



OMI is Solving the Industry's Need for Bandwidth and Capacity and is now Scaling to DDR5

Pekon Gupta
(SMART Modular Technologies)
Aug 2, 2022

Phases of Successful OMI Adoption



Flash Memory Summit



SAP HANA



MemVerge
Applications




Software



Enablement



- Modifications to BIOS
- OS Device drivers
- Management plugins
- Library / API support



Integrated Development Environment (IDE)

<https://opencapi.github.io/oc-accel-doc/>

Hardware




Integration

- Electrical Connectivity (cables, backplanes)
- Mechanical support (Form-factor, CEM/EDSFF)
- Thermal requirements



SAMSUNG



Development

- Protocol standardization
- Silicon Development
- Multiple suppliers

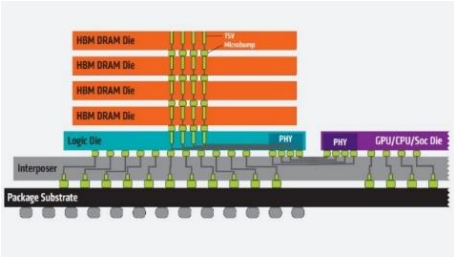
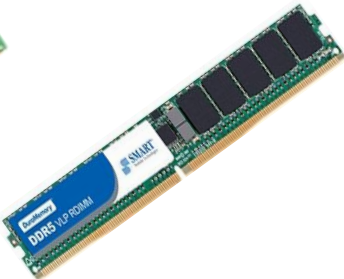
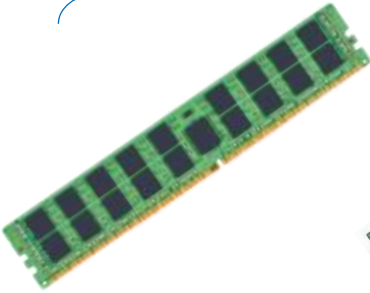


