



Flash Memory Summit

Top of Rack Shared Memory Pooling Using OMI

Presented by: Steven White, CTO
Scott Burns, VP R&D



Design Goals

- Create a centralized pool of memory
- Maintain compatibility with existing hardware
- Provide flexible connectivity options
- Sever the physical dependencies upon memory module types in servers




Design Obstacles

- **Cost**
 - Bill-of-Materials (BOM) cost must be reasonable
- **Density**
 - Maximize amount of memory per rack unit (RU) consumed
- **Performance**
 - Maintain throughput from memory to FPGA



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.

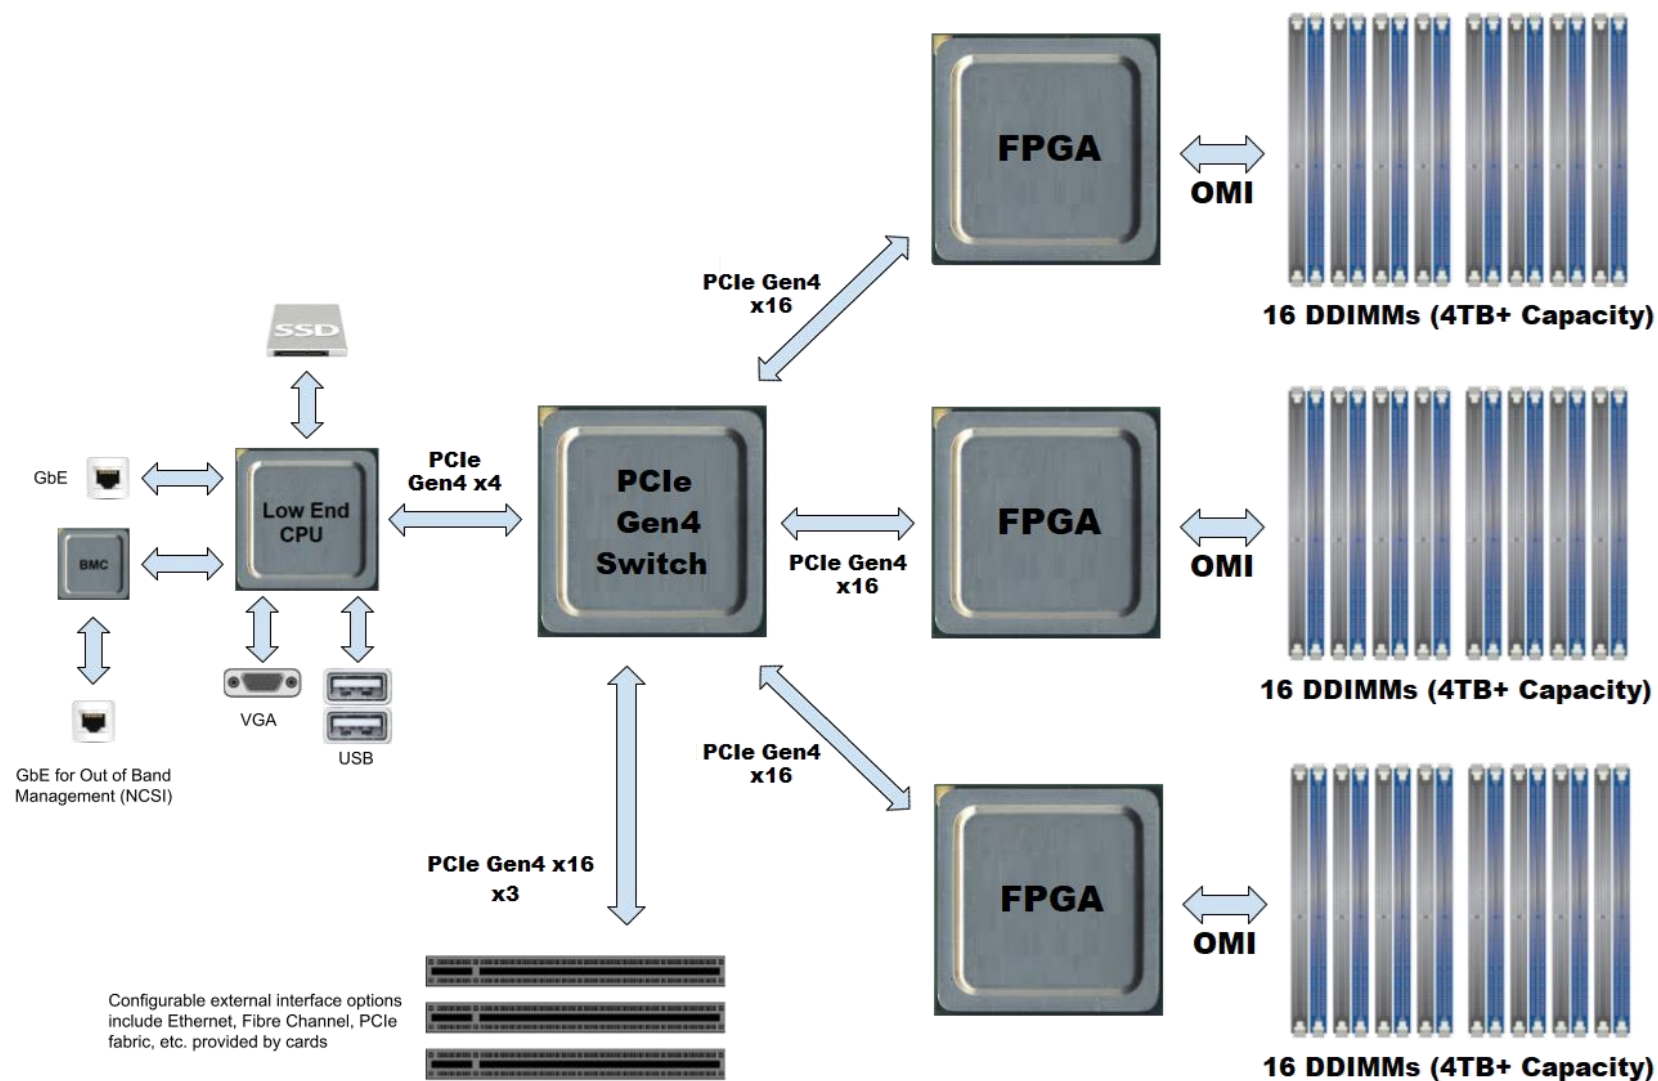


Performance

- At least comparable speeds
- OMI supports multiple lane configurations
 - x1, x4, x8
 - Enables capacity v/s speed flexibility for a given lane count
- Smaller resource requirements to implement

Specification	LRDIMM DDR4	DDR5	HBM2E(8-High)	OMI
Protocol	Parallel	Parallel	Parallel	Serial
Signalling	Single-Ended	Single-Ended	Single-Ended	Differential
I/O Type	Duplex	Duplex	Simplex	Simplex
Paths/Channel (Read/Write)	64	32	512R/512W	8R/8W
Data Transfer Rate	3,200MT/s	6,400MT/s	3,200MT/s	32,000MT/s
