



Custom AI Investor Event

June 17, 2025

Ashish Saran

SVP, Investor Relations

Forward-looking statements

Except for statements of historical fact, this presentation contains forward-looking statements (within the meaning of the federal securities laws) including, but not limited to, statements related to market trends and to the company's business and operations, business opportunities, growth strategy and expectations, and financial targets and plans, that involve risks and uncertainties. Words such as "anticipates," "expects," "intends," "plans," "projects," "believes," "seeks," "estimates," "can," "may," "will," "would" and similar expressions identify such forward-looking statements. These statements are not guarantees of results and should not be considered as an indication of future activity or future performance. Actual events or results may differ materially from those described in this presentation due to a number of risks and uncertainties.

For factors that could cause Marvell's results to vary from expectations, please see the risk factors identified in Marvell's Quarterly Report on Form 10-Q for the fiscal quarter ended May 3, 2025, as filed with the SEC on May 30, 2025, and Marvell's Annual Report on Form 10-K for the fiscal year ended February 1, 2025, as filed with the SEC on March 12, 2025, and other factors detailed from time to time in Marvell's filings with the SEC. The forward-looking statements in this presentation speak only as of the date of this presentation and Marvell undertakes no obligation to revise or update publicly any forward-looking statements.

Non-GAAP financial measures

- During this presentation, we may refer to certain financial measures on a U.S. non-GAAP basis.
- We believe that the presentation of non-GAAP financial measures provides important supplemental information to management and investors regarding financial and business trends relating to our financial condition and results of operations.
- While we use non-GAAP financial measures as a tool to enhance our understanding of certain aspects of our financial performance, we do not consider these measures to be a substitute for, or superior to, the information provided by GAAP financial measures.
- A reconciliation in accordance with SEC Regulation G for non-GAAP financial measures is available in the Investor Relations section of our website at <https://www.marvell.com>.

Agenda



Matt Murphy

Chairman and Chief Executive Officer



Chris Koopmans

Chief Operating Officer



Nick Kucharewski

SVP and GM, Cloud Platform



Sandeep Bharathi

Chief Development Officer



Will Chu

SVP and GM, Custom Cloud Solutions

Q&A



Custom AI Investor Event

Matt Murphy

Chairman and Chief Executive Officer

Cloud-optimized silicon

The era of cloud-optimized silicon

Software on x86
for everything

Heterogeneous
computing

Cloud-optimized
silicon

CPU

CPU

GPU

FPGA

Custom
CPU

Custom
ML

Custom
DPU


Custom
CPU

Custom
ML

Custom
DPU


AI infrastructure buildout drives incredible growth

Top 4 U.S. hyperscaler CAPEX




Source: public information

Emerging hyperscalers rising fast



AI infrastructure buildout drives incredible growth


Top 4 U.S. hyperscaler CAPEX



Source: public information


AI infrastructure buildout drives incredible growth

Total data center CAPEX



Source: public information, Dell' Oro


To reach \$1 trillion by 2028



Source: public information, Dell' Oro


Marvell® data center TAM

Data center TAM: AI Day April 2024




Source: 650 Group, CignalAI, Dell'Oro, LightCounting, and Marvell estimates

Marvell data center TAM




Compute is 29% larger, interconnect is 37% larger

Marvell data center TAM




Source: 650 Group, SignalAI, Dell'Oro, LightCounting, and Marvell estimates

Marvell data center TAM



Source: 650 Group, SignalAI, Dell'Oro, LightCounting, and Marvell estimates


Accelerated custom compute TAM = XPU



Custom XPU

Expanding opportunity with custom XPU attach

Accelerated custom compute TAM = XPU + XPU attach




Custom XPU

Custom XPU attach

Network interfaces, scale-up fabrics, security and host management co-processors, memory poolers and expanders, etc.


