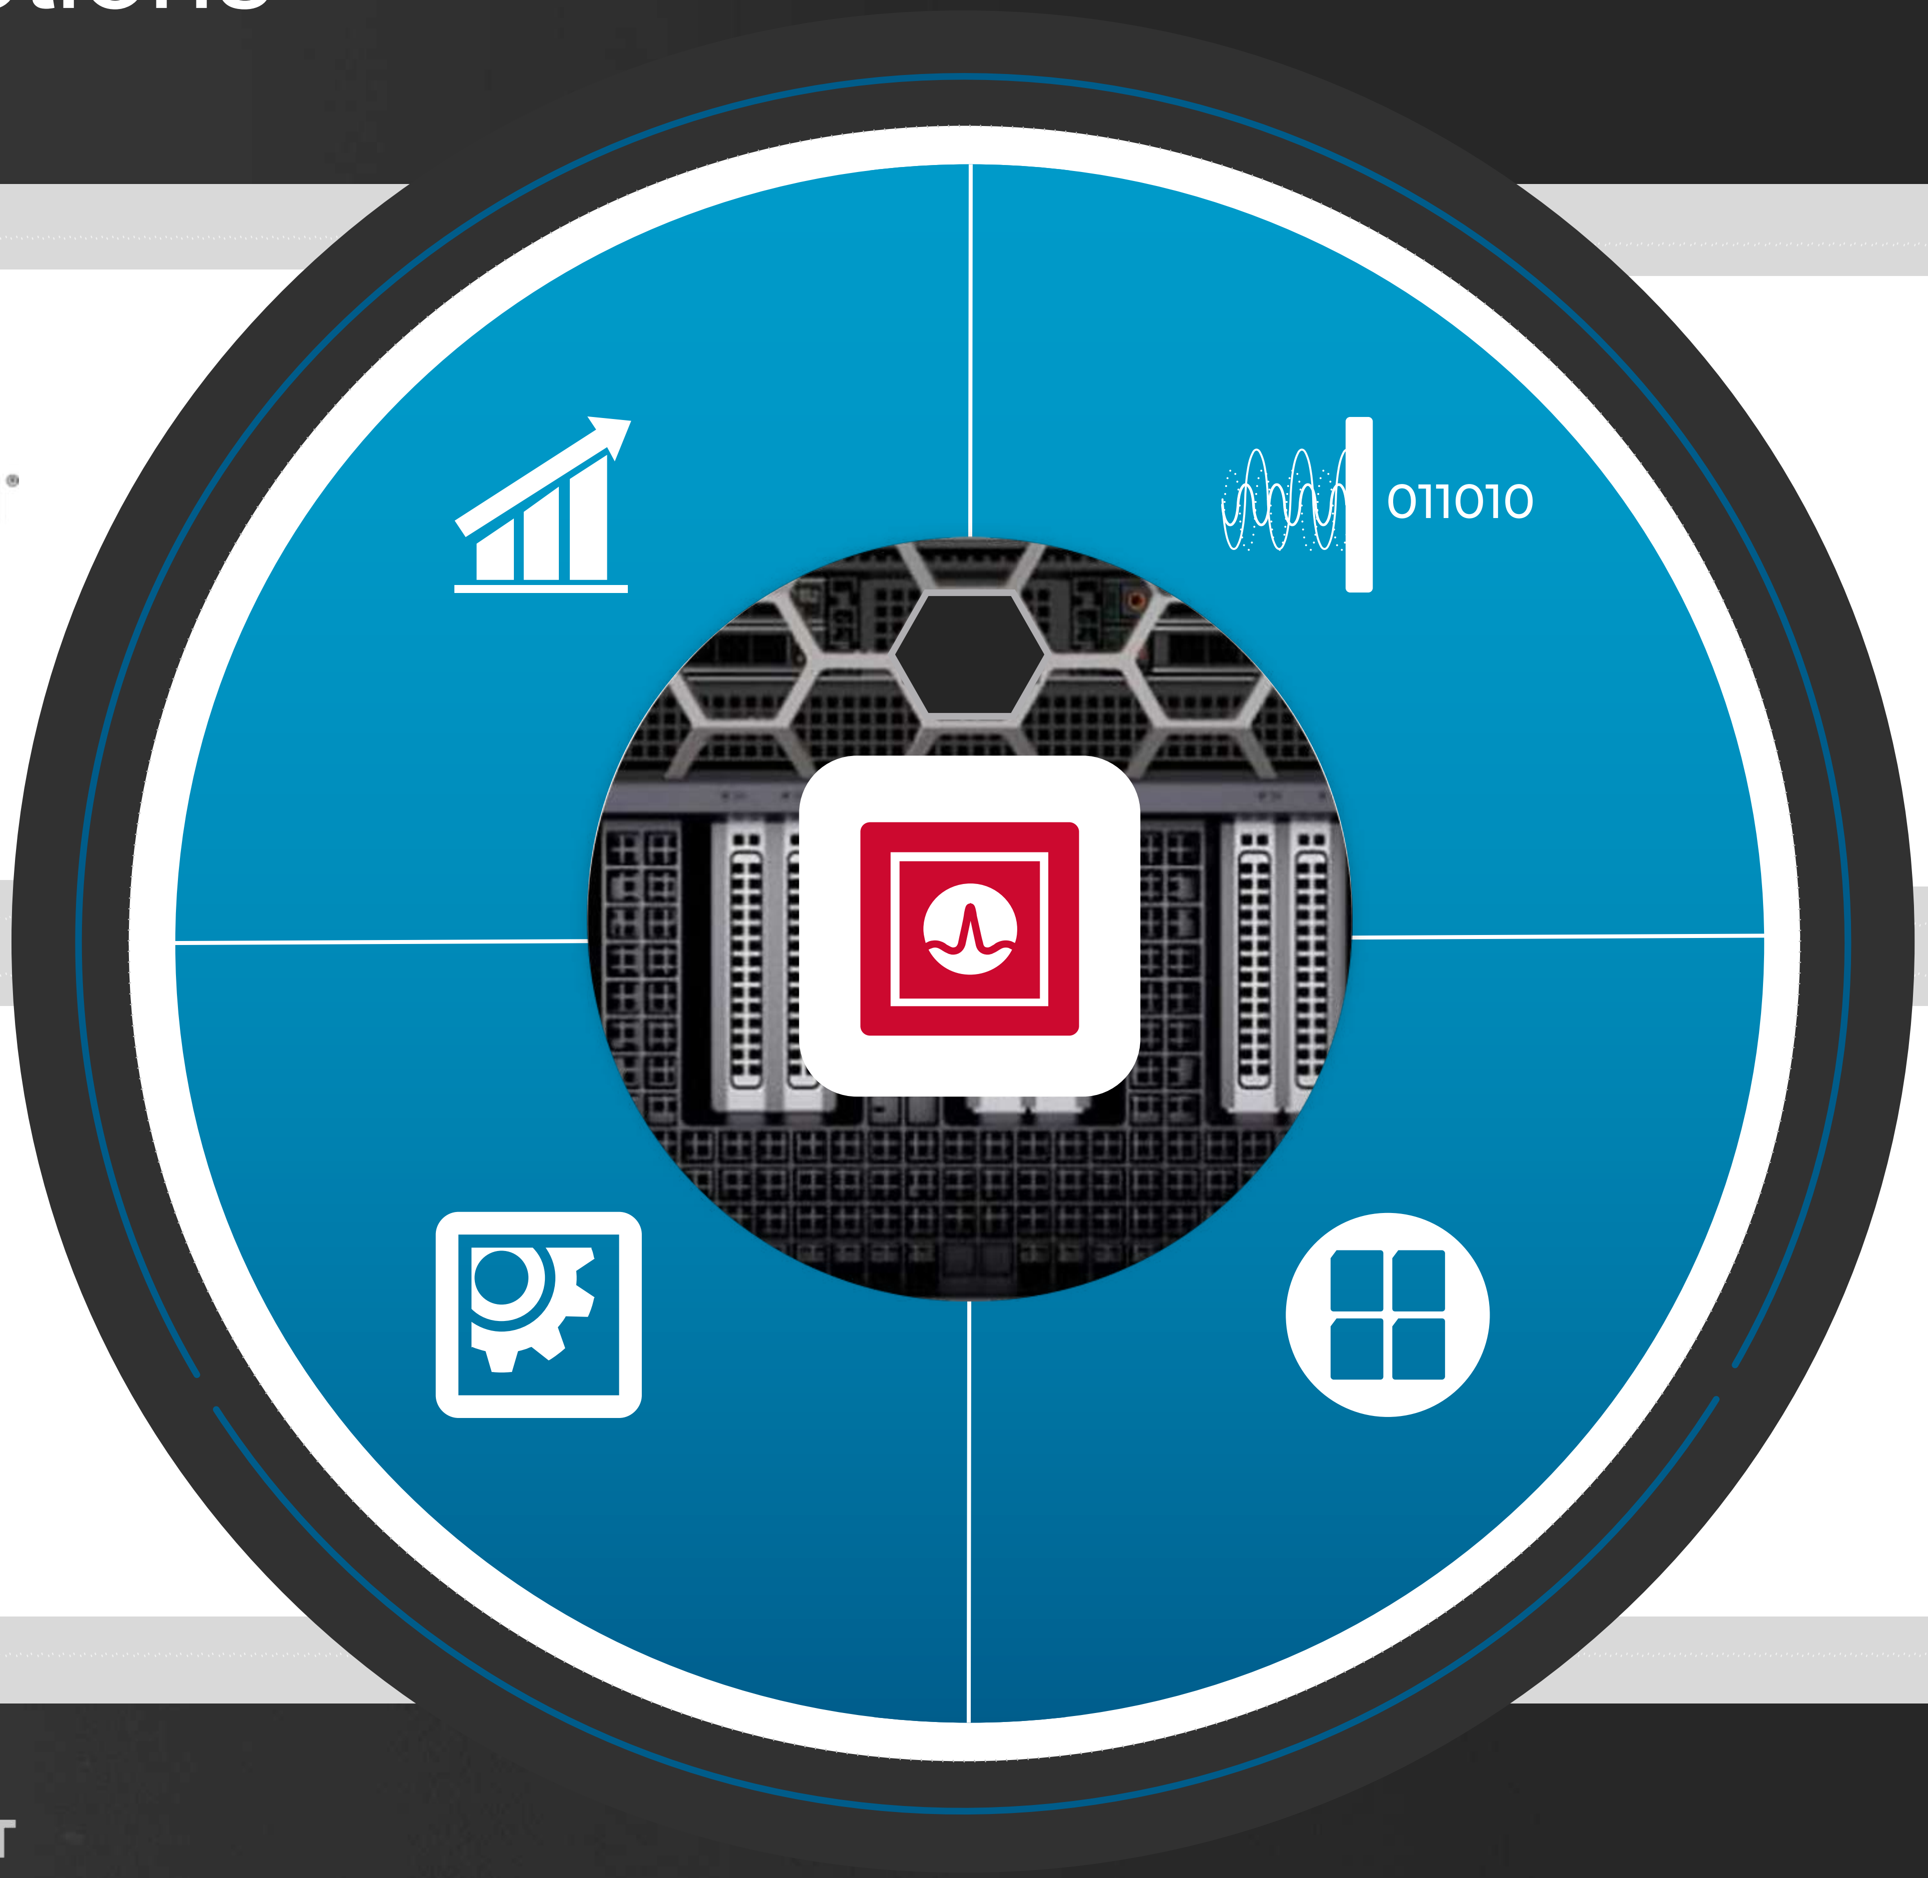

011010 NIC

a BROADCOM
intel NVIDIA

Storage & Memory

KIOXIA Micron
SAMSUNG SEAGATE
SOLIDIGM Western Digital

OPEN // SCALABLE // POWER EFFICIENT

AI Rack: Interconnect

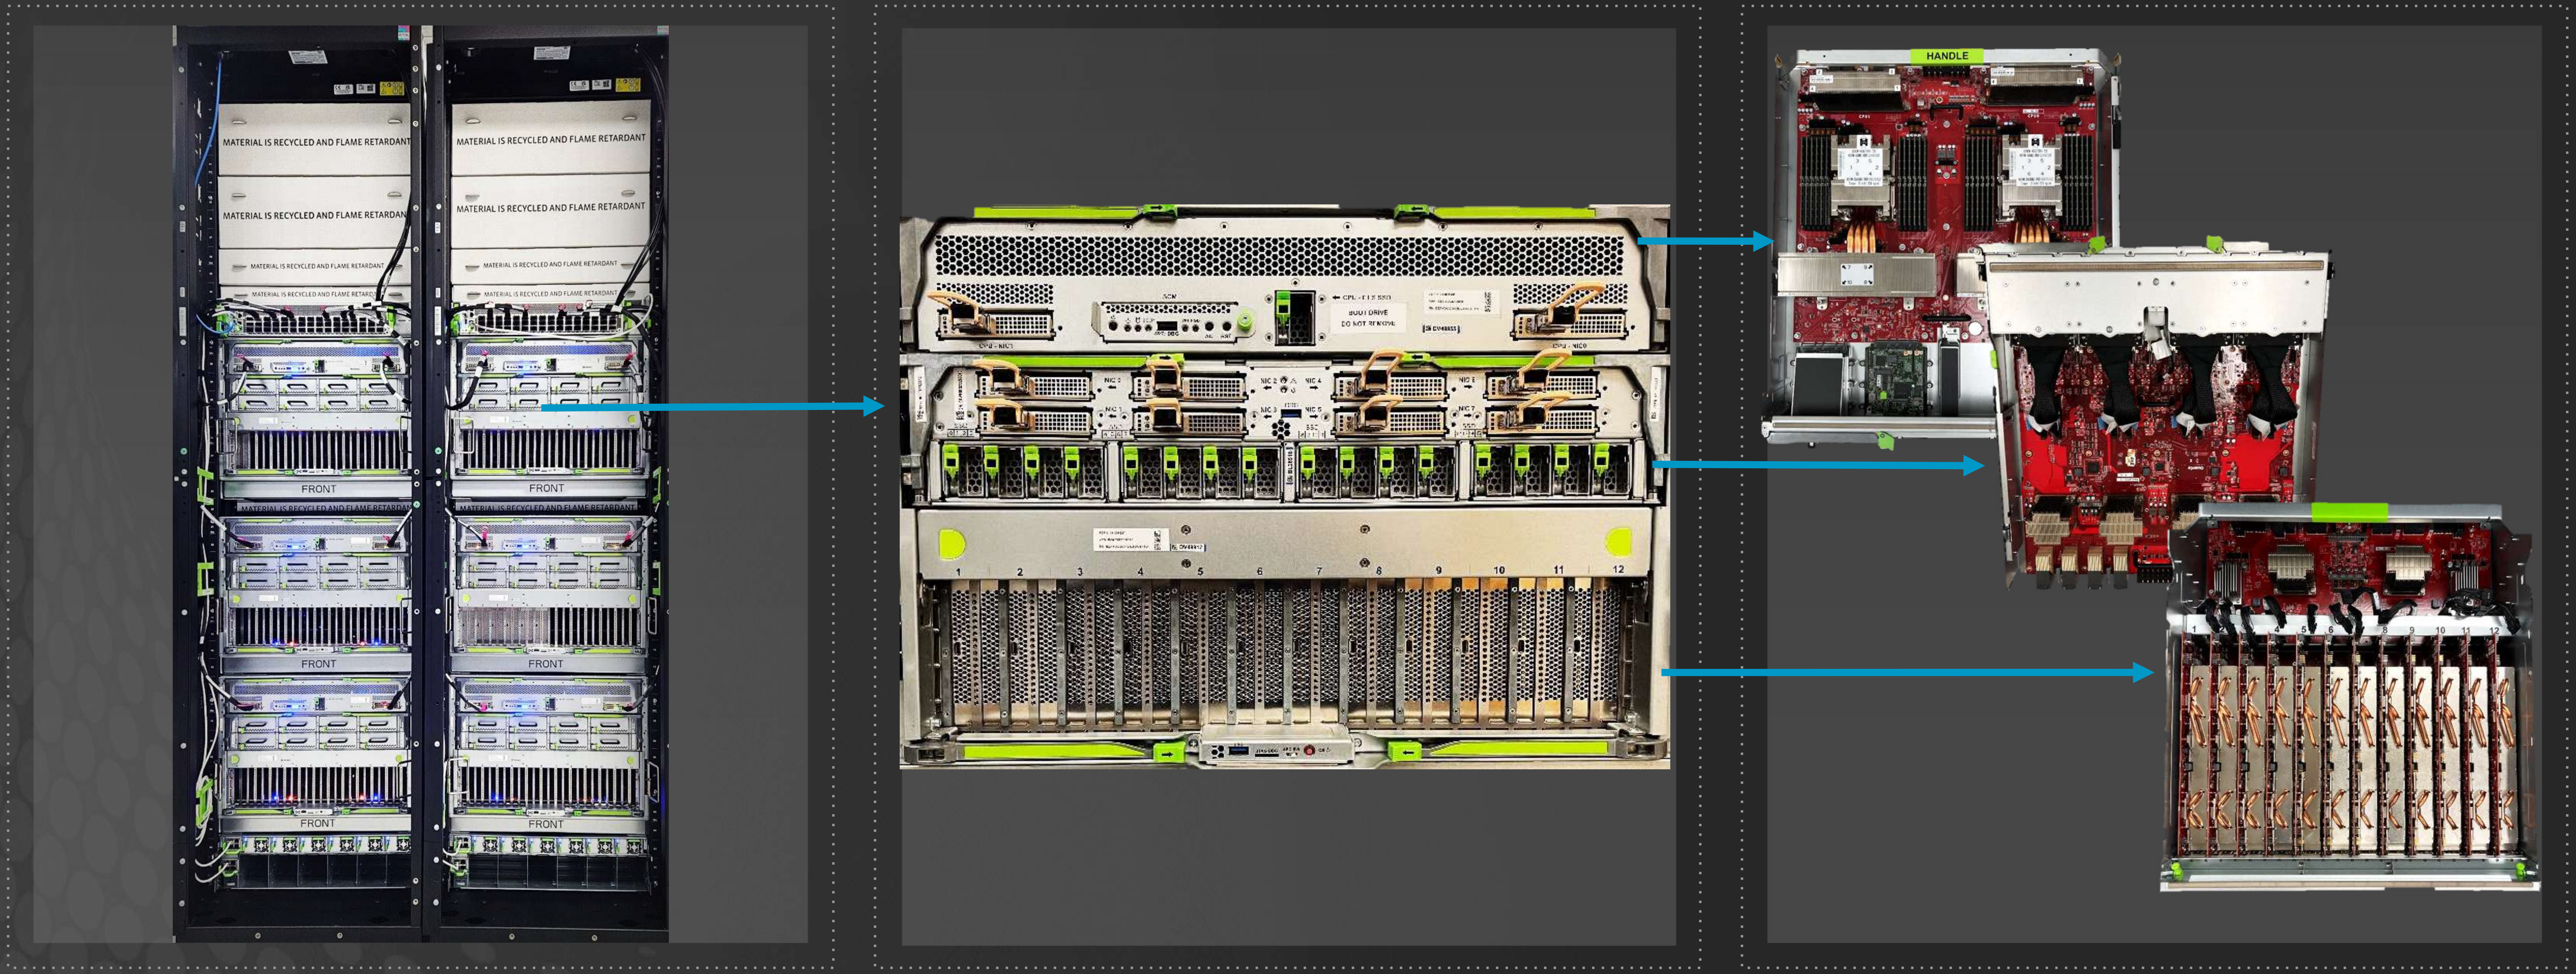

AI Server: Interconnect

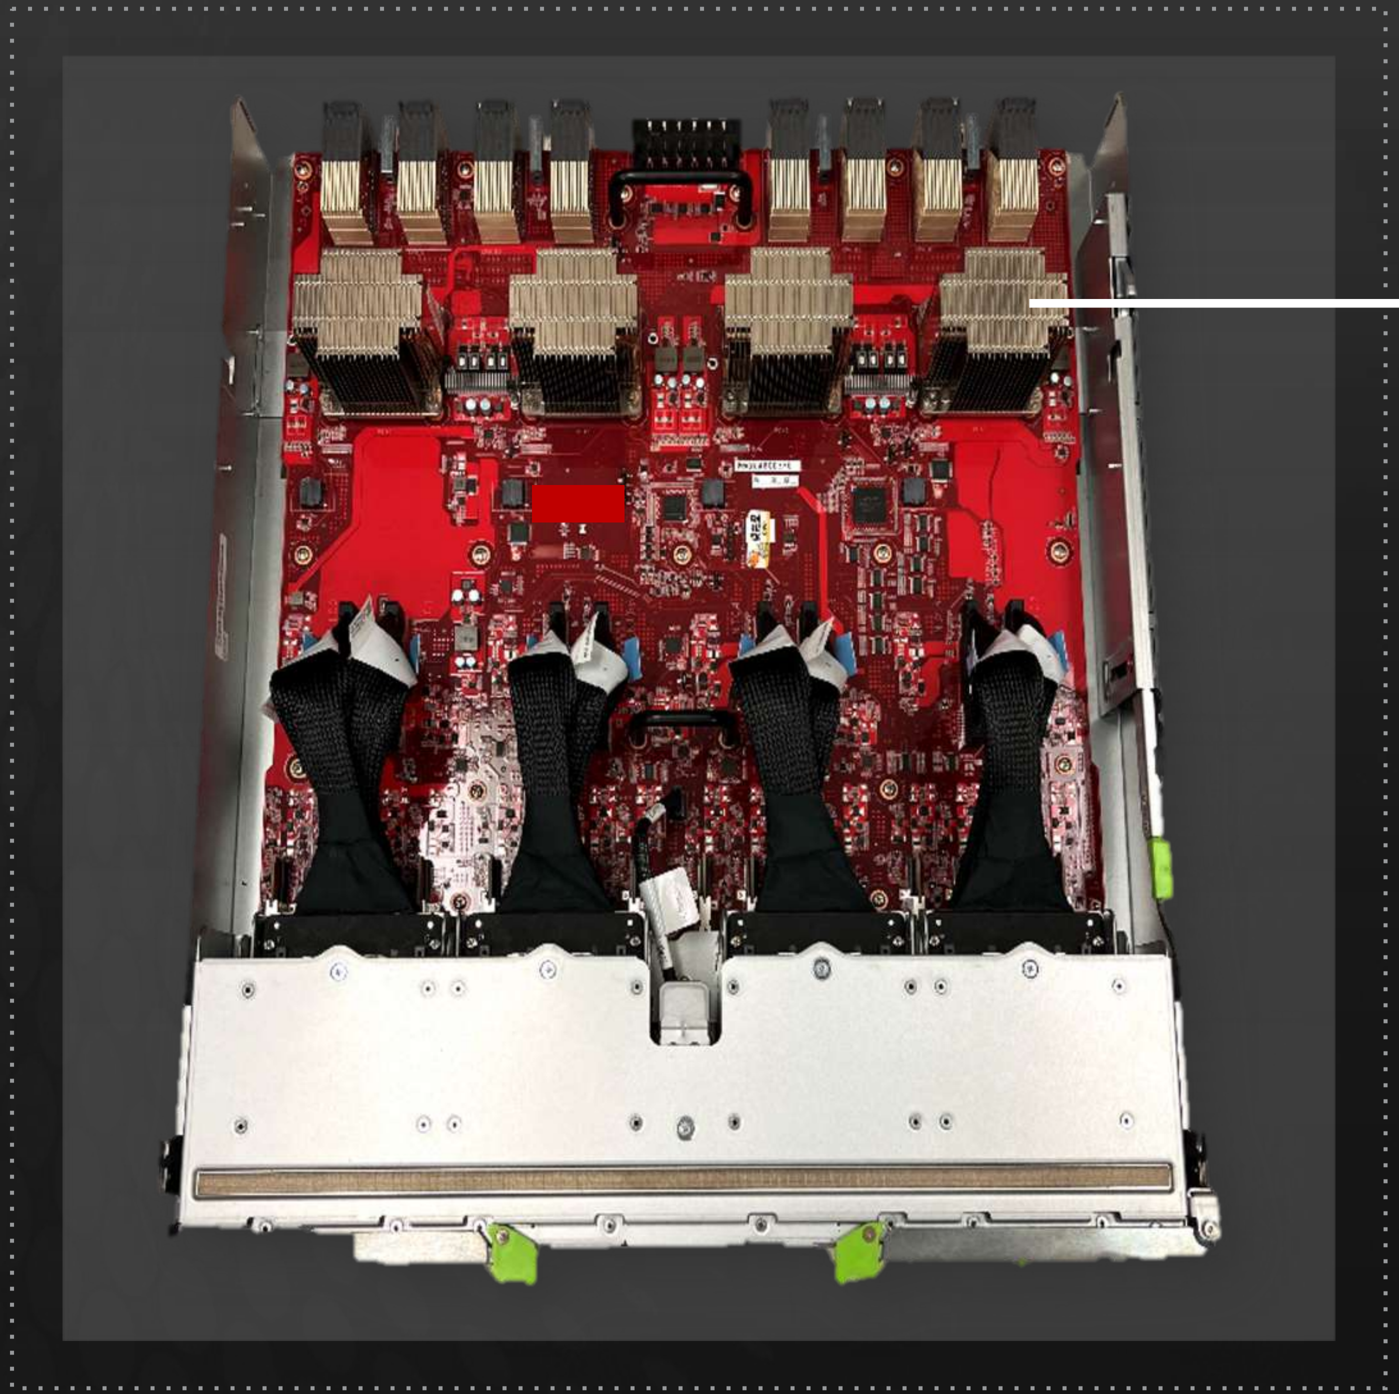

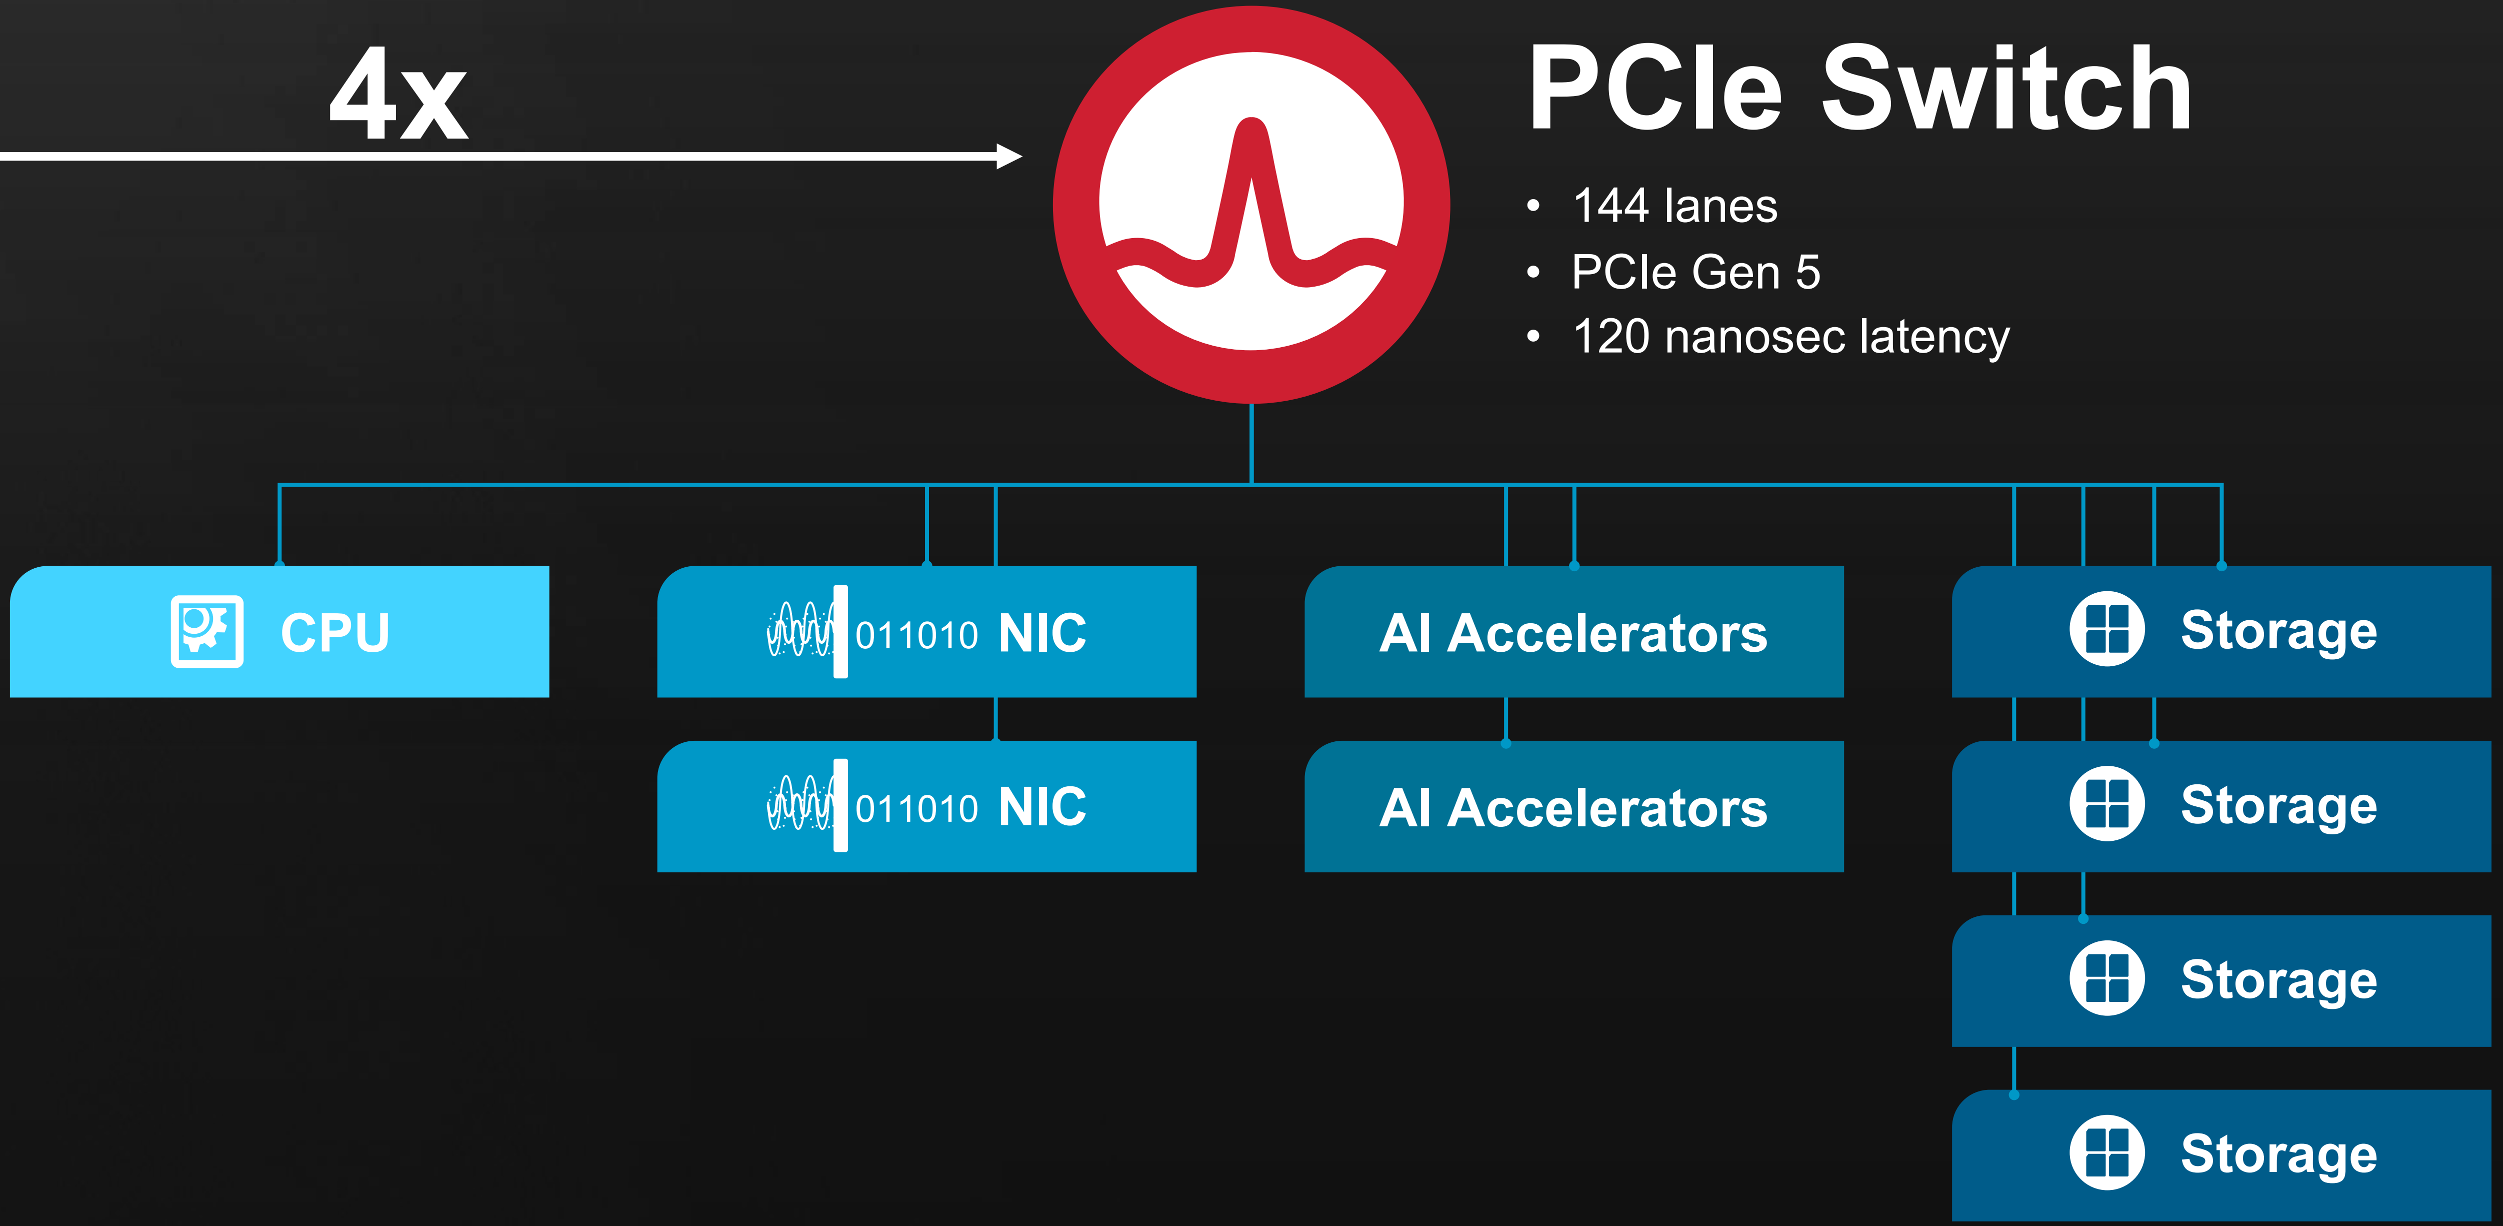

PCIe Switch

- 144 lanes
- PCIe Gen 5
- 120 nanosec latency

Most Trusted PCIe Device

OPEN // SCALABLE // POWER EFFICIENT

AI Server: PCIe Interconnect

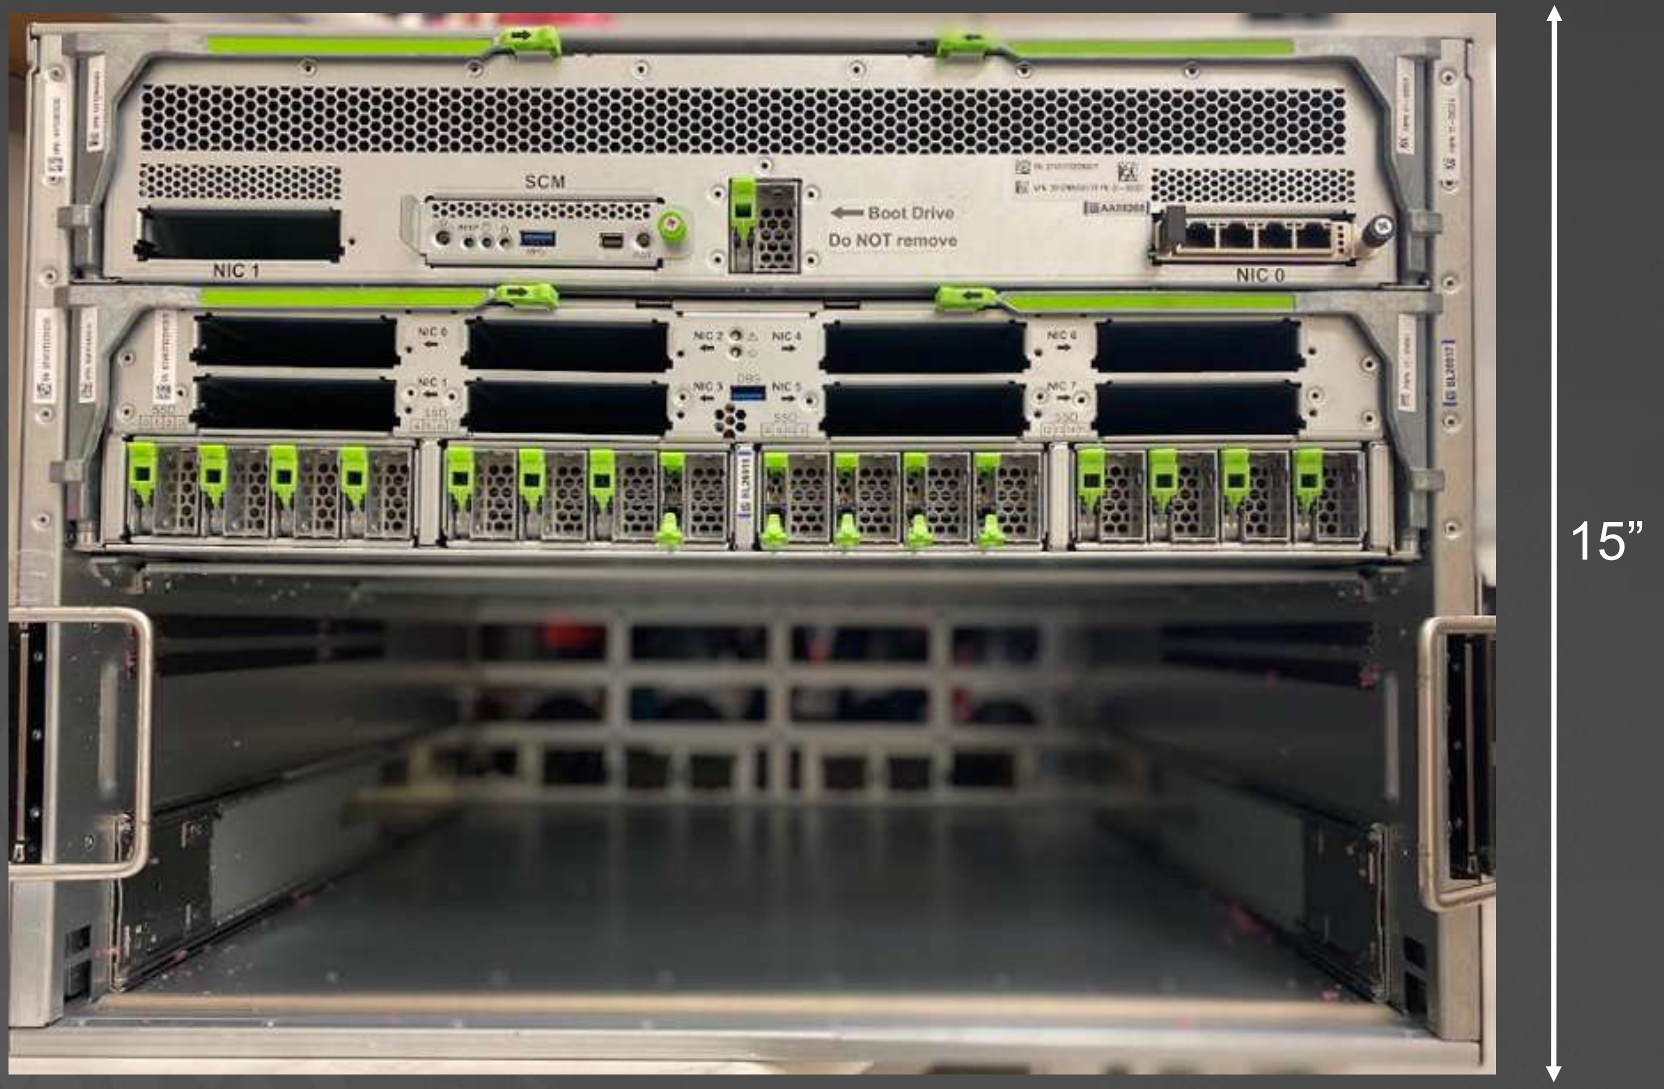

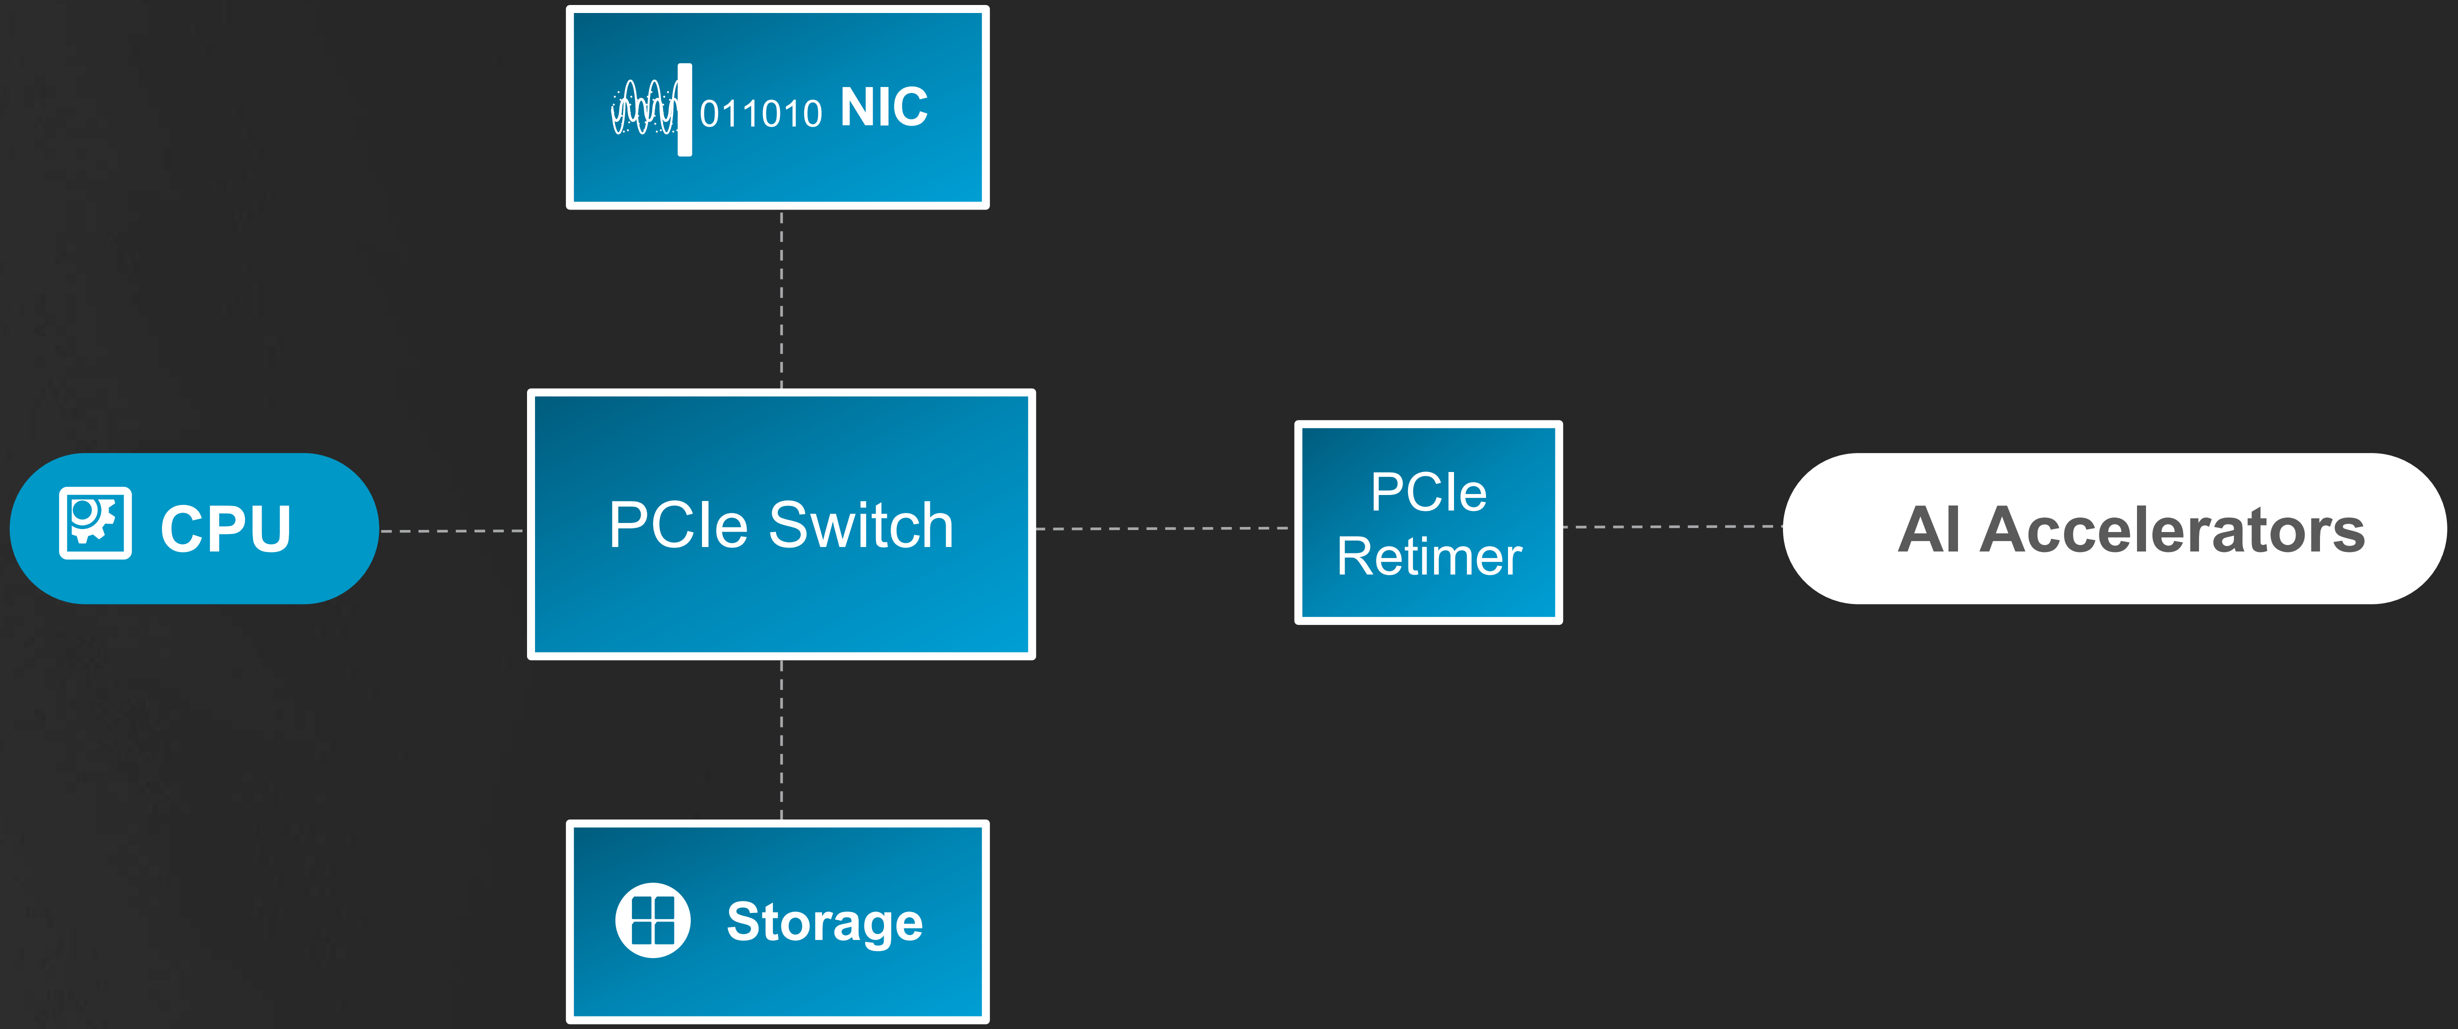

