

# TORmem's Disaggregated Memory Solutions

Presented by: Scott Burns, VP of R&D

- Founded in 2021
- Primary goal: Enhance existing datacenter servers and enable new software potential for non-specialized gear by providing massive pools of high-performance memory

**Thao Nguyen, CEO**



**Steven White, CTO**






# Why now?

# CPU Interface Speed




Flash Memory Summit

 **I/O BANDWIDTH DOUBLES  
EVERY 3 YEARS**





# Network Interface Speeds





# Recent Innovations and Demand

## Specialized Optimizations

- Memory tiering support in the mainline Linux kernel (v5.15)
- Compute Express Link (CXL)
- Open Memory Interface (OMI)

## Memory Hungry Applications

- In-memory databases
- AI/ML modeling
- Massive simulations
- Realtime analytics



# Difficulties to Overcome



# Design Obstacles


- **Cost**
  - Bill-of-Materials (BOM) cost must be reasonable
- **Density**
  - Maximize amount of memory per rack unit (RU) consumed
- **Management, Monitoring and Maintenance**
  - Health monitoring and system uptime is paramount
- **Integration Flexibility**
  - Compatibility with common existing interfaces, but flexible enough for future ones





# Cost and Density

- A density goal of 16 memory modules per FPGA memory controller was determined to be optimal.
- The FPGA pin count requirement for 16 parallel interface memory modules pushed FPGA pricing to unacceptable levels.
- Serially-connected memory provided by the OMI interface enables the desired memory density while also reducing the cost of the FPGA due to its lower pin count requirements.





# Management, Monitoring and Maintenance

- The M4000 and M5000 chassis designs are designed for rapid in-rack maintenance
  - All active components are field replaceable
  - Front sleds are hot-swap capable
  - Removable rear sled includes the entire backplane PCB
  - Power supplies are hot-swap capable
- Redundant power supplies
- Health Monitoring
  - Memory module health is constantly monitored by the FPGAs
  - Fan speeds are monitored for indications of failure
  - Temperature is monitored through the chassis and at each memory module
  - Power supplies provide PMBus data
  - All monitoring data is provided by the Management Controller in the chassis
- System management is handled by an integrated Management Controller
  - Provides API access to the Memory Controllers



# Integration Flexibility

- TORmem's memory appliance designs provide interface flexibility
  - Interface cards are installed in standard PCIe Gen4 or Gen5 x16 slots
  - Existing network switches are supported
  - Highest performance is with a PCIe fabric
    - TORmem is working on an integrated PCIe switch and a custom fiber connectivity solution for PCIe fabrics
- Planned CXL compatibility
  - CXL and OMI are complementary in TORmem's designs
  - Upcoming CPUs can take advantage of CXL while earlier generation processors use PCIe



# Products



# TORmem's Product Roadmap

- M1000

- Dual AMD EPYC CPUs
- 32x RDIMMs (up to 8TB)

- M4000

- 3x PCIe Gen4 Memory Controllers
- 48x OMI DDIMMs (12TB+)
- 3x PCIe Gen4 x16 for external connectivity
- Integrated PCIe Gen4 100 lane switch

- M5000

- 4x PCIe Gen5 Memory Controllers
- 64x OMI DDIMMs (16TB+)
- 4x PCIe Gen5 x16 for external connectivity
- Integrated PCIe Gen5 144 lane switch (CXL support planned)

# M1000



Flash Memory Summit

- Provides a baseline for benchmarking custom designs
- Up to 8 TB of local memory
- Includes TORMem's memory management software solutions
- Applications can run locally
- Network-based communication