Expanding opportunity with custom XPU attach

XPU attach becoming increasing portion of custom TAM



Source: 650 Group, Dell'Oro, and Marvell estimates

Expanding multi-generational relationships



Custom opportunity explosion

PIPELINE

>50

Opportunities

>10

Customers

\$75B

Lifetime revenue

Opportunities

1/3 custom XPU

2/3 custom XPU attach

Each XPU

Multi-\$B lifetime

18 - 24 months

Each XPU attach

Several-\$100M lifetime

2 - 4 years

Why we win: differentiating with full-service custom



**System
architecture**



Design IP



**Silicon
services**




Packaging



**Manufacturing
and logistics**


Leading-edge technology drives **multi-gen wins**

Process




**Nanometers
to angstroms**

IP



**World's first 448G electrical
and optical SerDes**

Packaging



**Advanced
packaging**



EDA in the cloud



Custom and
networking silicon



EDA in the cloud



Custom and
networking silicon



“Building a cost and power efficient cloud at the scale that only AWS can deliver begins with leading-edge semiconductors designed to meet the demanding infrastructure needs of our customers. Our expanded collaboration with Marvell enables us to deploy our comprehensive semiconductor portfolio and specialized networking hardware to advance our mission to provide the industry’s most robust and scalable cloud and AI services to our customers.”

Matt Garman, CEO at AWS






Rani Borkar

Corporate Vice President, Azure Hardware Systems
and Infrastructure, Microsoft




Video content

Marvell's **data center** revenue




Driving to 20% share in custom compute and attach




Source: 650 Group, Dell'Oro, and Marvell estimates

On track for 20% share in overall data center



Source: 650 Group, CignalAI, Dell'Oro, LightCounting, and Marvell estimates

On track for 20% share in overall data center



Market share based on Marvell fiscal year
Source: 650 Group, SignalAI, Dell'Oro, LightCounting, and Marvell estimates

Larger TAM

\$75B → \$94B

New markets

Custom XPU → Custom XPU attach

Increasing share

~10% → ~13% ... 20%

More sockets won

3 → 18

More customers

4 → >10

>50 opportunities

\$75B




Custom AI Investor Event

Chris Koopmans


Chief Operating Officer

Data center CAPEX to reach \$1 trillion by 2028




Source: Company public information, Dell'Oro

2028 data center infrastructure market




Source: 650 Group, SignalAI, Dell'Oro, LightCounting, and Marvell estimates

Breaking down accelerated computing TAM



Source: 650 Group, Dell'Oro, and Marvell estimates


Breaking down accelerated computing TAM



Source: 650 Group, Dell'Oro, and Marvell estimates


Custom silicon on track for 25% share of market by 2028

Accelerated compute




Source: 650 Group, Dell'Oro, and Marvell estimates

Custom accelerated compute



Marvell custom silicon TAM

Marvell TAM




Source: 650 Group, Dell'Oro, and Marvell estimates



Why
custom?