011010 NIC

CPU

PCIe Switch

PCIe Retimer

AI Accelerators

Storage

End-to-End AI PCIe Connectivity

OPEN // SCALABLE // POWER EFFICIENT

AI Server: PCIe Retimer

Retimer	Competition	Broadcom
Technology	7 / 16nm	**5nm**
SerDes	Third Party	**In House**
Power	1x	**35% Less**
Reach	1x	**40% More**

Source: Broadcom Internal Estimates

End-to-End AI PCIe Connectivity

OPEN // SCALABLE // POWER EFFICIENT

PCIe Switching Interconnect Leadership

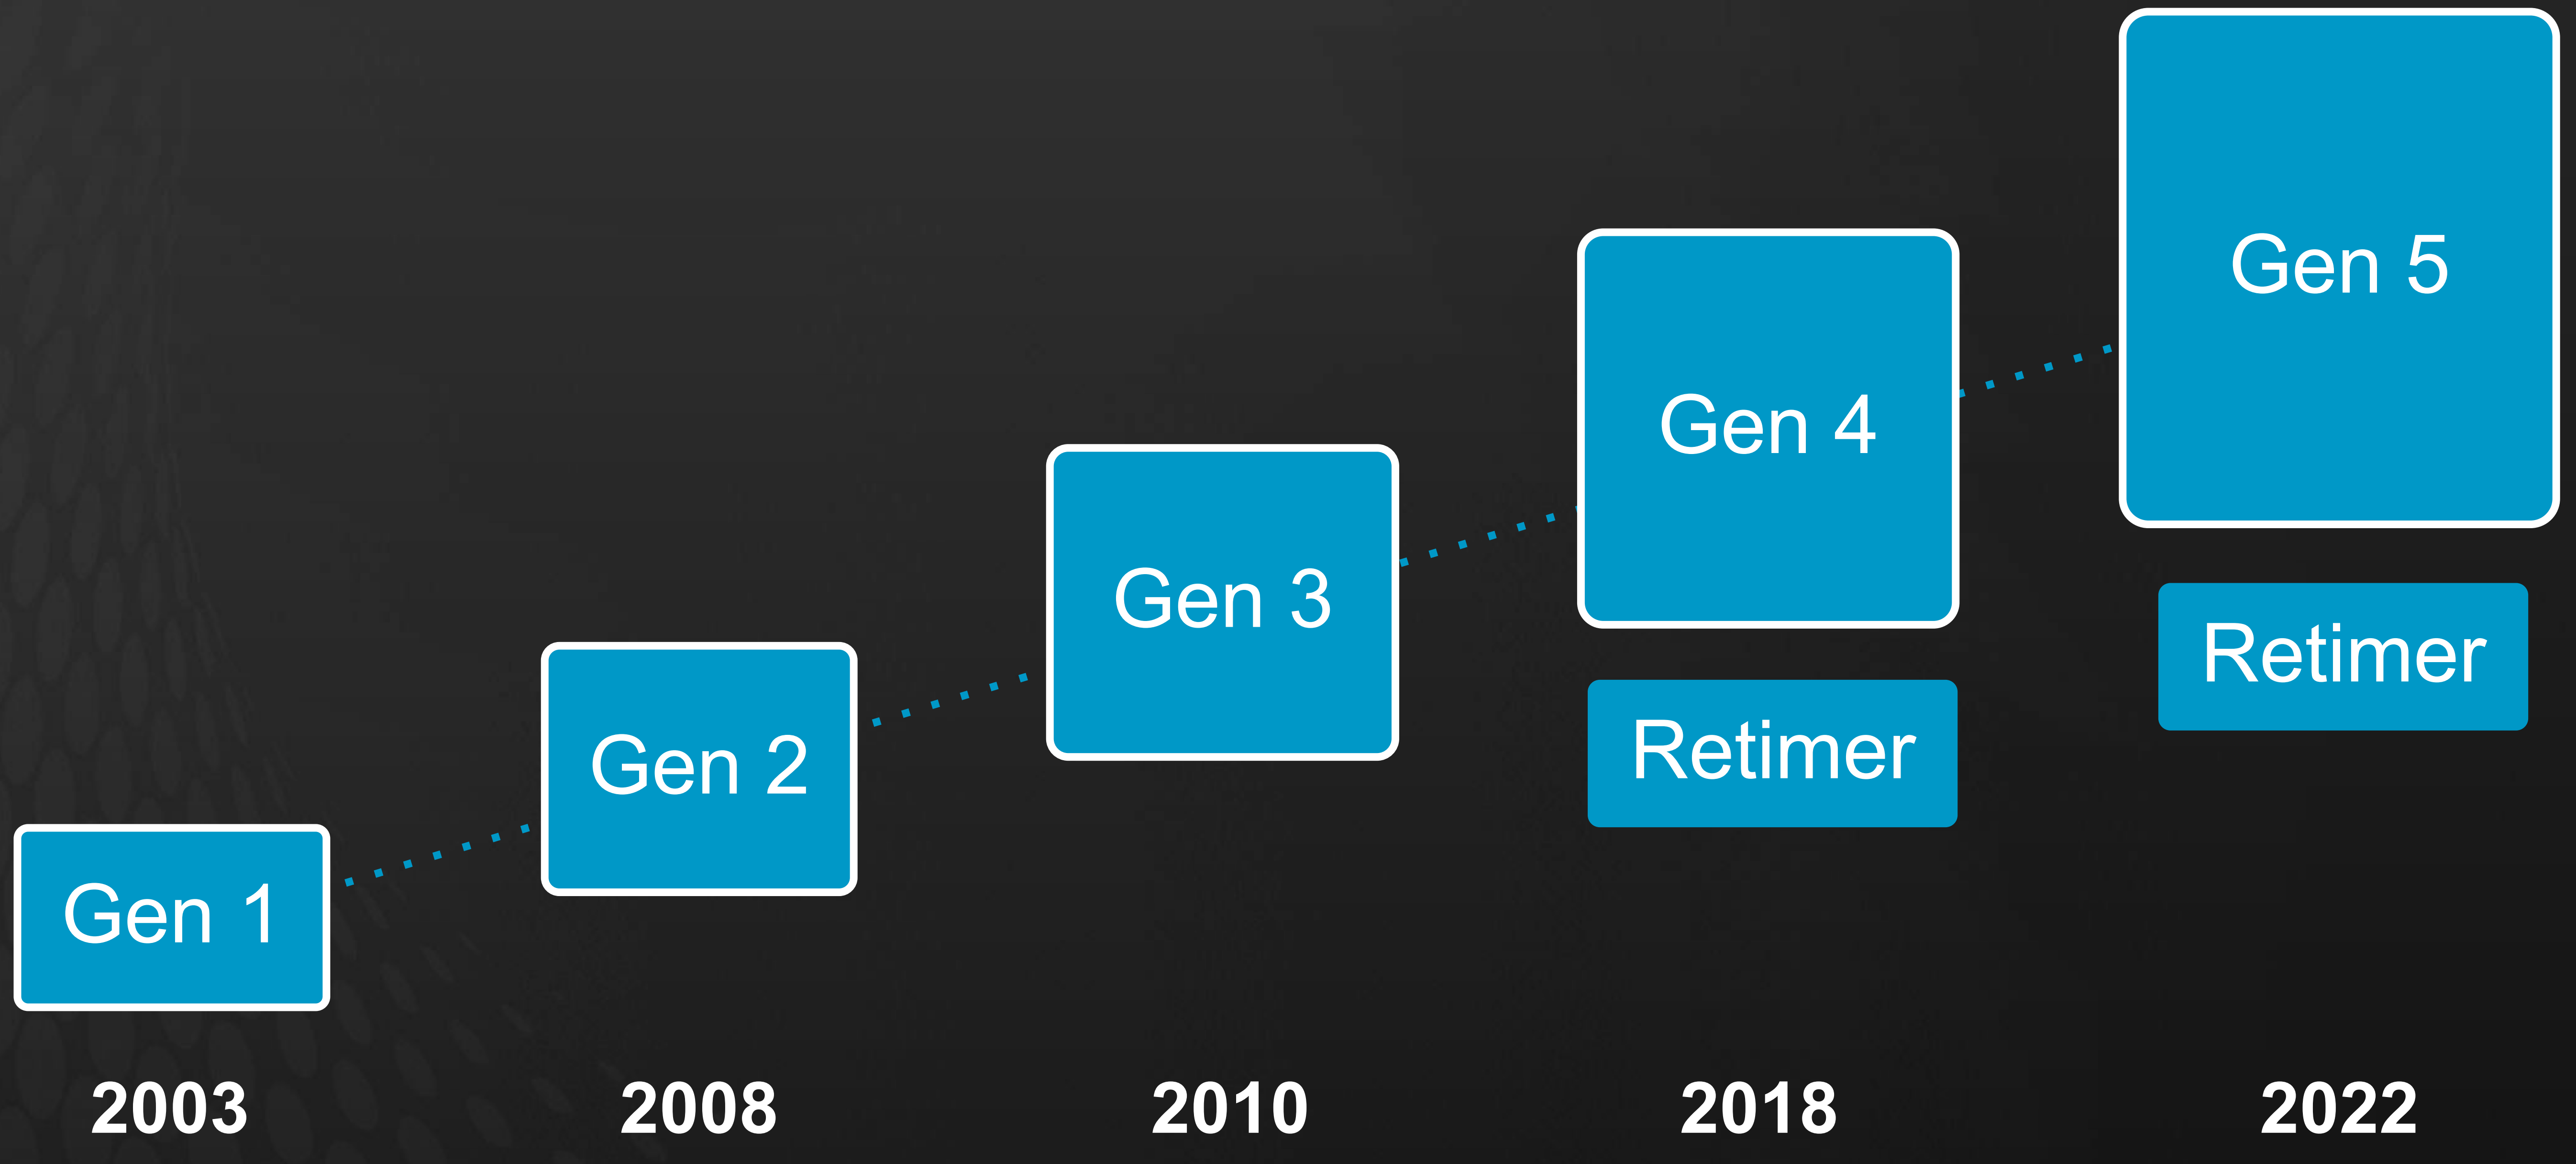

1B+ PCIe ports shipped

- PCIe Gen 5 Leadership
- Open: PCIe Golden Node
- Scale: 38% more lanes
- Power Efficient: 50% Less

