

A Platform for CXL Memory Services



Bernie Wu
VP Bus Dev & Alliances
MemVerge

EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA



MemVerge: Pioneering Big Memory Software since 2017



**Built World's 1st
Software-Defined
Memory Product:**

Memory Machine

**Founders were creators
of \$B product lines**



Raised \$93M of Investment



Memory Machine Customers



Tencent Cloud



Alibaba Cloud



CXL Road Map: CXL 3.0



CXL 3.0 Spec Feature Summary CXL Compute Express Link

Paper Release ! →

Features	CXL 1.0 / 1.1	CXL 2.0	CXL 3.0
Release date	2019	2020	1H 2022
Max link rate	32GTs	32GTs	64GTs
Flit 68 byte (up to 32 GTs)	✓	✓	✓
Flit 256 byte (up to 64 GTs)			✓
Type 1, Type 2 and Type 3 Devices	✓	✓	✓
Memory Pooling w/ MLDs		✓	✓
Global Persistent Flush		✓	✓
CXL IDE		✓	✓
Switching (Single-level)		✓	✓
Switching (Multi-level)			✓
Direct memory access for peer-to-peer			✓
Enhanced coherency (256 byte flit)			✓
Memory sharing (256 byte flit)			2025+
Multiple Type 1/Type 2 devices per root port			
Fabric capabilities (256 byte flit)			✓

2H23
To
1H24

1H24
To
2H24

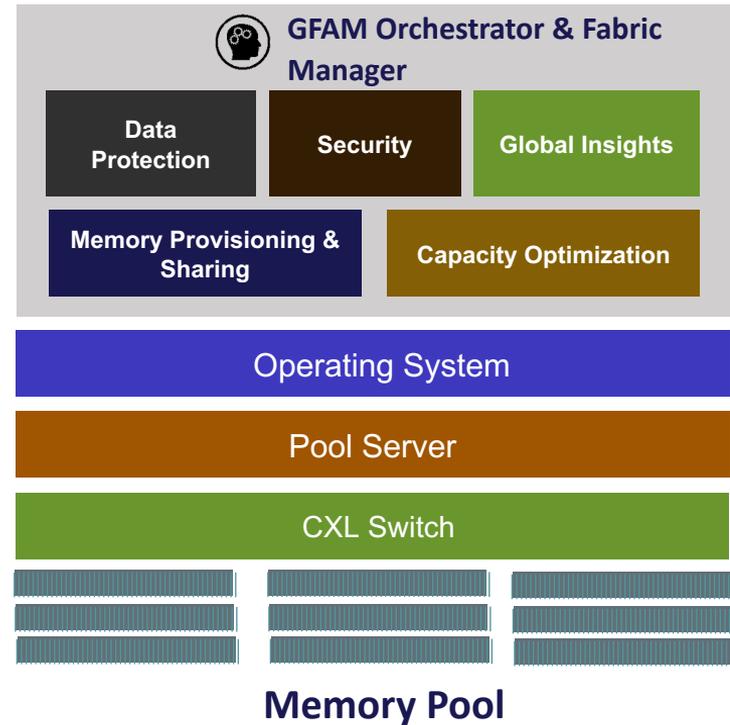
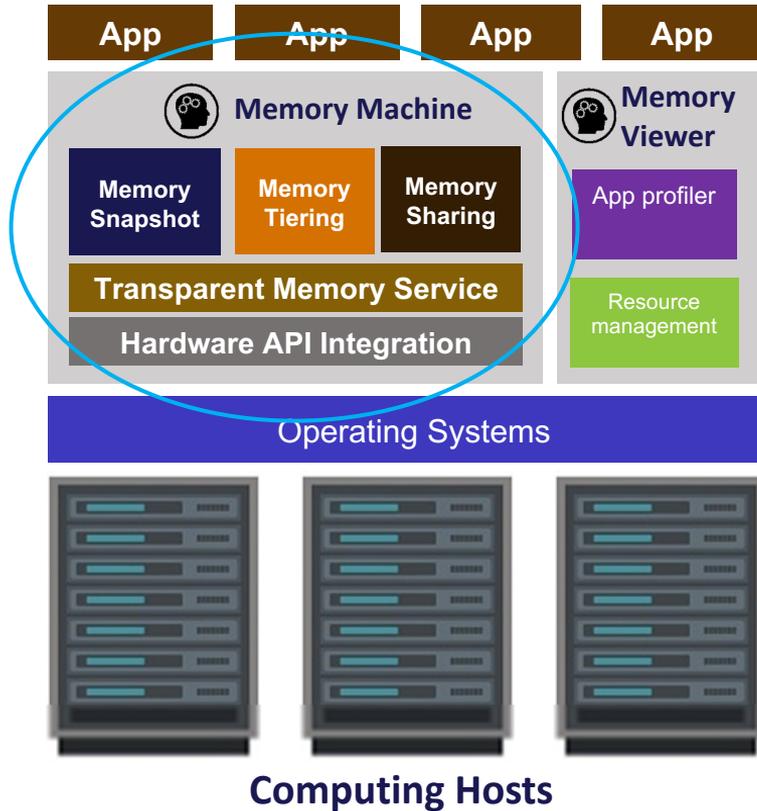
2025+