**Diversity drives specialization,
which drives customization**

Race to the **million-XPU** cluster speeding up



Model innovation continues

Large language model



Content
writing



Live
chatbots



Search
Analyze
Conclude

Reasoning model



Code
writing




Complex
math




Puzzle
solving

Rise of inference-optimized infrastructure


**One size
does not fit all**



Training




Chain of thought
inference




Inference

As AI workloads diversify, so must the infrastructure

Diversification creates more opportunities



Specialization drives diversity of compute core design



Multiple points of optimization

Core architectures


cores

Level of precision

Ratio of cores to memory


Programmable, powerful, purpose-built

And diversity of memory, packaging, connectivity



Beyond one chip, advanced packaging powers scale

And diversity of the entire scaled-up platform design



Optimized infrastructure for specific workloads

AI infrastructure platform options



NVLink Fusion
announced May 18

UALink
announced June 11

XPU

Custom

Custom

Custom

Custom

**XPU
attach**

General-purpose
platform

Third-party
platform

Standards-
based platform


Full-custom
platform

AI infrastructure platform options




NVLink Fusion
announced May 18

UALink
announced June 11



Growing TAM supports further **pipeline expansion**



18 current sockets
targeting **20%** share

>50 opportunities
driving **\$75B**

Source: 650 Group, Dell'Oro, and Marvell estimates



Custom AI Investor Event

Nick Kucharewski

SVP and GM, Cloud Platform

Marvell Custom Cloud Solutions



30 years of experience

First-time-right silicon design



Comprehensive portfolio


Silicon-proven core enabling IP



Full-lifecycle partnership

Definition through deployment





Choosing a custom silicon partner

Pace of innovation in AI cloud infrastructure



New key enabling technologies
every year

... applications demand system innovation
every < 2 years

... a 1-day delay means
millions \$ in lost revenue



**Full-service
custom**



System
architecture



Design IP



Packaging



Silicon
services



Manufacturing
and logistics

Physical design services



Silicon
services



Manufacturing
and logistics

Manufacturing



Manufacturing
and logistics

Marvell: Full-Lifecycle Custom Cloud



**System
architecture**



Design IP



**Silicon
services**



Packaging



**Manufacturing
and logistics**

Unique differentiated technology, not available on the open market

Projects begin years ahead of time



**System
architecture**

New technology development

Incubation 2-4 years before project start



Evaluation and assessment

Concept systems enable customer exploration



Cooperative system definition

Iterate on power, performance, cost tradeoffs



Key enabling core tech



Design IP



SerDes



Die-to-die



Compute
cores



Packet
processing



Compute
fabric



Multi-die



Security



Ethernet
MAC



Packet
fabric



Specialized
SW/FW



Optics



Compress /
decompress



Custom
HBM



Dense
SRAM



Hardened
Arm

Marvell

Ecosystem

arm

cadence

synopsys

Start-ups

Running ahead of the customer project



**Silicon
services**

5nm

3nm

2nm

16/14Å



Process enablement



IP
migration



Test
silicon



Model
creation



EDA
toolflow

Design execution



Logic
design



Physical
design



DFT
and DFM



Power
tuning



Yield
enhancement

The key to scaling beyond Moore's Law




Packaging



Multi-die

Marvell custom
and TSMC CoWoS




HBM

Marvell custom
and JEDEC




**Advanced
substrate**

Leadership
roadmap



CPC

Co-packaged copper



CPO

Co-packaged optics

Delivering quality at massive scale



**Manufacturing
and logistics**

Scale

Leading fabless manufacturer



Quality

Billions of device-hours over multiple decades



Readiness

Operations engineering for volume delivery



**Product
engineering**

**Test
engineering**

**Quality
assurance**

**Reliability
assessment**



**Full-service
custom**



System
architecture



Design IP



Packaging



Silicon
services



Manufacturing
and logistics

Physical design services



Silicon
services



Manufacturing
and logistics

Manufacturing



Manufacturing
and logistics

Running faster and faster

~2x every ~2 years

Process technology



3nm, 2nm, 16/14Å...

Die-to-die interconnect



16 Gbps, 32 Gbps, 64 Gbps...

High-speed SerDes



224 Gbps, 448 Gbps...

Advanced packaging



2.5D, 3.5D, advanced...

Networking protocols



Ethernet, NVLink, UALink...



10x
every
~2 years

Leading systems need access to the leading tech, every generation

Marvell Full-Lifecycle Custom Cloud

✓ **Partnership** for system design

✓ **Exclusive** enabling technology

✓ **Experience** at volume scale



CPU



Accelerator



Fabric



Interconnect



NIC



Switch



Custom AI Investor Event

Sandeep Bharathi

Chief Development Officer



Ken Chang

SVP, Analog and Mixed
Signal Engineering



Mark Kuemerle

VP, Technology,
Custom Cloud Solutions



Radha Nagarajan

SVP and CTO,
Optical Engineering



Mayank Mayukh

Senior Distinguished Engineer,
Advanced Packaging

AI compute cycle **accelerates**

General-purpose compute

2 years




Accelerated infrastructure

< 1 year



AI compute doubling less than a year


AI accelerates **single-die** to **multi-die** transition



Source: Adapted from TSMC slide from IEDM conference.

