Hyperscale Composable Memory Systems with

Dynamically Adjusting Compressed Tier

Presenter:
Angelos Arelakis, CTO
Nilesh Shah, VP Business Development
ZeroPoint Technologies

FMS the Future of Memory and Storage

Challenge

Hyperscalers spending significant $$ on software based compression

CPU cycles used for compression

Meta **4.6%**

Google Cloud **3%**

Hyperscaler requirement:
hardware compression is a **MUST-have**

OCP Hyperscale CXL Tiered Memory Expander Spec

FMS

Opportunity|Add Compressed CXL Memory Tier

Deploy New Tiers: Ordinary + Compressed DRAM memory on CXL

	Benefit	Spec Component
1	**Reduction** in total cost of ownership	Standardization, Hardware accelerated, Lossless Compressed Memory Tier
2	Energy **Efficiency, Sustainability**	Transparent Hardware accelerated Compression
3	**Preserve** Software Investments	Support for legacy Compression Algorithms

Game Changer

OPEN Compute Project

Hyperscale CXL Tiered Memory Expander Specification

Revision 1

Version 1.0

Base Specification Template v1.2

Effective October 27, 2023

Removes barriers, enables diversity of Hyperscalers + Enterprise Customers

FMS

Requirement: Hyperscale CXL Tiered Memory Expander Spec

Preserving Software Investment, without compromising performance

- Support for Legacy Compression Algorithms
- Future Proof , supports Algorithm Innovation
- Stringent Latency & Bandwidth Specifications

Parameter	Specification
Latency Uncompressed Access	90 to 150ns
Latency Compressed Access	250ns to <1us
Bandwidth Efficiency Read only/ Write only	80% / 75%

Stringent Requirements

https://www.opencompute.org/documents/hyperscale-cxl-tiered-memory-expander-for-ocp-base-specification-1-pdf