Switch + Retimer

First to Market for 20 Years … 5 Successful Generations

PCIe Switching Interconnect Leadership

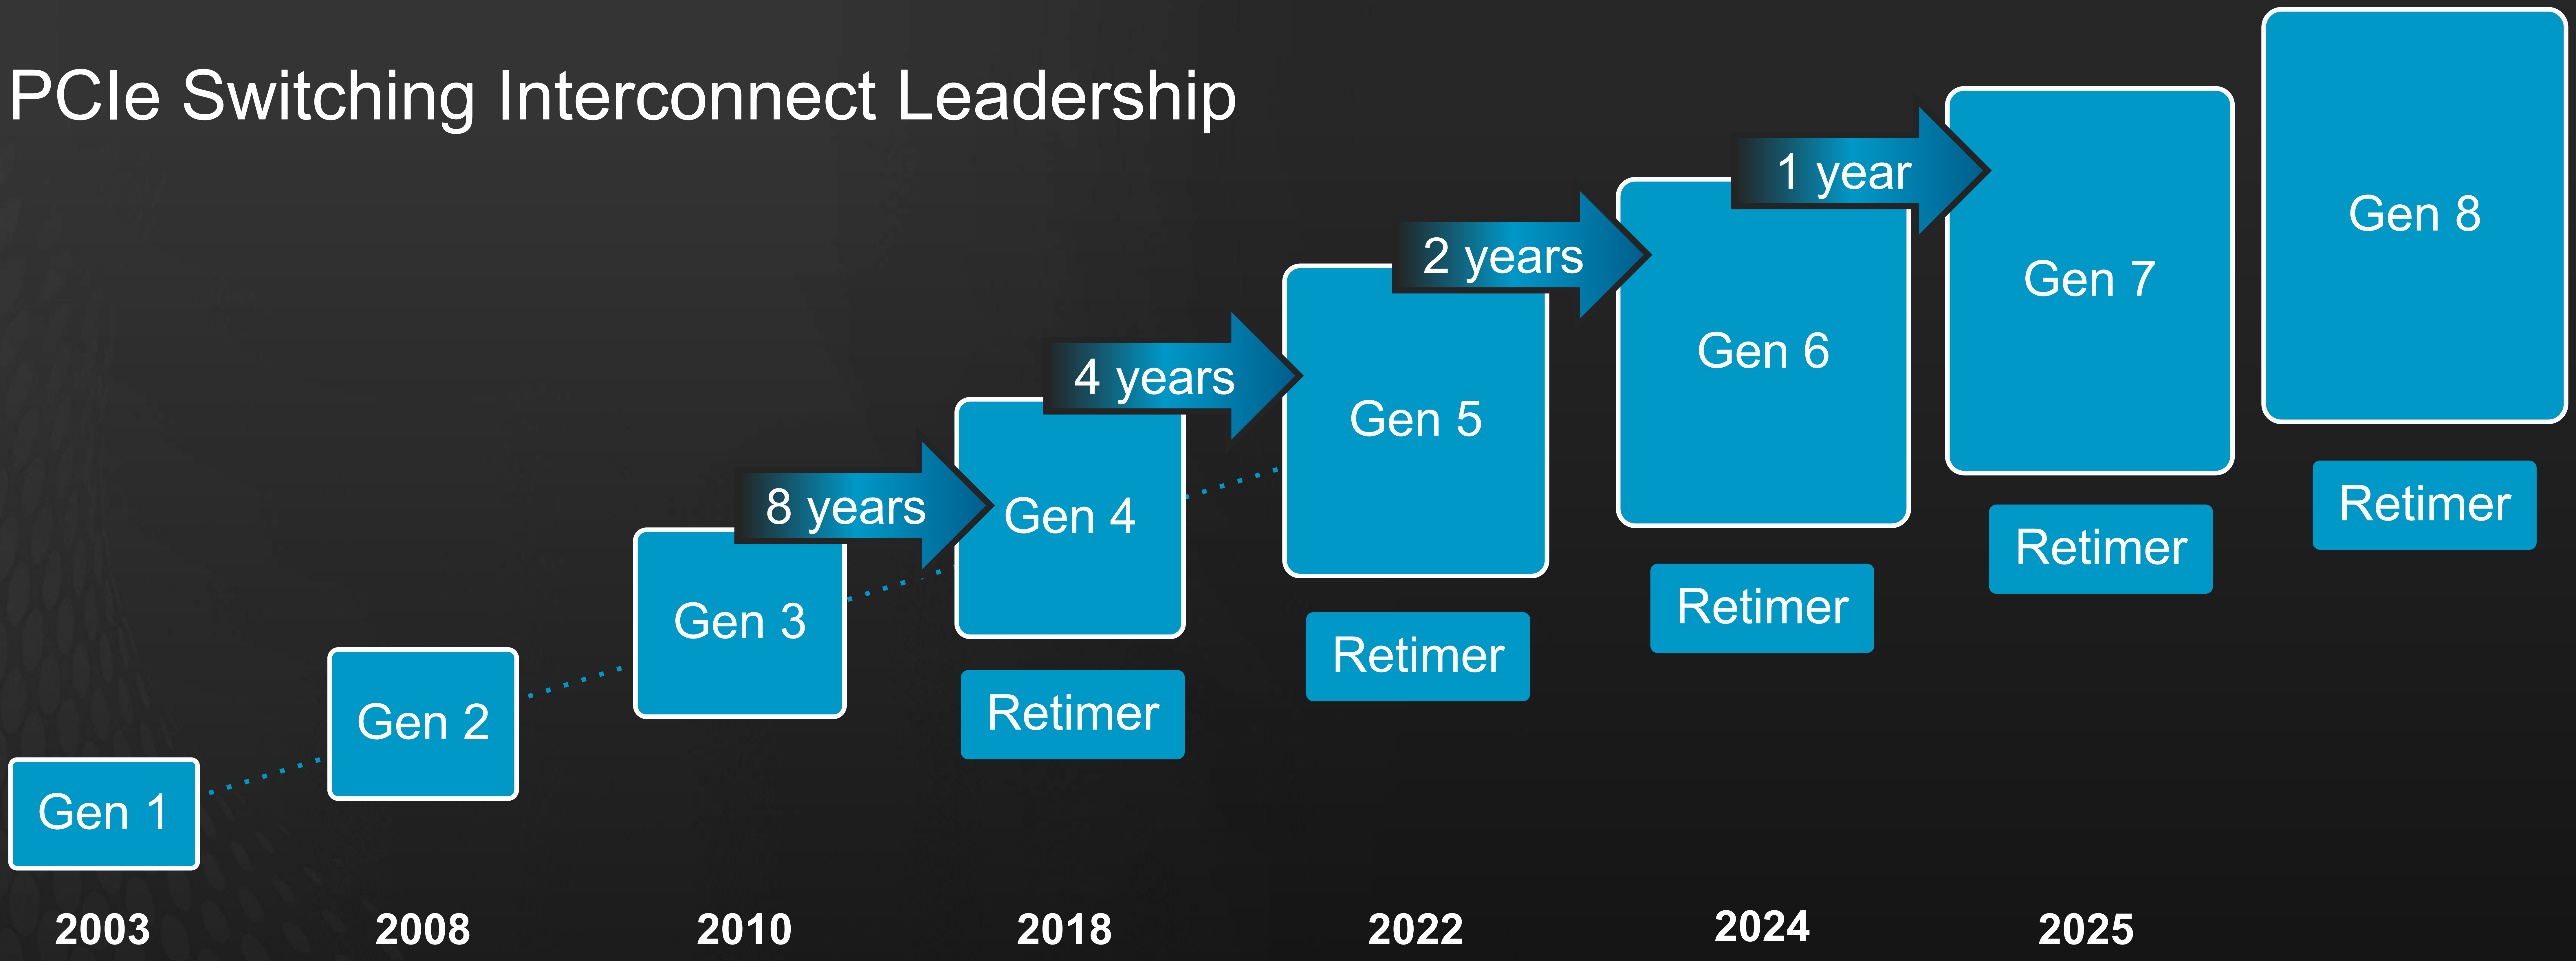

Gen 1

Gen 2

Gen 3

8 years

Gen 4

Retimer

4 years

Gen 5

Retimer

2 years

Gen 6

Retimer

1 year

Gen 7

Retimer

Gen 8

Retimer

2003

2008

2010

2018

2022

2024

2025

Accelerating PCIe Switch + Retimer Introduction

AI Server: Scale Up

PCIe Gen 7 Switch

- Accelerator Fabric Link Latency: ~120 nanosecond
- Scale: Hundreds of Nodes

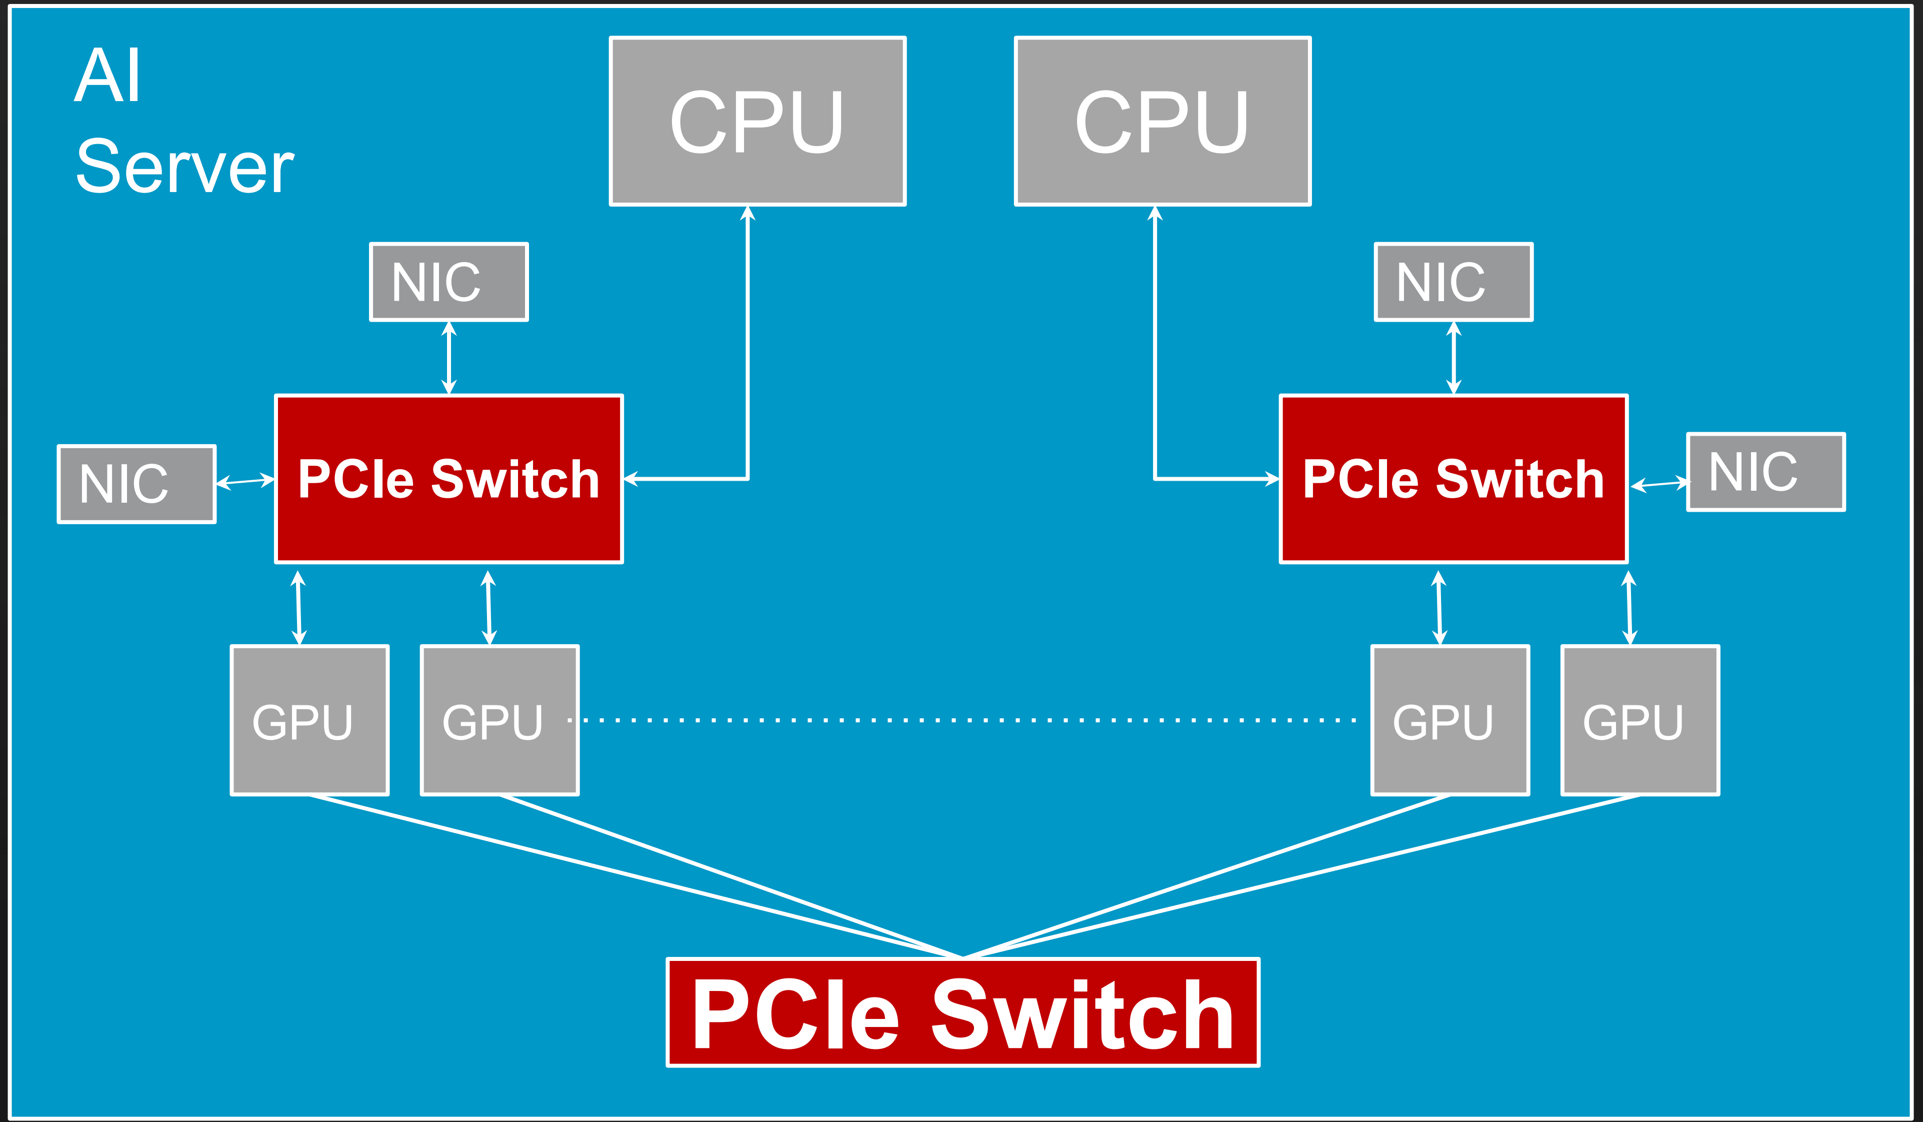

Scale Up

AMD/Broadcom Collaborate on Scale Up Fabrics

OPEN // SCALABLE // POWER EFFICIENT

AI Server Connectivity

Key Points

PCIe is key to AI Server **OPEN** interconnect

Decades of market leadership in **SCALABLE** PCIe Switches and Retimers

Broadcom end to end PCIe delivers 40% greater reach with 50% lower **POWER**

Source: Broadcom Internal Estimates

OPEN // SCALABLE // POWER EFFICIENT

BROADCOM

ENABLING AI Infrastructure

Optical Interconnects

Near Margalit, Ph. D.

General Manager
Optical Systems Division
Broadcom

OPEN // SCALABLE // POWER EFFICIENT

The Optical Interconnect for AI

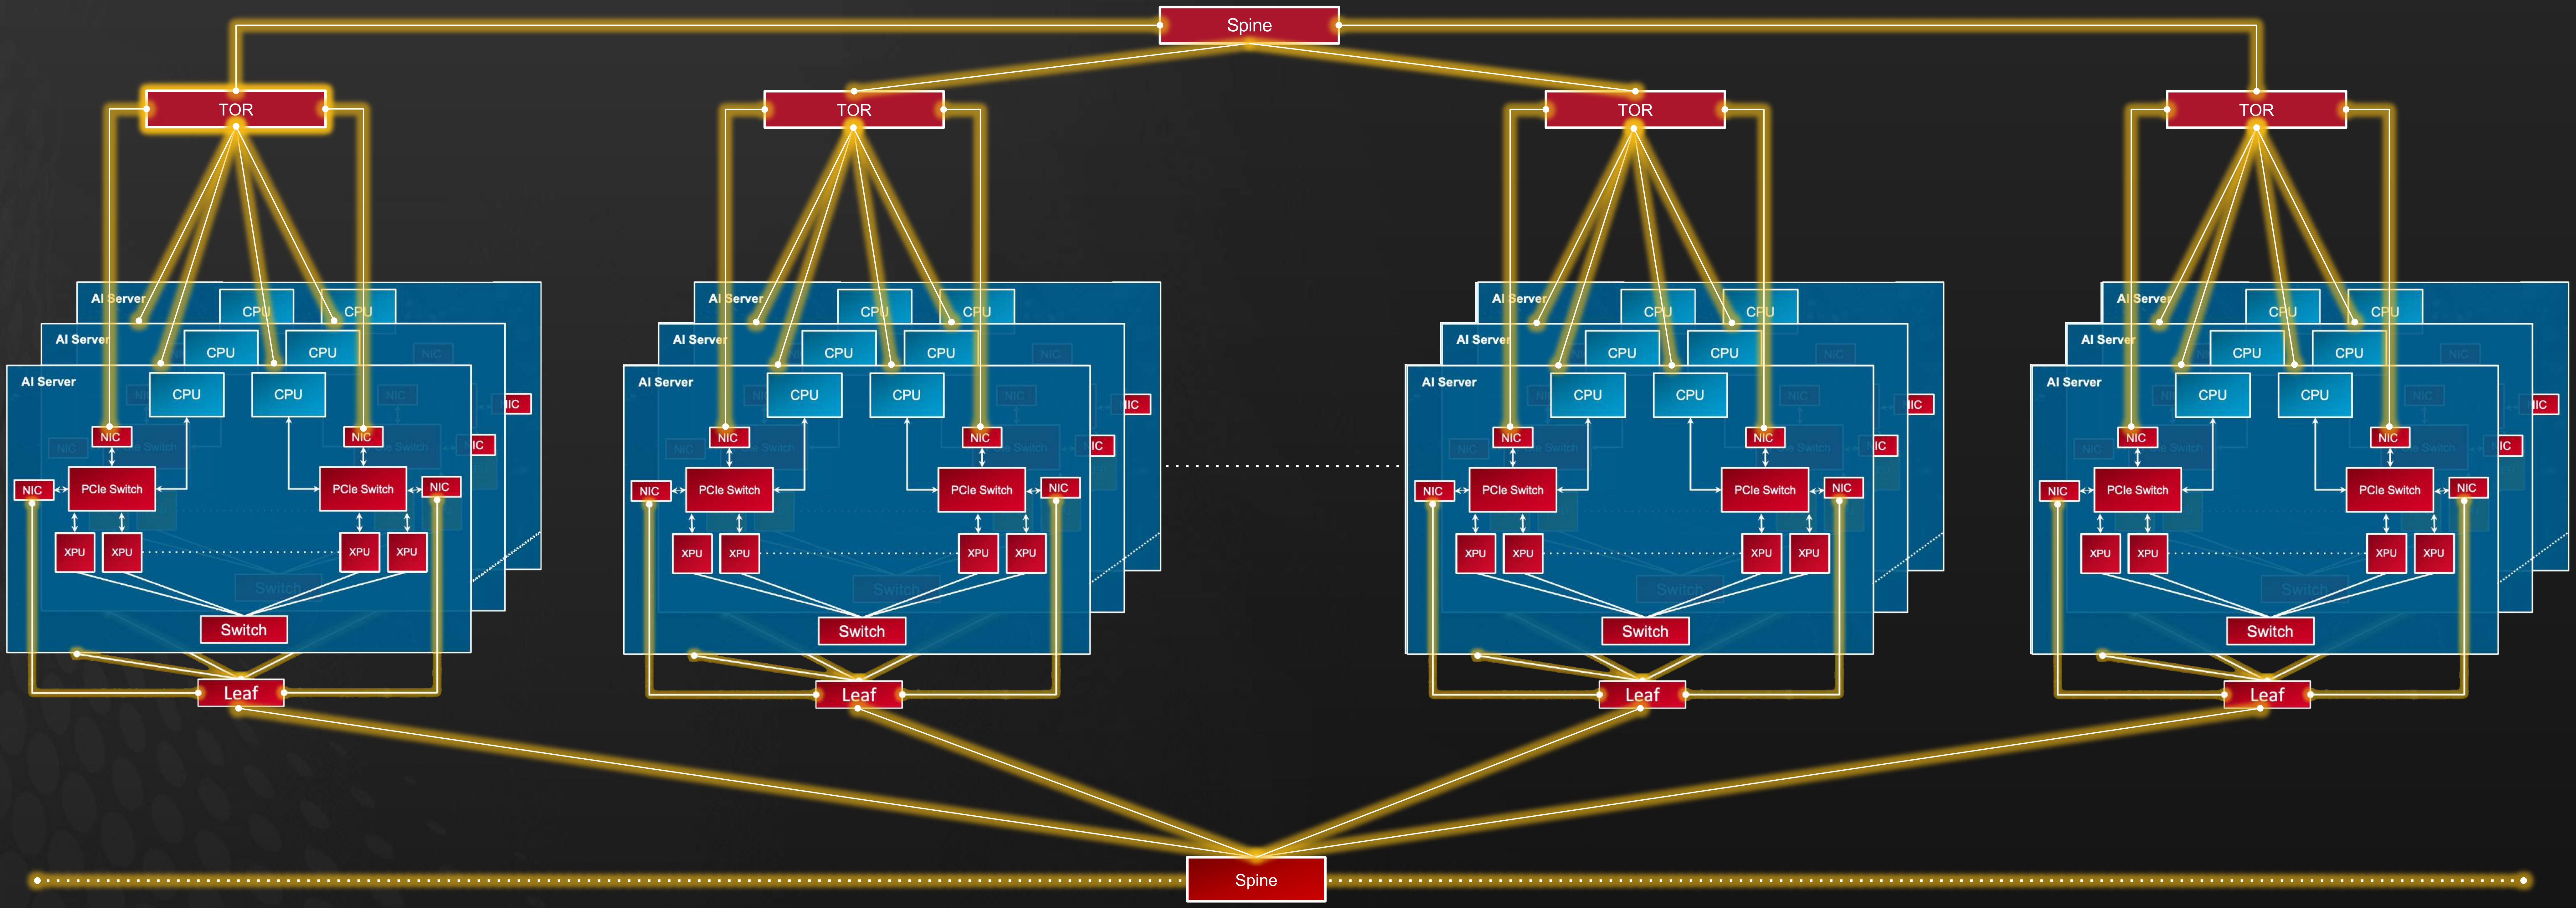

OPEN // SCALABLE // POWER EFFICIENT

Optical Interconnect for AI

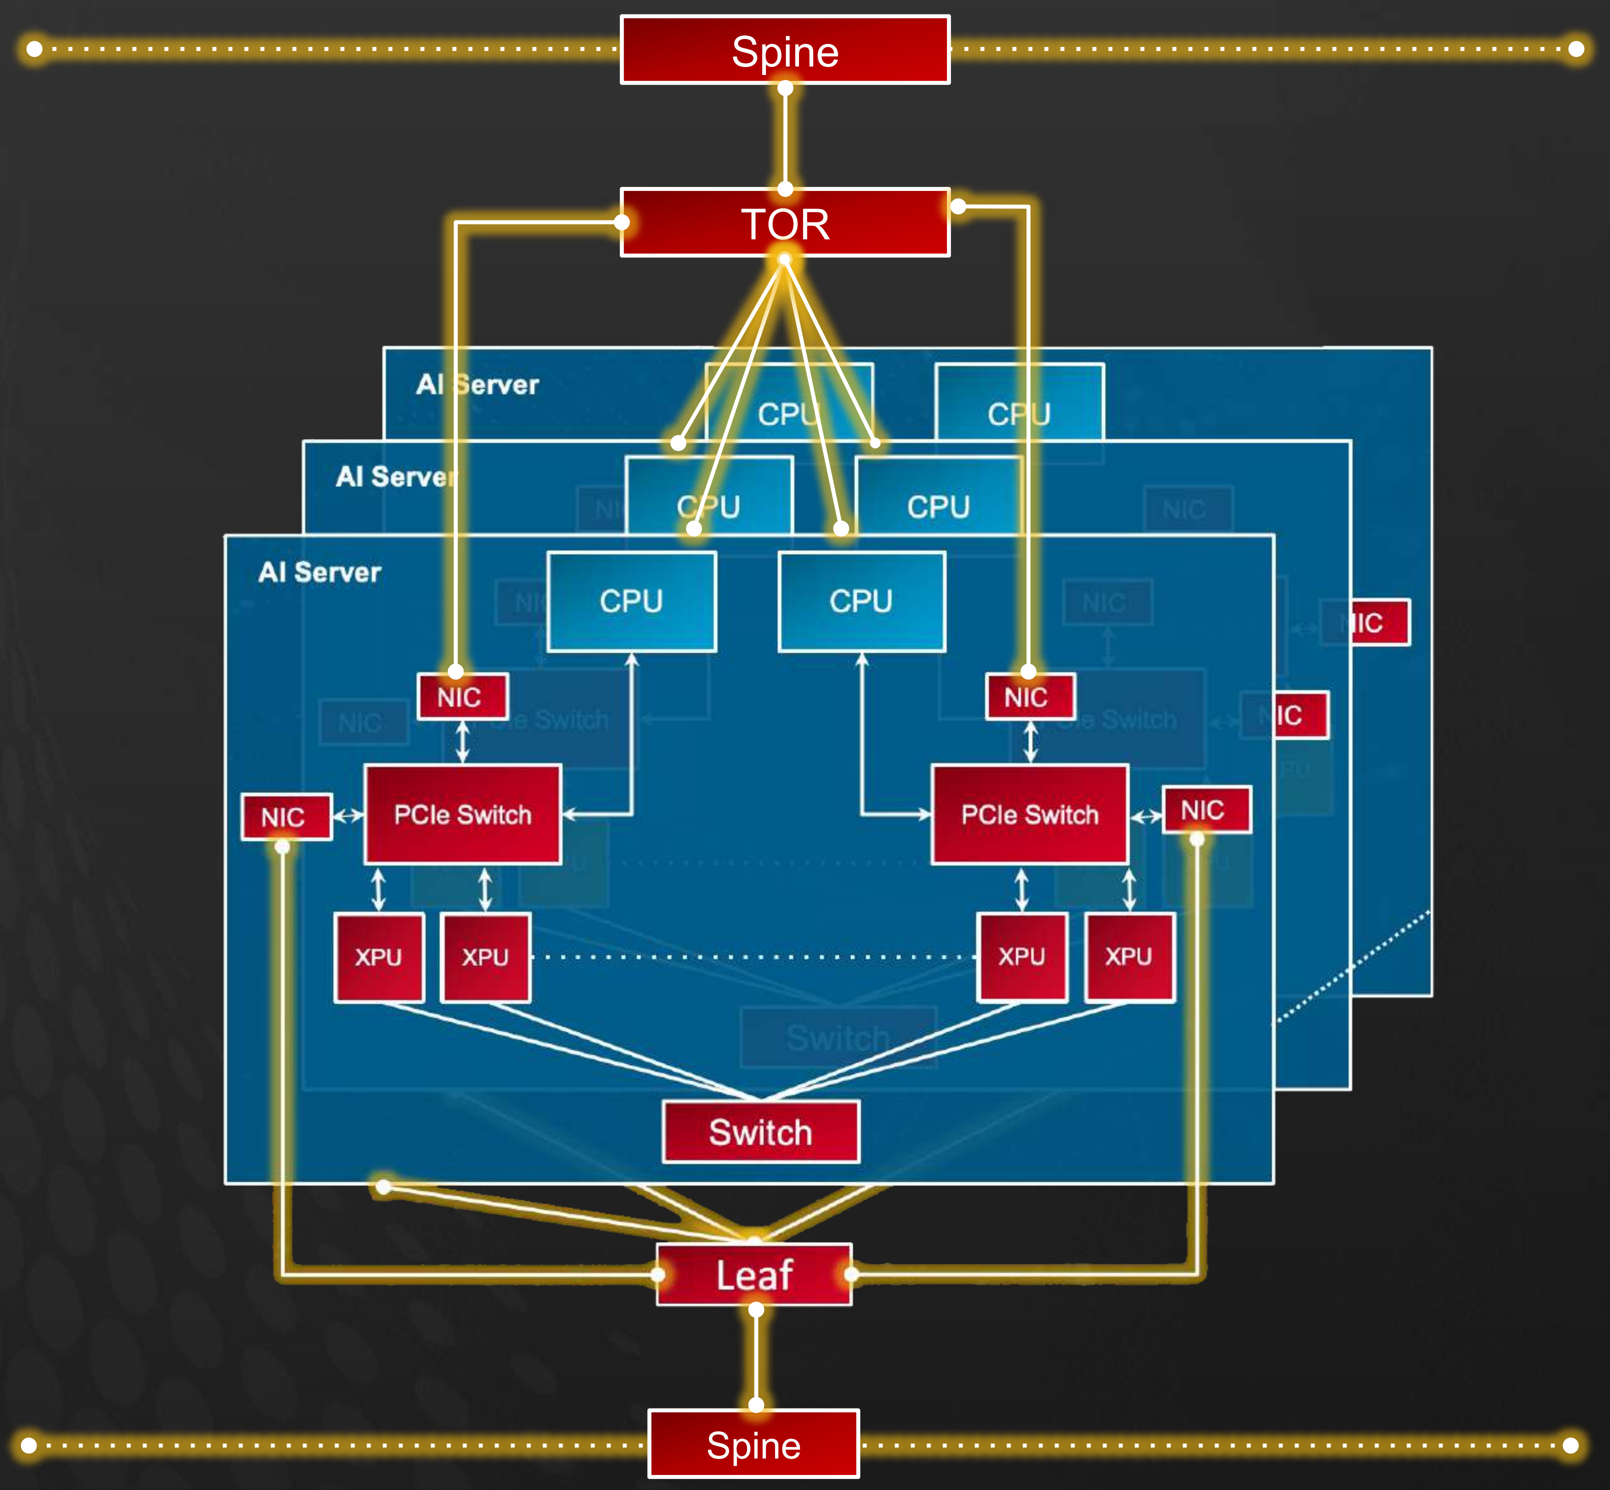

#1		**VCSEL**	Multi-mode, short reach
#1	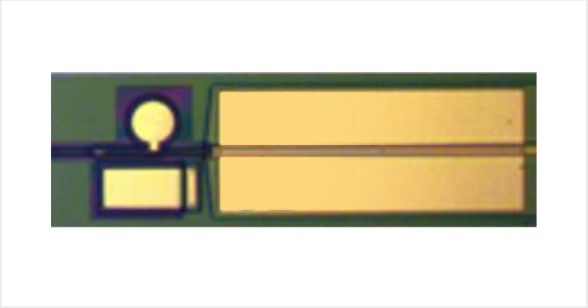	**EML**	Single-mode, long reach
1st	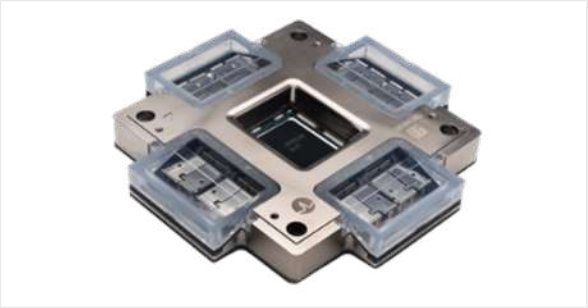	**CPO** (SiPho)	Power and cost leadership