Channel Bandwidth (R+W)	25.6Gbytes/s	25.6Gbytes/s	400Gbytes/s	64Gbytes/s
Latency	41.5ns		60.4ns	45.5ns
Channels / Processor Die	8 (EPYC Rome IO)		5 (Nvidia Ampere)	16 (POWER10)
Processor Die Size	416mm ²		826mm ²	602mm ²
Driver Area / Channel	7.8mm ²	3.9mm ²	11.4mm ²	2.2mm ²
Bandwidth / mm ²	3.3GBytes/s/mm ²	6.6GBytes/s/mm ²	35GBytes/s/mm ²	29.6GBytes/s/mm ²
Max Capacity / Channel	64GB	256GB	16GB	256GB
Connection	Multi Drop	Multi Drop	Point-to-Point	Point-to-Point
Data Resilience	Parity	Parity	Parity	CRC

M4000



M4000



Flash Memory Summit

PCIe Gen4 Backplane


External Interfaces

Management Node

Memory Controller Nodes

FPGAs


DDIMMs



M4000




Flash Memory Summit





Path to CXL – M5000

The modular nature of the M5000 design makes it straightforward to upgrade to CXL as components become available. The FPGA in the Memory Controller Nodes supports CXL with a change to its programming and the PCIe Gen5 backplane can be swapped out for a CXL compatible equivalent with no other modifications necessary. The FPGA's OMI interface is complementary to CXL in this application.





TORMEM

ONE MEMORY FOR ALL™

Thank you!