© 2025 Marvell. All rights reserved.

AI accelerates **single-die** to **multi-die** transition




Trillion+ transistors requires rack-scale integration


Source: Adapted from TSMC slide from IEDM conference.

© 2025 Marvell. All rights reserved.


Proven Marvell technology stack



Multi-generational execution track record



Process technology **leadership**




**Most advanced
process nodes**



Validated critical IP ahead of product development

Marvell IP across process technologies



CMOS


Silicon germanium

Silicon photonics





Foundries

Process-agnostic. Foundry-ready. Future-proof.




Anatomy of a modern **XPU**



Industry-leading customized IP portfolio

Anatomy of a modern **XPU**



SerDes



Die-to-die



Custom
HBM




Custom
SRAM





Silicon
photonics



Anatomy of a modern XPU



Modern XPU and XPU attach



Modern **XPU** and **XPU attach**



Customization to achieve the highest performance per watt

Why is SerDes so difficult?



Figures of merit

Reach

Power

Latency


**Bit
error rate**

Marvell has multi-generational SerDes leadership


Marvell IP leadership



SerDes



**Custom
HBM**



**Co-packaged
optics**



**Advanced
packaging**



Die-to-die



**Custom
SRAM**



**Silicon
photonics**

Focused investment enables first-to-market advantage




Custom AI Investor Event

Ken Chang

SVP, Analog and Mixed Signal Engineering

Marvell SerDes leadership in IEEE spotlight






Marvell SerDes recognized by IEEE

Source: ISSCC, VLSI, OFC, industry SerDes papers


© 2025 Marvell. All rights reserved.

Marvell SerDes leadership in IEEE spotlight

	ISSCC 2022	ISSCC 2024	VLSI 2024	ISSCC 2025	VLSI 2025	OFC 2025 Demo	
Electrical SerDes	 5nm 112G XSR	 5nm 112G LR	 3nm 112G XSR	 5nm 224G LR	 3nm 128G LR	 3nm 224G LR	
Optical SerDes		 5nm 200 G/lane optical gearbox	 5nm 100 G/lane optical retimer	 5nm 200 G/lane optical retimer	 5nm 100 G/lane optical retimer	 3nm 200 G/lane optical retimer	 3nm 400 G/lane

Enabling next-generation XPU platforms

Marvell 3nm 224Gbps SerDes at OFC



Enabling next-generation XPU platforms

Marvell die-to-die (D2D)

Today


10-20 Tbps / mm
~ 0.3pJ/bit

Tomorrow

30+ Tbps / mm
~ 0.2pJ/bit

Future

50+ Tbps / mm
<0.1pJ/bit






Custom AI Investor Event

Mark Kuemerle

VP, Technology, Custom Cloud Solutions

Custom dense SRAM




On-chip embedded SRAM – customization is key

Marvell custom SRAM

Key metrics for embedded memories


	Marvell vs. industry
Area at same bandwidth	50% lower
Standby power	66% lower
Bandwidth at same area	17x higher

Bandwidth / mm² of memory options



Comparison of 1 Mb instance of third-party dense memories vs. Marvell
Marvell IP provides **17x** bandwidth-density improvement


Custom HBM



Custom HBM – high capacity with minimal overhead


Enhancing XPU's with custom HBM architecture

Standard HBM No chiplets




- 1X useful compute area
- No compute area on HBM
- 1X power

Marvell Custom HBM With I/O chiplets



- **1.7X** useful compute area
- **More compute** area on HBM
- **75% lower** memory I/O power

Package-integrated voltage regulation (PIVR)