Broadcom Optics Leadership

AT&T
2.5G EML and DML

AT&T
10G EML and DML

AVAGO TECHNOLOGIES
50G EML and DML

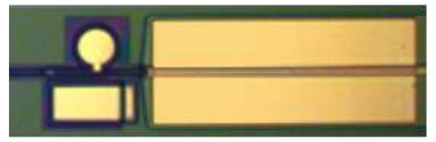

100G EML

50M+
High-Speed Communication Lasers Shipped Annually

10 Trillion
Device Service Hours in the field with < 1 FIT

hp
1G VCSEL

AVAGO TECHNOLOGIES
10G VCSEL

AVAGO TECHNOLOGIES
25G VCSEL

50G VCSEL

1998 2004 2013 2019 2023 2025

OPEN // SCALABLE // POWER EFFICIENT

Broadcom Optics Leadership

AT&T
2.5G EML and DML

AT&T
10G EML and DML

AVAGO TECHNOLOGIES
50G EML and DML

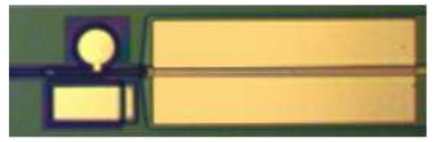
100G EML

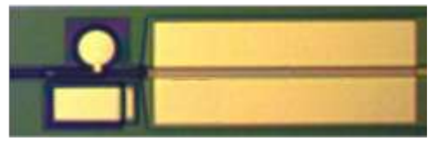
200G EML

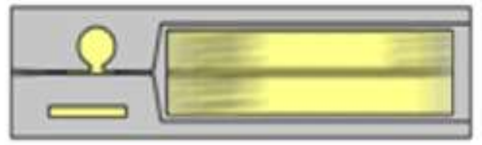
Next Gen EML

50M+
High-Speed Communication Lasers Shipped Annually

10 Trillion
Device Service Hours in the field with < 1 FIT


Bailly CPO

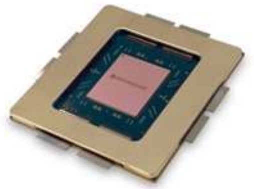
Next Gen CPO

hp
1G VCSEL

AVAGO TECHNOLOGIES
10G VCSEL

AVAGO TECHNOLOGIES
25G VCSEL


50G VCSEL


100G VCSEL


200G VCSEL

1998 | 2004 | 2013 | 2019 | 2023 | 2025

Why Co-Packaged Optics (CPO)?

Lowest Cost/Bit
Massive reduction in components and interconnects

Power/Performance Leadership
Eliminate electrical interconnect power

Increased Reliability
Silicon integration with pluggable Laser

POWER COMPARISON

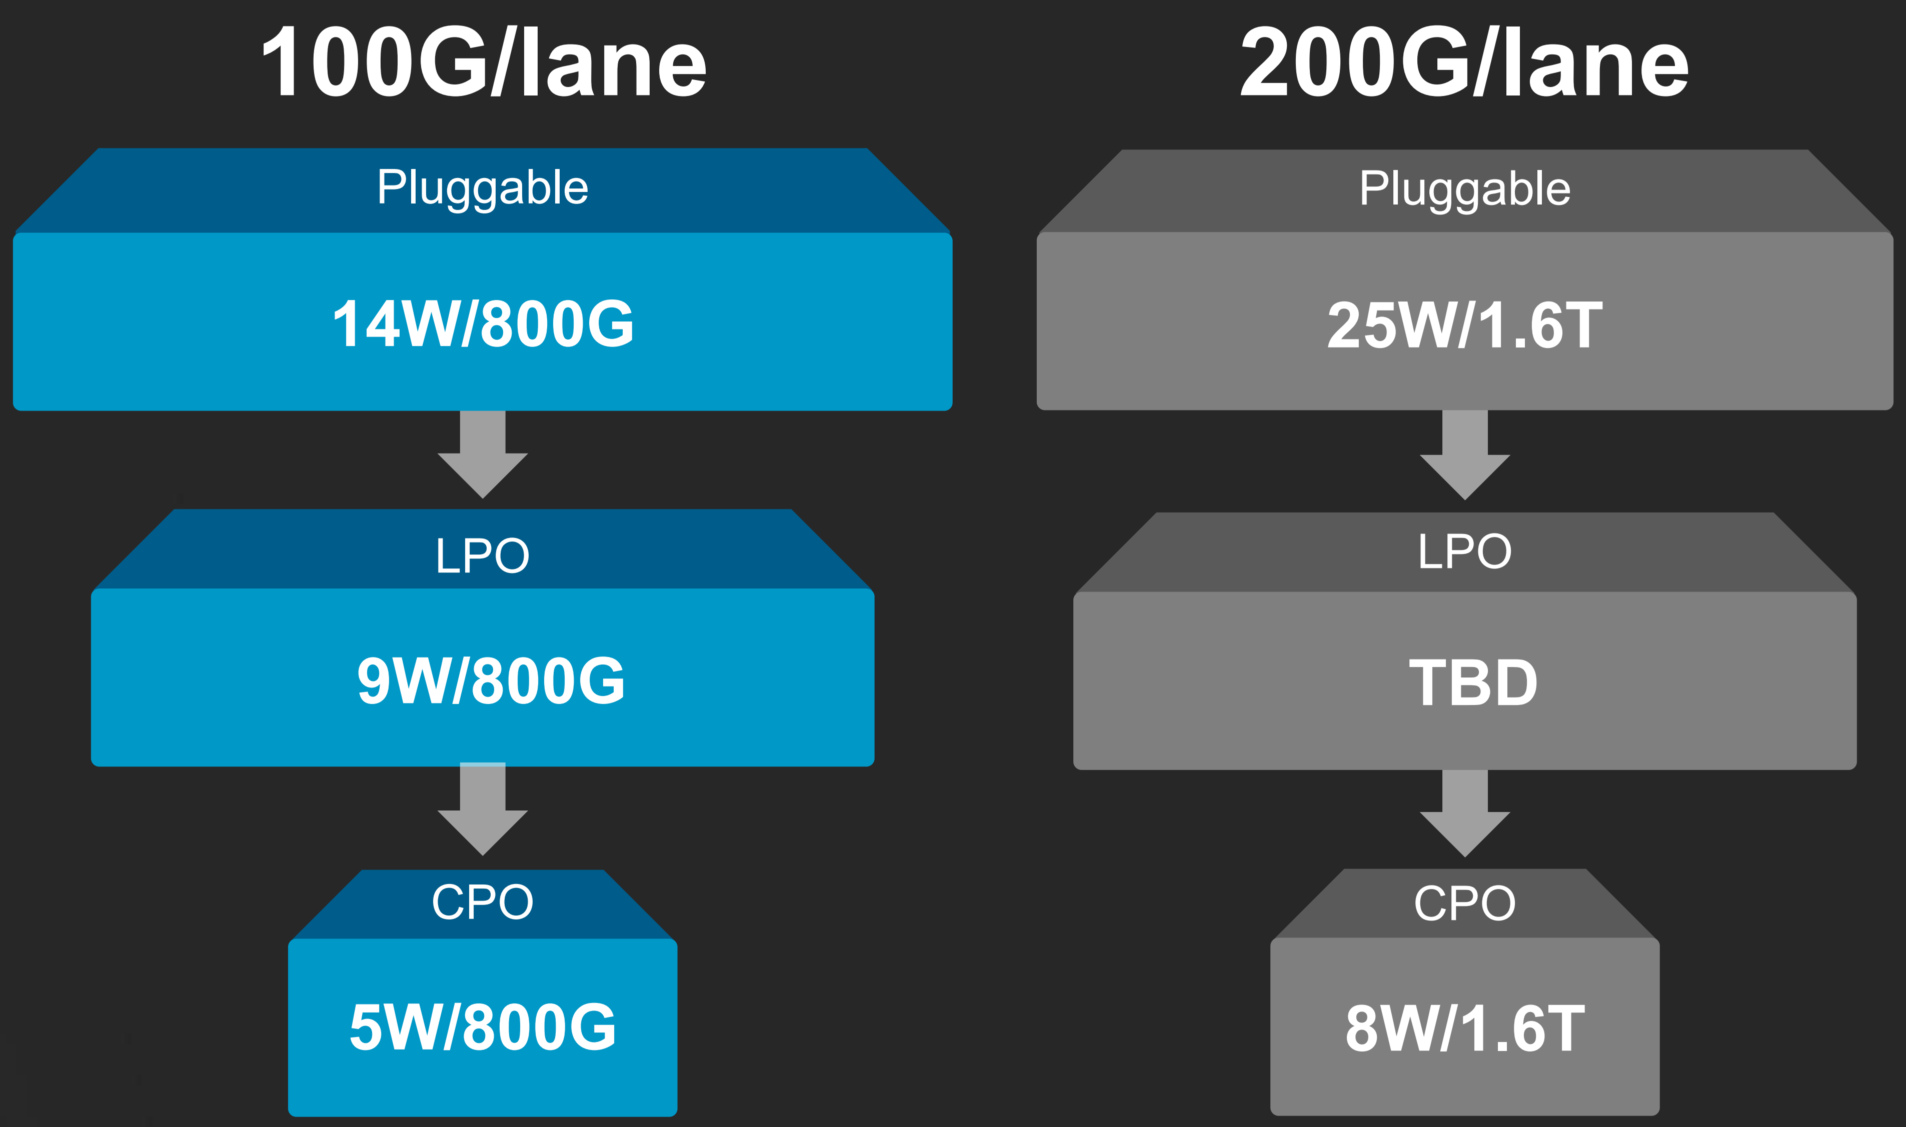

	100G/lane	200G/lane
Pluggable	14W/800G	25W/1.6T
LPO	9W/800G	TBD
CPO	5W/800G	8W/1.6T

The Journey to Integration at Scale

Conventional Module Design

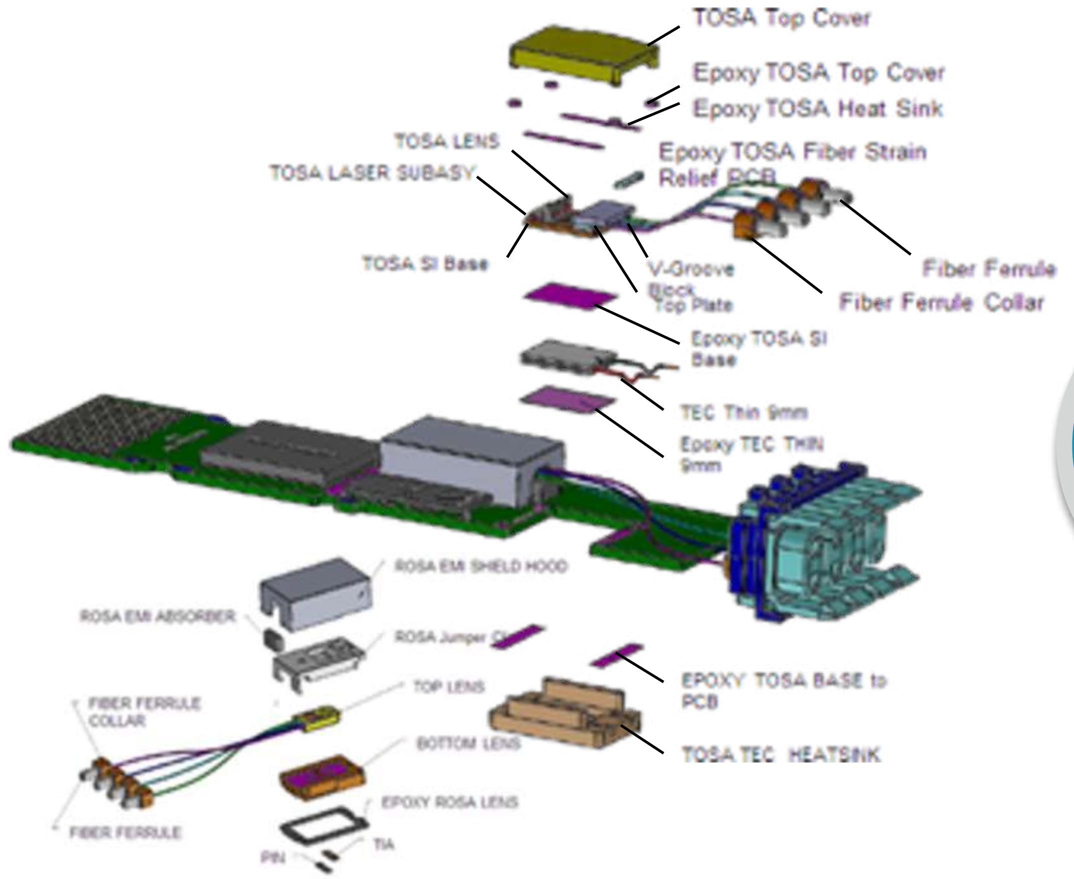

Engineering and manufacturing limits to scale

Integrated SiPh Module Design

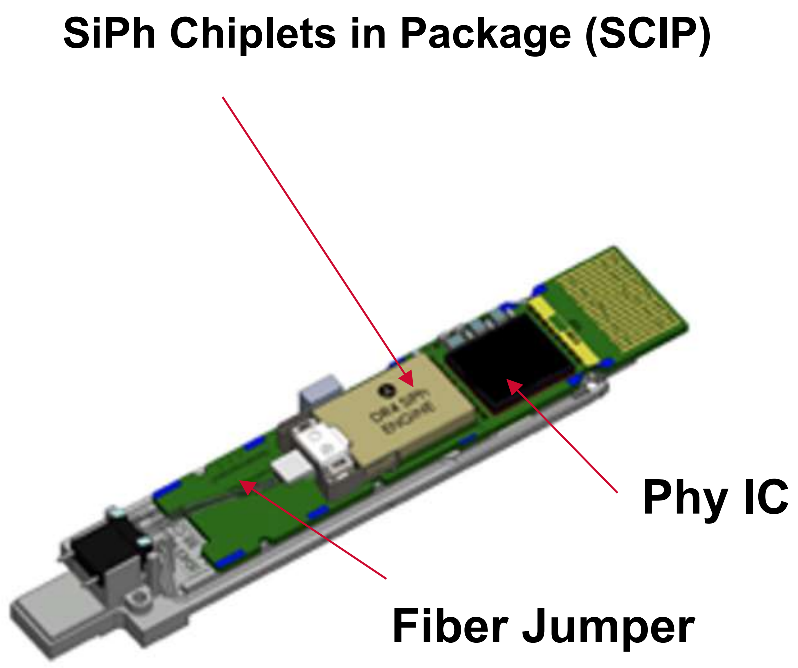

30% fewer piece parts

Module Integration = First step to improved scale

Co-Packaged Optics

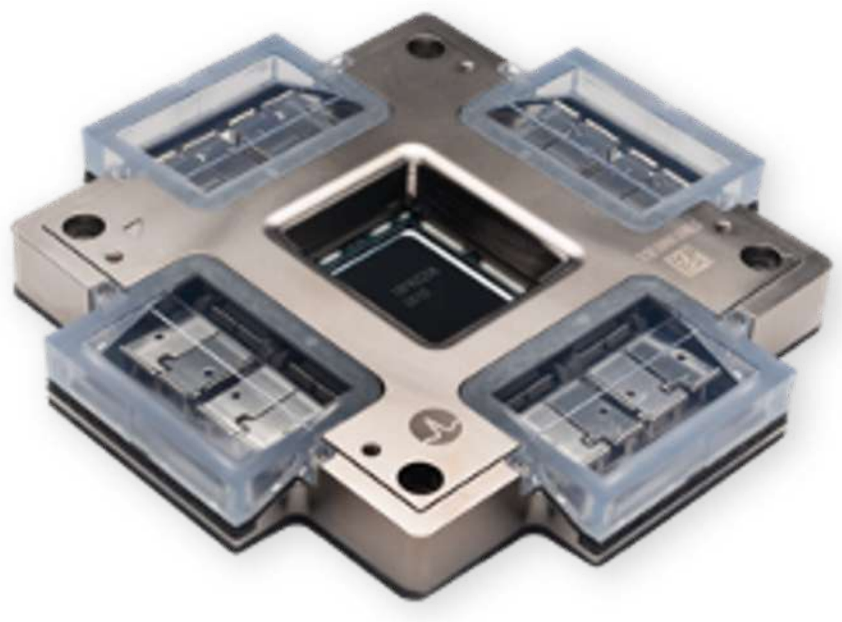

Highly Integrated Optical Engines (up to 6.4T)

