

Landscaping the Future for Robust CXL Memory Development

Thomas Won Ha Choi
Director of DRAM Product Planning, SK hynix

EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

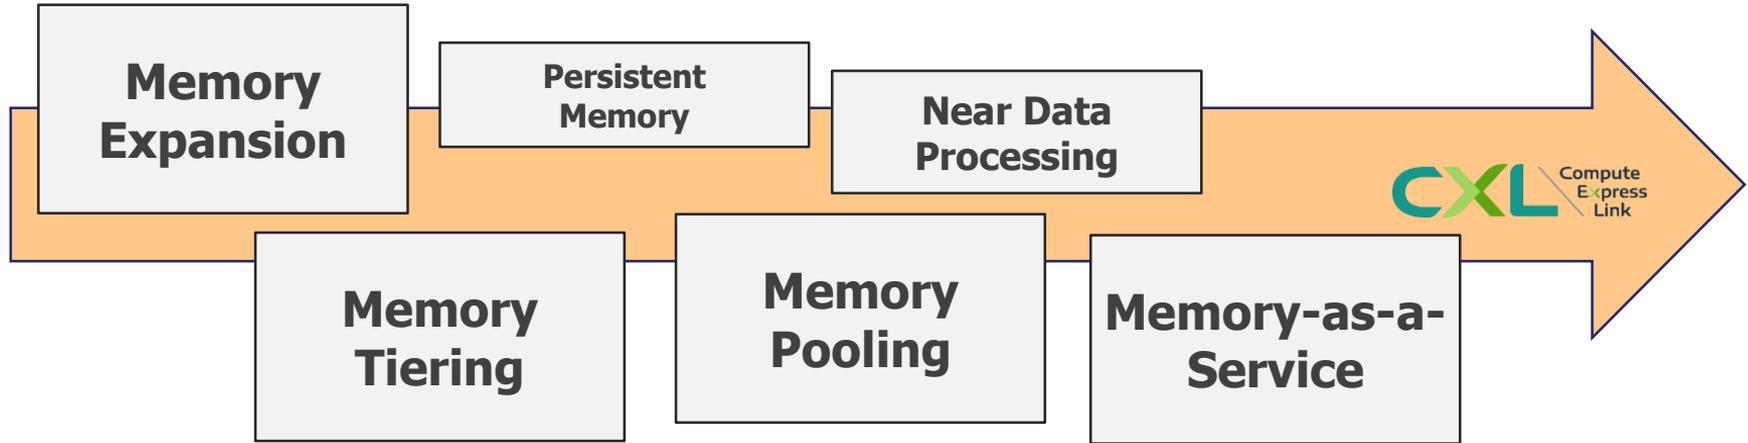
OCTOBER 18-20, 2022
SAN JOSE, CA



CXL™: Great Expectations



- Exciting opportunities to enhance memory system with emerging use cases!
- Datacenter industry players are teaming up to introduce first CXL memory product very soon.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.



SK hynix's Vision in CXL Memory

- Overcome current memory system's limitations to scale-out of bandwidth and capacity.
- With the new CXL memory hierarchy, aim to optimize data efficiency and datacenter TCO.
- More details on latest demos of SK hynix CXL PoCs/prototypes at SK hynix Booth (#B25) & Expo Hall Talk Video!

Bandwidth Memory Expansion (BME)

Assume 12CH/CPU

Get 13~16th Mem. CH

Better Interleaving

Better Loaded Latency

Capacity Memory Expansion (CME)

2DPC Era near the end

Allow Capacity Scaling

Improve RAS & W/GB

Expansion w/ Tiering

Memory Media Differentiation

Media Agnostic Nature

Standardized + Custom

Various use cases

Mult. Design Tradeoffs

Memory Controller Differentiation

Beyond Common IPs

Better RAS/Security

Lower-power

Add computational

Memory-as-a-Service (MaaS)

Thru Tiering & Pooling

Choose memory types

Build pool appliances

Composable Arch.



