



MemVerge

OCP 2022 Wrap Up – The Latest and Greatest in SSDs, CXL, Accelerators, & Other Big Trends for Datacenter & Cloud

G2M

RESEARCH

Multi-Vendor Webinar Tuesday Nov 15, 2022

### G2M Research Introduction and Ground Rules

#### Mike Heumann Managing Partner, G2M Research





11/15/2022

Copyright © 2022 G2M Research. All Rights Reserved.

#### Webinar Agenda



- **10:00-10:05** Ground Rules and Webinar Topic Introduction (G2M Research)
- **10:06-10:13** Kioxia Presentation (Sam Bhattarai, Director Prod Mgmt)
- **10:14-10:21** Astera Labs Presentation (Ahmad Danesh, Sr. Dir, Prod Mgmt)
- **10:22-10:29** MemVerge Presentation (Frank Berry, VP of Marketing)
- **10:30-10:42** Panel Discussion Question #1
- **10:43-10:43** Audience Survey #1
- **10:44-10:56** Panel Discussion Question #2
- **10:57-10:57** Audience Survey #2
- **10:58-11:10** Panel Discussion Question #3
- **11:11-11:20** Audience Q&A (10 minutes)

#### The OCP Summit – Impact Beyond Cloud-Scale Deployments and Products?

- As a tradeshow, the OCP Global Summit has become more than just a conference focused on cloud-scale hardware and software
- The 2022 OCP Global Summit had over 60 exhibitors, nearly 2/3 the number of exhibitors at Flash Memory Summit 2022
- The OCP Global Summit has become a show where storage, compute, and networking companies introduce new technology for a variety of use cases.





#### Hot Storage Tech Areas at OCP 2022



- Several vendors showcased their NVMe SSDs
  - These included new drives with greater capacity, higher speeds, and a variety of flash-optimized form factors
  - A number of system-level storage solutions were shown both by vendors and hyperscale companies
- CXL memory and switching solutions were the focus both on the expo floor and in the Day 3 CXL Forum
  - There were a number of vendors showing CXL memory expanders, switches, and management software
  - The CXL Forum included presentations from 18 companies and the CXL exec board



#### **CXL<sup>™</sup> Forum** at OCP Summit

MemVerge

	arm		AsteraLabs		Elastics,	<u> </u>	MARVELL	
🔊 Meta	Micr	on	E Micro	osoft 🕅	MONTAGE			
Rambus	SAMSU	NG	FIELD DA	Uber	<b>vm</b> wa	are	XCONN	









### ΚΙΟΧΙΑ

------

Sam Bhattarai Director, Product Mgmt, Data Center Division, SSD BU <u>www.kioxia.com</u>





Ahmad Danesh Sr. Director, Product Management <u>www.asteralabs.com</u>





Frank Berry Vice President of Marketing www.memverge.com





Mike Heumann Principal Analyst www.g2minc.com



Copyright © 2022 G2M Research. All Rights Reserved.

11/15/2022

#### **KIOXIA SSD Portfolio**



#### **KIOXIA CM SSD Series**

CM7



Enterprise-class NVMe<sup>™</sup> SSD with the highest performance on the market, suitable for servers and storage arrays requiring consistency and high availability

Overview	Features	Benefit
<ul> <li>PCle® 5.0 1x4 or 2x2</li> <li>2.5" U.3 and U.2 support</li> <li>EDSFF E3.S support</li> </ul>	<ul> <li>1/3 DWPD @ 5 years warranty</li> <li>NVMe 1.4, NVMe-MI 1.1</li> <li>SR-IOV, CMB, SGL</li> </ul>	<ul> <li>The highest performance of any TLC SSD today</li> <li>Robust and highly reliable multi- generation platform</li> </ul>
<ul> <li>1600GB to 30720GB drive configurations</li> <li>Power Loss Protection</li> </ul>	<ul><li>included as standard</li><li>Dual-port support</li><li>SIE, SED and FIPS</li></ul>	<ul> <li>Advanced feature sets make this product ideal for cutting edge use cases</li> </ul>