FIRST to Productize High-Density Photonic Integration

CPO Applications: Ethernet Switch, XPU Interconnect

Ethernet Switch

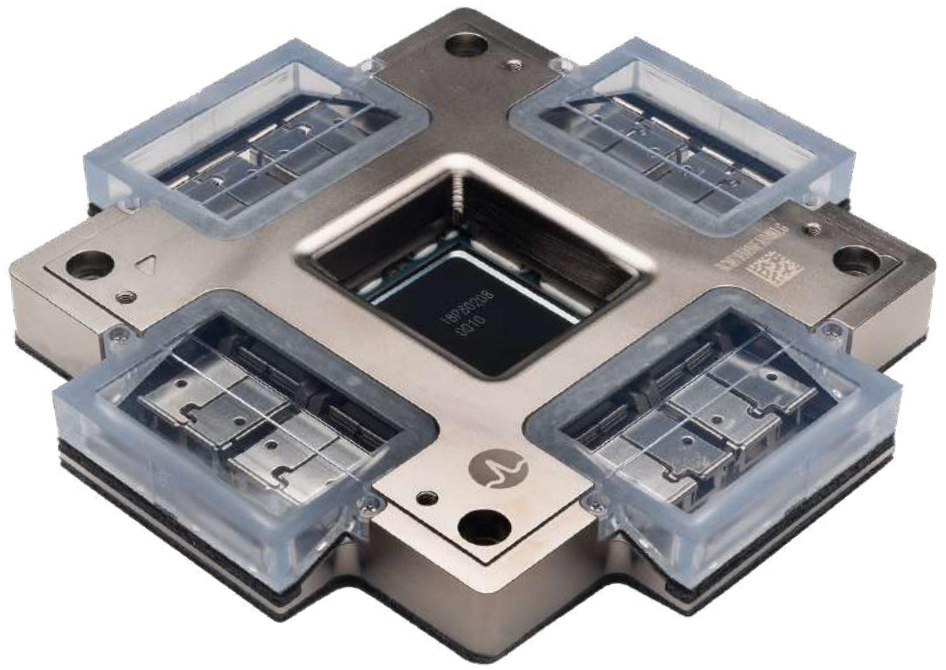

More than 50Tbps optical connectivity co-packaged with an Ethernet Switch

XPU Interconnect

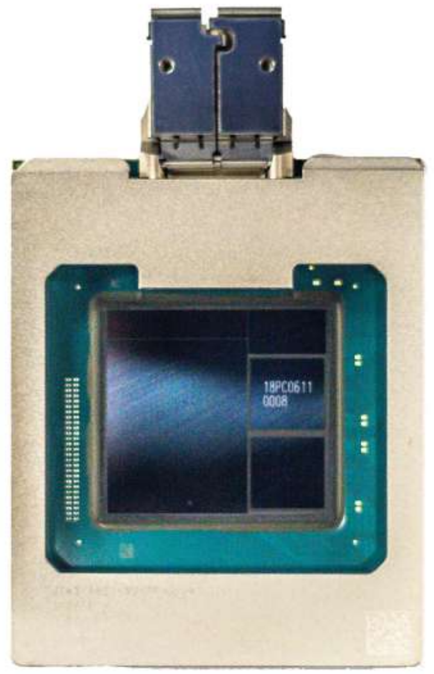

Up to 50Tbps optical connectivity co-packaged with a XPU

OPEN // SCALABLE // POWER EFFICIENT

51.2T TH5-Bailly CPO available today

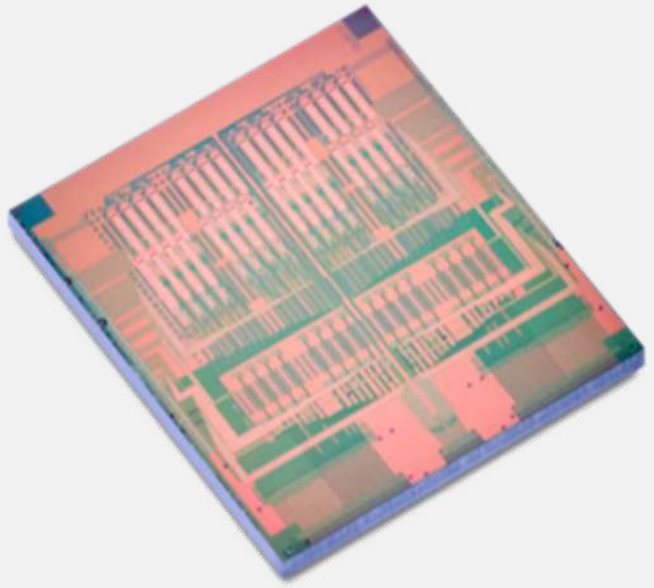

64 Channel FR4 PIC with Integrated Mux/Demux

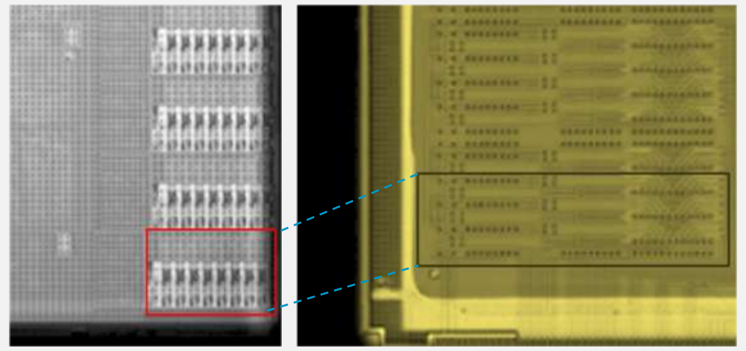

64 Channel EIC with CMOS Driver and TIA

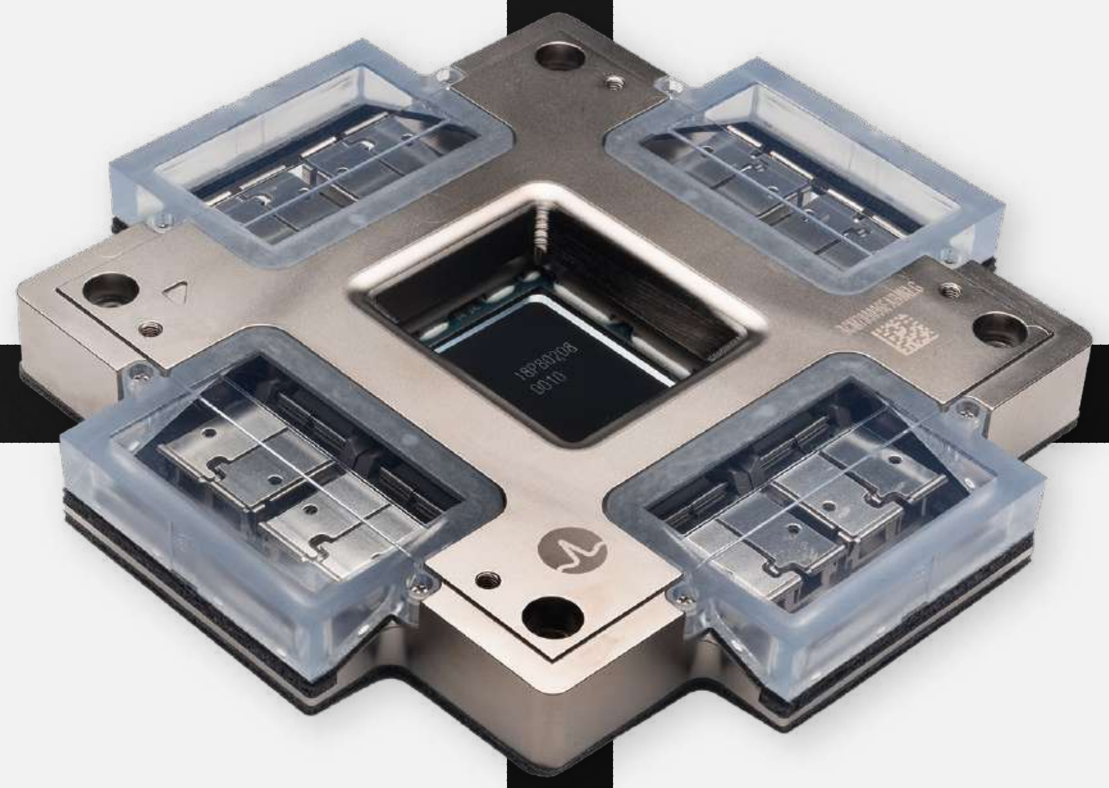

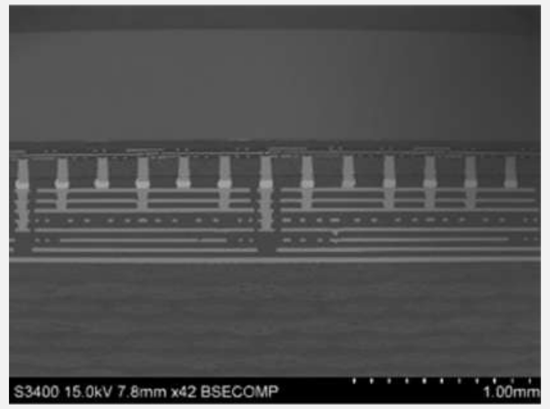

Advanced Packaging

ODM System Integration

Optical Interconnects

Key Points

Industry leadership on **100G → 200G VCSEL** and **EML**

First commercial **CPO** with pluggable lasers

Unmatched **>70%** power reduction* and **>30%** cost savings

*Source: Broadcom Internal Estimates

ENABLING AI Infrastructure

Foundational Technology for AI Interconnect

Vijay Janapaty

General Manager
Physical Layer Products
Broadcom

OPEN // SCALABLE // POWER EFFICIENT

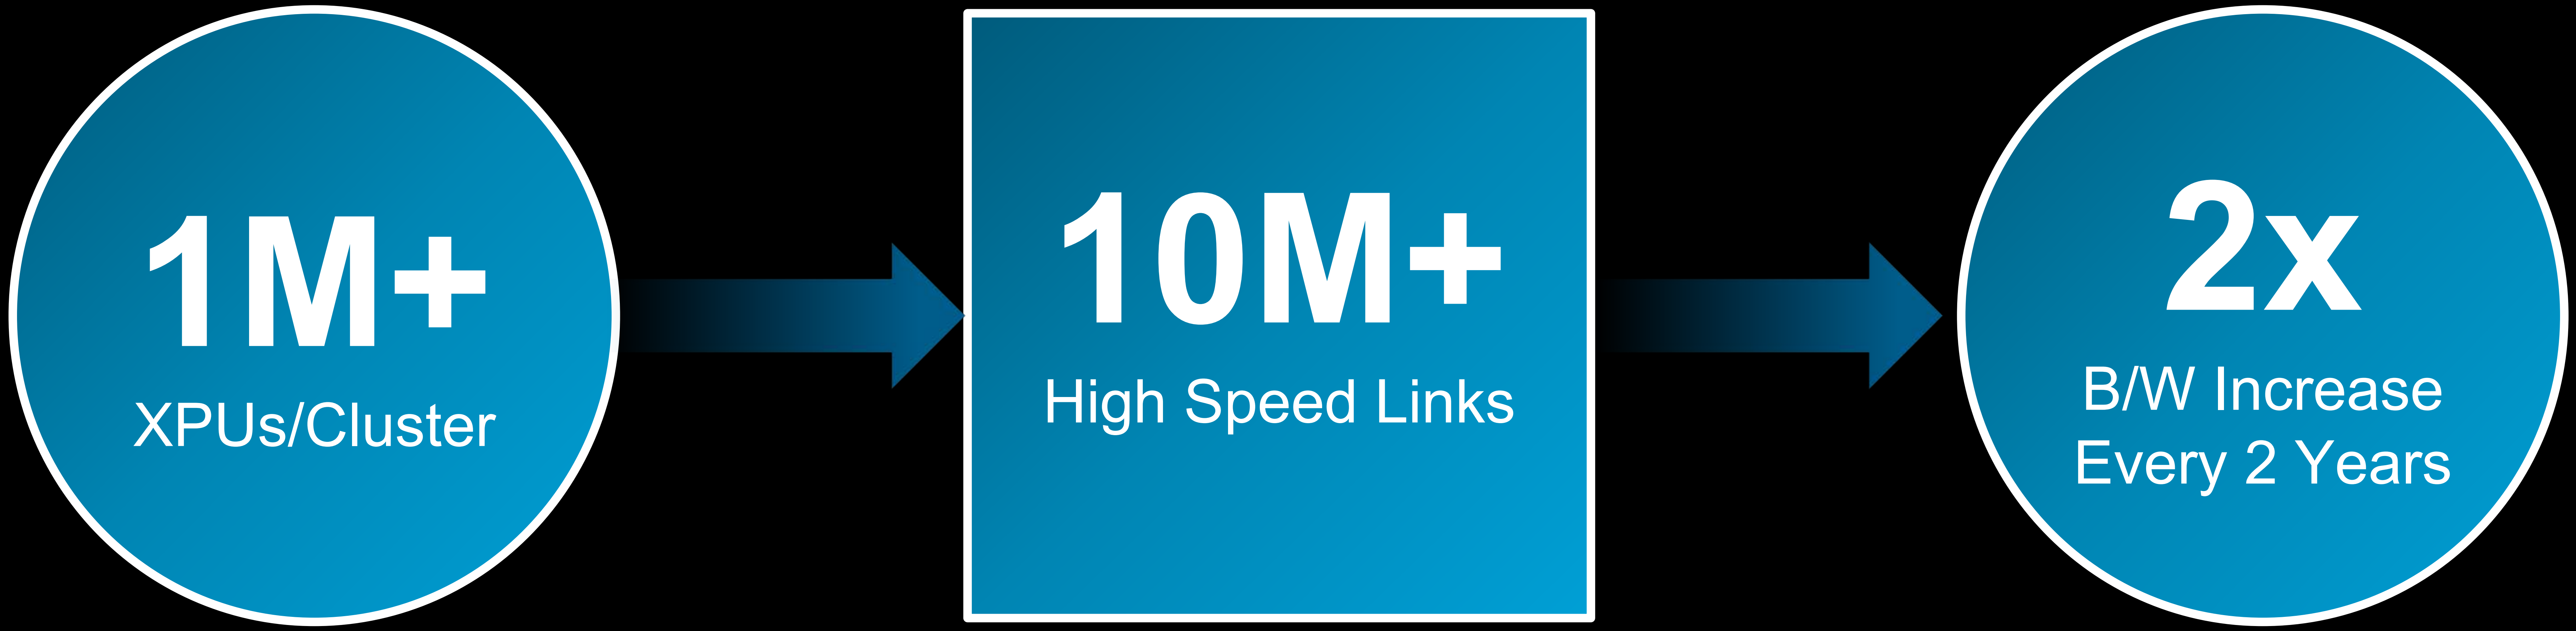

1M+
XPUs/Cluster

10M+
High Speed Links

2x
B/W Increase
Every 2 Years

High Speed Links : 2nd Highest Power & Cost in AI Cluster

10M+

High Speed Links

Backplane /DAC

Reach: < 5m
Lowest Power
Lowest Cost

SerDes

Optics

Reach: 10m–10 Km
Highest Power
Highest Cost

SerDes

DSPs

Source: Broadcom Internal Estimates

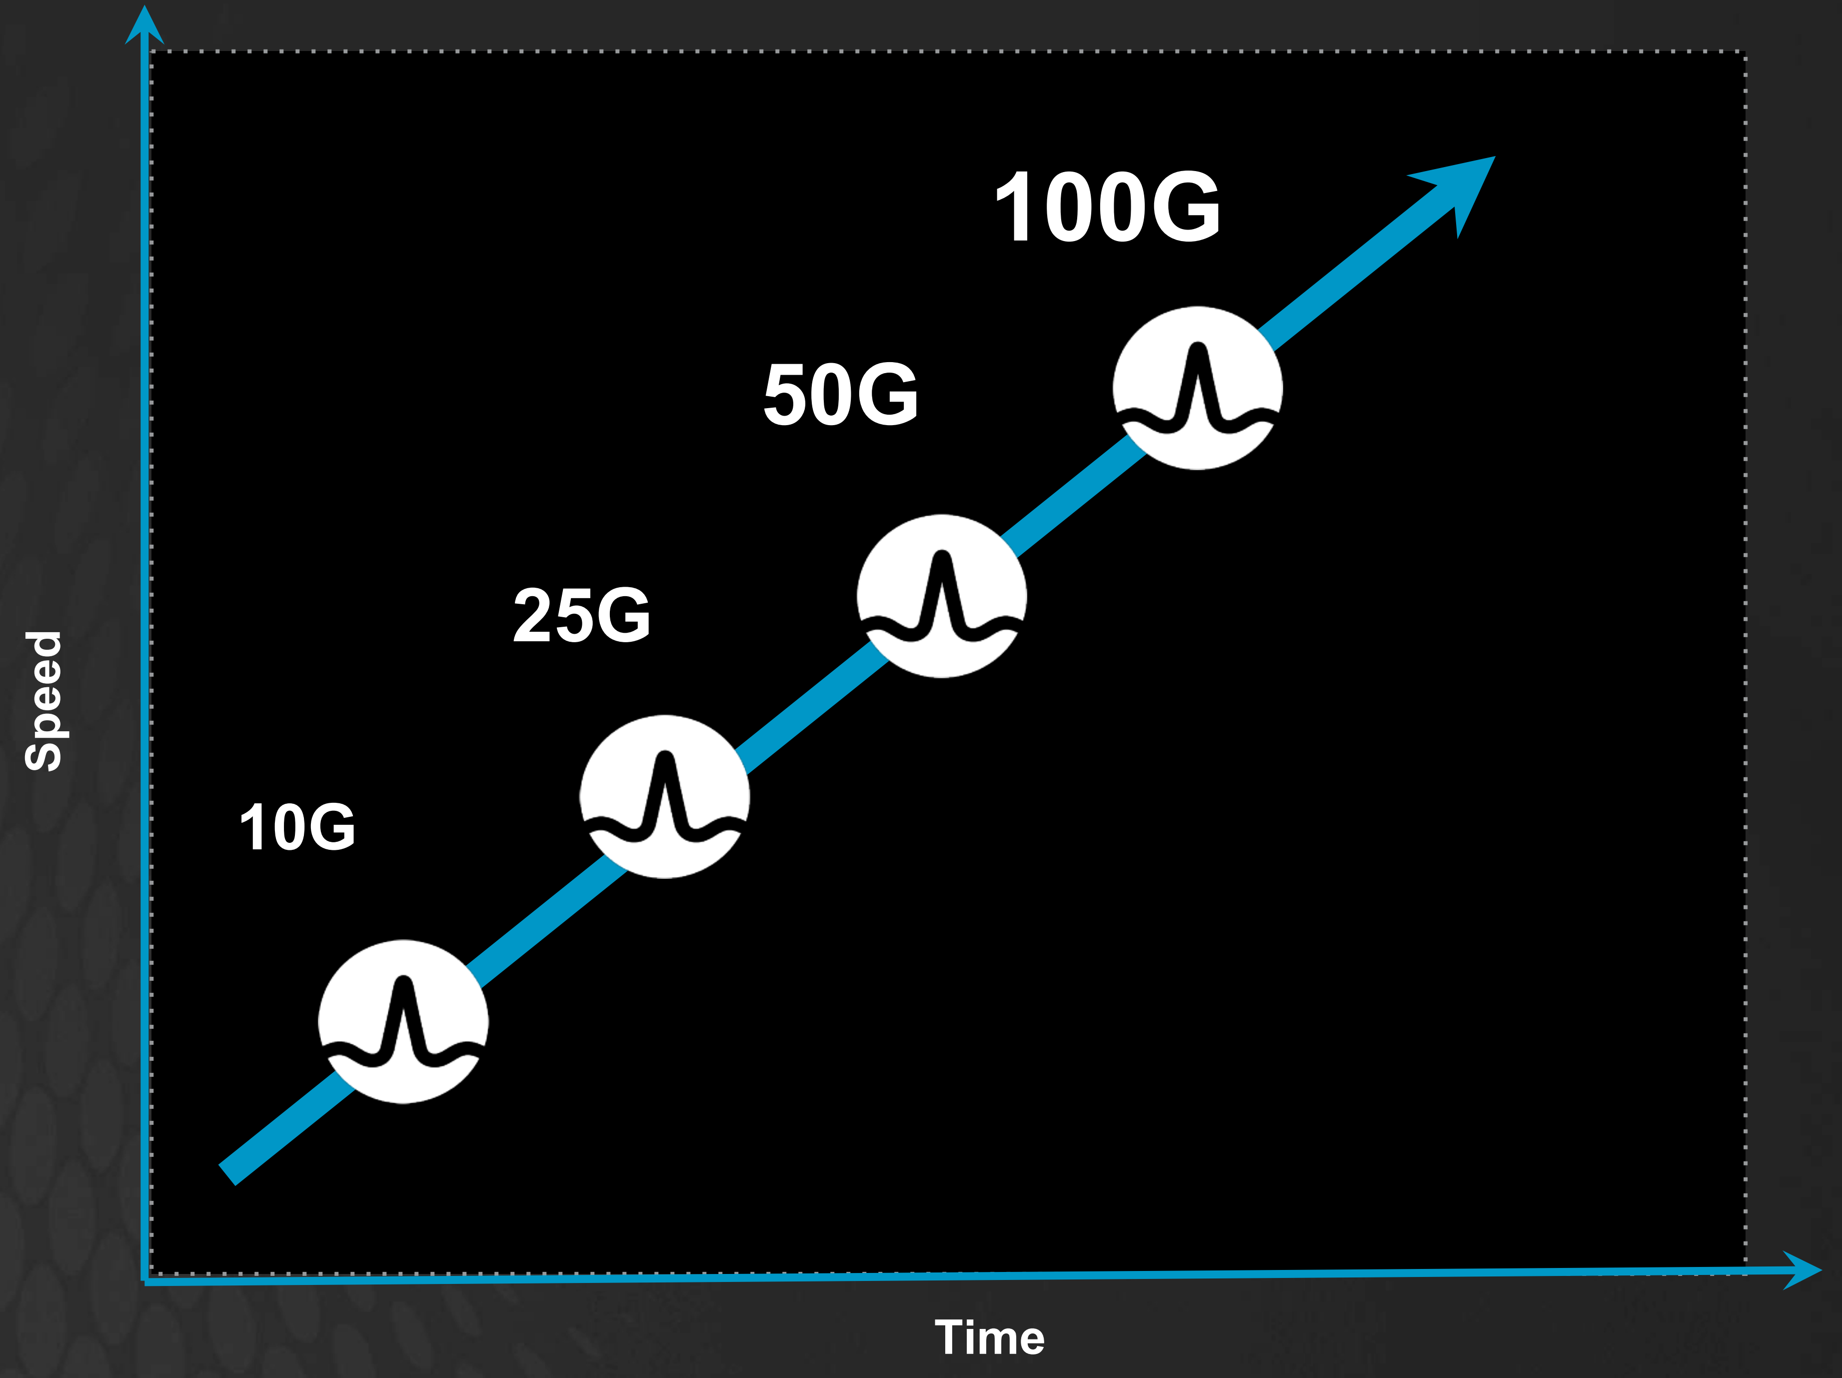

- High-Speed **Analog-to-Digital Conversions**
- High-Speed Custom **Digital-Signal Processing**
- Leading-Edge **Process Technology Nodes**
- **Deployment** in 100s of Systems

#1 SerDes for 4+ Generations

Industry-Leading 5 nm 112 Gb/s SerDes: Peregrine

45+ dB Backplanes

Higher Fanout

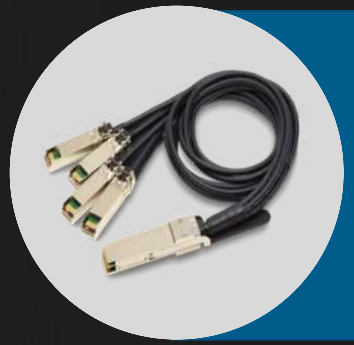

4+m DAC Cables

Lowest Power and Cost for Intra-Rack & Inter-Rack

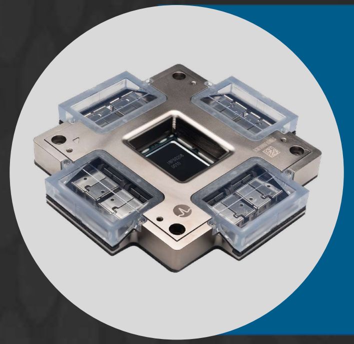

New

Drives CPO Natively

50+% Power Reduction

New

Drives LDO Natively

33+% Power Reduction

• • • Switches XPUs

Same SerDes Benefits

Across All Broadcom Products

NICs DSPs • • •

OPEN // SCALABLE // POWER EFFICIENT

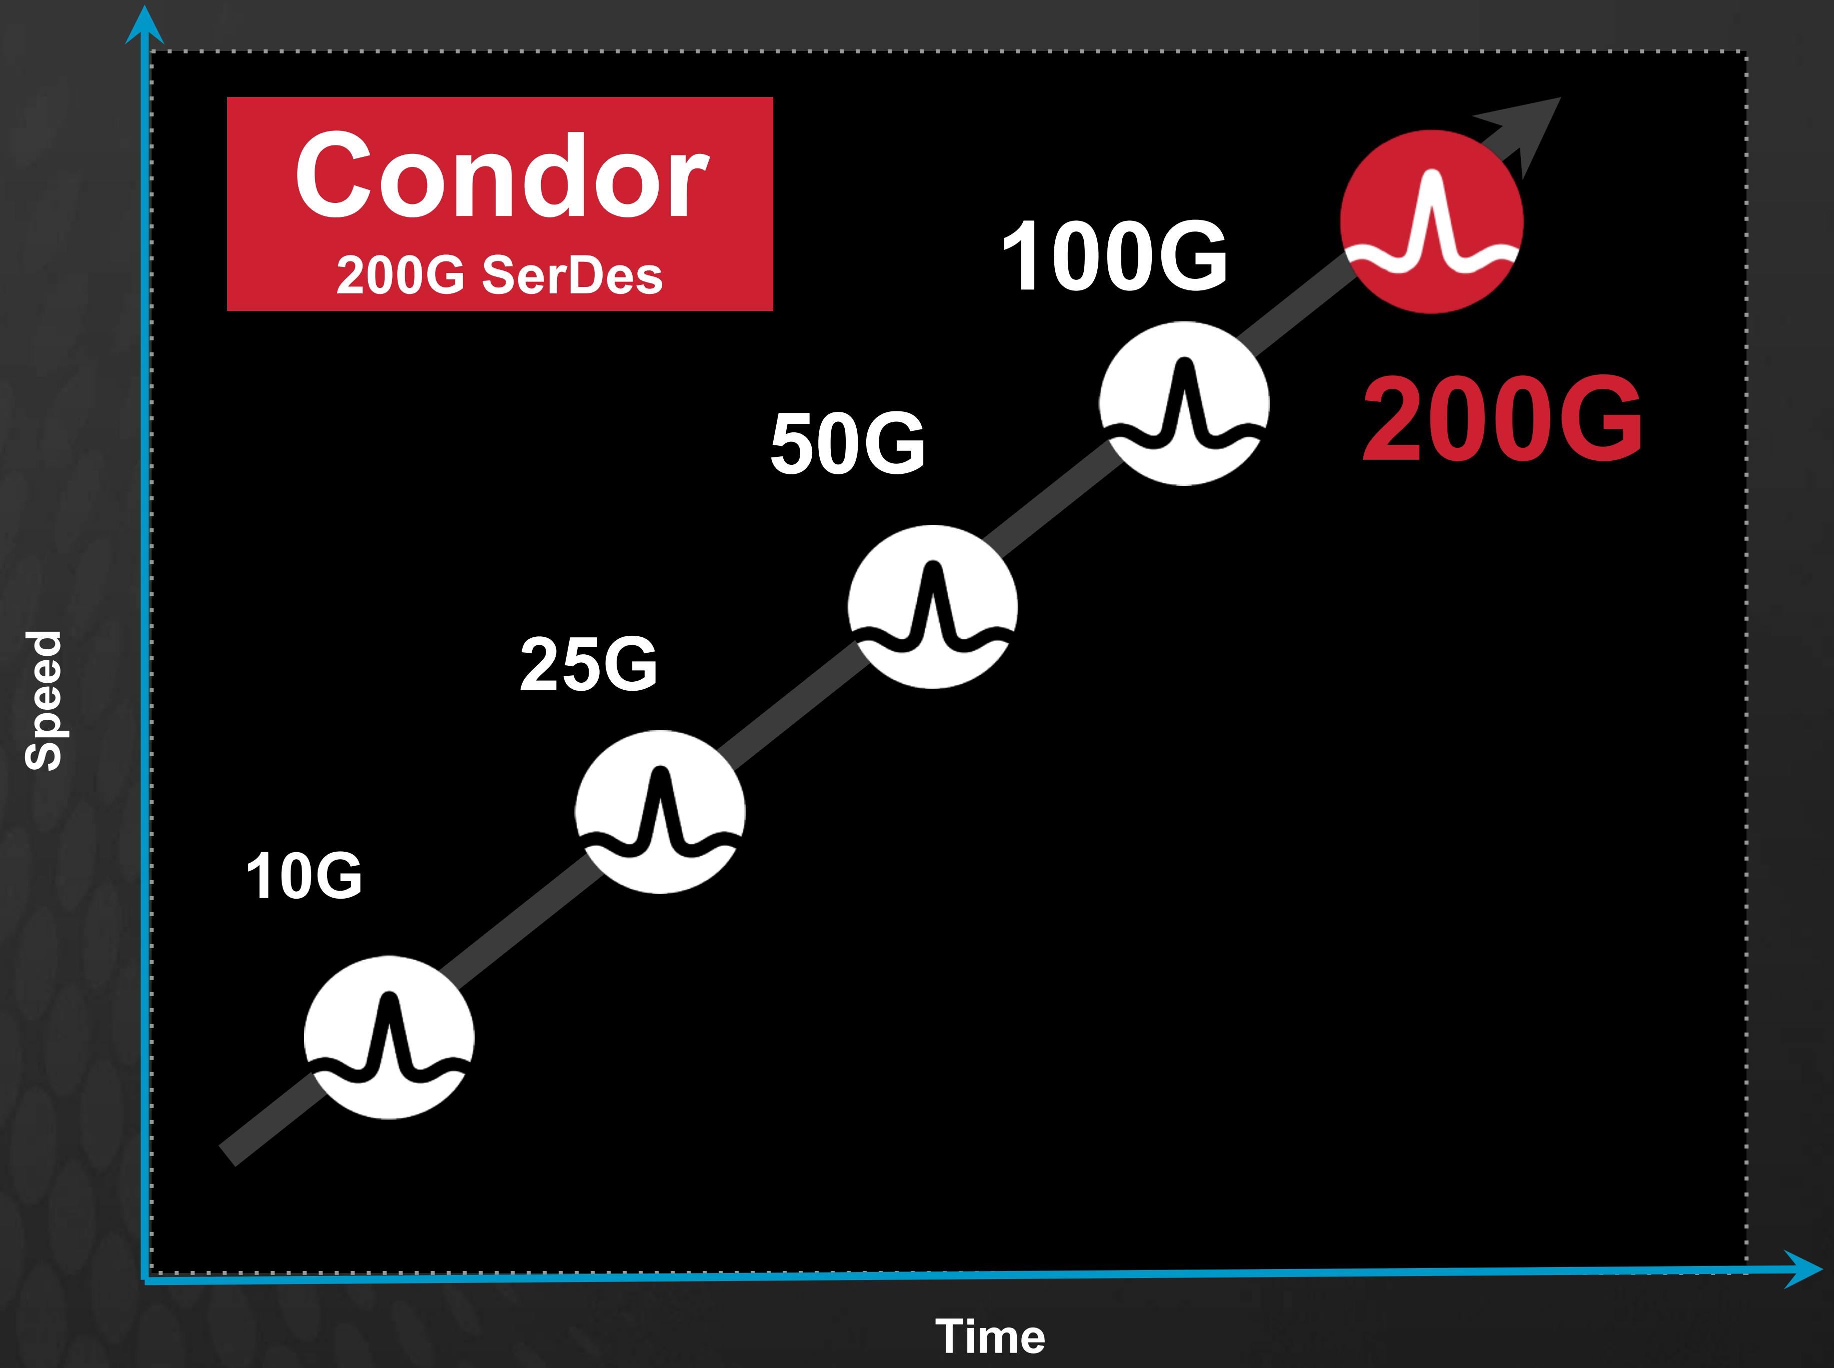


45+ dB Backplanes

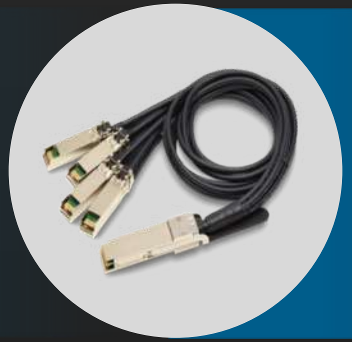
2+m DAC Cables

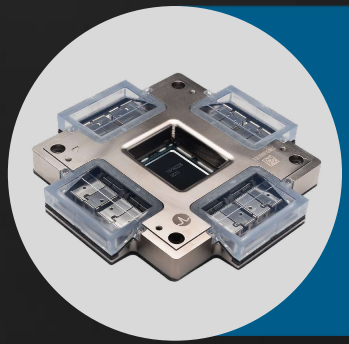
Drives CPO

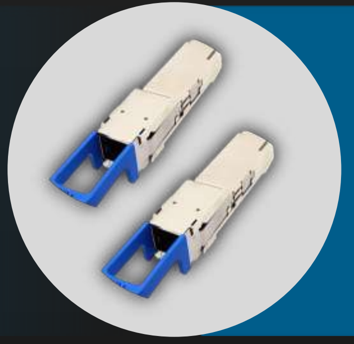
Drives LDO

Available Now

Industry’s Best 200 Gb/s SerDes in 3 nm

Industry's First: CMOS DSP + Driver + TIA for 100 G/λ PAM4 Optics

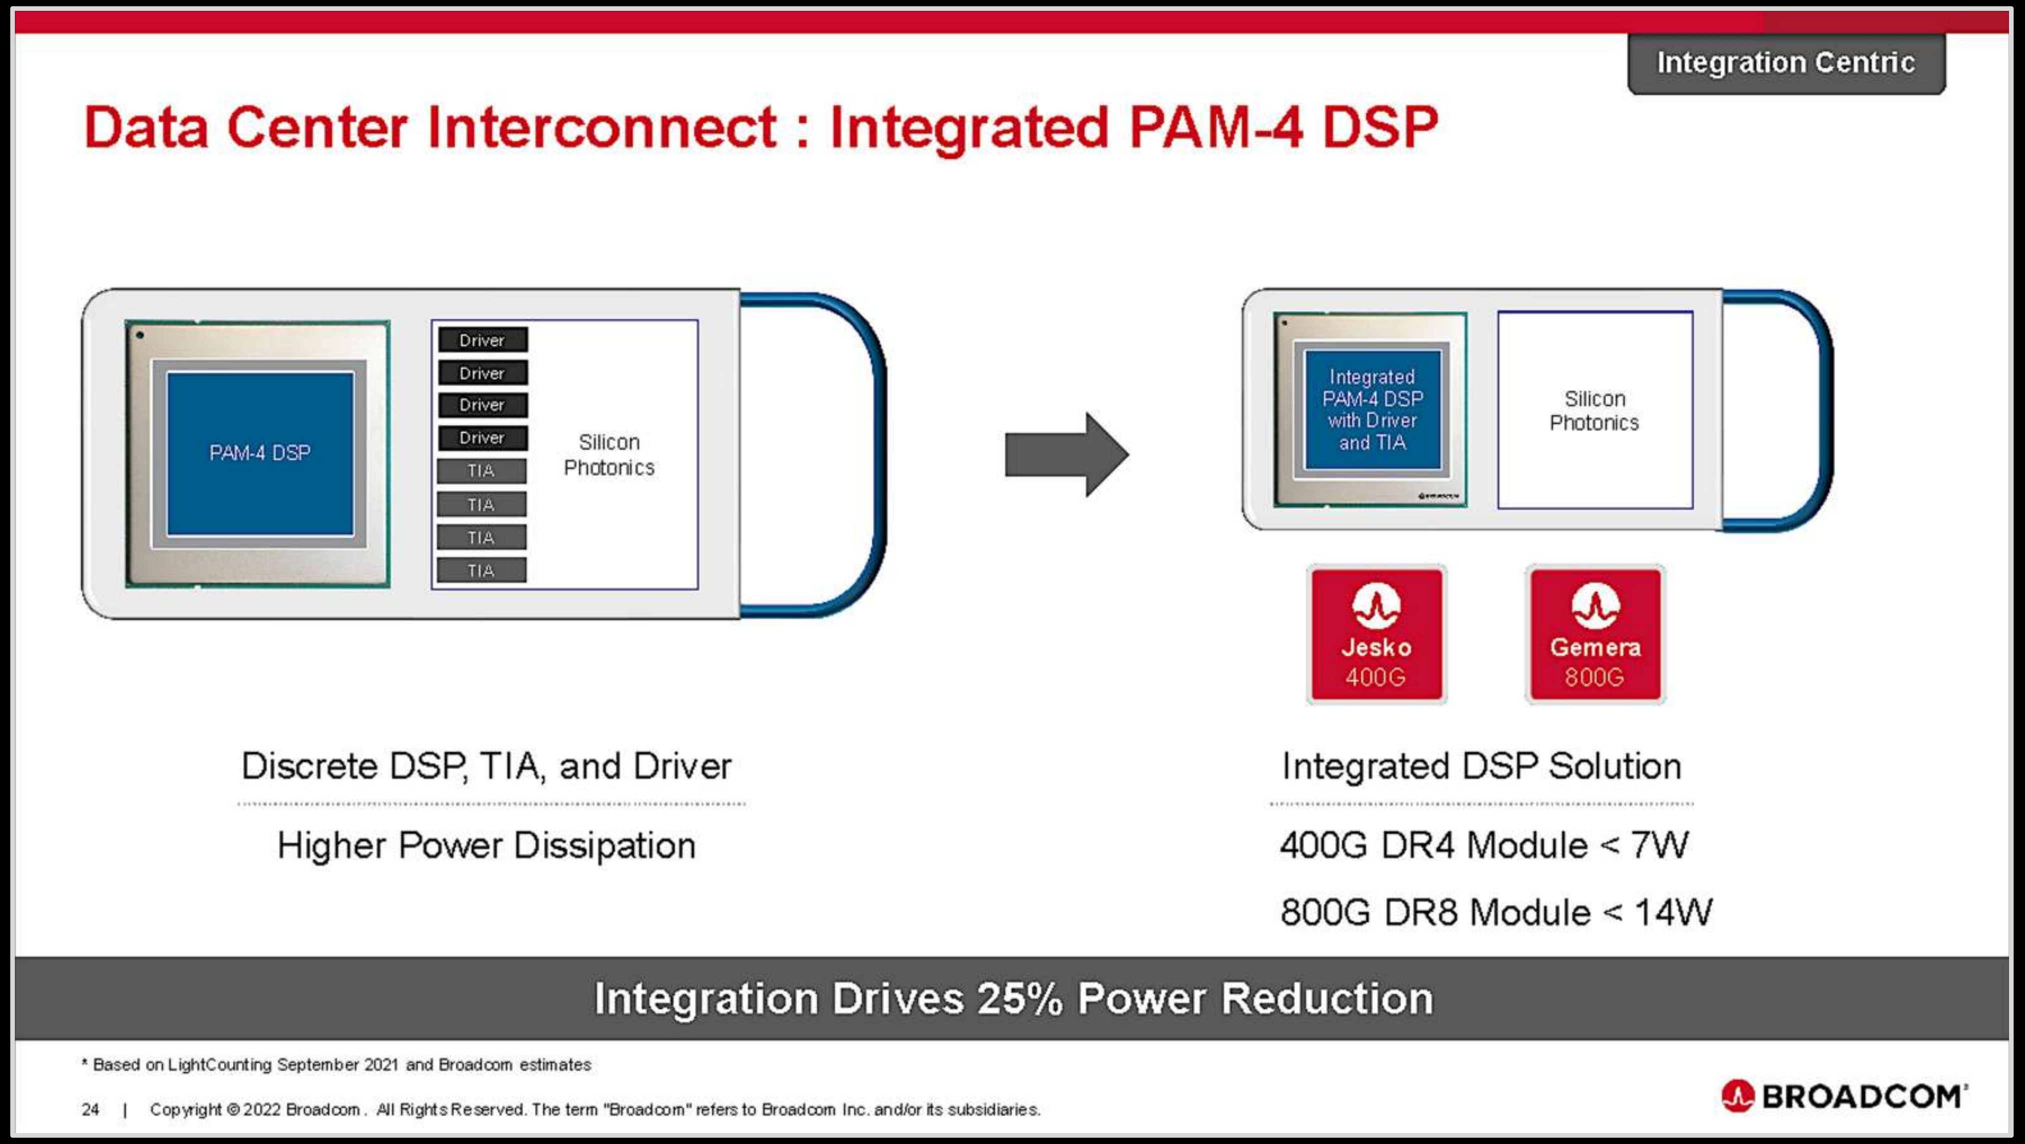

Integration Centric

Data Center Interconnect : Integrated PAM-4 DSP

Discrete DSP, TIA, and Driver

Higher Power Dissipation

Integrated DSP Solution

400G DR4 Module < 7W

800G DR8 Module < 14W

Integration Drives 25% Power Reduction

* Based on LightCounting September 2021 and Broadcom estimates

24 | Copyright © 2022 Broadcom. All Rights Reserved. The term "Broadcom" refers to Broadcom Inc. and/or its subsidiaries.