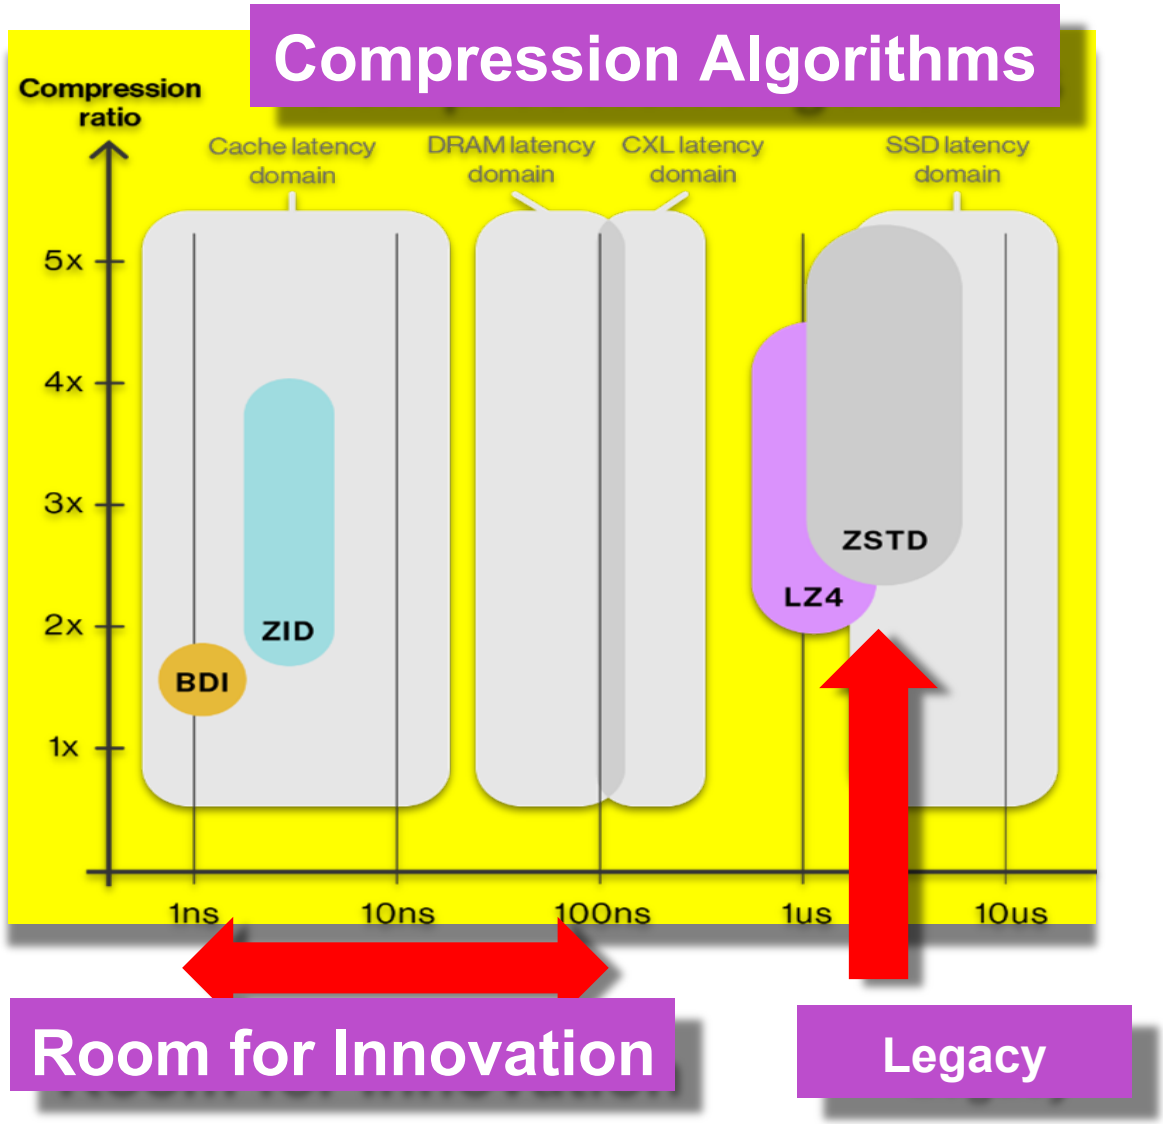

Hyperscale CXL Tiered Memory Expander Spec

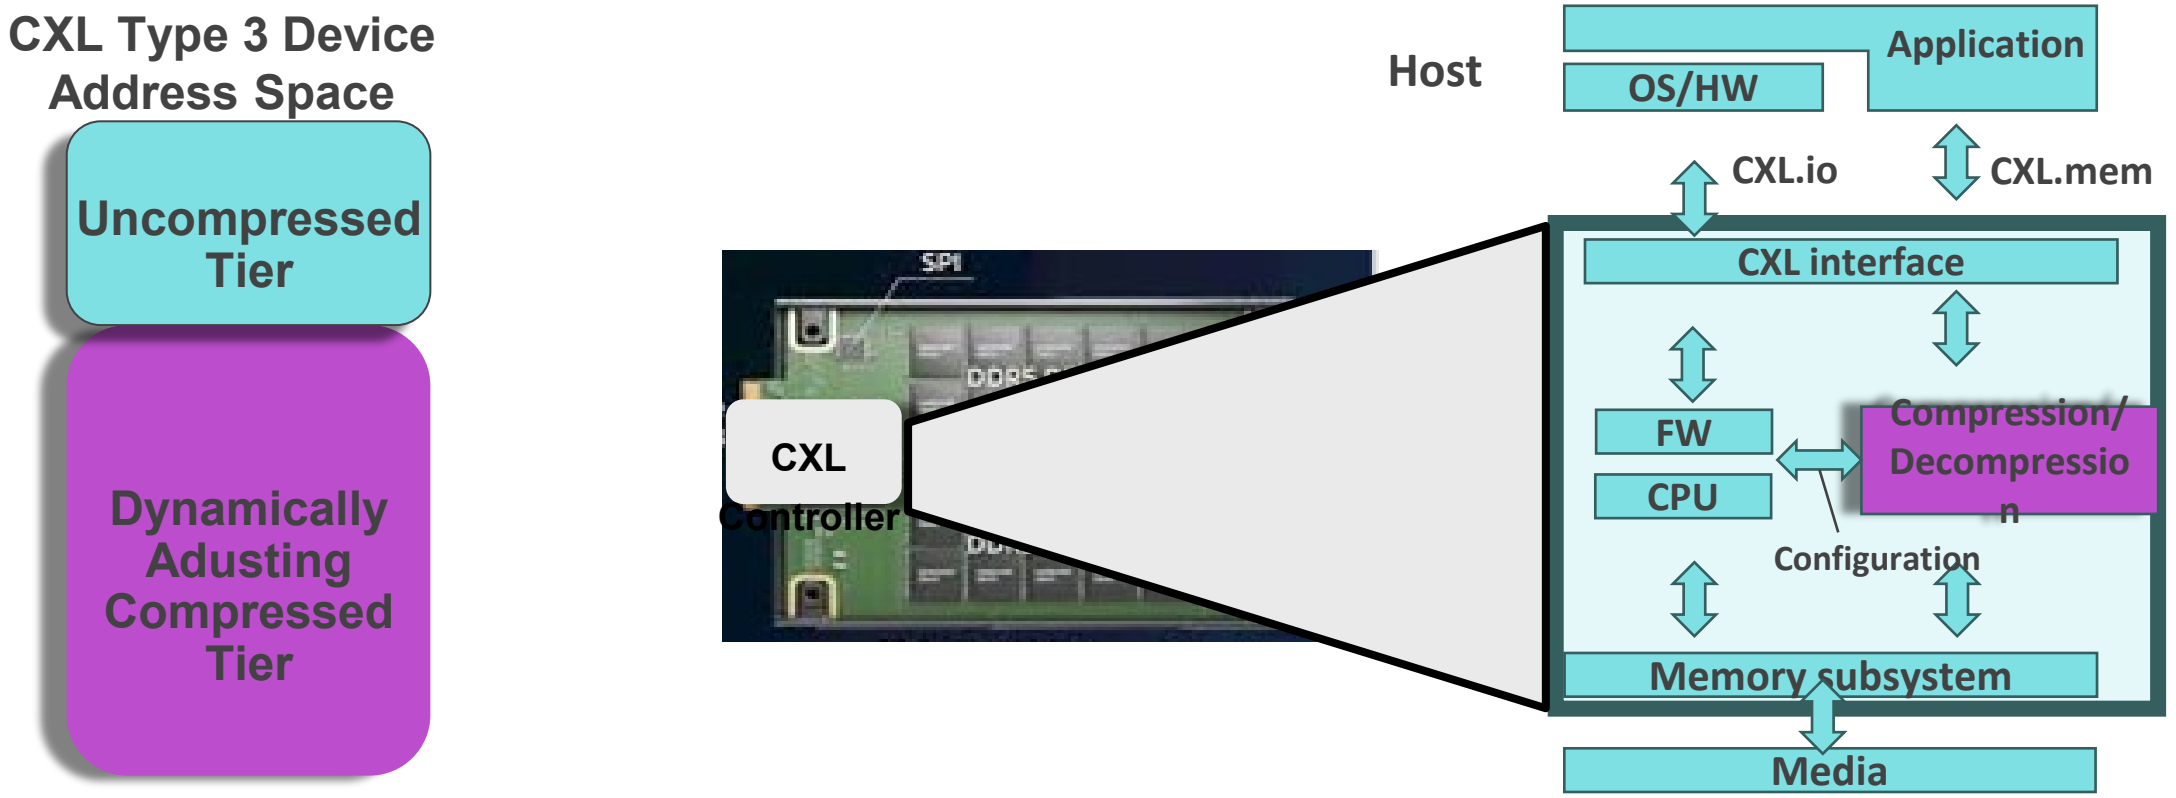

CXL Controller with integrated Hardware Acceleration:
(De)Compression+ Compaction + Transparent memory management

The ZeroPoint IP Solution Results

- OCP Spec compliant Hardware Accelerated CXL memory **(De)Compression + Compaction + Transparent** memory management IP block
- **2-4x** transparent (de)compression major Datacenter workloads
- Protocol: CXL 2.0, 3.0, 3.1
- Compression Algorithms: LZ4, ZID (proprietary)
- Portable: AXI4, CHI, Leading process node support