OCP
GLOBAL
SUMMIT

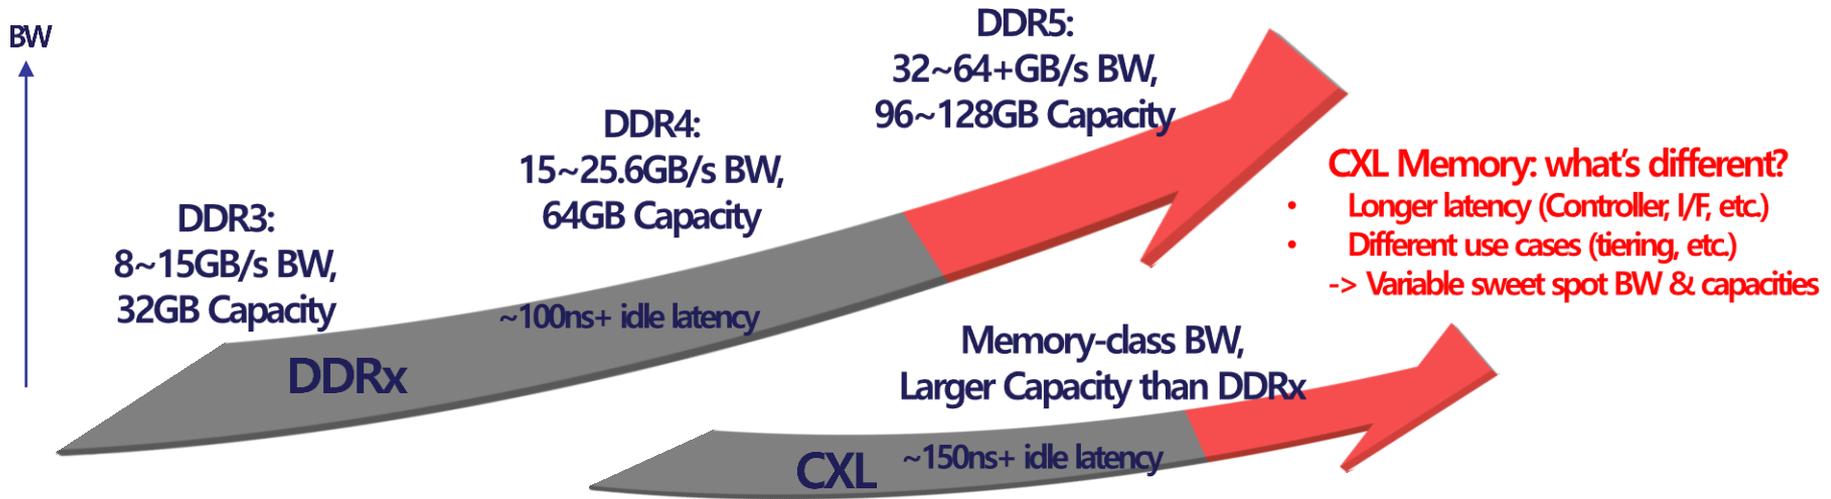
OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.



Memory Requirements for CXL

- Traditionally for local DDRx memory: Scale-up BW & capacity, keep power & cost at the same level.
- CXL memory requirement: slightly different from traditional due to different components and use cases.





Newly Introduced Memory Features for CXL

- Migration of memory controller features into CXL controller.
- Media-specific features also introduced to optimize for CXL.
- Any specifics related to module mechanicals also must be considered.

CXL Memory Module: mechanical specifics related to integrating all components within the module

CXL Controller/Bridge:

- *Enhanced RAS Features, ECC*
- *Security Features*
- *Low-power Features*
- *FW Management Related*
- *Computation & Acceleration Features*
- *Vendor-Optimized Features*

CXL Memory Media:

- *Custom Features to optimize for different design tradeoff points (BW, power, capacity, cost, etc.)*
- *Variations of feature modified from standardized DRAM*



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Challenges in Enabling CXL Memory



- How to meet CXL memory requirements?
- How to evaluate and validate CXL memory features? Development process may be new to the industry.

Challenge Points in meeting Requirements

- Achieve Reasonable BW amid longer latency from CXL
- Achieve 2X+ module capacity vs. DDR5 with TCO awareness
- Maintain power amid more components in the module
- Maintain server-class RAS & security with TCO awareness

Challenge Points in Evaluation / Validation

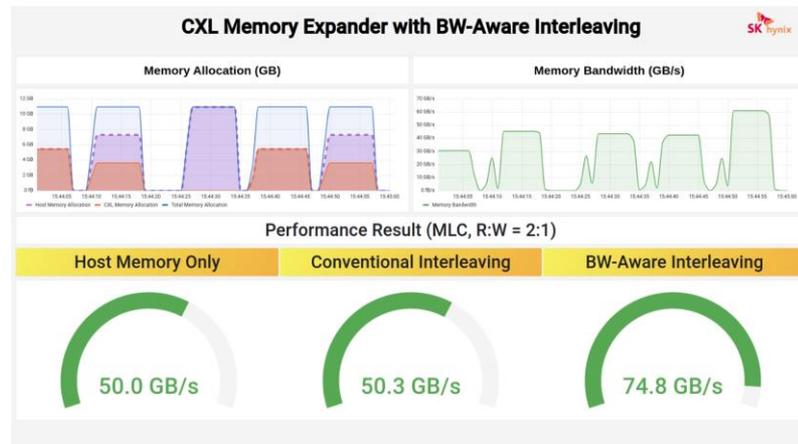
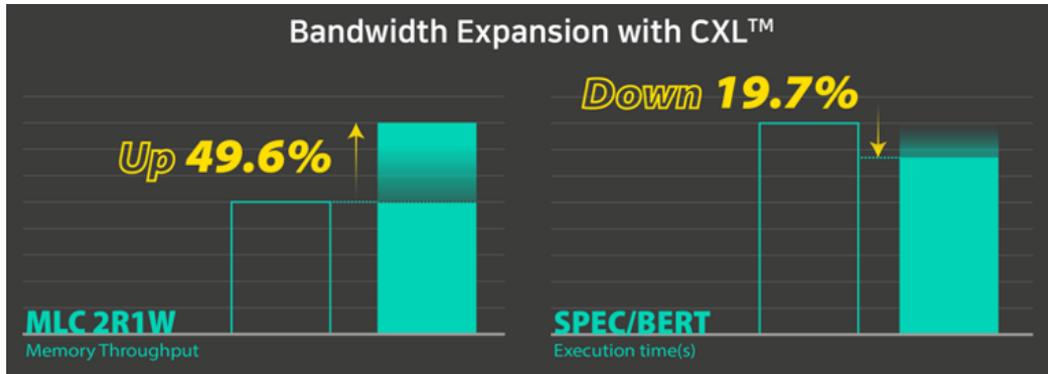
- Offloading traditional host-related features to CXL controller
- Consistent level in RAS & security vs. those in local memory
- Getting used to CXL memory evaluation environments
- Who is responsible for each validation items?