BROADCOM

Promise from 4-19-2022 Teach-in: CMOS DSP + Driver + TIA

7 nm **Jesko**

400G DR4 Modules
- < 7W Power

7 nm **Gemera**

800G DR8/SR8 Modules
- < 14W Power

5 nm **Cygnus**

800G DR8/SR8 Modules
- **< 11W Power**

In High-Volume Production

Broadcom iDSPs ⇒ Lowest Power/Cost MM/SM Pluggable Modules

Best-Performing 200 G/λ DSPs in the Industry

1.6T DR8 Modules

- 16x100G to 8x200G
- 212 Gb/s FEC Bypass
- 226 Gb/s w/ Inner FEC
- Integrated Driver

1.6T DR8 Modules

- 8x200G to 8x200G
- 212 Gb/s FEC Bypass
- 226 Gb/s w/ Inner FEC
- Integrated Driver

Integrated DSPs for 1.6T Generation

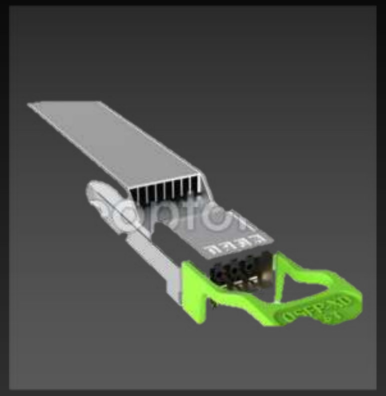

MMF

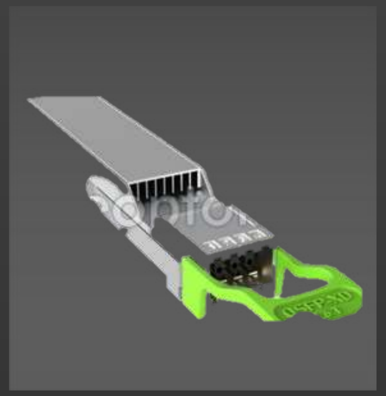

SMF

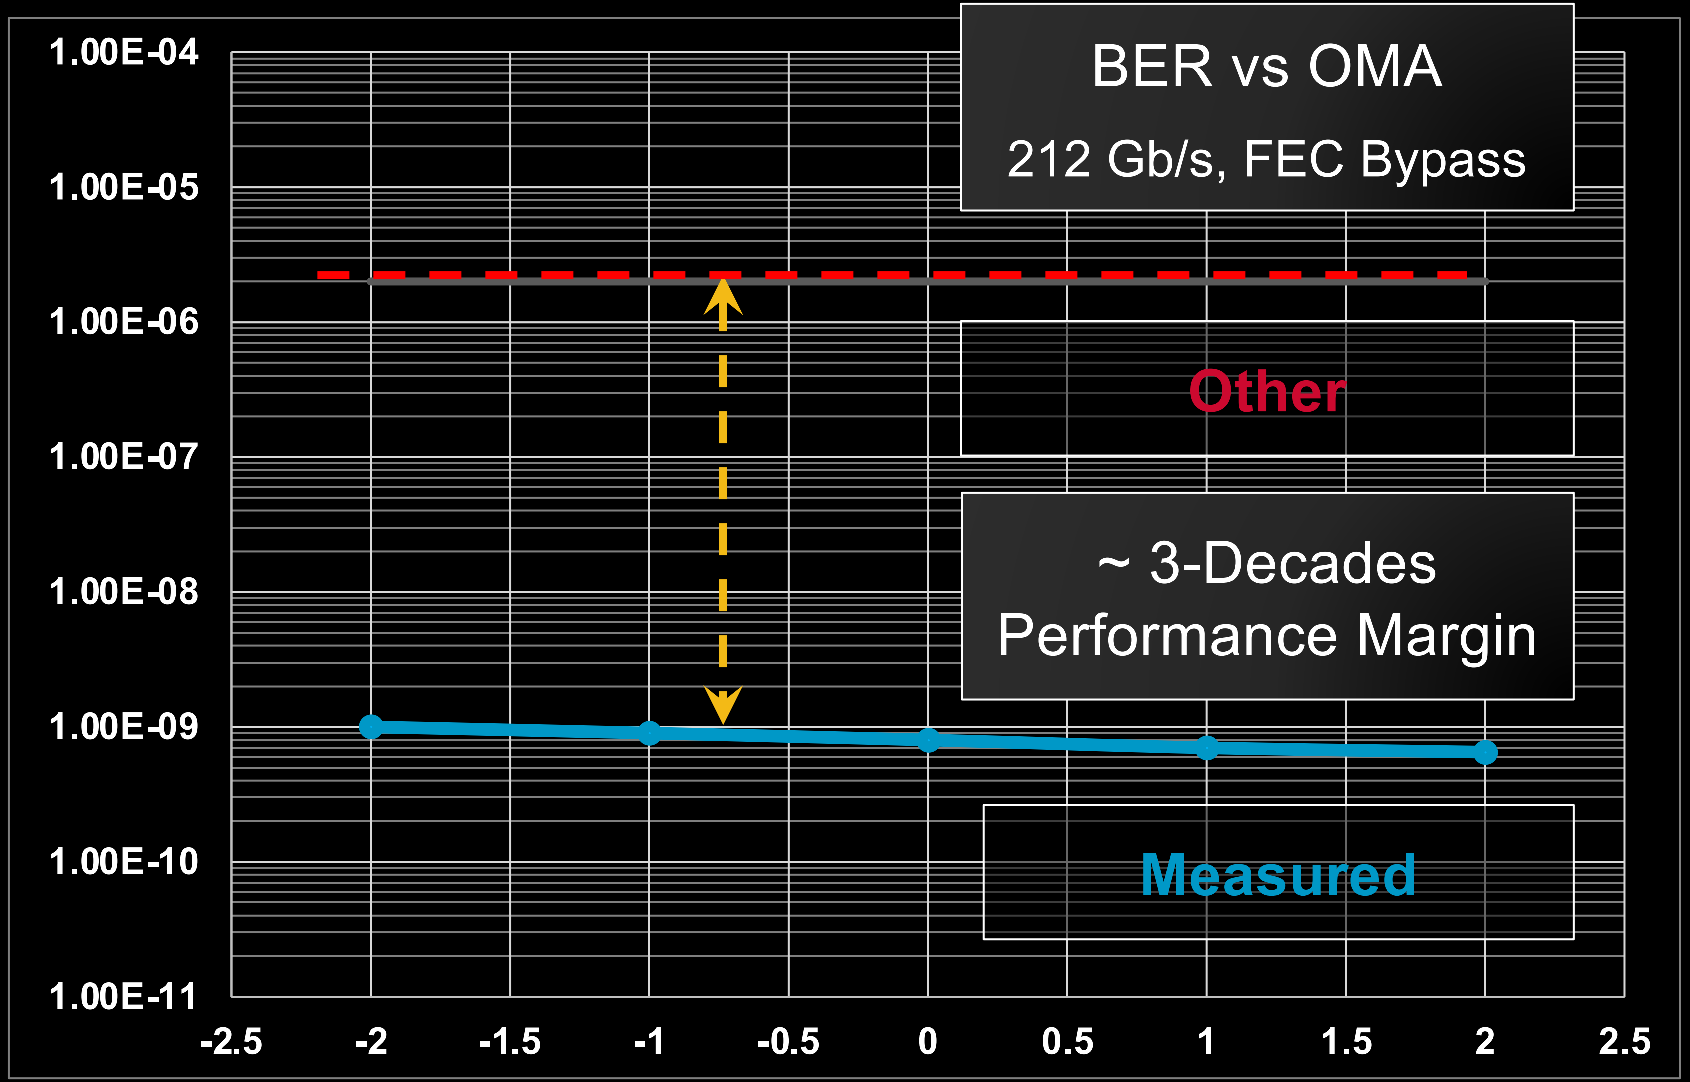

Performance Result from Partner 1.6T Module

200G iDSPs ⇒ Performance, Power and Latency Leadership

Foundation Technology

Key Points

100 Gb/s **BEST-IN-CLASS** Innovative Connectivity

200 Gb/s **LEADS** Open Ecosystem

DIFFERENTIATION from Analog + DSP Expertise, R&D Scale

OPEN // SCALABLE // POWER EFFICIENT

BROADCOM®

ENABLING AI Infrastructure

Custom Accelerators (XPUs)

Frank Ostojic

General Manager
Custom AI Product Division
Broadcom

OPEN // SCALABLE // POWER EFFICIENT

#1 in custom silicon for 10 years

Consumer AI Requires Custom Accelerators (XPUs)

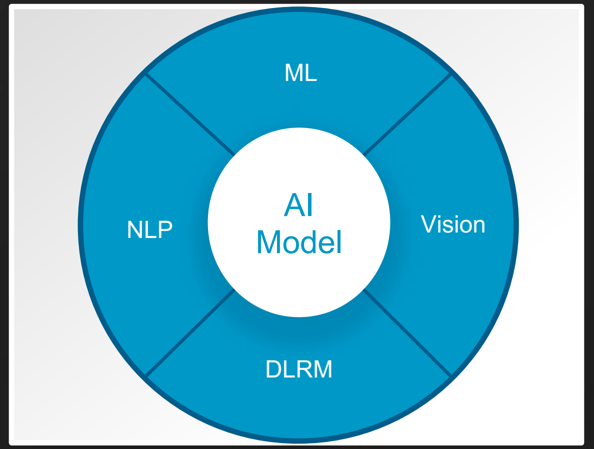

Architecture optimized for internal workloads

Larger volumes → lower power & cost

Memory & I/O architecture customized for AI workloads

Industry-Leading IP Portfolio for Differentiated XPUs

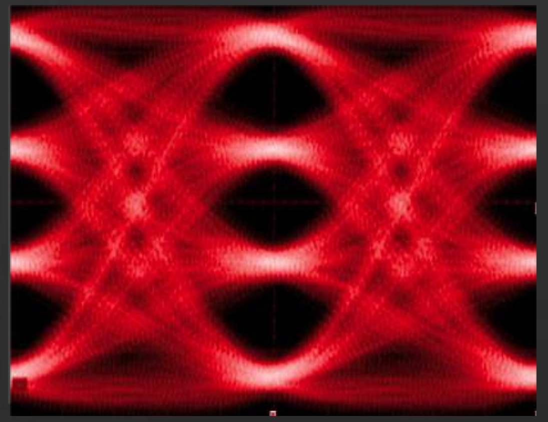
SerDes IP


AI-Optimized NICs IP

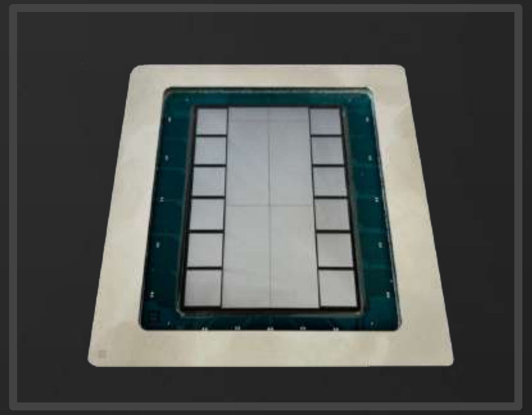
Advanced Pkg

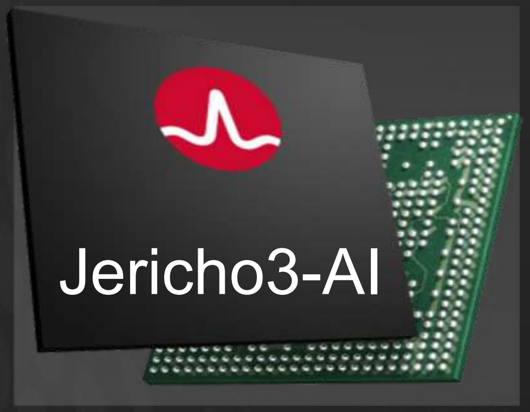

Switching IP Cores

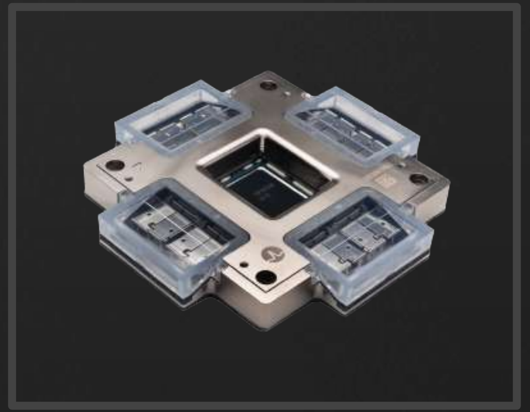
Co-Packaged Optics

Buffer Memory IP

$3B+

R&D Investment

XPUs

Optimized Architecture + Lowest Power → Best Performance/TCO

Broadcom IP Enables XPU Critical Components

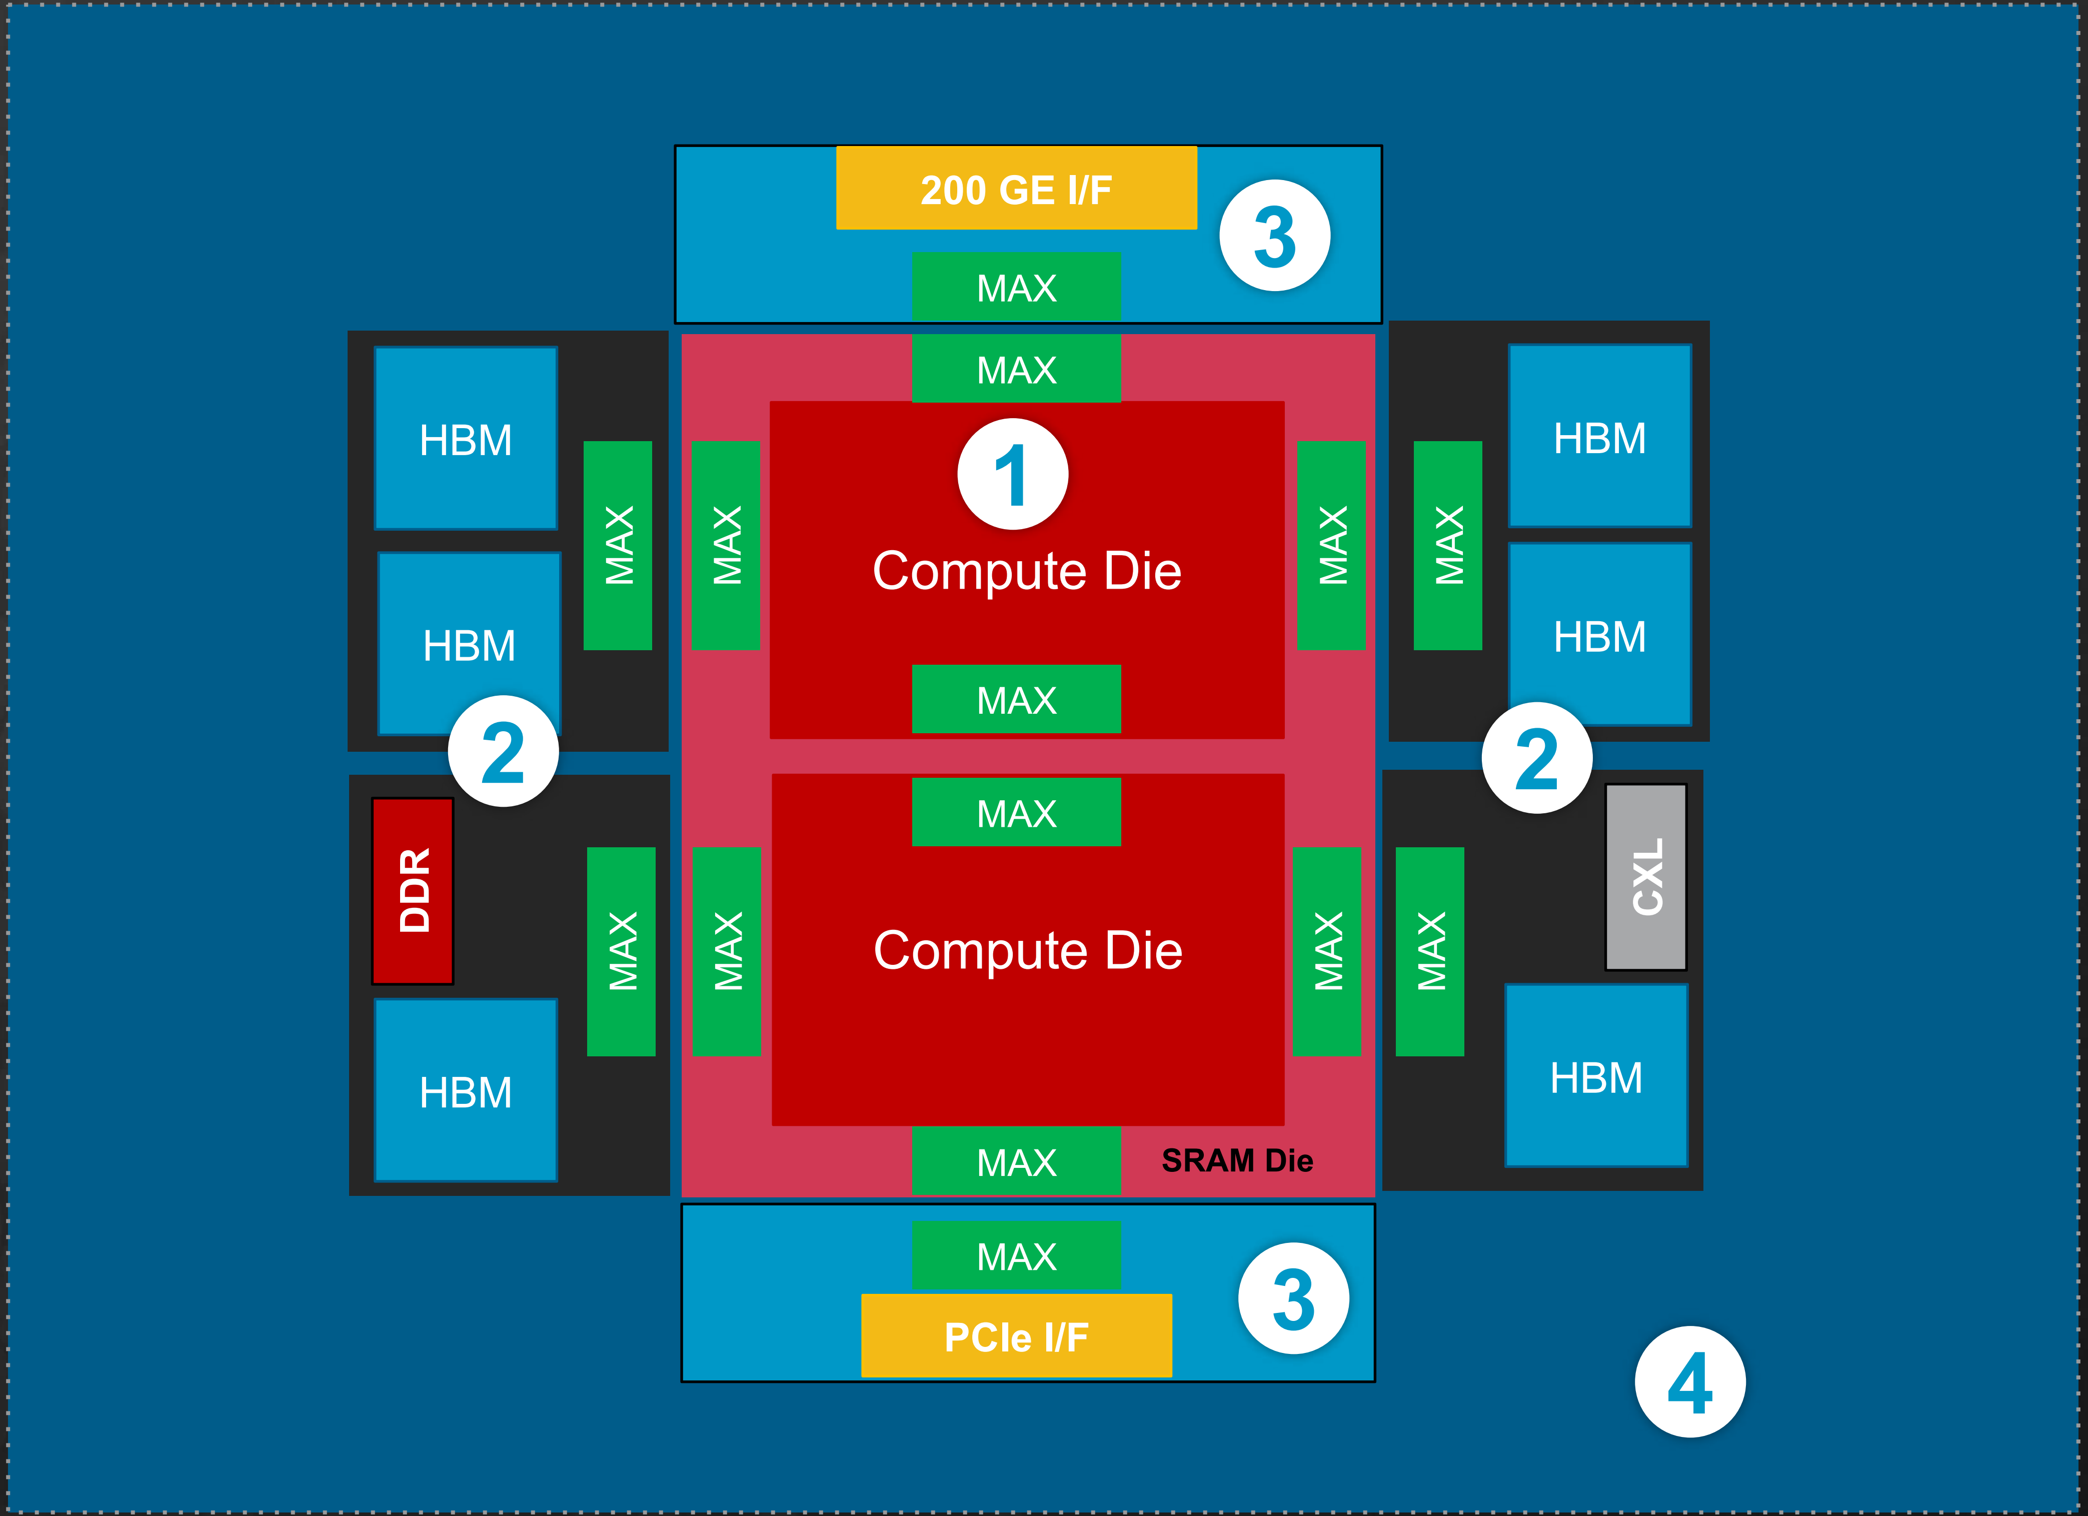

1 Compute

Processing Unit Architecture (Customer Owned)
Design Flow & Performance Optimization (Broadcom Owned)

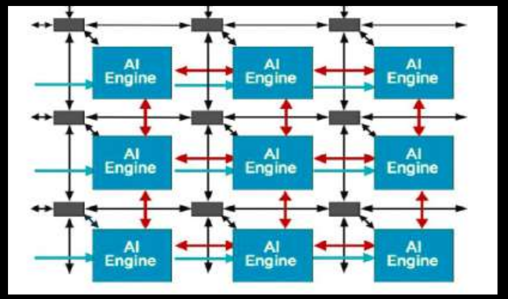

2 Memory

HBM PHY, Integration & Performance (Broadcom Owned)

3 Network IO

Architecture & Implementation (Broadcom Owned)
Full Solution Chiplets (Hardware, Firmware, & Software)

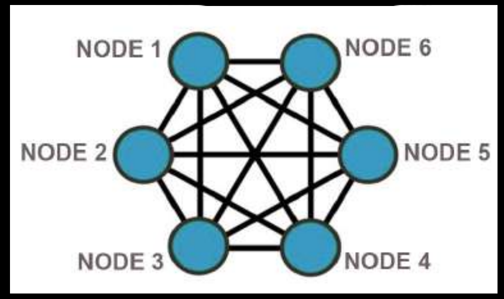

4 Package

2.5D, 3D, & Silicon Photonics Architecture & Implementation (Broadcom Owned); Vertical Integration Advantage