Not supported

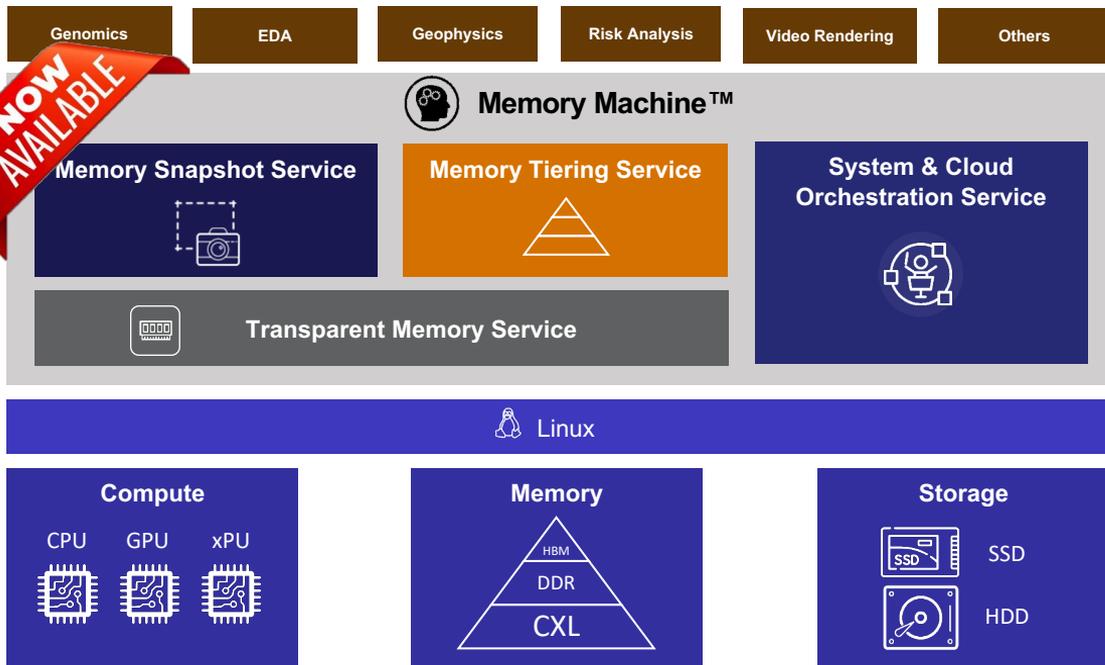
✓ Supported

Market Availability

The ❤️ of Big Memory in the CXL era is Software



Memory Machine Available Now



Memory Capacity Expansion

- Software-defined Memory Pool with intelligent Auto-tiering
- No application change required