#### **KIOXIA CD SSD Series**

CD8



Datacenter-class NVMe SSD leveraged from KIOXIA's robust enterprise architecture with an excellent performance to power ratio, suitable for servers and datacenters requiring read-intensive performance, and low latency

Overview	Features	Benefit
<ul> <li>PCIe 4.0x4, 5.0x4 support</li> <li>2.5-inch form factor, supporting U.2</li> </ul>	<ul> <li>1 &amp; 3 DWPD @ 5 years warranty</li> <li>Very high random write performance</li> </ul>	One of the best performing single port drives on the market
<ul> <li>800GB to 15360GB drive configurations</li> <li>Power Loss Protection</li> </ul>	<ul> <li>Leverages KIOXIA's enterprise architecture</li> <li>SIE, SED options to safely encrypt data</li> </ul>	<ul> <li>Performance comparable to PCIe 4.0 generation Enterprise NVMe SSDs</li> <li>Robust and highly reliable multi- generation platform</li> </ul>

KIOXIA

#### **KIOXIA XD SSD Series**

## XD7P



Datacenter-class NVMe SSD leveraged from KIOXIA's robust enterprise architecture with a balance of performance to power ratio, designed to optimize storage density, suitable for servers and datacenters requiring read-intensive, low latency performance

Overview	Features	Benefit
PCIe 4.0x4 support	<ul> <li>1 DWPD @ 5 years warranty</li> </ul>	<ul> <li>One of the best performing single port</li> </ul>
E1.S form factor – 9.5mm/ 15mm options	OCP v2.0 support	drives on the market
• 1920GB to 7680GB	Nearly 1.5x sequential write and 2.0x random	<ul> <li>Excellent performance, QoS and Reliability</li> </ul>
BICS 112L Flash	read/write performance	Proprietary KIOXIA     architecture
Power Loss Protection	<ul> <li>SED / non-SED options to safely encrypt data</li> </ul>	

#### ΚΙΟΧΙΑ

#### Combination of CXL<sup>™</sup> & FLASH will offer increased capacity with competitive performance



KIOXIA <u>https://www.youtube.com/watch?v=-POD\_pBV3Xk&list=PLsf8NUp2sz\_iqrlzDVOuwl1Z47MlojoX&index=20</u>

#### **KIOXIA Views on Flash Native Memory Hierarchy**



- ✓ BiCS FLASH<sup>™</sup> as Capacity Storage
   ✓ Leading GB density
  - XL-FLASH<sup>™</sup> as Storage Class Memory
     ✓ Low latency & high endurance flash
- ✓ FLASH for Memory Expansion with CXL™
   ✓ Cost efficient "capacity memory"

#### Flash Memory continues to expand its applications

#### **XL-FLASH<sup>™</sup>** in the Memory Tiers

XL-FLASH will cover from CXL<sup>™</sup> memory tier to high performance SSD. KIOXIA is designing XL-FLASH based CXL memory with SRAM/DRAM tiers.



#### KIOXIA

#### What can be done with XL-FLASH<sup>™</sup>?

Application performance with XL-FLASH would be as fast as DRAM in **parallel random read** even in the most random-read-heavy algorithms such as graph-processing.







Connectivity Solutions for Intelligent Systems

#### Taking CXL from Promise to Reality

Ahmad Danesh Sr. Director, Product Management

#### **Leo Memory Connectivity Platform**

Purpose-built Memory Expansion/Pooling Controllers and Solutions for Cloud Servers

#### Leo Smart Memory Controllers

Leo E-Series



✓ Memory Expansion



- ✓ Memory Pooling
   ✓ Memory Sharing
  - ✓ Memory Expansion

#### **Aurora A-Series Hardware Solutions**



#### Leo Memory Connectivity Platform

CXL-attached memory expansion and pooling for cloud servers



#### In Pre-Production Sampling Phase



#### Leading the Server Industry with CXL Solutions

Unleashing AI/ML Performance and Reducing TCO for Cloud Computing





#### **Taking CXL from Promise to Reality**

Demonstrated CXL Memory Expansion and industry's first CXL Memory Pooling at OCP Summit











1111 1111 1111

### The Dawn of CXL<sup>™</sup>

Frank Berry Vice-President of Marketing MemVerge

All and and a start where a share a share the second the start where the second and the second and the second as a second as a



#### The Past 50 Years





#### MemVerge: Pioneering Big Memory Software

Founders created products with \$B run-rates



## XtremIO

Built Memory Machine™: World's 1st Big Memory software



#### **Memory Machine customers**





### **Memory Machine**<sup>™</sup>

Big Memory software that provides a pool of tiered, composable memory with storagelike in-memory data services.