OCP Spec Compliant, portable IP

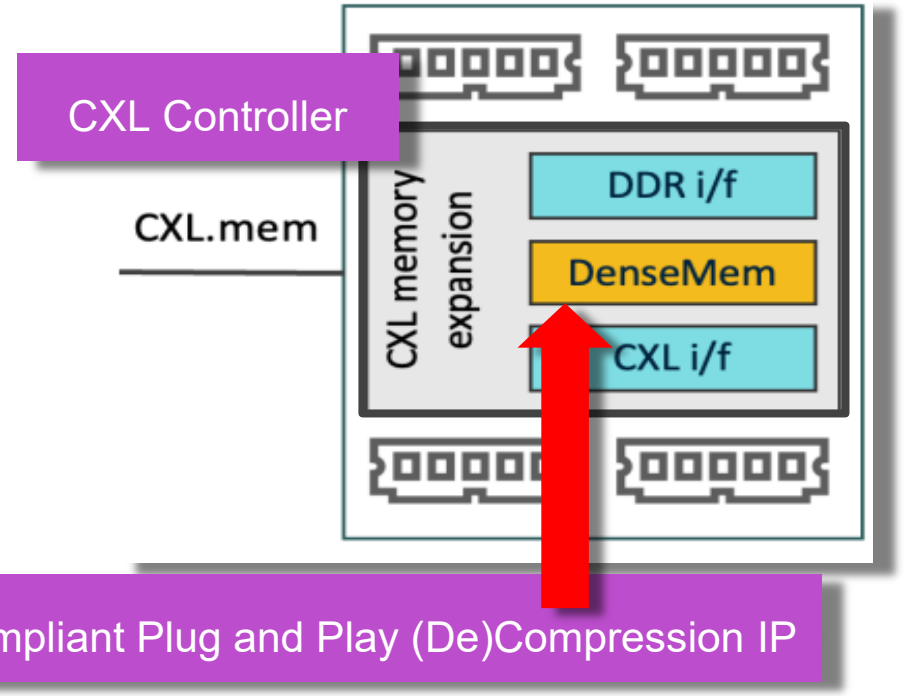

IP Solution Performance Characteristics	Value
Compression Ratio	2-3X
Block cache (SRAM) hit latency	<30ns
Cache line in uncompressed region latency	<90ns
Cache line in an uncompressed block latency	<150ns
Cache line in a compressed block latency	<250ns
Tail latency [cache line in a compressed block]	<1us
Decompress bandwidth[4x 1867MT/s]	>46G/s

