Marvell Structera: Overcoming Computing's Memory Challenges with Optimized CXL Devices

Khurram Malik

Director, Product Marketing, Marvell

MEMORY FABRIC FORUM

OCP GLOBAL SUMMIT

OCT 15-17, 2024
SAN JOSE, CA

Forward-looking statements

Except for statements of historical fact, this presentation contains forward-looking statements (within the meaning of the federal securities laws) including statements related to future revenue, future earnings, and the success of our product releases 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.

Forward-looking statements are only predictions and are subject to risks, uncertainties and assumptions that are difficult to predict, including those described in the “Risk Factors” section of our Annual Reports on Form 10-K, Quarterly Reports on Form 10-Q and other documents filed by us from time to time with the SEC. Forward-looking statements speak only as of the date they are made. You are cautioned not to put undue reliance on forward-looking statements, and no person assumes any obligation to update or revise any such forward-looking statements, whether as a result of new information, future events or otherwise.

Data center server market trends

CPU core growth outpacing **memory bandwidth** growth

Constrained CPU pin-count limiting **memory capacity**

Sustainability and CAPEX driving **memory recycling** initiatives

CXL use cases

Expander

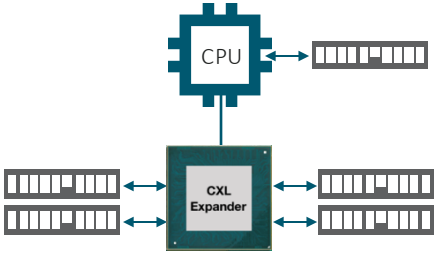

Pooling

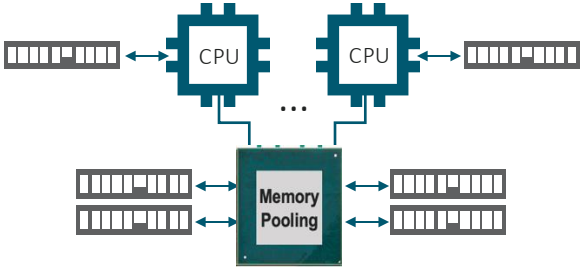

Accelerator

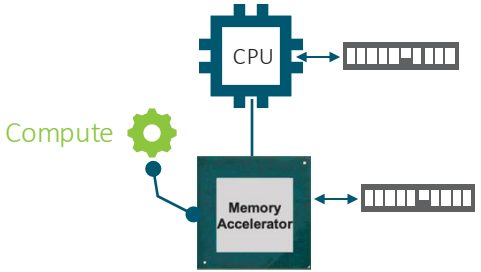

Optimizing memory-intensive workloads using CXL

CXL addresses **memory-intensive applications**

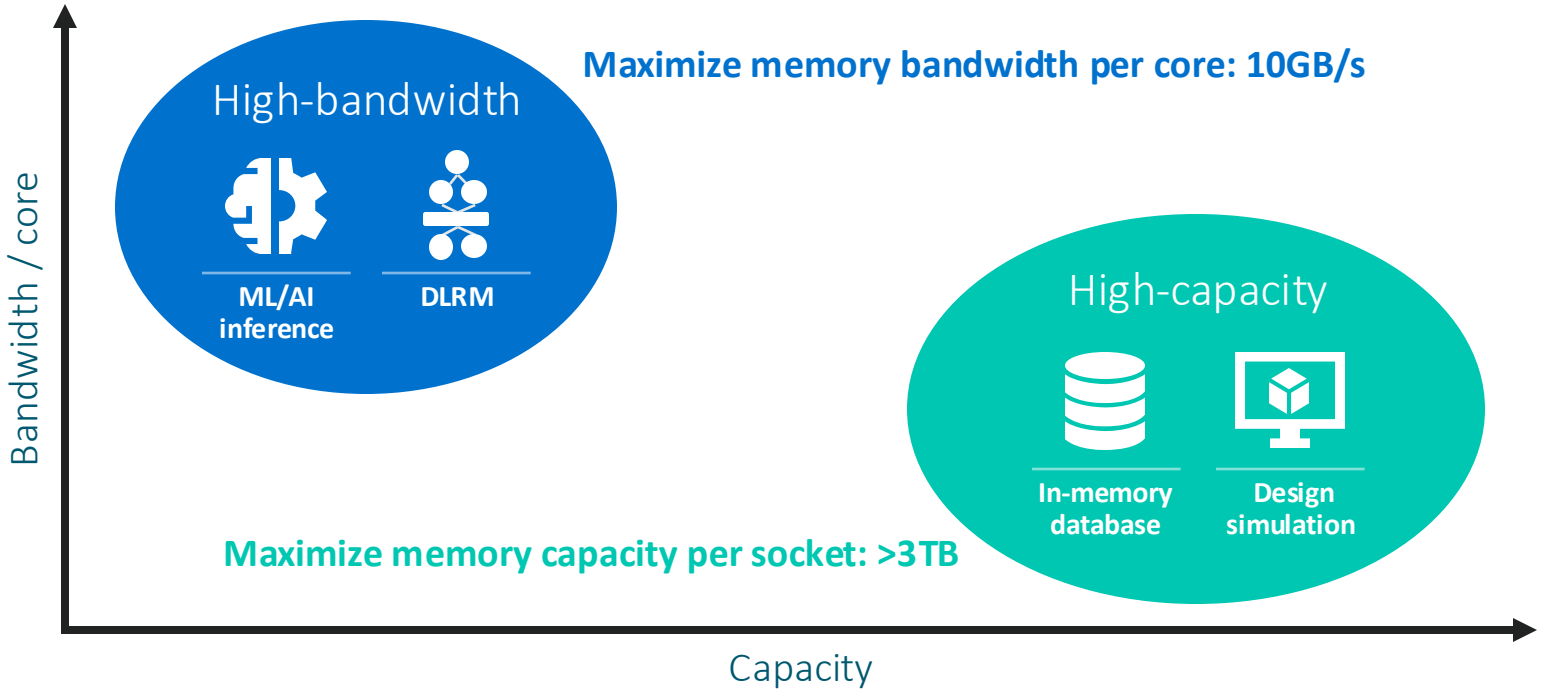

Introducing the Marvell® **Structera™** CXL product line

Structera™ A
Near-memory compute accelerators

Structera™ X
Memory-expansion controllers

Applications

High memory bandwidth

High memory capacity

Structera A 2504:
CXL 2.0 DDR5 4-channel accelerator

- CXL 2.0 / PCIe 5.0 x16 port controller
- 200 GB/s memory bandwidth
  - **4 x DDR5-6400 memory channels**
  - Support for up to two DIMMs per channel
- **16 Arm® Neoverse® V2 (Demeter) cores at 3.2 GHz**
- **Inline LZ4 compression / decompression**
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot
- **Built on industry-leading Marvell 5nm IP**
- Typical power consumption of <100W

Enables optimal compute and memory scaling

Deep learning recommendation model (DLRM) server

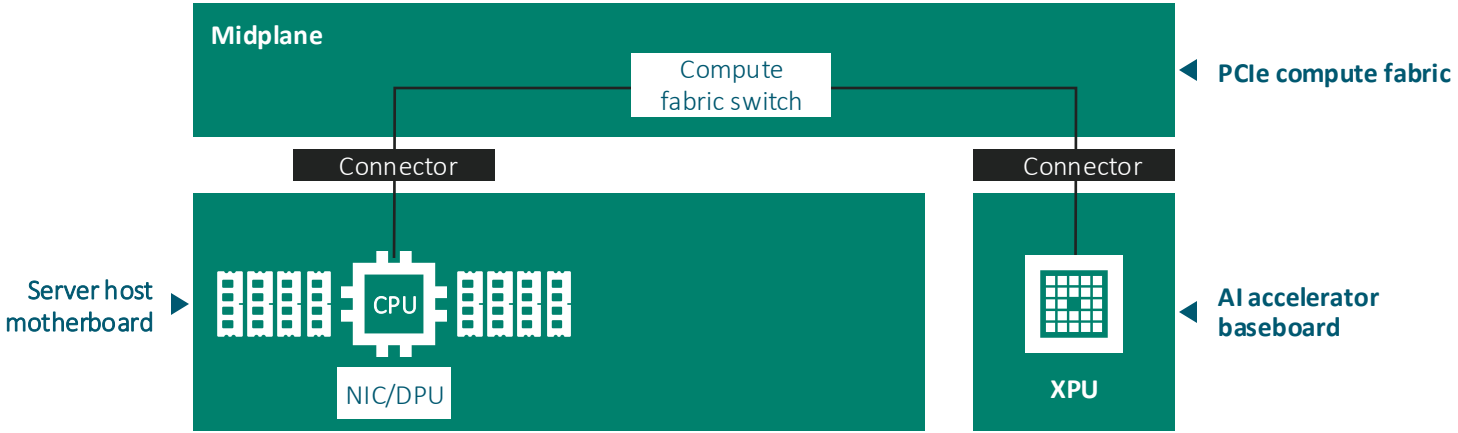

Server has **64 cores** with **400 GB/s of DRAM bandwidth** and consumes **400W power**