Accelerate Time-to-discovery

- Transparent checkpointing
- Roll-back, restore and clone anywhere any time

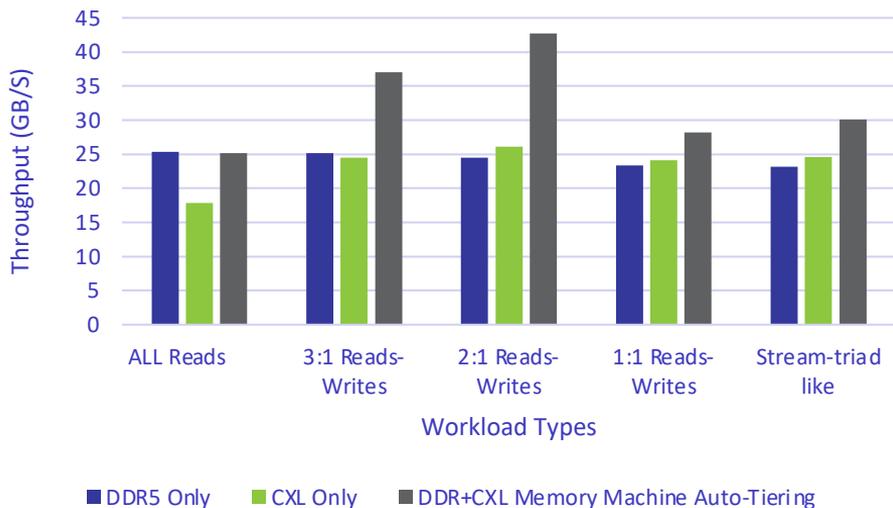
Reduce Cloud Cost by up to 70%

- Enable long-running applications to use low-cost Spot instances
- Integration with cloud automation and scheduler to auto-recover from CSP preemptions

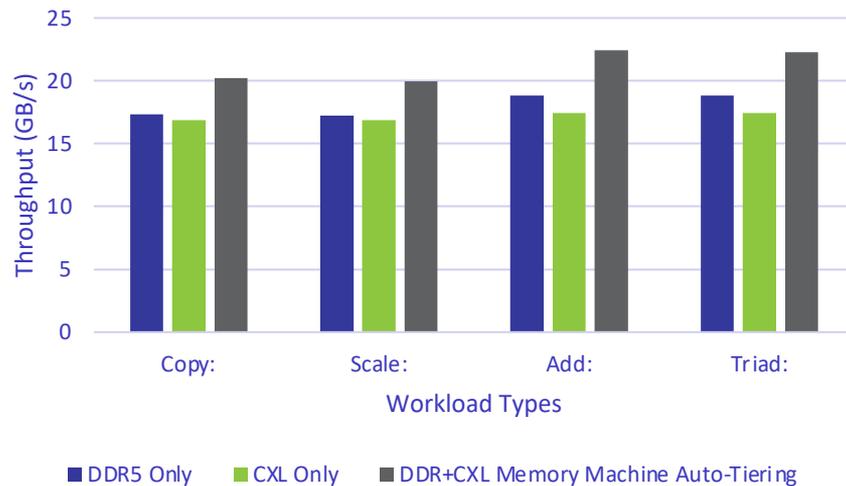


Early Results Running Memory Machine on CXL

MLC (Memory Latency Checker) Results



Stream Results



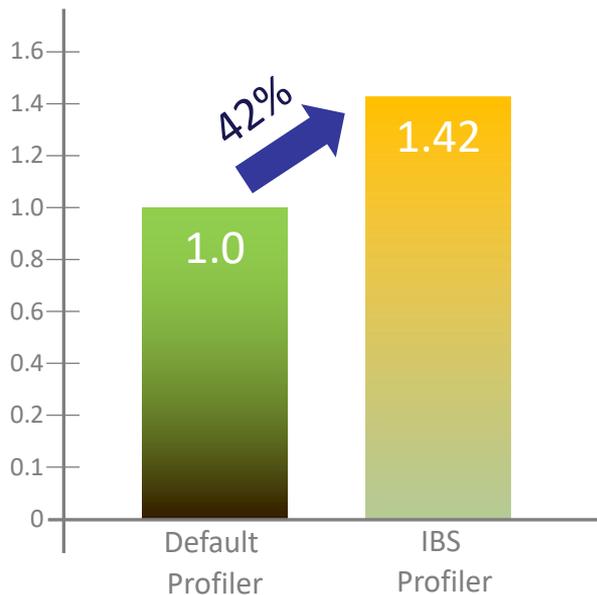
DDR5 Latency: 108ns

CXL Latency: 272ns

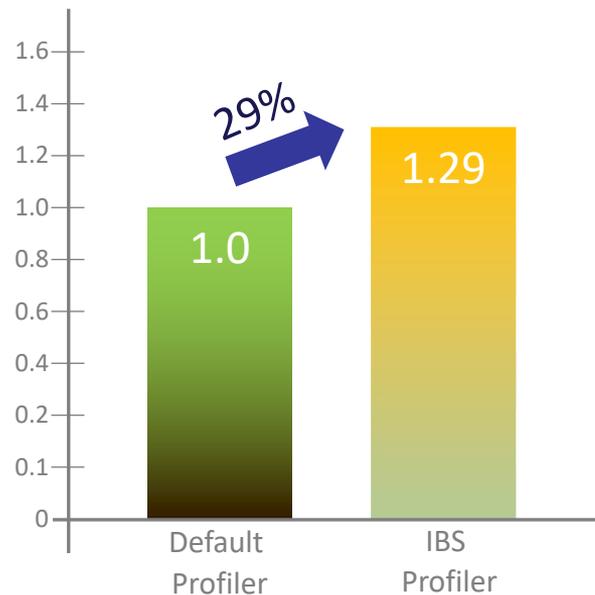
HW-assisted Memory Profiling Boosts Tiering Performance