Reduce Data Center TCO 20-25%

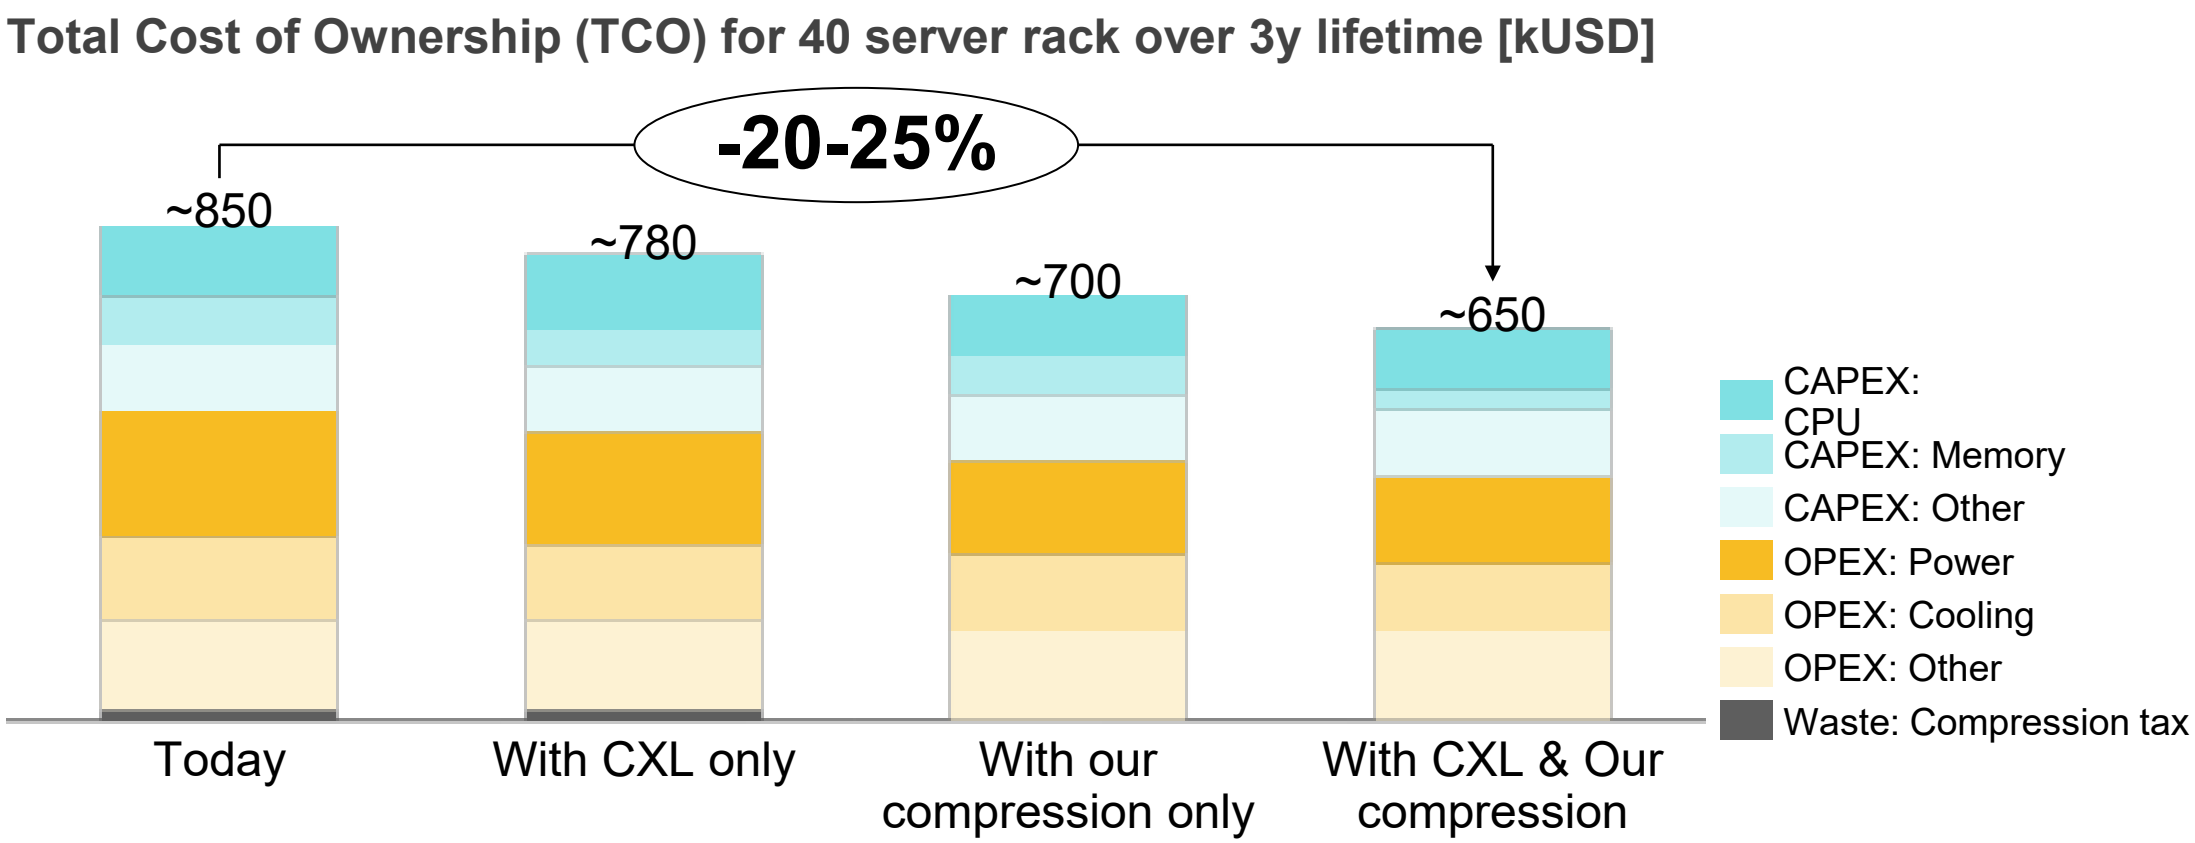

Demo

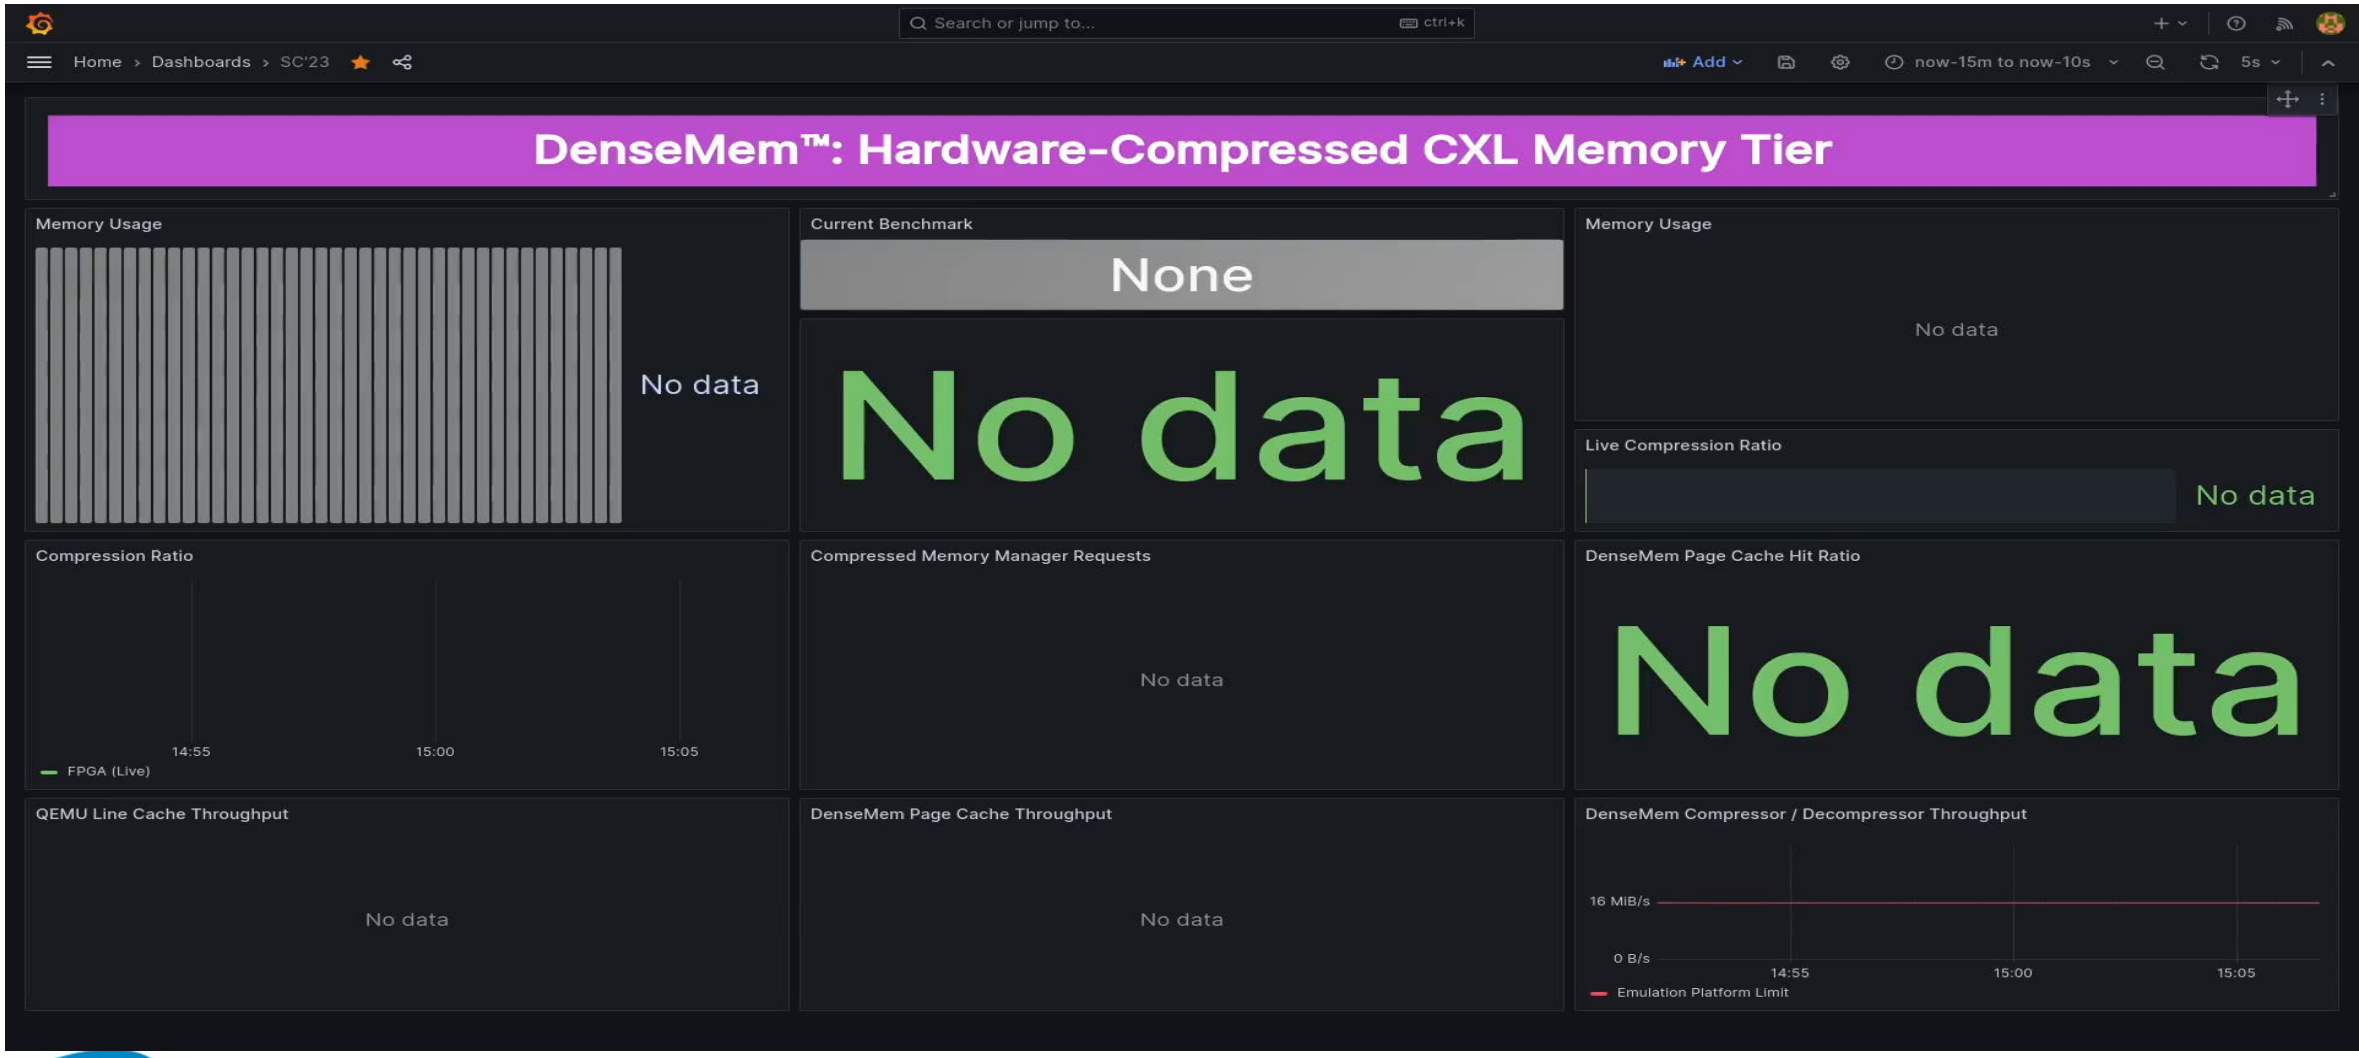
Search or jump to... ctrl+k
Home › Dashboards › SC'23
Add
now-15m to now-10s
5s
DenseMem™: Hardware-Compressed CXL Memory Tier
Memory Usage
No data
Current Benchmark
None
No data
Memory Usage
No data
Live Compression Ratio
No data
Compression Ratio
14:55
15:00
15:05
FPGA (Live)
Compressed Memory Manager Requests
No data
DenseMem Page Cache Hit Ratio
No data
QEMU Line Cache Throughput
No data
DenseMem Page Cache Throughput
No data
DenseMem Compressor / Decompressor Throughput
16 MiB/s
0 B/s
14:55
15:00
15:05
Emulation Platform Limit

