CXL Ready for Take-off

Charles Fan
Co-founder & CEO

MemVerge

CXL Enables Memory-Centric Computing

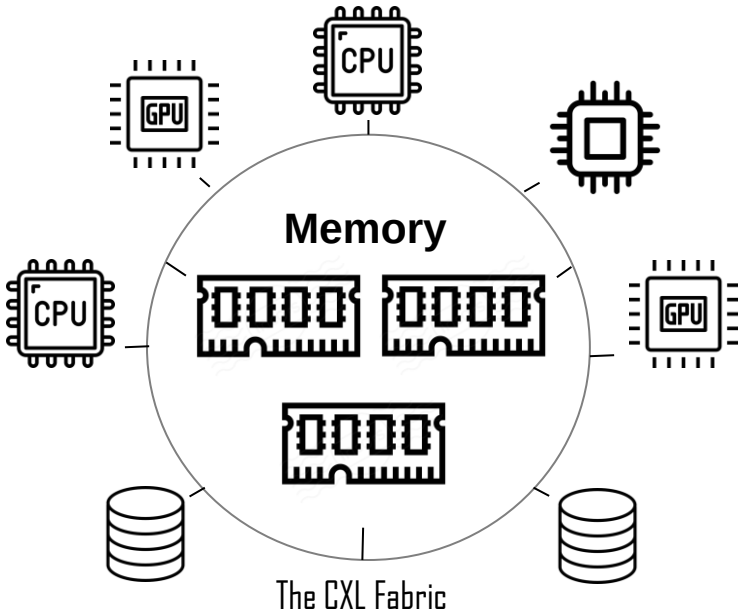

- Disaggregates memory from processors
- Enables
  - Memory expansion
  - Memory pooling
  - Memory sharing

Project Endless Memory

- 2 Intel Sapphire Rapids Servers
- 1 SK hynix Niagara Pooled Memory Development Platform (128GB Multi-host Single Logical Device)

No More Out-of-memory Erros

1. Host at ‘steady-state’ using Main Memory only

2. As high VM demand exhausts Main Memory, **Memory Machine** detects and dynamically provisions capacity from the memory pool