# M4000



Flash Memory Summit

PCIe Gen4 Backplane

External Interfaces

Management Node

Memory Controller Nodes

FPGAs


DDIMMs



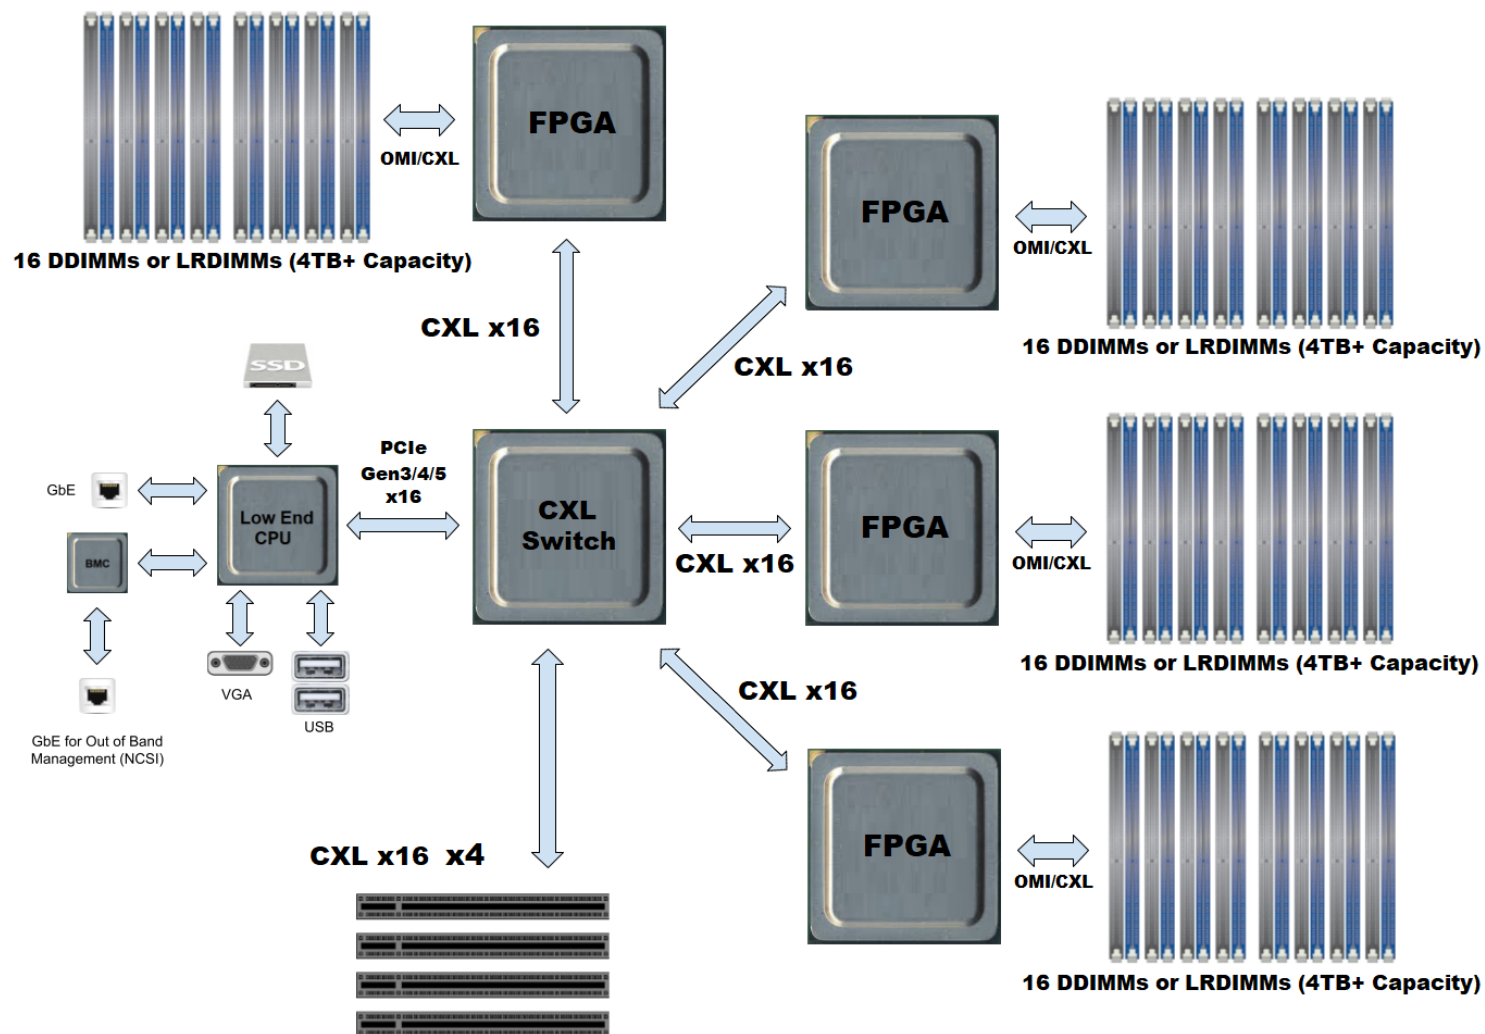
# M4000



Flash Memory Summit







Configurable external interface options include Ethernet, Fibre Channel, PCIe fabric, etc. provided by cards



# Development Efforts

- Collaboration with IBM, AMD and Smart Modular Technology for OMI
  - TORmem is an early adopter of OMI technology
  - TORmem has contributed reference designs to aid in OMI adoption by others
  - Software utilities and hardware designs are open source
  - Experiences are incorporated into TORmem's own designs
- Collaboration with AMD
  - TORmem is actively engaged with the FGPA division of AMD (formerly Xilinx)
  - Early adopter of emerging FPGA technology
  - Design reviews and roadmap planning for new FPGA offerings
- Engineering Support from Samtec
  - Excellent response time to address connector requirements
  - Guidance to preserve signal integrity at extremely high speeds
  - Solutions for integrated optical technology

- Memory is no longer tied to the server's physical requirements
- Older servers can use new memory technology since it isn't directly connected
- Applications normally limited to costly HPC hardware can become feasible for "standard" hardware
- Integrated PCIe switching minimizes the rack units consumed
- Servers can be purchased with small amounts of local memory
- Disaggregated memory provides the final piece for a fully configurable solution



# ***TORMEM***

**One Memory for ALL™**

# Thank you!