Performance : Datacenter workloads

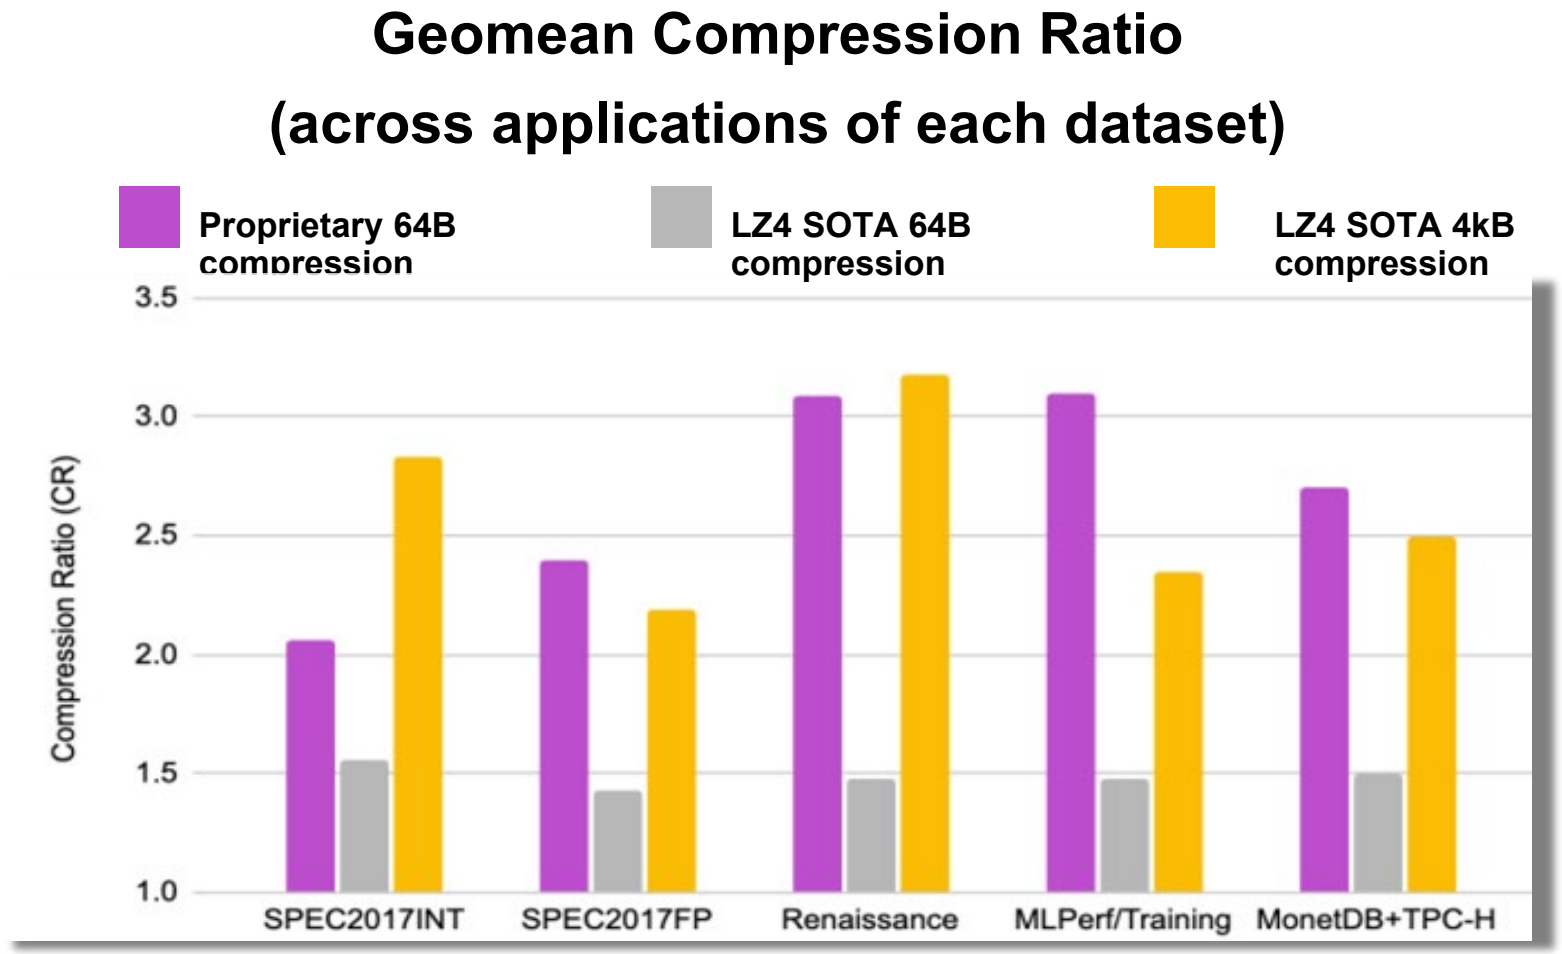

APIs

- [.io] Telemetry: Remaining actual free space (not DPA)
- [.io] Config: 4 capacity watermarks {High, …, Critically Low}
- [.io] MSI/-X interrupts: The IP raises interrupts to the Host as the watermarks are crossed. Internally it reacts:
  - write back-pressure delay mechanism (effectively delay write credits)
  - tune compressor to favor CR over bandwidth
  - defragmentation
- [.io] Command: Freeing capacity (aka overwrite DPA pages with 0s)

Host SW leverage IP Solution APIs to manage migrations

Call To Action

- **Summary**
  - OCP Spec Compliant IP Solution
  - Portable across process nodes, AXI/CHI interface
  - Performance verified

- **Call To Action**
  - **Controller manufacturers**: Collaboratively address Hyperscale OCP requirements
  - **ISVs**: Host software integration, joint workload demo