6.25 GB/s per CPU core and **1 W per GB/s**

Scaling compute and memory bandwidth per core critical

DLRM server example with **Structera A** CXL accelerator

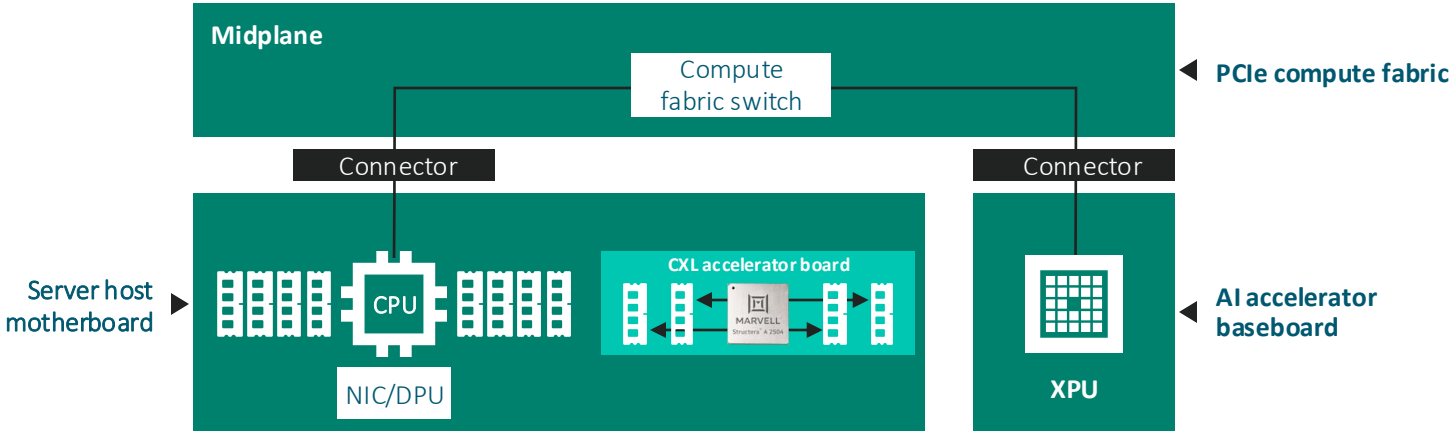

Server has **80 cores** with **600 GB/s of DRAM bandwidth** and consumes **500W power**

7.5 GB/s per CPU core and **0.83W per GB/s**

Increases number of cores by 25% and memory bandwidth by 50%

DLRM server example with
2 Structera A CXL accelerators

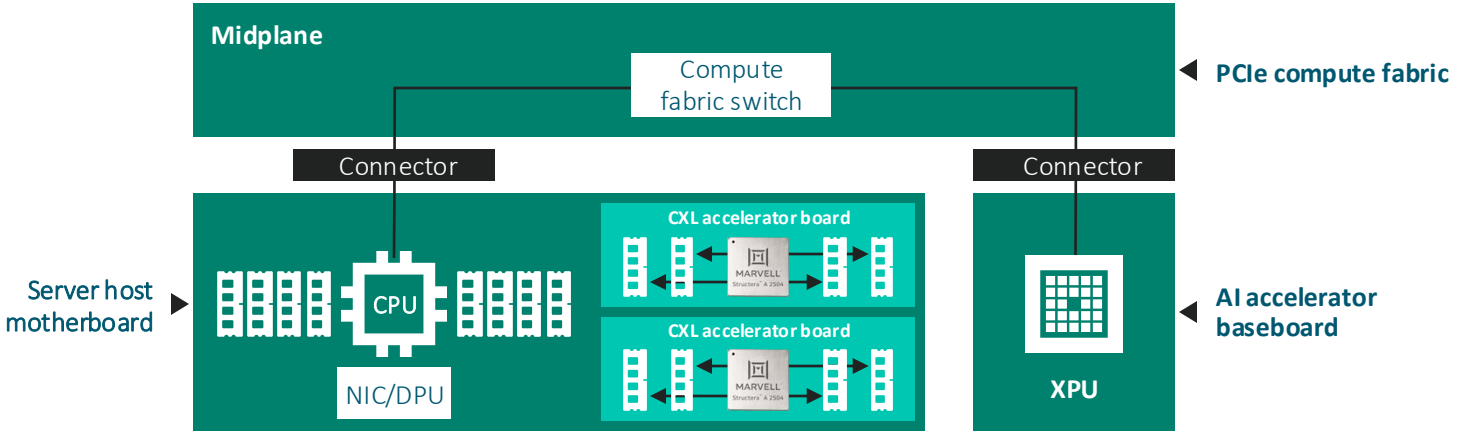

Server has **96 cores** with **800 GB/s of DRAM bandwidth** and consumes **600W power**