3. As VM demand diminishes, **Memory Machine** releases the memory back to the pool and returns to ‘steady state’

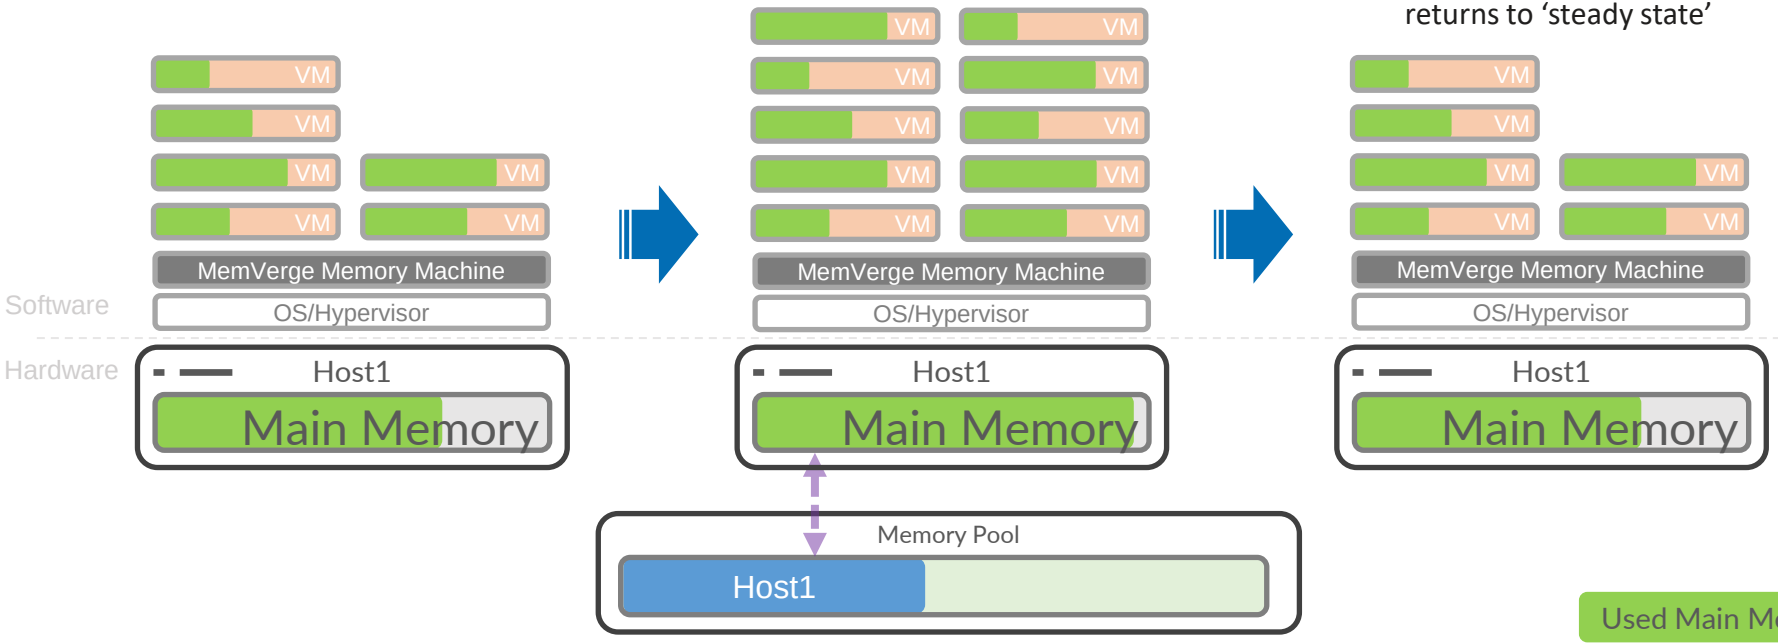

The I/O Wall in Distributed Applications

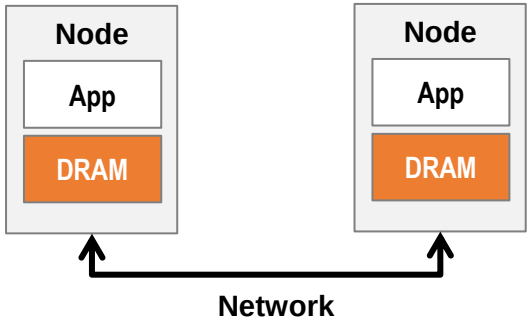

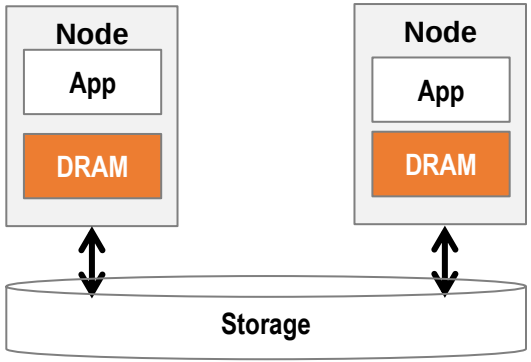

Serialization & Deserialization

Data Copies

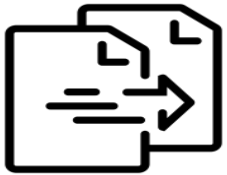

Media Speed

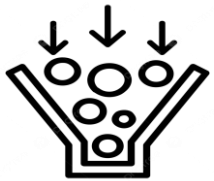

CXL Makes Cross-node Shared Memory Possible!

- Multiple nodes have direct memory access to the same CXL memory region
- Cache Coherence is part of CXL 3.0 hardware specification
- Software cache coherence can be implemented on top of CXL 2.0 hardware