<https://openmemoryinterface.org/omi-products/>





Comparing Memory Interfaces

Parallel Interfaces



Serial Interfaces



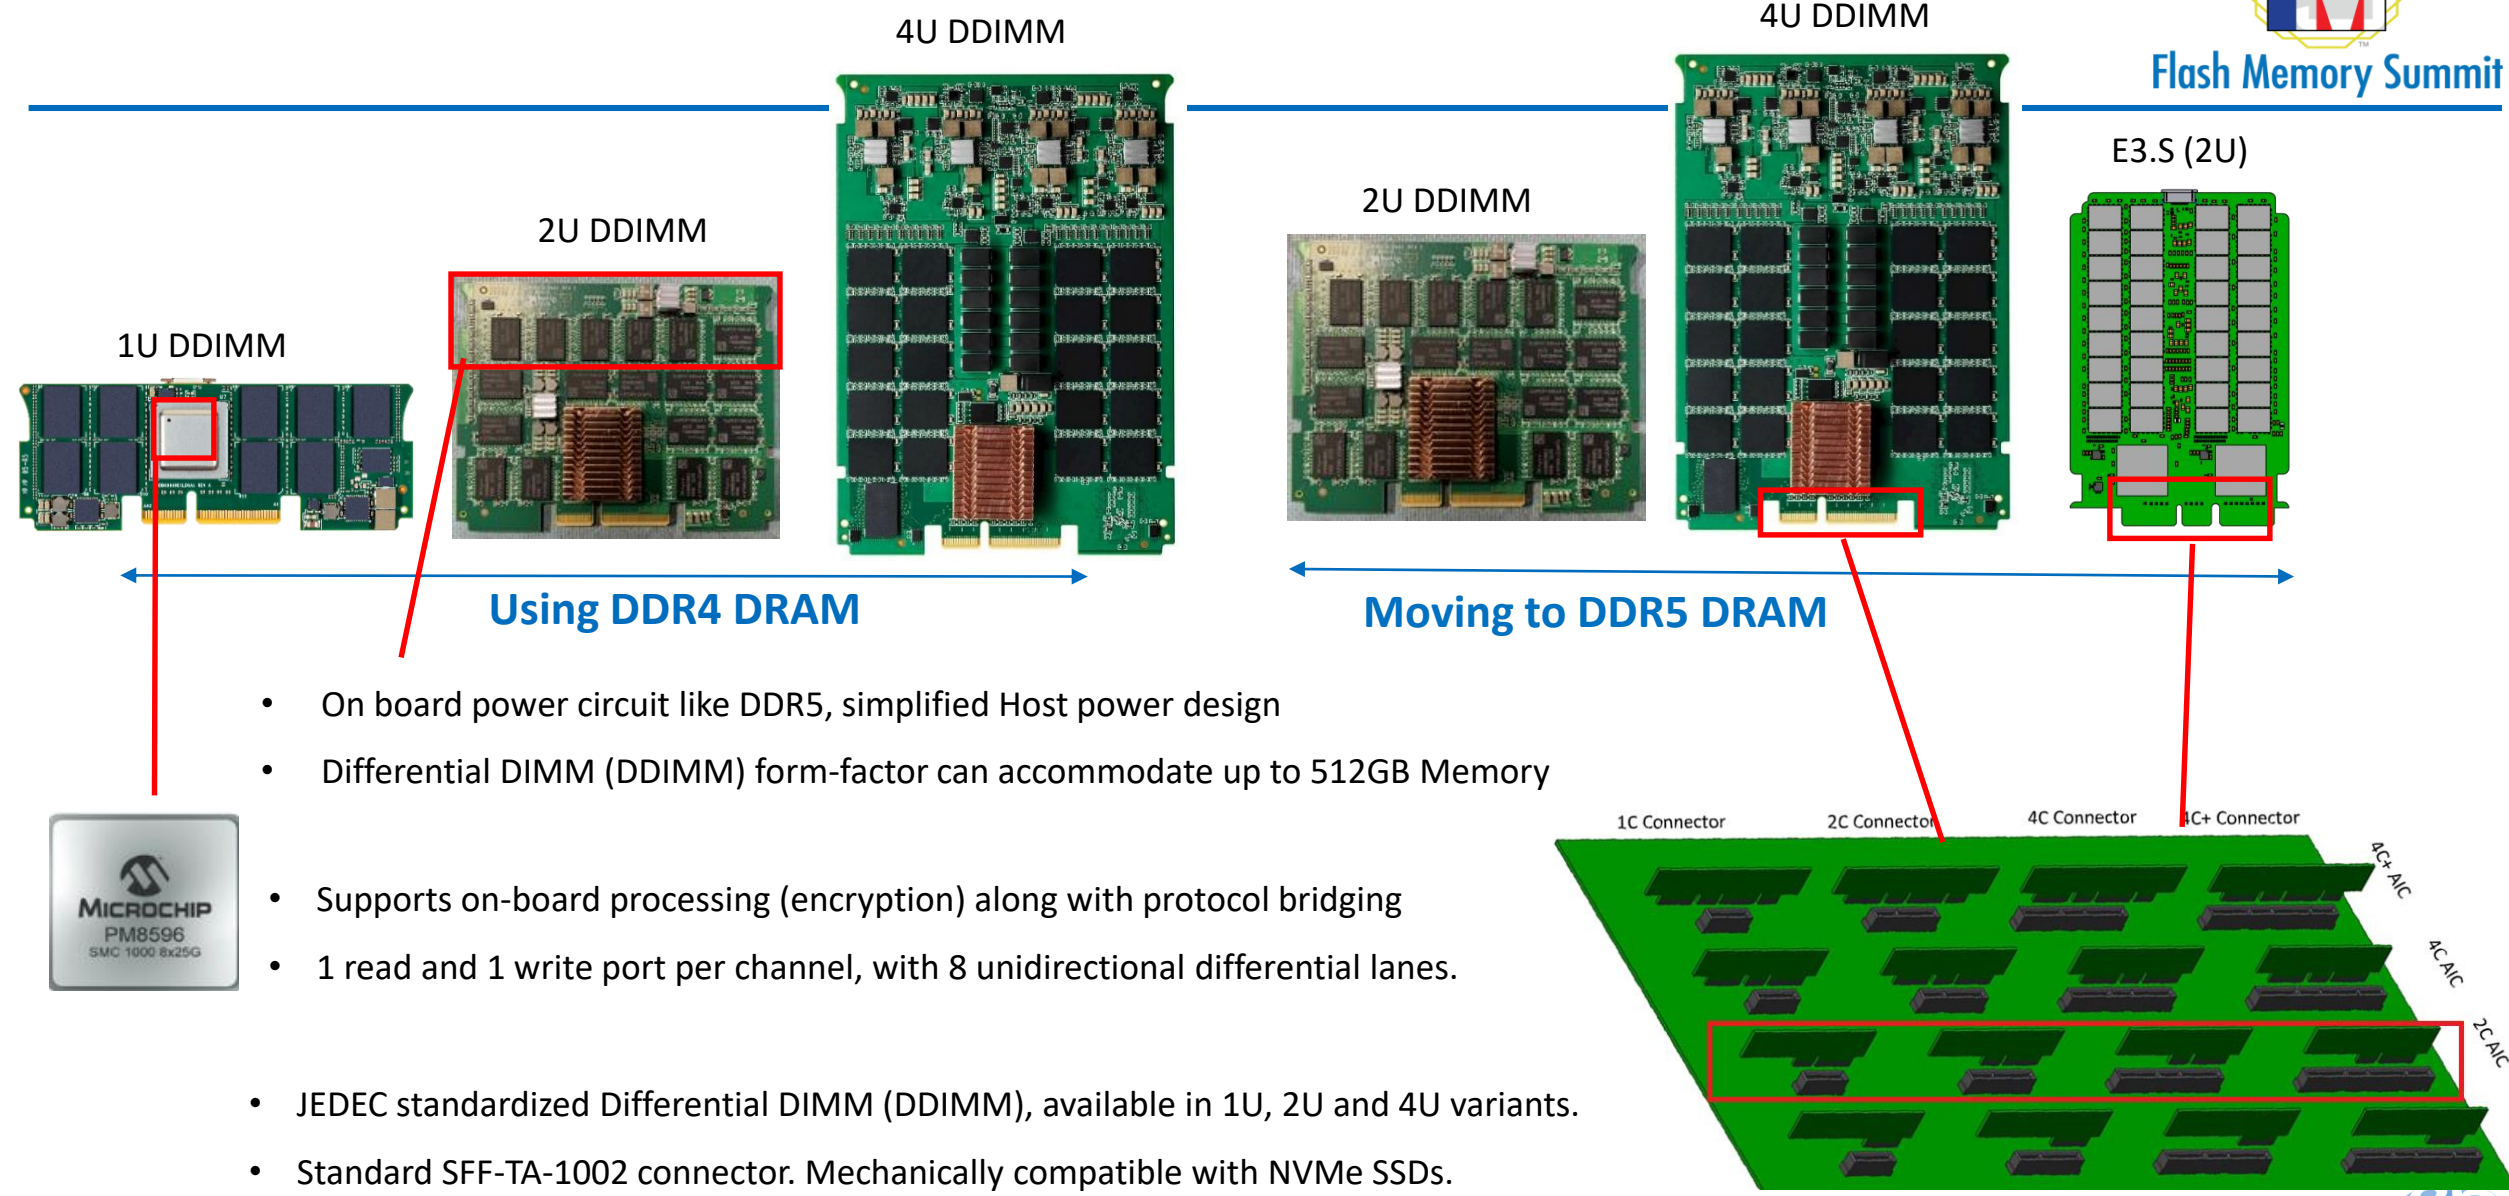
Attributes	DDR4 (LRDIMM)	DDR5 (LRDIMM)	HBM2E (8 High)	OMI (w/ DDR4 DRAM) ^[1]	CXL (w/DDR5 DRAM)
Channel Bandwidth (R+W)	25.6GB/s x (8 Ch)	TBD x (8 Ch)	400GB/s	64GB/s x (8 ch)	TBD x (8 Ch)
Latency	41.5ns	35-45ns	60.4ns	45.5ns	<100n (estimated)
I/O Pin count ^[2]	Parallel (288 pins)	Parallel (288 pins)	Parallel (1024)	Serial (84 pins)	Serial (84 pins)
Driver Area/Channel	7.8mm2	TBD	11.4mm ²	2.2mm2	TBD
Connection	Multi-Drop	Multi-Drop	Point-to-Point	Point-to-Point	Point-to-Point
Data Resilience	Parity	Parity	Parity	CRC	CRC