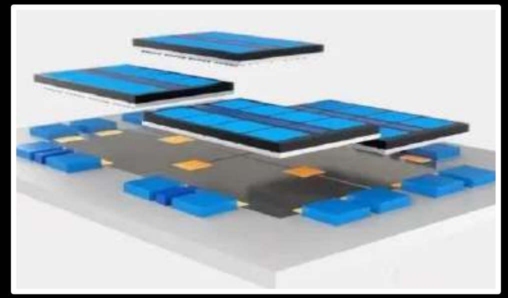

OPEN // SCALABLE // POWER EFFICIENT

Delivering Multi-Generational Custom AI Accelerators

over the Past Decade

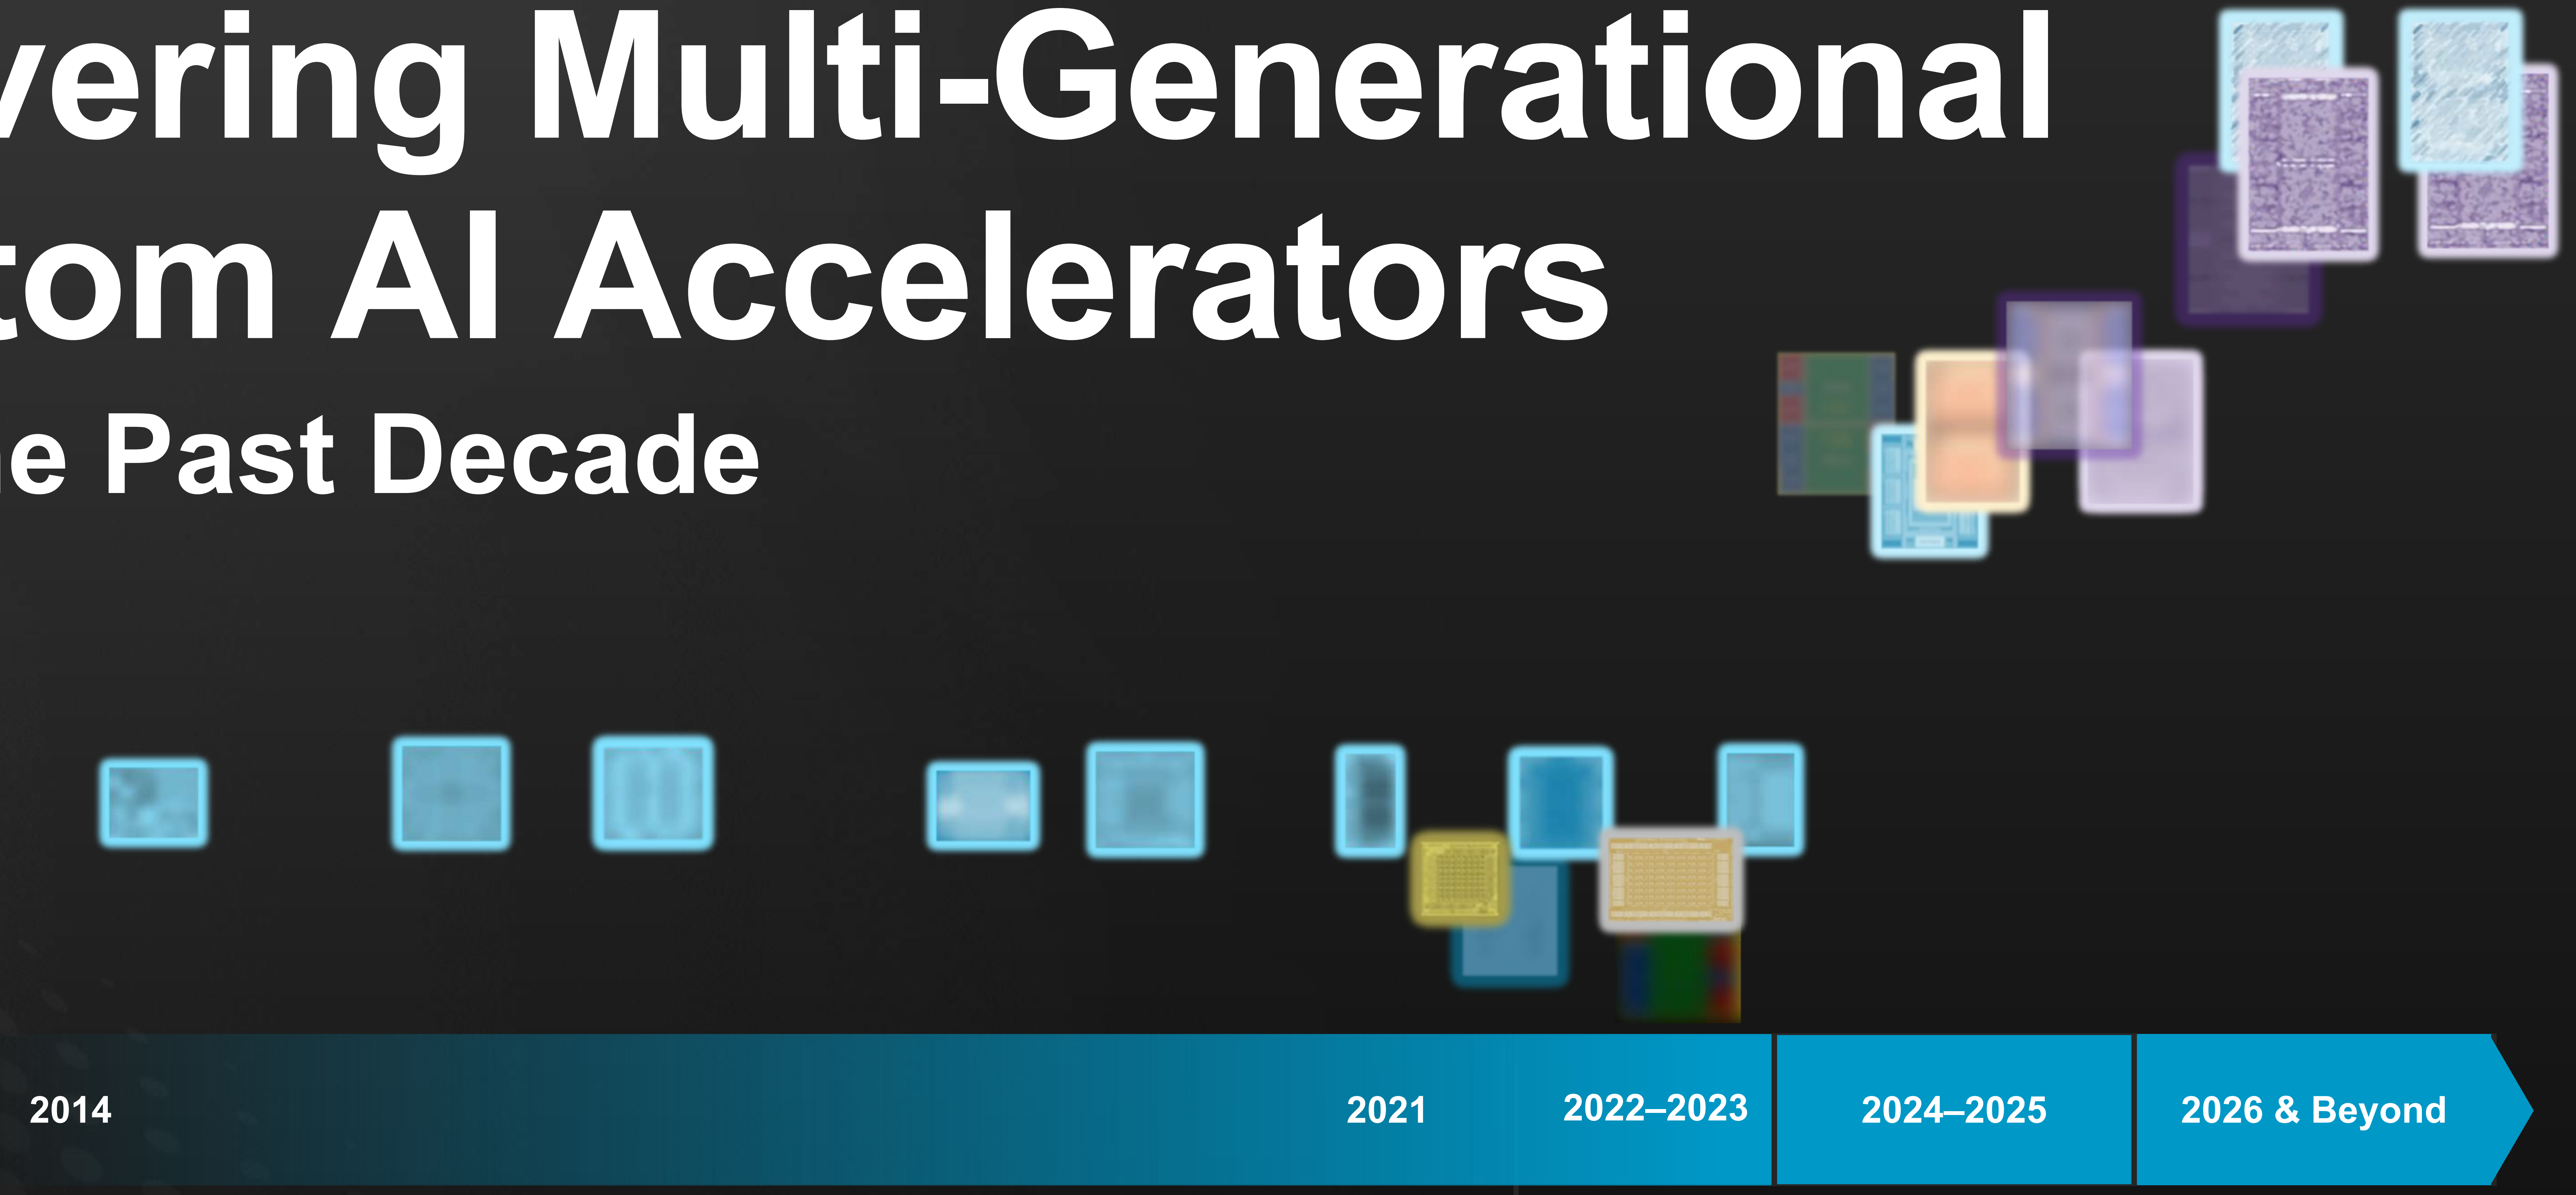

OPEN // SCALABLE // POWER EFFICIENT

Deep, Strategic, Multi-Year Partnership with Consumer AI Customers

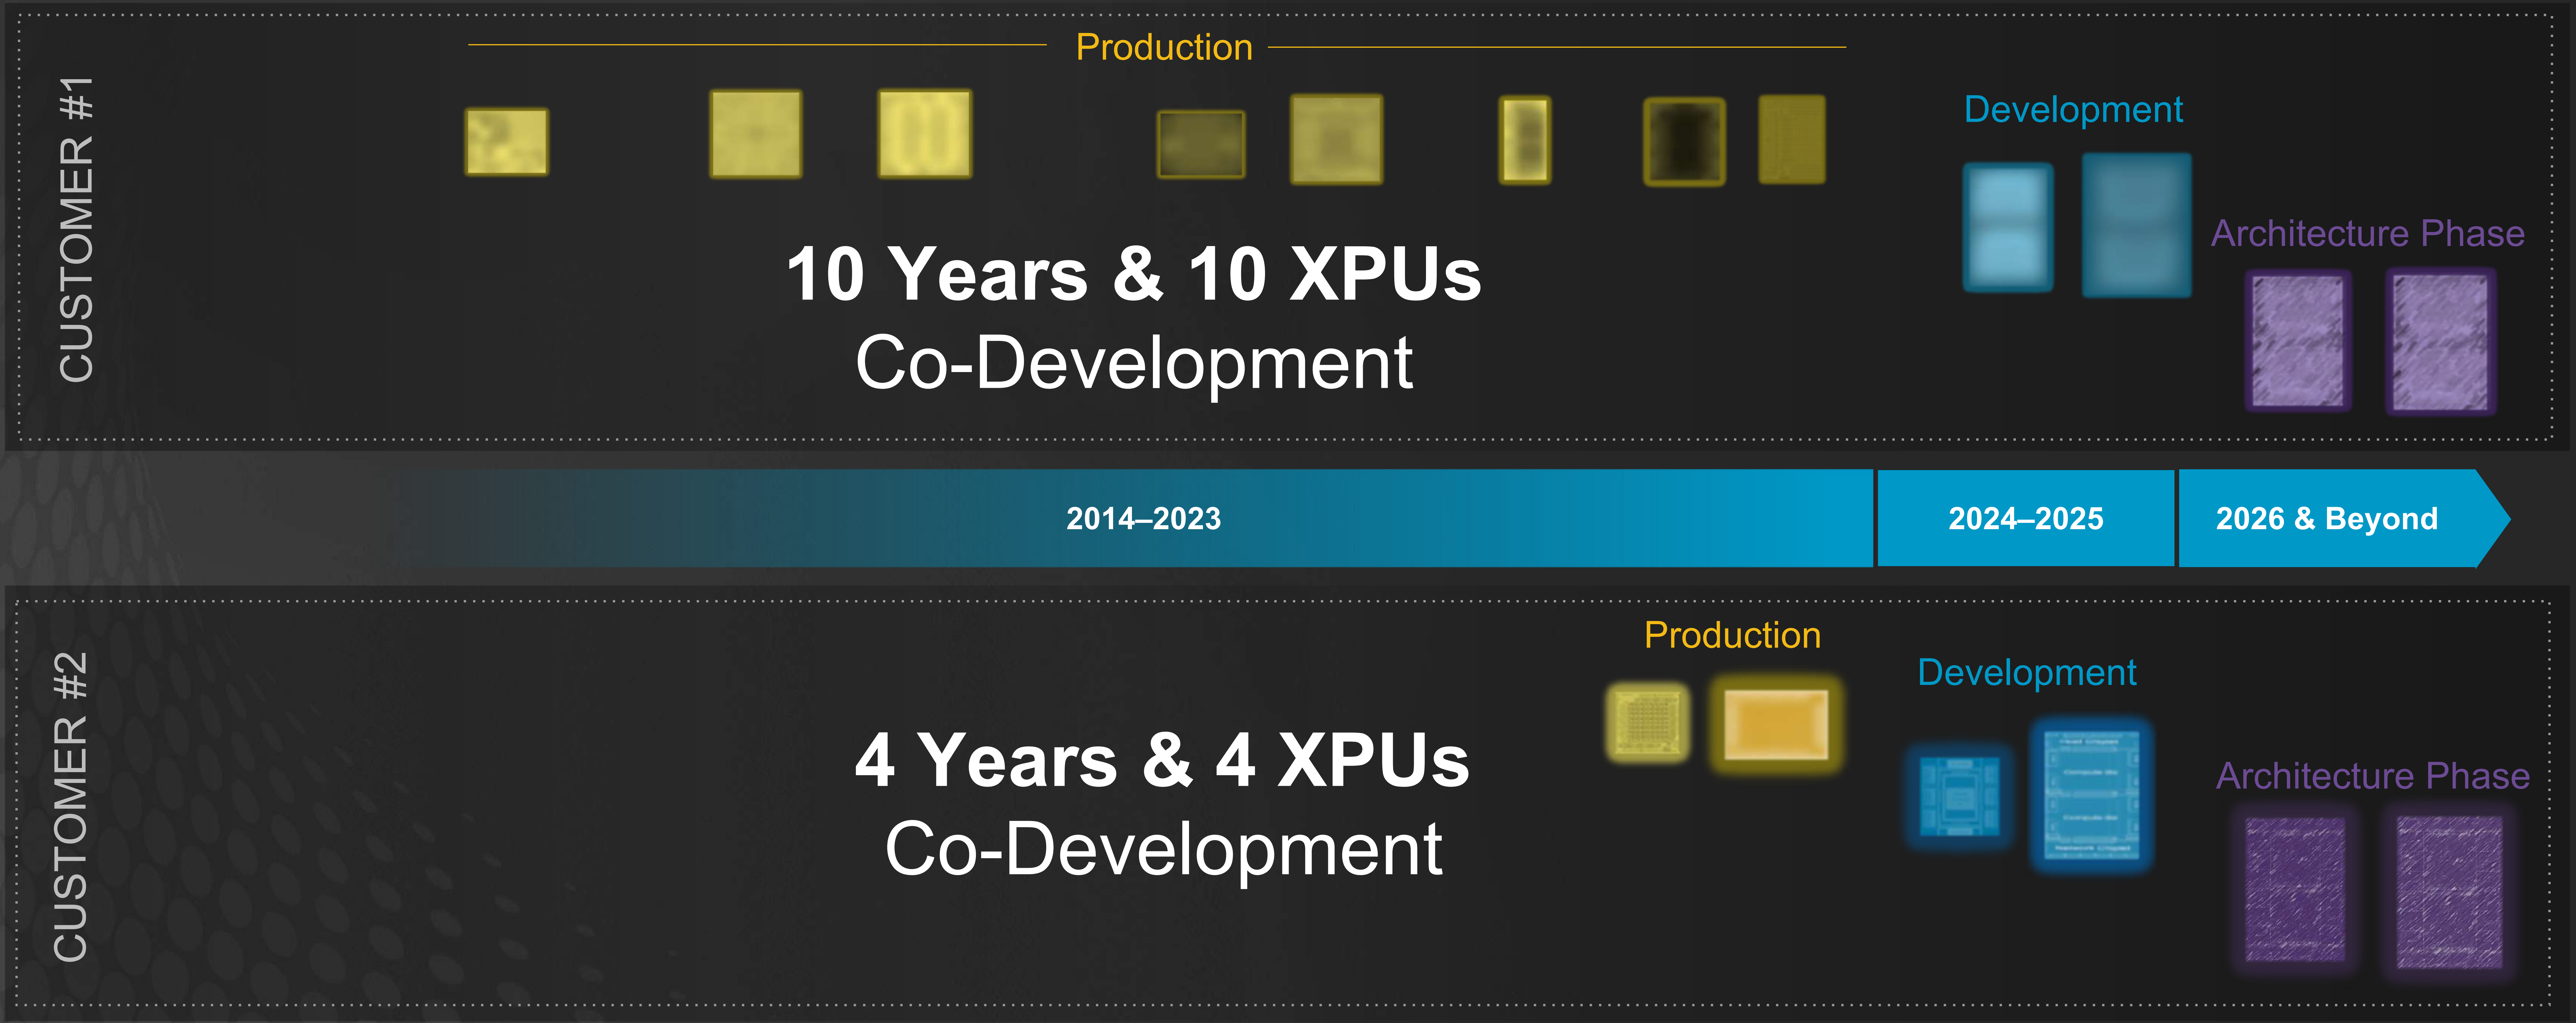

CUSTOMER #1

Production

10 Years & 10 XPUs
Co-Development

Development

Architecture Phase

2014–2023

2024–2025

2026 & Beyond

CUSTOMER #2

Production

4 Years & 4 XPUs
Co-Development

Development

Architecture Phase

OPEN // SCALABLE // POWER EFFICIENT

Industry's Fastest Time-to-Production with Annual Cadence

XPU Platform Engineering Years prior to Design Start → **7–9 Month Development** → **FAB** → **3-Month Production Ramp**

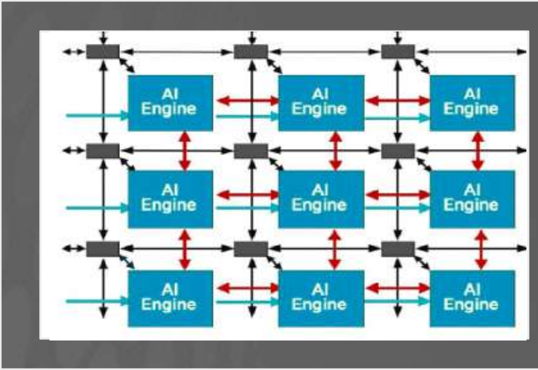

Automated Design Flow

AI Design Flow Optimized over 10 Years of XPU experience

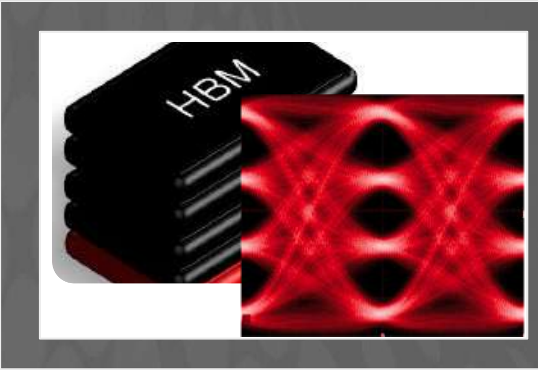

AI IP

Plug & Play Ready
- ✓ Silicon
- ✓ Firmware
- ✓ Software-API

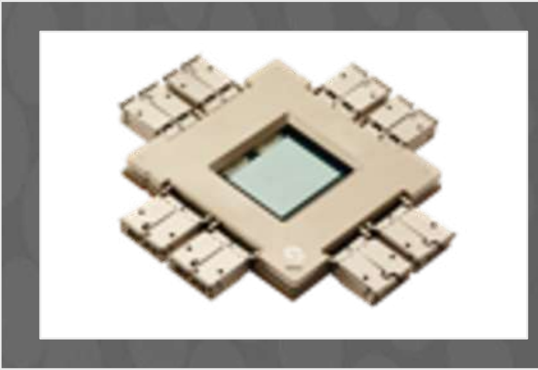

SOC Package

Pre-Qualified
- ✓ Thermal
- ✓ Mechanical
- ✓ Electrical

Co-Development with Consumer AI Customer

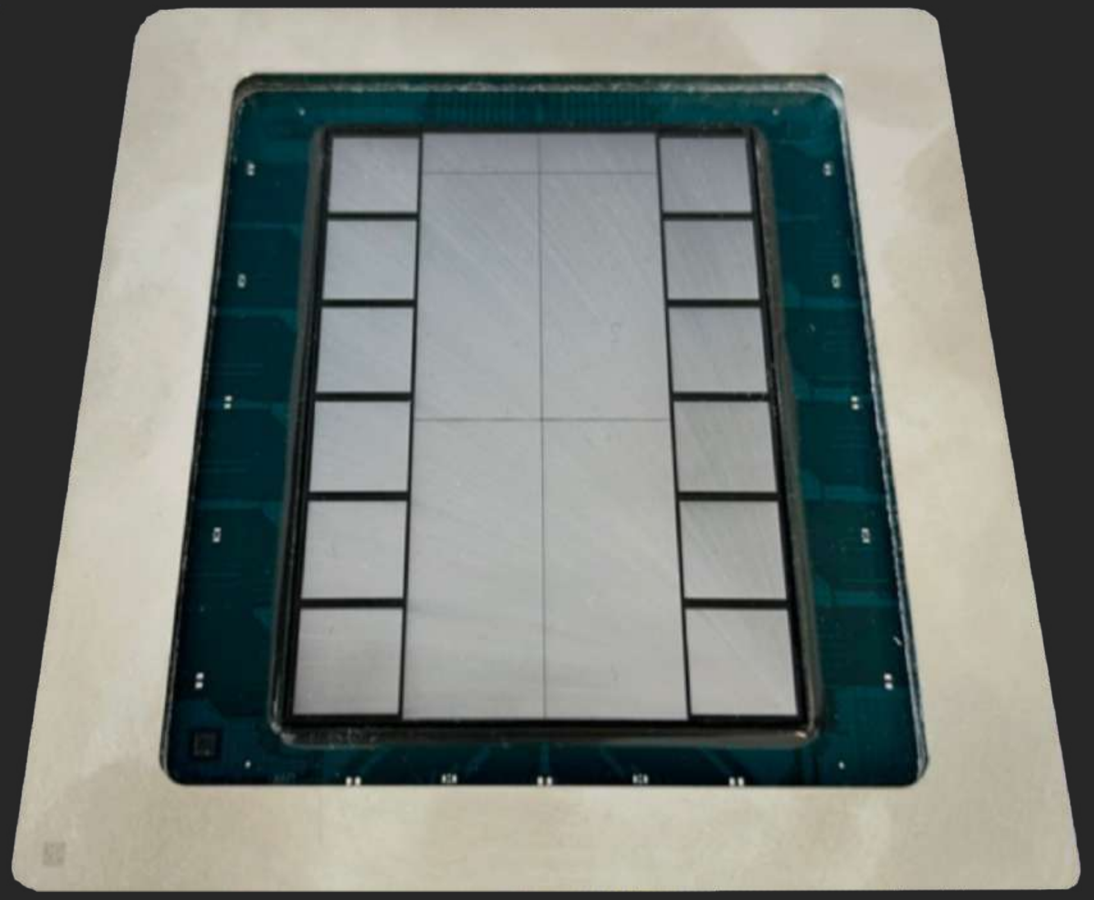

OPEN // SCALABLE // POWER EFFICIENT

Deep & Strategic Collaboration During Architectural Phase

INVESTMENT PERIOD
2023-2025

ARCHITECTURE PHASE
2026 & Beyond

Early IP investments years before XPU design kick-off across critical components

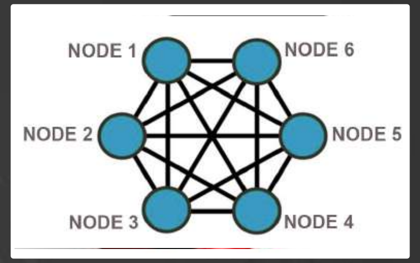

Concurrent technology, IP & packaging investments for next two generations of XPUs