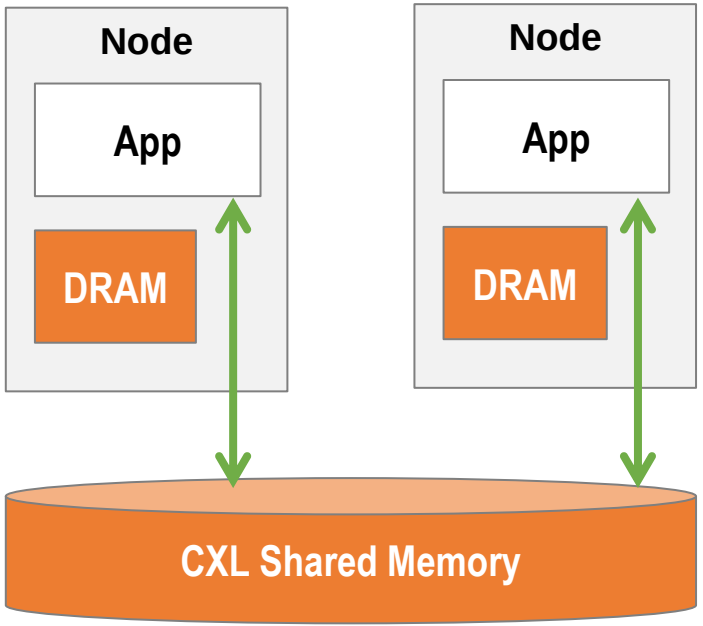

Introducing Project Gismo

Global I/O-free Shared Memory Objects

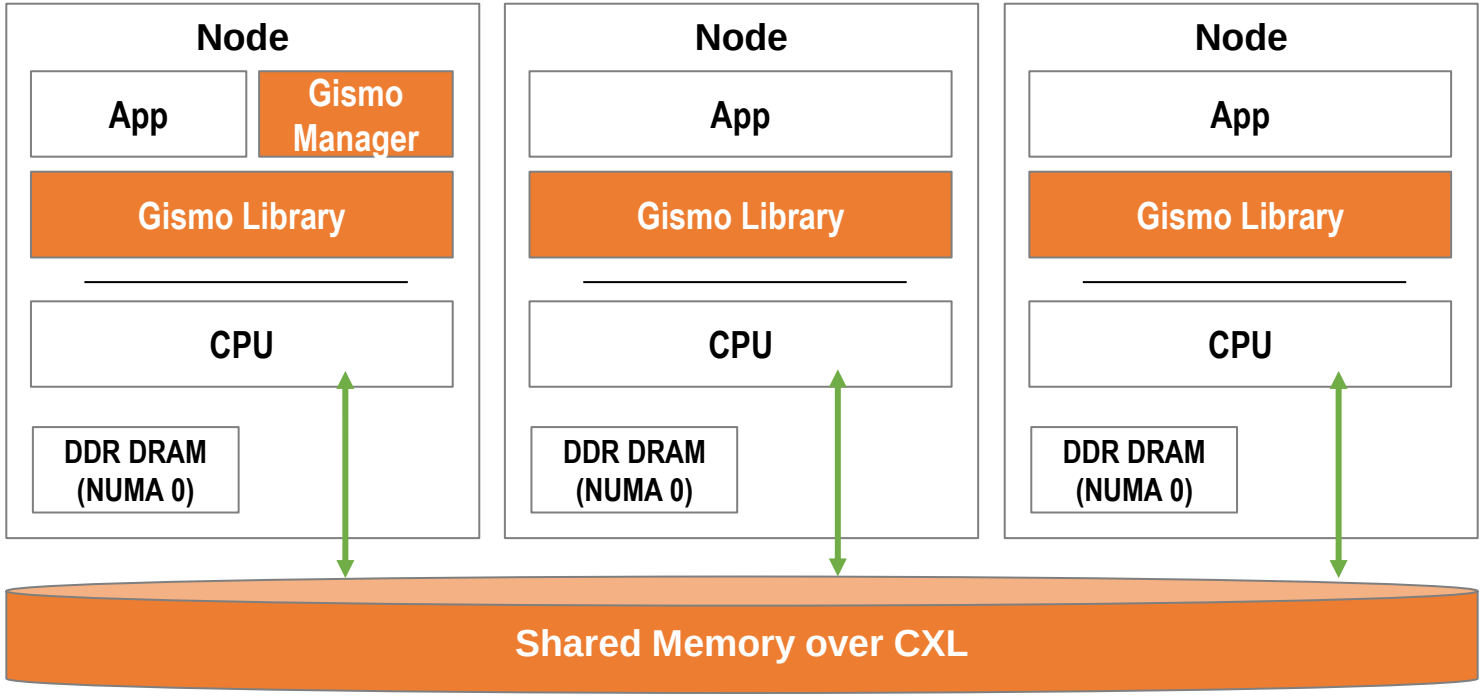

Use Case Example: Ray+Gismo

- **IO-free:** No more object serialization and transfer over network for remote object access
- **Zero Copy:** No more duplicate object copies on different nodes
- **No Spilling:** Reduce object spilling and data skewing because each node has access to the whole memory pool