#### Accelerates Time-to-Discovery By eliminating IO to storage













### **Memory Machine™ Cloud Edition**



System & Cloud Orchestration Service - integrated with cloud automation and scheduler services so that applications can auto-recover from Spot preemptions.

#### Allows long-running apps to run on low-cost Spot instances

Auto Save On Frequent checkpointing **Spot Reclaimed** 

Coordinate with cloud services and schedulers Restore AppCapsule on new Spot



### The Dawn of CXL<sup>™</sup>

### Big Memory Market Dynamics

With the work work and the added and the and the work work the work the

MemVerge

HE AND A RELE

### Growth of Data in Memory Led by Al

**2,600,000x** Growth of EDA data since 1984

> Cerberus WSE-2 The Largest Chip Ever Built

46,225	mm <sup>2</sup> silicon
2.6	Trillion transistors
850,000	Al optimized cores
40	Gigabytes on chip memory
20	Petabyte/s memory bandwidth
220	Petabit/s fabric bandwidth
1.2	Terabit/s ingest bandwidth
7nm	Process technology at TSMC

**1,000x** Greater cell data since 2009



#### 1,000x

Larger AI model data in last 2 years



Source: VentureBeat

Source: Analytical Biosciences

Source: VentureBeat



### Memory is worth its weight in gold

#### Oz. of Gold

~\$2,000



~\$2,000







### **Memory Biggest Percent of Compute Cost**

### 52.0% at 576GB

\$14,985 for DRAM = \$28,787 total price

#### 70.3% at 1.1TB

#### \$29,970 for DRAM = \$42,631 total price



Compelling performance and high scalability Enterprise server powered by 3rd generation Intel<sup>®</sup> Xeon<sup>®</sup> Scalable processors designed to optimize workload performance and data center density. Starting at \$8,618.99 Dell Business Credit As low as \$864 /mo.\* | Apply for credit End Of Year Savings: 53% Off

Add to Cart: \$28,787.26

Order Code pe\_r650\_tm\_vi\_vpsb

PowerEdge R650 Rack Server

8GB RDIMM, 3200MT/s, Single Rank	Included in price
Qty - 8 + \$276.65 /ea.	
16GB RDIMM, 3200MT/s, Dual Rank	\$412.86 /ea.
32GB RDIMM, 3200MT/s, Dual Rank, 16Gb BASE x8	\$779.23 /ea.
64GB RDIMM, 3200MT/s, Dual Rank	Included in price
Oty - 8 + \$1,596,50 /ea.	

https://www.dell.com/en-us/work/shop/servers-storage-and-networking/poweredge-r650-rack-server/spd/poweredge-r650/pe\_r650\_tm\_vi\_vpsb



#### Help Me Choose

 8GB RDIMM, 3200MT/s, Single Rank
 16 + \$276.65 /ea.
 16GB RDIMM, 3200MT/s, Dual Rank
 32GB RDIMM, 3200MT/s, Dual Rank, 16Gb BASE x8
 \$779.23 /ea.
 64GB RDIMM, 3200MT/s, Dual Rank
 Included in price Qty - 16 + \$1,596.50 /ea.

PowerEdge	R650	Rack	Server	

#### Compelling performance and high scalability

Enterprise server powered by 3rd generation  $Intel^{\odot}$  Xeon $^{\odot}$  Scalable processors designed to optimize workload performance and data center density.