Typical
accelerator

PIVR
on underside

85% lower IR power

15% lower total product power

60% lower power noise





Custom AI Investor Event

Radha Nagarajan

SVP and CTO, Optical Engineering

Silicon photonics is critical for co-packaged optics




- Optical system-on-chip
- High speed, long reach
- High volume CMOS fabs
- Complex electronics and photonics integration

Highly integrated light engines for all interconnects

Marvell silicon photonics at scale

100G

Field-proven



2017

400G


Shipping



2020

800G

Sampling



2023

1.6T

Development



2026

Multiple generations of field-deployed silicon photonics

Marvell breaking **500 Gbps** data rate barrier

Data rate and analog
bandwidth (GHz) scaling


480
Gbps
Today

224
Gbps

112
Gbps

450 Gbps

Optical demo
OFC March 2025
enabled by Marvell
high-speed analog
components



Enables	Application	Single optical lane	Lanes	Format
1.6T ZR	Between DC	1.6T	1	Coherent
12.8T CPO	Inside DC	448G	32	PAM4

2X analog bandwidth every 2 years

Co-packaged optics to custom silicon interconnect



3D integrated silicon photonics light engine

Electronics ▶

Photonics ▶

Substrate ▶

Integration on reverse side ▶




6.4T per engine

32 x 224G
optical lanes

25.6T 4 engines

Custom silicon
interconnect



Marvell high-density custom silicon interconnects for AI




Custom AI Investor Event

Mayank Mayukh

Senior Distinguished Engineer, Advanced Packaging

Marvell advanced packaging **strategy**


Robust ecosystem for faster cadence between custom packaging generations



Marvell advanced packaging portfolio

Road to packaging a trillion transistors together


2D package



1X

Yesterday


2.5D package



4X


Today

3.5D package



8X


4D package



16X


Tomorrow

4.5D package



Marvell advanced packaging **toolbox**

Fungible toolbox for generational portfolio




Integrated optics and copper

3D-stacked silicon chiplets

Bridge or RDL interposer

Advanced substrates





Custom AI Investor Event

Will Chu

SVP and GM, Custom Cloud Solutions



+

averasem!


Compute

ASIC




Custom Cloud Solutions

Strong design win momentum



Typical system architecture



**Custom
XPU**



**Custom
XPU attach**



Providing unique customer value: typical XPU



Rack-scale enablement and optimization



Critical IPs



SerDes



Die-to-die




Dense SRAM



Hardened
Arm



Design-to-spec and co-development



Marvell custom and TSMC CoWoS



Faster time to market

Providing unique customer value: typical **XPU attach**



Production FW/SW, boards, and ecosystem



Critical IPs



Compress /
decompress



Compute
fabric



Security




SerDes




Design-to-spec integrating customer IP

Unique value across a single complete program



Unmatched combination of technology, expertise and scale

Unique value driving multi-generational partnerships



Concurrent engagement enables a new chip every year

Investing ahead in **breakthrough** technologies

Marvell Introduces Optimized
Integrated Power Solutions
to Boost Performance, Efficiency,
and ROI of Accelerated Infrastructure

June 17, 2025

Marvell Develops **Industry's First 2nm
Custom SRAM** for Next-Generation
AI Infrastructure Silicon

June 17, 2025

Marvell Delivers **Advanced Packaging
Platform** for Custom AI Accelerators

May 9, 2025

Marvell Demonstrates
Industry's Leading 2nm Silicon
for Accelerated Infrastructure

March 3, 2025

Marvell Announces **Breakthrough
Co-Packaged Optics** Architecture
for Custom AI Accelerators

January 6, 2025

Marvell Announces **Breakthrough
Custom HBM** Compute Architecture
to Optimize Cloud AI Accelerators

December 10, 2024

Rapidly expanding pipeline



Engaged in every opportunity and well-positioned for growth



Custom AI Investor Event



Custom AI Investor Event

Q&A



Custom AI Investor Event

Thank you



Custom AI Investor Event

June 17, 2025