Normalized MVBench Throughput
48GB Dataset w/ 6GB DRAM Tier



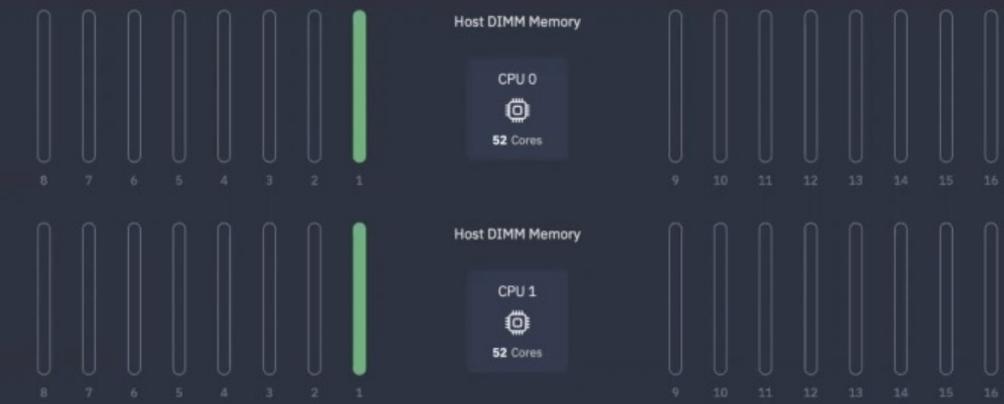
Normalized MVBench Throughput
48GB Dataset w/ 9GB DRAM Tier



Note: This demonstration used pre-production hardware, firmware, and devices. Performance varies by use, configuration and other factors

Overall

Total Memory Capacity: 125.63 GB



SE02.ENG.MEMVERGE.LOCAL

App Direct Mode

Application Memory Heatmap

System Topology

Process Monitor

All | Insight Groups: Top 10 Memory Consumers KVM

Insight Group Settings

Multi-selection

Duration: 1 hour

Start

Stop

Name	PID	DRAM Usage	CPU	Start Time	User	Monitoring Status	Report
gnome-shell	3512	133.88 MB	0.09%	09:01:05 May 18, 2022	gdm	1.00 hours	(45 KB)
mmctl	2052	91.86 MB	0.02%	09:01:02 May 18, 2022	mvmm	1.00 hours	(42 KB)
grafana-server	5469	81.23 MB	0.24%	09:06:09 May 18, 2022		1.00 hours	(45 KB)
prometheus	117975	66.18 MB	0.12%	20:47:07 May 18, 2022	nfsnobody	1.00 hours	(43 KB)
mmagent	2055	61.09 MB	0.12%	09:01:02 May 18, 2022	mvmm		
dockerd-current	5032	52.54 MB	0.25%	09:05:25 May 18, 2022	root		
mvmv	78301	49.88 MB	2.27%	14:00:38 Jul 25, 2022	root		
Xorg	3375	33.34 MB	0%	09:01:04 May 18, 2022	root		
gsd-color	3698	31.45 MB	0.22%	09:01:06 May 18, 2022	gdm	1.00 hours	(44 KB)
docker-containerd-current	5053	31.02 MB	0.12%	09:05:25 May 18, 2022	root		

Process Monitor - gnome-shell (3512)