Starting at \$8,618.99 Dell Business Credit

As low as \$864 /mo.^ | Apply for credit End Of Year Savings: 53% Off

View Full Specs
Order Code pe_r650_tm_vi_vpsb

https://www.dell.com/en-us/work/shop/servers-storage-and-networking/poweredge-r650-rack-server/spd/poweredge-r650/pe\_r650\_tm\_vi\_vpsb



### **Memory Dedicated to Processors & Apps**





### Adding DIMMs May not Add Bandwidth

(1) 128GB DIMM per Channel

(2) 128GB DIMMs per Channel





### **Result: Memory is Stranded in the Cloud**

- •Up to 25% is *stranded* (no free CPU cores but unsold memory)
- Customers also
   overprovision memory
- Median VM: **45%** untouched memory





### **Result: Memory is Stranded in VMware**

ALCONTRACTOR AND ALCONTRACTOR	Monitor Configure Permissions	VMs Datastores Networks	Update Manager		
Issues Performance Task	s & Events Hardware Status				
44 Overview	The graph refreshes every 20 s	econds.	Chart Options		View: Default 🔻 C .
Advanced	.50000000 [			426GB C	onsumed
	40000000				
	30000000 -				
	G				
				-	
	10000000			31GB Act	ive (7%)
			~		
	0 1.36 PM 1.41 Performance Chart Legend	PM 1:48 PM 1:51 PM	4 1.56 PM 2.01 PM 2.06 Tim	PM 2:11 PM 2:16 PM 2: 9	21 PM 2:28 PM 2:31 PM
	0 1:36 PM 1:41 Performance Chart Legend Key Object	PM 1:46 PM 1:51 Ph Measurement	A 1.56 PM 2:01 PM 2:06 Tim Rollup	PM 2:11 PM 2:16 PM 2:1 PM Units Latest	21 PM 2:28 PM 2:31 PM Maximum Minimum Average
	0 1:36 PM 1:41 Performance Chart Legend Key Object	PM 1:46 PM 1:51 Ph Measurement Active	A 1.56 PM 2:01 PM 2:06 Tim Rollup Average	PM 2:11 PM 2:18 PM 2:1 ie Units Latest KB 31219344	21 PM 2:28 PM 2:31 PM Maximum Minimum Average 4 1034468 27302728 33058
	0 1:36 PM 1:41 Performance Chart Legend Key Object	PM 1:46 PM 1:51 Ph Measurement Active Ballooned memory	A 1:56 PM 2:01 PM 2:06 Tim Rollup Average Average	PM 2:11 PM 2:10 PM 2:1 e Units Latest KB 31219344 KB 0	Maximum Minimum Average 4 41034468 27302728 33058 0 0 0

Results may vary

ılı.

Data from VMware customer production environments tends to shows a lot of stranded memory



Active memory < 25-40% of consumed/physical memory is good fit for memory tiering

Source: VMware



#### **Result: Gap in Memory Bandwidth per Core**



MemVerge

### The Dawn of CXL<sup>™</sup>

### CXL 101

We have a total where we all the well the well the best of the well and the well an