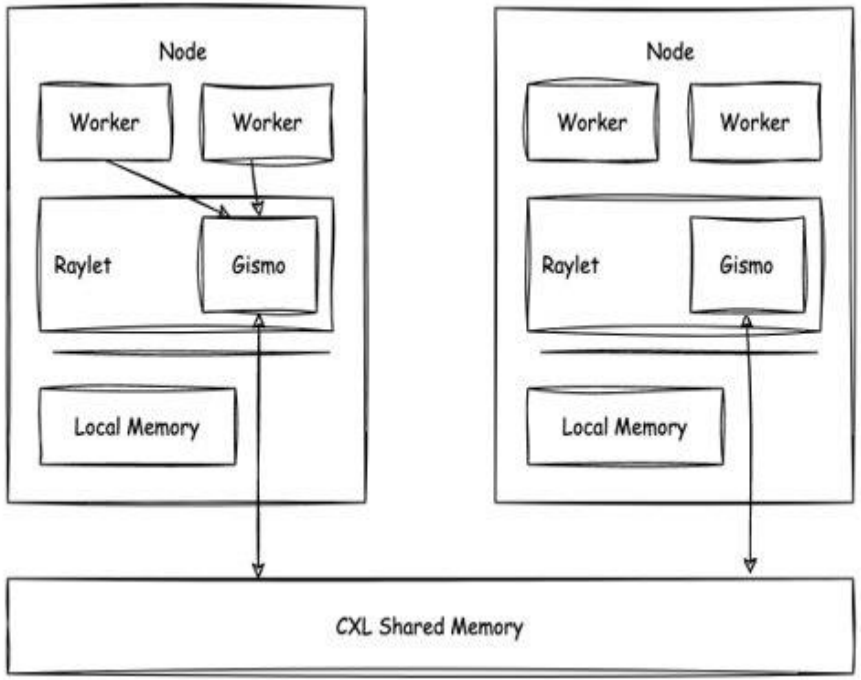

Shuffle Benchmark Results

	Baseline Ray	Ray + Gismo
Local Get 1GB object	0.4 sec	0.4 sec **CXL shared memory as fast as local memory**
Remote Get 1GB object	2.7 sec	0.4 sec **675% faster**
Shuffle 50GB 4 nodes, each 4 cores, 128 GB object store	515 sec	185 sec **280% faster**