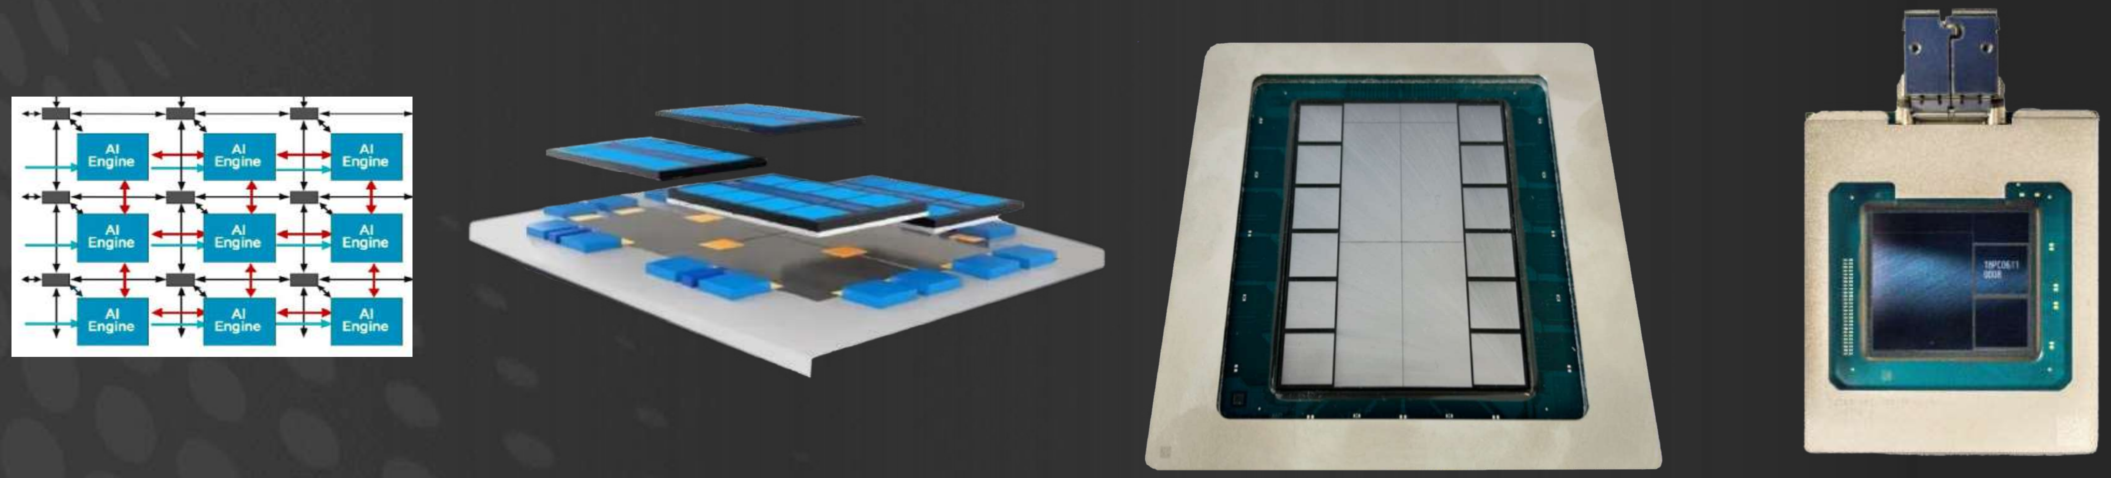

OPEN // SCALABLE // POWER EFFICIENT

Execution Excellence in Delivering to Consumer-AI Requirements

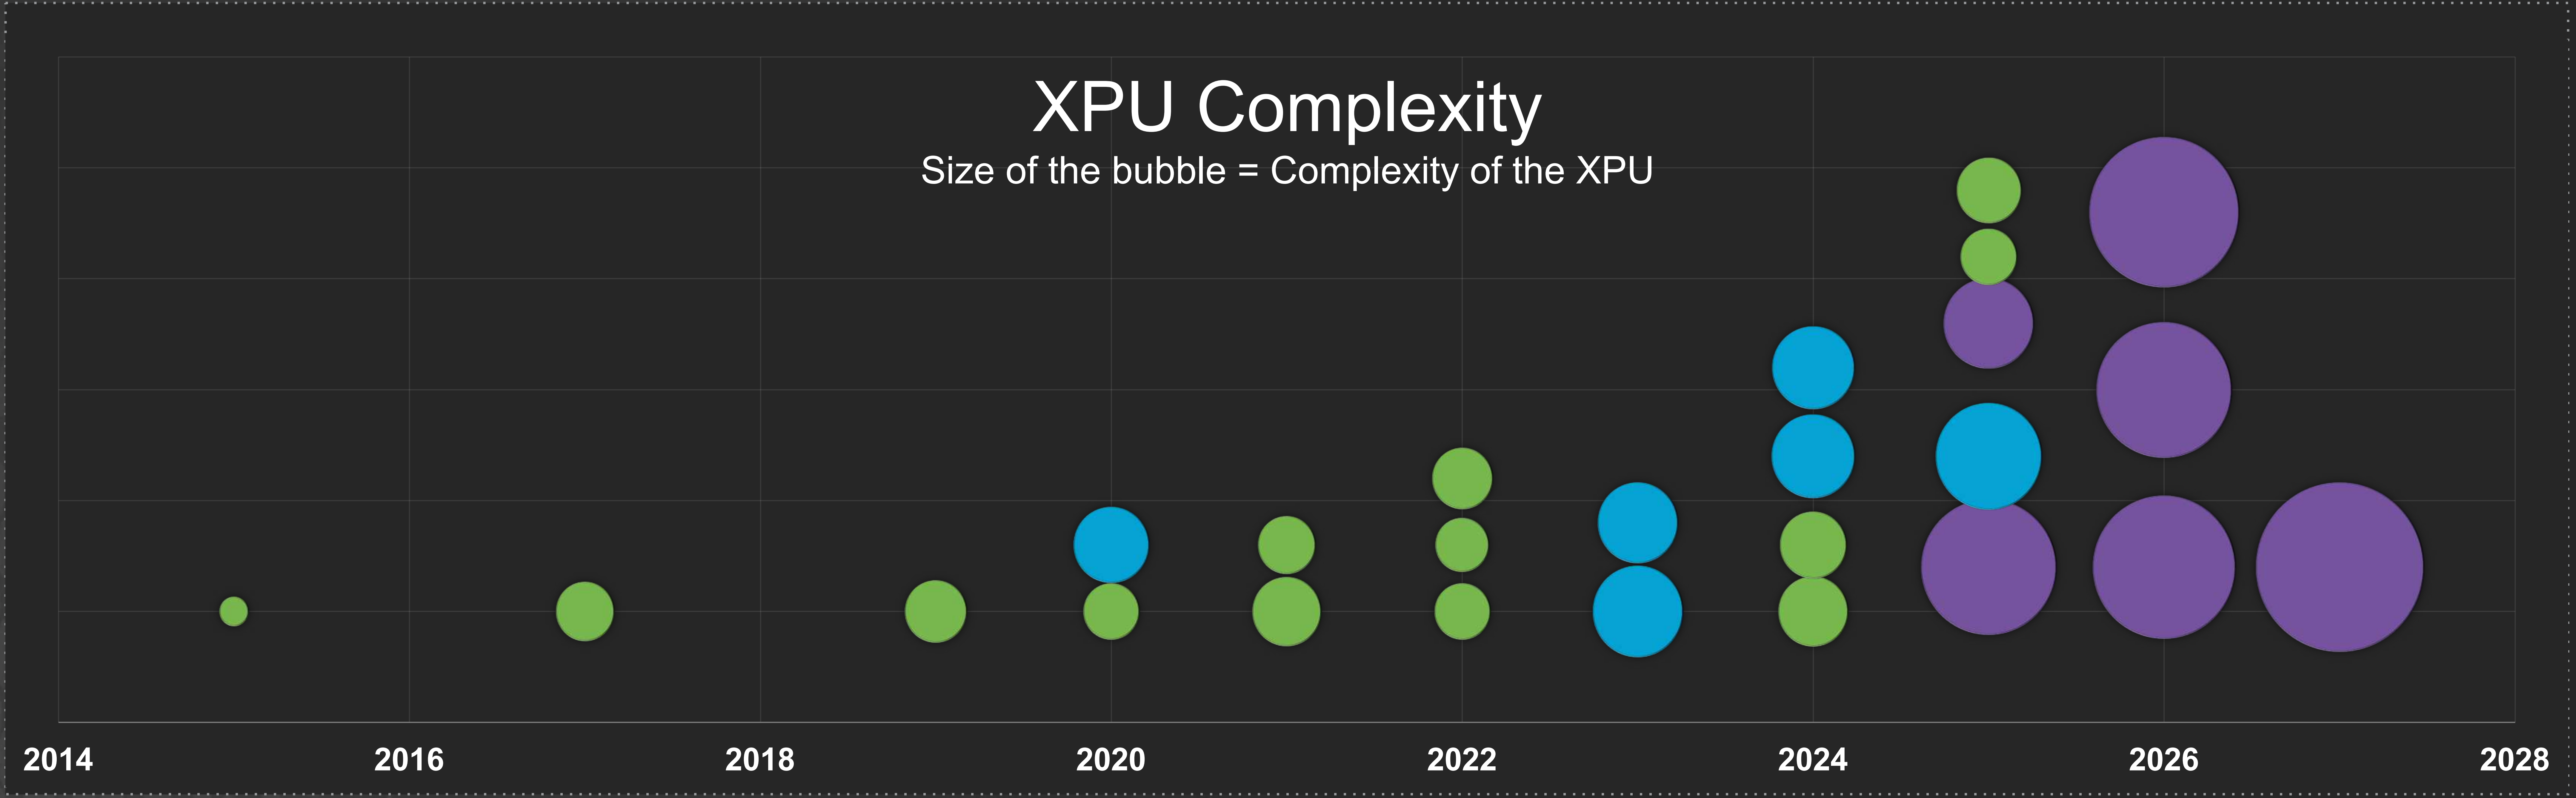

F (Complexity) = {Compute Performance, Network Bandwidth, Memory Bandwidth, Power Delivery, Thermal Integrity, Mechanical Reliability}

OPEN // SCALABLE // POWER EFFICIENT

Custom AI Accelerators

Key Points

- Focused on Consumer-AI Custom Accelerators
- **$3B+** Investment Scale in Differentiation **XPU IP**
- 10-year Engineering Excellence in High-Complexity XPUs
- 3-year Deep Architectural Advantage in Co-Development

OPEN // SCALABLE // POWER EFFICIENT

BROADCOM®

ENABLING AI Infrastructure

Charlie Kawwas, Ph. D.
President, Broadcom

OPEN // SCALABLE // POWER EFFICIENT

Consumer AI

Search, Social Networking, …

Enterprise AI

Public Cloud, On-Prem, …

AI Accelerators (XPU)

Custom

Merchant

AI Connectivity

Merchant

XPUs → Multi-Generational Strategic Engagements

	2024	2025	2026
Customer #1	Production	✓	✓
Customer #2	Production	✓	✓
Customer #3 *New*	Ramp → Production	✓	✓

Broadest & Industry-Leading AI Networking Portfolio

Ethernet


Broadcom (AVGO) to Deploy World's First 51.2 Tbps Switches

source: Yahoo Finance

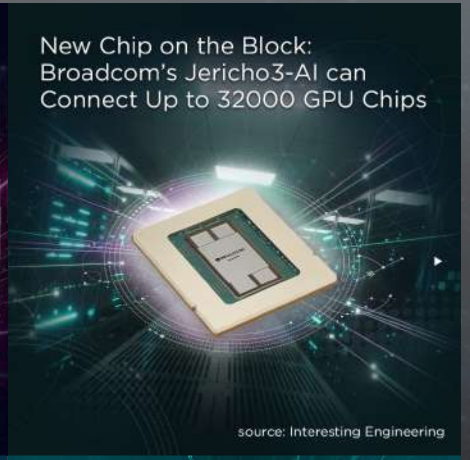
New Chip on the Block: Broadcom's Jericho3-AI can Connect Up to 32000 GPU Chips

source: Interesting Engineering


Broadcom Qumran3D 5nm 25.6T Switch Router Chip

Source: ServeTheHome

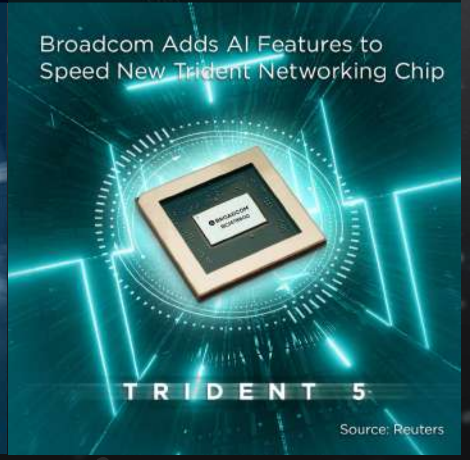
Broadcom Adds AI Features to Speed New Trident Networking Chip

Source: Reuters

PCIe


Industry's First End-to-End PCIe/CXL Connectivity Portfolio

Optics

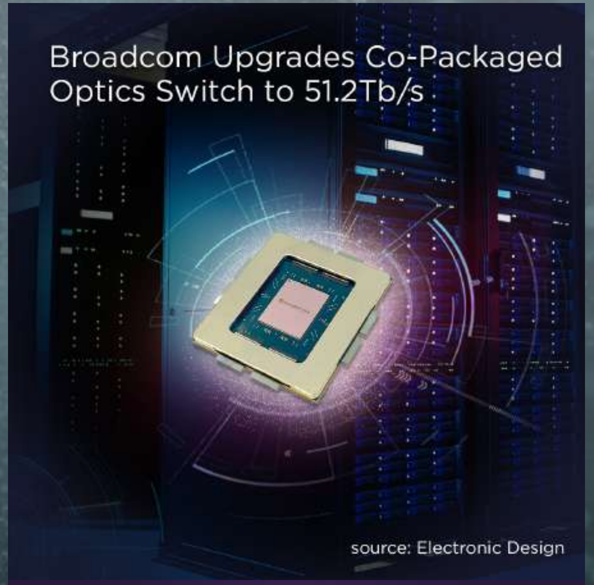
Broadcom Upgrades Co-Packaged Optics Switch to 51.2Tb/s

source: Electronic Design


World Leader in AI Infrastructure Optics

Foundational Technology

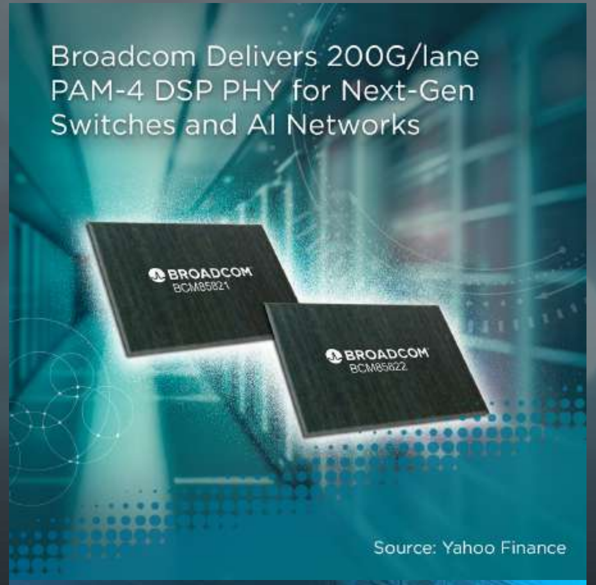
Broadcom Delivers 200G/lane PAM-4 DSP PHY for Next-Gen Switches and AI Networks

Source: Yahoo Finance

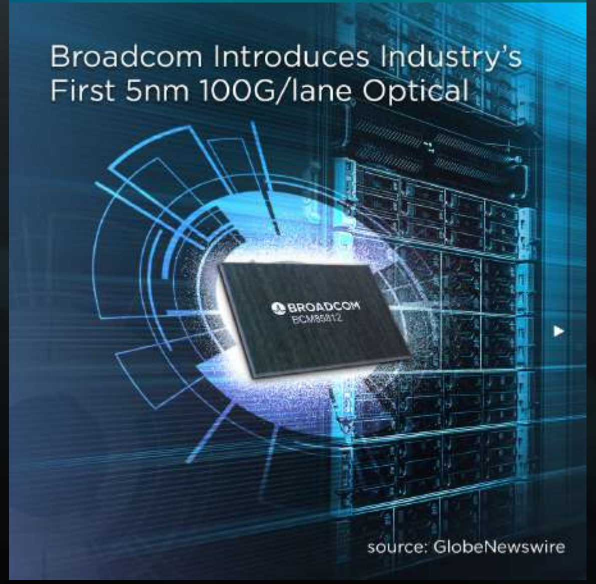
Broadcom Introduces Industry's First 5nm 100G/lane Optical

source: GlobeNewswire

OPEN // SCALABLE // POWER EFFICIENT

… Enabling AI Infrastructure

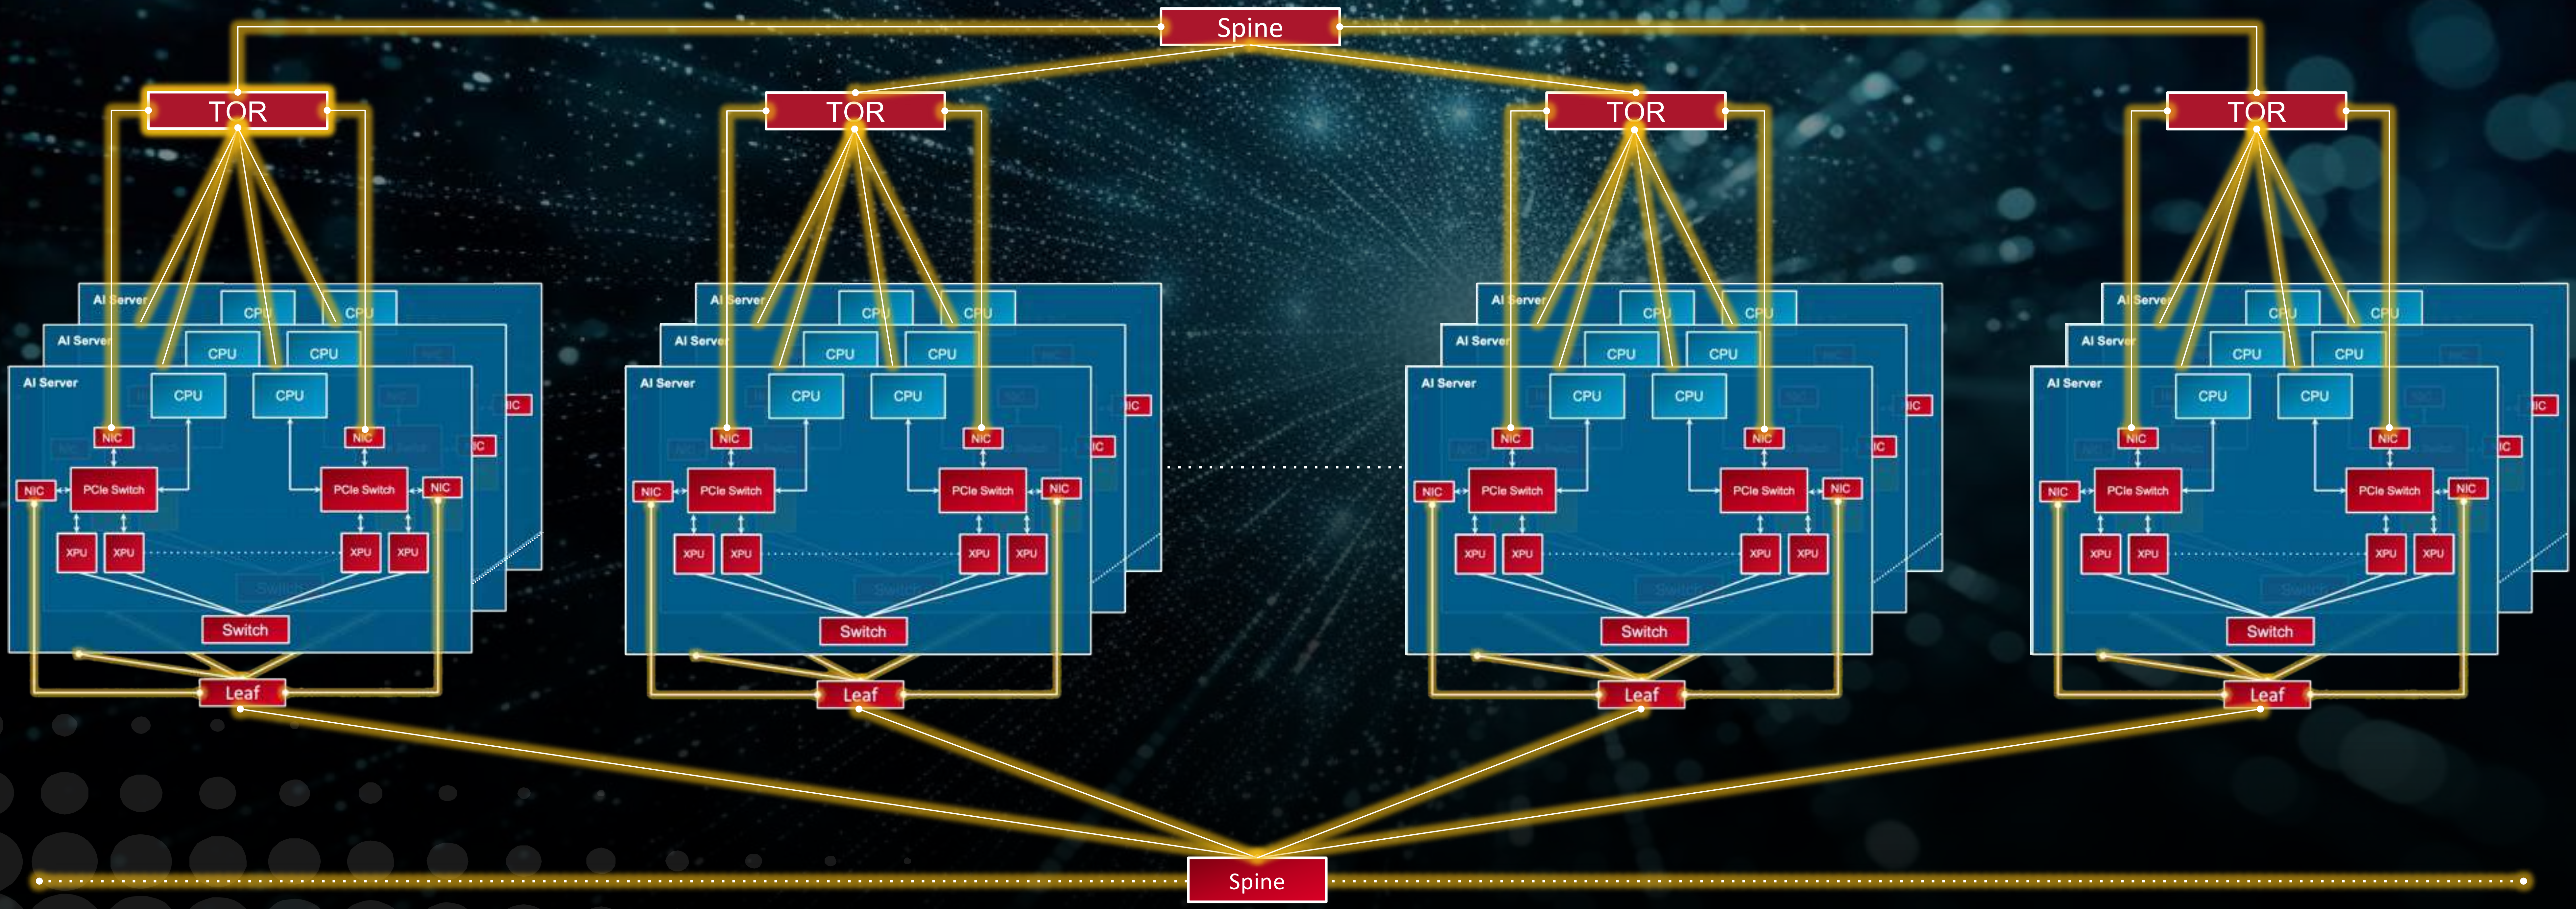

OPEN // SCALABLE // POWER EFFICIENT

Open Scalable Power Efficient

Sustainable Technology Leadership

Ethernet PCIe Optics SerDes DSP Custom

Investment Scale **Innovation Leadership** **Execution Excellence**

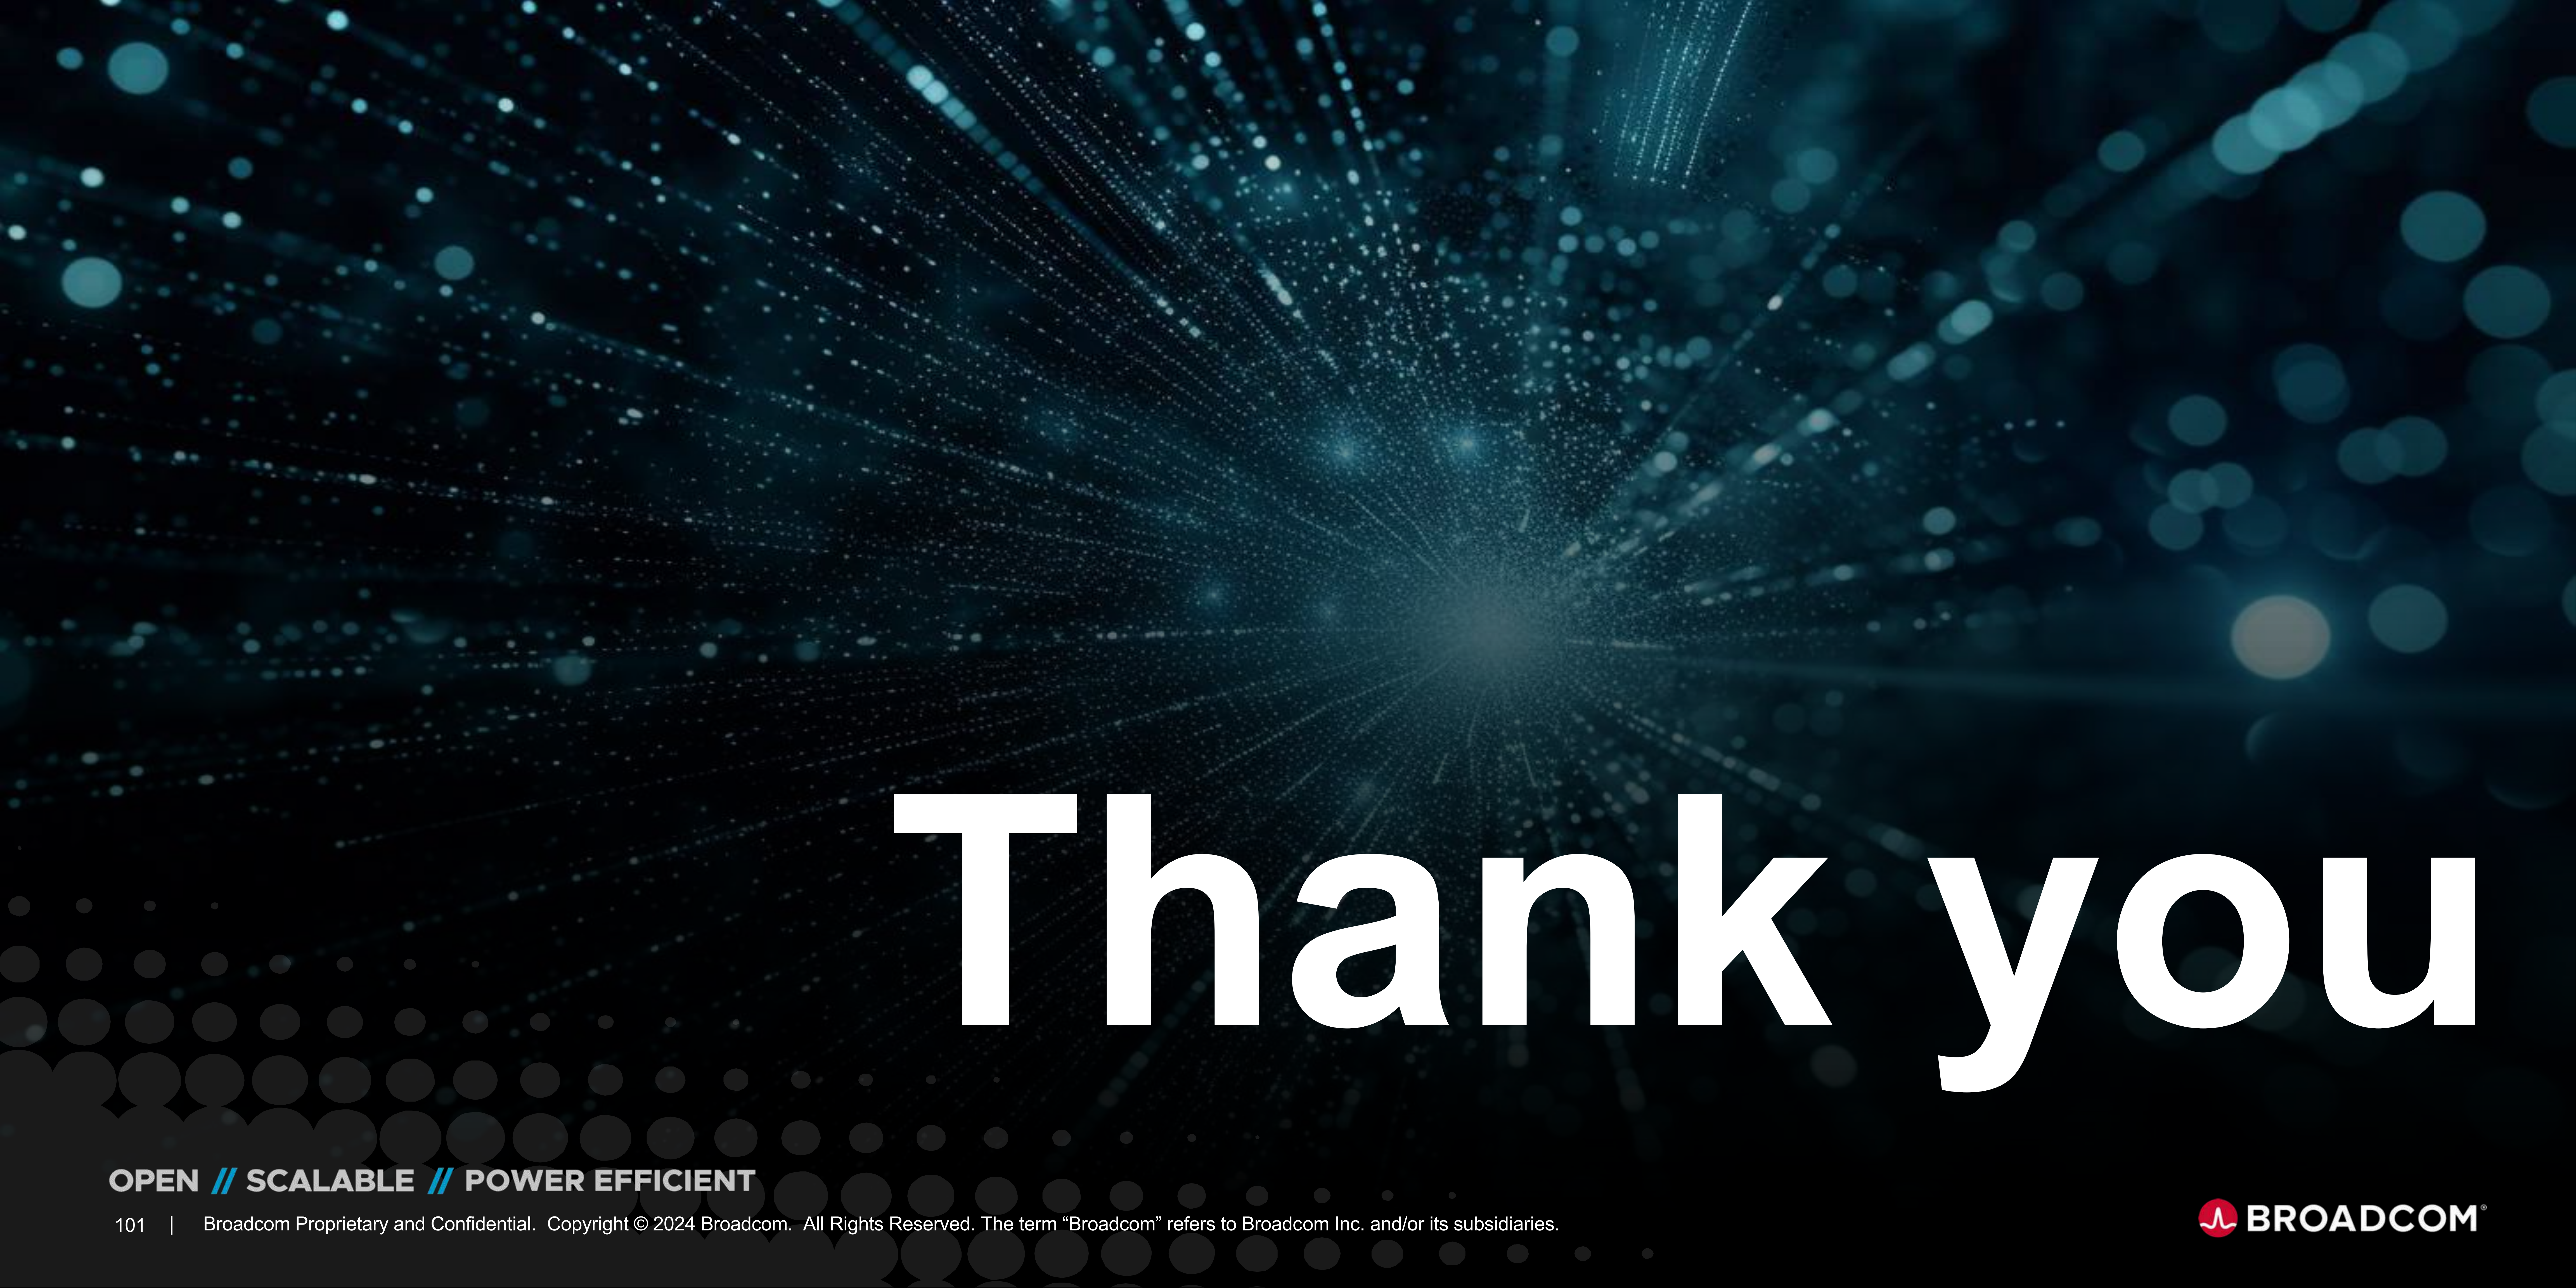

Thank you

BROADCOM®

ENABLING AI Infrastructure

OPEN // SCALABLE // POWER EFFICIENT