### **CXL: The Backbone of Future Servers**

Internal and External memory expansion





### **The Storage Analogy**



MemVerge

### **Broad CXL Ecosystem**



MemVerge

### **CXL 1.1 Memory Expansion on PCIe**

#### **Motherboard**

#### Expanding memory using PCIe slots





#### Expanding memory using DIMM slots





#### Memory Expansion Overcomes the DDR Bandwidth Bottleneck

Bandwidth for pool of DDR & CXL memory scales beyond the capabilities of DDR



DDR5 Latency: 108ns, CXL Latency: 272ns

MemVerge

### **CXL 2.0 Memory Pooling**

**Expansion Cards** 

**Memory Pooling Appliance** 



### **Tiered, Pooled Memory Reduces Cost**

#### 2TB DRAM

#### 2.5TB Memory (DRAM + Tier 2)

Host cor	figuration DRAM only	List Price <sup>1</sup>	Host conf	iguration DRAM + PMem	List Price <sup>1</sup>
CPU	2 x 6348 (28C @ 2.6GHz)	\$6,144*	CPU	2 x 6348 (28C @ 2.6GHz)	\$6,144*
DRAM	2048GB (32 x 64GB)	\$58,368	DRAM	512GB (16 x 32GB)	\$14,592
Slower Memory	-	-	Slower Memory	2048GB (16 X 128GB)	\$17,121
Storage capacity	15 X 3.84TB <u>vSAS</u> (57TB)	\$27,000	Storage capacity	15 X 3.84TB <u>vSAS</u> (57TB)	\$27,000
Cache drives	3 X P5800X 400GB	\$6,780*	Cache drives	3 X P5800X 400GB	\$6,780*
Software	vSAN + vSphere ENT+	\$28,563	Software	vSAN + vSphere ENT+	\$28,563
Total		\$126,855*	Total		\$100,200*

#### 20% Platform level Cost Reduction

\*Estimated

1 https://www.dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r750/pe r750 14794 vi vp?configurationid=b605e5ac-c8b9-4578-b0e2-7d9b15772b04 Your costs and results may vary

MemVerge

Source: VMware

### **CXL 3.0 Switched Fabric**











### **Peta-Scale Memory**





### **Global CXL Fabric Management**



🔲 MemVerge Memory Machine 🗙 👹 Memory Viewer 🛛 🗙 🕂	
← → C ☆ ① localhost:18080/#/	0 🖈 🏞 🗖 🥥 ፤
MemoryViewer FEDORA	Version v1.0.0 () Memory Viewer
System Topology 🔲 Process Monitor	
Overall	Bandwidth CPU 소
Total Memory Capacity: 125.63 GB	DRAM Bandwidth Real Time 🗸
Host DIMM Memory B 7 6 5 4 3 2 1 Host DIMM Memory 9 10 11 12 13 14 15 1 Host DIMM Memory	Co CEN So CEN
CPU 1 EVALUATION CONTROL OF A	CXL Memory Bandwidth Real Time T7 GB/6 14 GB/6 14 GB/6 15 GB/6 5.6 GB/6 2.8 GB/6 2.8 GB/6
	00.22 00.23 00.24 00.25 00.26

CXL Memory 04

Capacity: 96.00 GB

CXL Memory 03

Capacity: 64.00 GB

CXL Memory 02

Capacity: 512.00 GB

CXL Memory 01

Capacity: 512.00 GB

#### SE02.ENG.MEMVERGE.LOCAL App Direct Mo

System Topology Process Monitor

All | Insight Groups: 🖓 Top 10 Memory Consumers 🖓 KVM

3	Multi-selection											Dura	tion: 1 hour				Stop
	Name	\$ Q	PID	٩	DRAM Usage	÷ (	CPU 🗘		itart Time	÷ •	User	۹	Monitoring Status	T F	Report		
	gnome-shell		3512		133.88 MB	(	0.09%	C	9:01:05 May 18, 2022		gdm						•
	mmetl		2052		91.86 MB	(	0.02%	C	9:01:02 May 18, 2022		mvmm		📕 1.00 hours		L (42 H	(B)	0
	grafana-server		5469		81.23 MB	(	0.24%	C	9:06:09 May 18, 2022				1.00 hours		L (45)	(B)	0
	prometheus		117975		66.18 MB	¢	0.12%	2	0:47:07 May 18, 2022		nfsnobody		1.00 hours		L (43 H	(B)	0
	mmagent		2055		61.09 MB	(	0.12%	C	9:01:02 May 18, 2022		mvmm						
	dockerd-current		5032		52.54 MB	(	0.25%	C	9:05:25 May 18, 2022		root						
	mvmv		78301		49.88 MB	1	2.27%	1	4:00:38 Jul 25, 2022		root						
	Xorg		3375		33.34 MB	(	)%	C	9:01:04 May 18, 2022		root						
	gsd-color		3698		31.45 MB	(	0.22%	C	9:01:06 May 18, 2022		gdm		1.00 hours		L (44 )	(B)	•
	docker-containerd-current		5053		31.02 MB		0.12%	c	9:05:25 May 18, 2022		root						