* Running in emulation environment

Not STAC benchmarks

OOM

IO

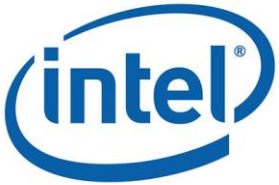
intel

MONTAGE

SAMSUNG

SK hynix

AMD

AsteraLabs

Micron

XCONN TECH

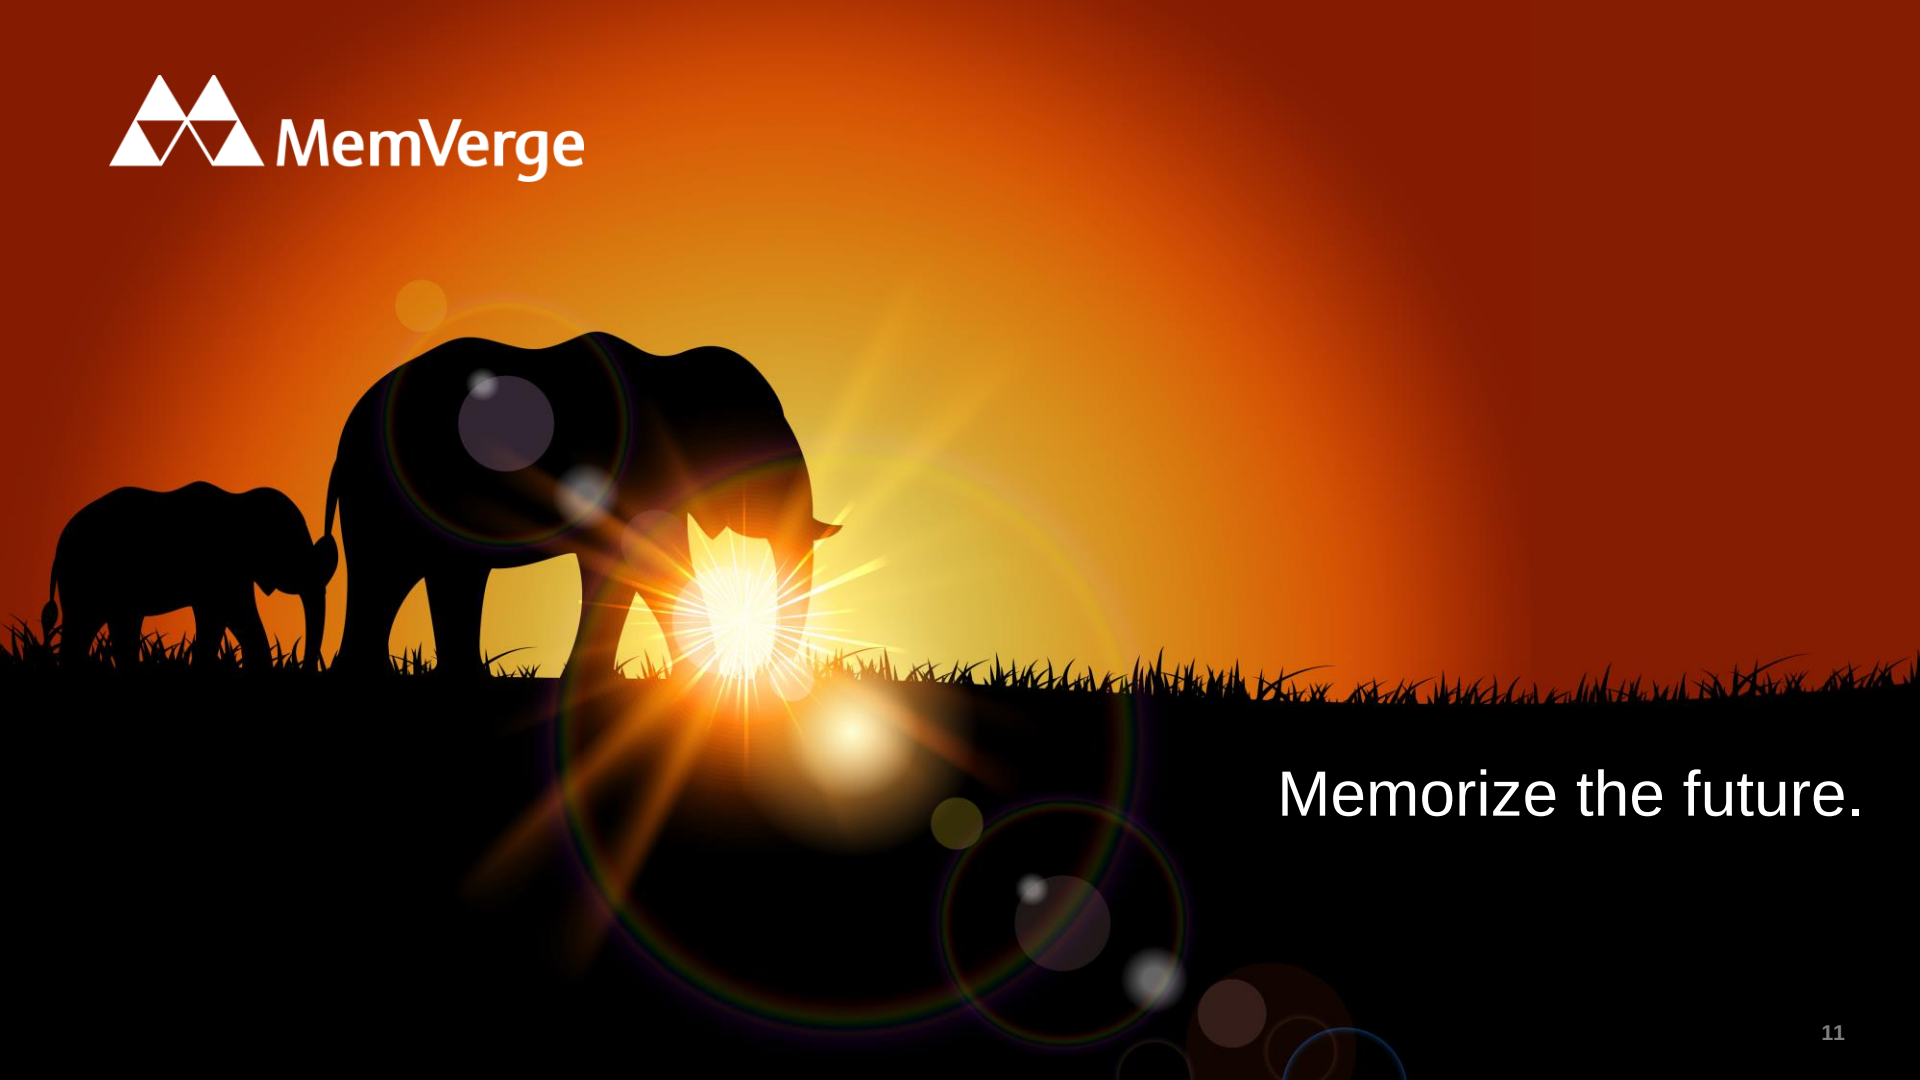
MemVerge
Memorize the future.