Started at: 10:39:36 Jul 06, 2022

Ended at: 11:39:26 Jul 06, 2022

Memory Used

Peak 134.01 MB

Average 133.94 MB

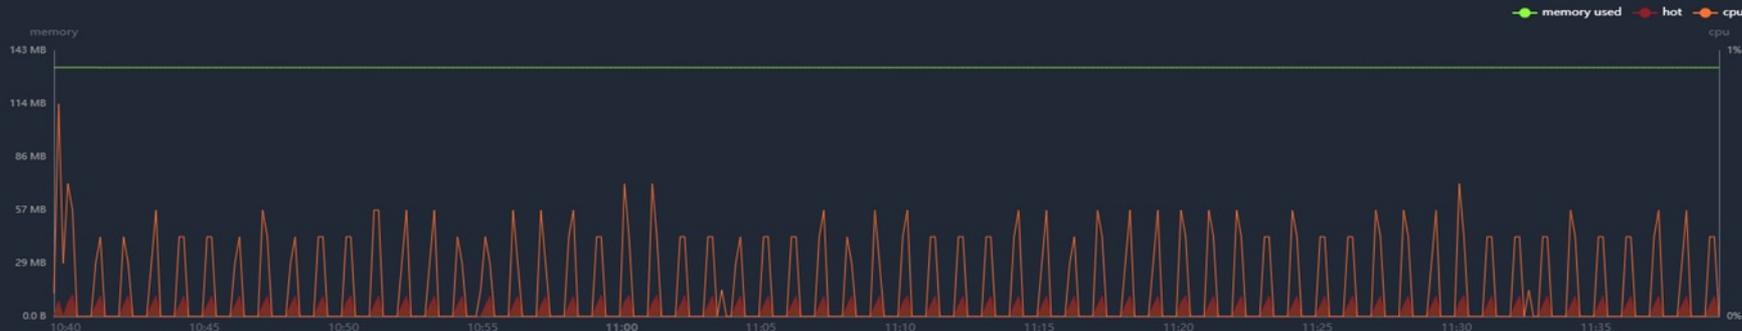
Standard Deviation 11.78 KB | 0.01 %

Memory Hot Size

Peak 11.64 MB | 8.69 %

Average 2.87 MB | 2.15 %

Standard Deviation 4.20 MB | 146.28 %

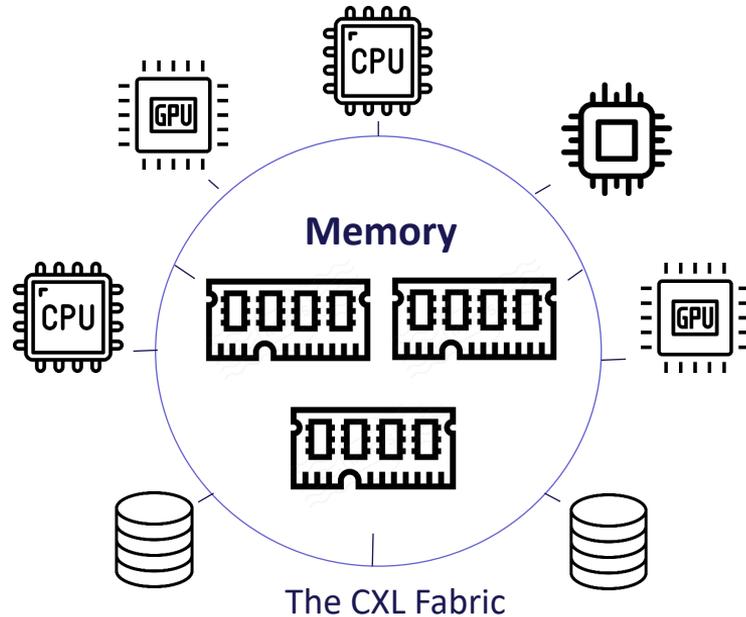


Examples of CXL Memory Services that can be enabled with Software



- Stranded Memory problem- automatic tiering/dynamic profiling
- Underutilized/Overallocated Memory within VMs
- Enable Memory Composable infrastructure
- Increase Resilience
 - Self-healing for DIMM failures
 - Transparent checkpoint/recovery and data protection with high-speed memory snapshots
- Enable Shared Memory Collaboration through shared memory objects or snapshots
- Consolidate distributed compute memory caches

CXL + Software Enables Memory-Centric Data Center



- Memory Expansion with Tiering
 - Capacity Expansion
 - Bandwidth Expansion
- Memory Pooling
 - Dynamic Elastic Memory
 - Composable Infrastructure
- Memory Sharing
 - Low Latency Pub/Sub
 - Shared Storage Cache
- For the first time in history, software becomes critical for memory!

Summary



- CXL is the game changer that will enable Big Memory
- Software will be at the heart of a CXL-based Big Memory solution
- MemVerge looking to partner with rest of industry to deliver solutions to Memory tiering, pooling, and sharing as well as advanced memory services.
- MemVerge software already supports CXL. MemVerge intends to deliver the best CXL middleware software solution to customers.





Thank you

EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