Process Monitor - gnome-shell (3512)

Started all 10.392.36 Jul 06, 2022 Ended all 11.392.26 Jul 06, 2022

#### 👌 Application Memory Heatmap

級 Insight G

#### Join the CXL Forum Community

- Receive news about technology, products, and CXL Forum events
- Attend CXL Forum events
- Get early access to CXL hardware and software for PoC
- Participate in an advisory board and help us shape the future of CXL software
- https://memverge.com/cxl-forum-community-news-subscribe/

## arm

## AMDA



#### Access over 50 CXL presentations <u>CXL Forum at FMS</u> <u>CXL Forum on Wall Street</u> <u>CXL Forum at OCP Summit</u>

#### Get an Endless Memory t-shirt

#### Send:

Name Title Organization Mail to address Size

#### To:

Frank.berry@memverge.com

Sorry, US only.



### The Dawn of CXL<sup>™</sup>

### Thank You

- Allow and the work work and and and the and the and the work and the allow



## Panel Questions and Audience Surveys

#### Audience Survey Question #1



- We are only buying NVMe flash storage devices going forward: 18%
- We buy mostly NVMe storage devices, and only use SAS/SATA devices for cost-conscious storage (nearline, cold storage, etc.): 45%
- We buy mostly SAS/SATA storage devices, and only use NVMe where it is needed for its performance:
- We only buy SAS/SATA-based flash storage devices: 9%
- We let our storage solution vendors choose what they think is best: 0%
- Don't know/no opinion: 27%

0%



The OCP Summit is specifically focused on cloud-scale technology deployment, but more and more vendors see it as a general-purpose tradeshow to introduce products and reach their target audiences. What lessons can the larger enterprise datacenter ecosystem learn from the OCP Summit?

- Sam Bhattaria- Kioxia
- Frank Berry MemVerge
- Ahmad Danesh Astera Labs

#### Audience Survey Question #2



What do you believe is the approach for your company to deploy CXLbased memory infrastructure and networking (pick all that apply):

- We are planning on deploying CXL networked memory in 2023: 0%
- We are planning on deploying CXL local expansion memory in 2023: 0%
- We plan to start evaluating CXL networked memory in 2023: 22%
- We plan to start evaluating CXL local memory expansion in 2023: 22%
- We are waiting for more clarity (product intros, OEM server and product timelines, software support) before starting CXL evaluations: 22%
- We currently have no plans to evaluate CXL in the near future: 11%
- Our compute infrastructure is cloud-based, so we don't care: 0%
- No opinion/don't know:

56%

#### Panel Question #2



CXL brings the promise of disaggregated, software-defined memory to computing, similar to how NVMe and NVMe-oF did the same thing for flash storage five years ago. Do you think that CXL will push changes in how NVMe storage is deployed in the datacenter and in the cloud?

- Frank Berry MemVerge
- Ahmad Danesh Astera Labs
- Sam Bhattaria- Kioxia

#### Audience Survey Question #3



What do you see as the biggest hurdles to the large-scale deployment of CXL-based memory solutions (pick all that apply):

• I don't understand the use cases where CXL-based memory could help: 0%

<ul> <li>Lack of technology knowledge about CXL:</li> </ul>	10%
<ul> <li>Pack of product knowledge about CXL:</li> </ul>	20%
<ul> <li>Lack of support/active participation from server vendors:</li> </ul>	0%
<ul> <li>Lack of support/active participation from "big memory" app vendors:</li> </ul>	20%
<ul> <li>Lack of management frameworks for CXL:</li> </ul>	20%
<ul> <li>Cloud-based computing ("it's not really my problem"):</li> </ul>	10%
<ul> <li>Don't know/no opinion:</li> </ul>	40%



With CXL and NVMe being the dominant storage/memory connects in the datacenter, and with them both running over PCI Express, how do you see these standards merging and/or becoming complimentary over time?

- Ahmad Danesh Astera Labs
- Sam Bhattarai Kioxia
- Frank Berry MemVerge



# Audience Q&A





#### Effective Marketing & Communications with Quantifiable Results