8.33 GB/s per CPU core and **0.75W per GB/s**

Increases number of cores by 50% and doubles memory bandwidth

Structera X 2404:
CXL 2.0 DDR4 4-channel expander

- CXL 2.0 / PCIe 5.0 1x16-port or 2x8-port controller
- 100 GB/s memory bandwidth
  - **4 x DDR4-3200 memory channels**
- **Support for up to three DIMMs per channel**
- **Support for >6TB of DRAM memory capacity**
- **Inline LZ4 compression / decompression**
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot
- **Built on industry-leading Marvell 5nm IP**
- Typical power consumption of <30W

Enables recycling of DDR4 DIMMs to increase server memory capacity

Structera X 2404 enables **DDR4 memory recycling**

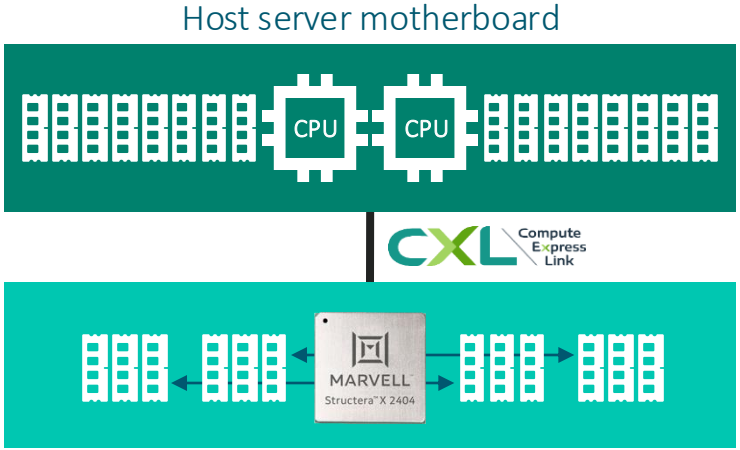

Recycle up to **12 DDR4 DIMMs** per expander (up to 6TB)

Increases server memory capacity with lower CAPEX and reduces e-waste

Structera X 2504:
CXL 2.0 DDR5 4-channel expander

- CXL 2.0 / PCIe 5.0 1x16-port
- 200 GB/s memory bandwidth
  - **4 x DDR5-6400 memory channels**
  - **Support for up to two DIMMs per channel**
- **Support for >4TB of DRAM memory capacity**
- **Inline LZ4 compression / decompression**
- Inline AES-XTS 256-bit encryption and decryption
- Embedded hardware security module and secure boot
- **Built on industry-leading Marvell 5nm IP**
- Typical power consumption of <30W

Enables high-capacity DDR5 memory servers

Structera X 2504 enables
DDR5 memory expansion

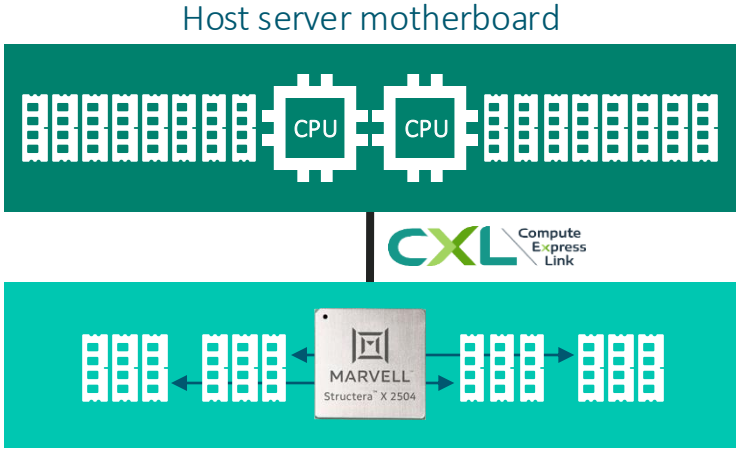

Up to 8 DDR5 DIMMs per expander (up to 4TB)

Increases server memory capacity with high-capacity DDR5 DIMMs

Structera product family overview summary

Structera™ A
Near-memory accelerators

Structera™ X
Memory-expansion controllers

CXL 2.0 / PCIe 5.0

4 x DDR
memory channels

Industry-first features
to enable CXL adoption

Conclusion

1. CXL to address data center server memory bandwidth, capacity and recycling trends
2. New Marvell Structera CXL product line optimized for memory-intensive cloud applications
3. Structera devices are 1st to support 4 memory channels, integrate compression and use 5nm
4. Structera A 2504 is industry's 1st CXL accelerator with Arm Neoverse v2 cores to optim
5. Structera X 2404 is industry's 1st CXL expander to support recycling of 12 DDR4 DIMMs

Thank you!

MEMORY FABRIC FORUM | OCP GLOBAL SUMMIT | OCT 15-17, 2024 SAN JOSE, CA

Open Discussion