[1] Sources: (a) White Paper on The [Future Of Low-Latency-Memory by Jim Handy and Tom Coughlin](#) of Objective Analysis, (b) <https://openmemoryinterface.org/about/>
[2] Form-factor dependent

OMI Hardware: Differential DIMM (DDIMM) and E3.S



Flash Memory Summit



Platform Integration: Proof of Low-Latency OMI


IBM's Power-10 Server



IBM's Power-10 E1080 system has multiple DDIMM slots, but “no” DDR4 or DDR5.

- Up to 16TB memory in a single node is enabled using OMI.
- Giving Over 400GB/s bandwidth per socket.
- With Memory Transparent Encryption, and Controller managed ECC/redundancy.

Search for “IBM Z16 Facility Tour!” on YouTube



IBM Power E1080 systems with DDR4 DDIMM Slots

Recommended Demo: [https://www.ibm.com/demos/it-infrastructure/product.html#15/1985;IBM Power Interactive Update/index.html](https://www.ibm.com/demos/it-infrastructure/product.html#15/1985;IBM_Power_Interactive_Update/index.html)

Key Takeaways

OMI is a building block for Memory Disaggregation.

OMI gives ***HBM2E like bandwidth, DDR4 like latency*** at $\frac{1}{4}$ the number of pins than DDR4 DIMM.

OMI decouples interface between Processor and Memory.

Works with DDR5, DDR6 and future Persistent or volatile technologies.

OMI works with existing eco-system of hardware form-factors and software frameworks.

References: <https://www.nextplatform.com/2021/09/08/this-is-what-the-most-powerful-server-in-the-world-looks-like/>
<https://objective-analysis.com/uploads/2021-04-18%20Objective%20Analysis%20White%20Paper%20%E2%80%93%20The%20Future%20of%20Low-Latency%20Memory.pdf>

- ❖ **Did you know: OMI Transaction and Data Link Layer IP code is available for “Free” without any royalty?**
 - ❖ Yes, OMI Transaction and Data Link code for both Host Controller and Memory Buffer IP can be found at
 - ❖ <https://openmemoryinterface.org/reference-designs-enablement/>
 - ❖ IP also has different code versions optimized for both FPGAs and ASIC implementations.


- ❖ **Did you know: AMD, IBM, SMART Modular and other OpenCAPI members are enabling**
 - ❖ FPGA based (Xilinx) PCIe Development Platform
 - ❖ Fully functional OMI Transaction and Data link layers, and Traffic generators
 - ❖ Supports OpenCAPI FMC+ to two OMI DDIMM Sockets Adapter

- ❖ **Did you know: OpenCAPI and OMI have two separate websites ?**

OMI (Open Memory Interface) is low latency optimized version of OpenCAPI.

OMI Website: <https://openmemoryinterface.org/>

OpenCAPI Website: <https://opencapi.org/>



References: <https://www.nextplatform.com/2021/09/08/this-is-what-the-most-powerful-server-in-the-world-looks-like/>

<https://objective-analysis.com/uploads/2021-04-18%20Objective%20Analysis%20White%20Paper%20%E2%80%93%20The%20Future%20of%20Low-Latency%20Memory.pdf>