Planning BW Scaling with Perf. Evaluations

- To-be: understand BW/latency tradeoff in CXL, and find performance enablers to allow BW scaling.
- Current findings: some promising results, but still room for improvements from various area (module, CPU, etc.)
- Utilizing SW accompaniments to achieve better interleaving may help to boost the performance significantly.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

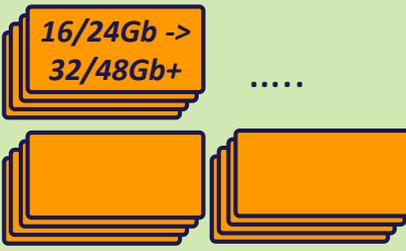
EMPOWERING OPEN.



Building Concepts for Capacity Scaling

- Many consider the capacity scaling as the #1 factor in adopting CXL memory, preparing for the 1DPC era.
- To-be: find enablers to allow CXL memory module capacity scaling with cost-effectiveness.
- Memory die density scaling is getting more difficult; critical to find a cost-effective stacking method.
- Moreover, maintaining BW to GB ratio suitable for CXL memory must be addressed.

Era (DPC: DIMMs Per CH)	DDRx Gen.	DRAM Speed
3DPC	DDR3~DDR4	~ DDR4-2133
2DPC	DDR4~DDR5	DDR4-2400 ~ DDR5-5600-ish
1DPC	DDR5~DDR6	DDR5-6400 ~



16/24Gb ->
32/48Gb+

.....

Paths to Capacity Scaling:
*Die density scaling or
Cost-effective Stacking*

Keep in mind that:
BW-to-GB ratio must be consistent

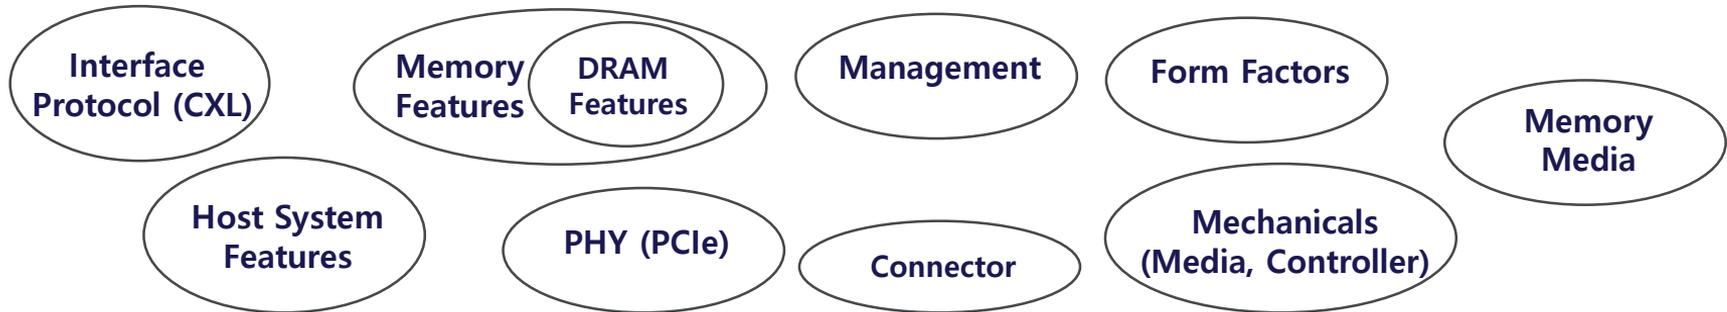
The diagram shows two stacks of orange rectangular memory modules. The top stack is labeled '16/24Gb -> 32/48Gb+'. To the right of the top stack are three dots '.....'. Below the top stack is another stack of orange rectangular memory modules.



Landscaping CXL Memory Enabling Process

- Enabling process of CXL memory may have to integrate from multiple literatures (DDRx w/ DIMM, PCIe w/ SSD).
- Utilizing standardization of CXL memory will be critical in expanding the volumes of CXL memory
 - Important mission is to allow standardization of major CXL memory features without stifling innovations
 - Recent announcement on MOU between CXL and JEDEC is a great step forward to a robust CXL memory standardization

Who validates what? How and whether to standardize? What existing literatures can be leveraged?



Exploring Sweet-Spot Application Use Cases



- Goal: Improve data efficiency by selectively offloading cold accesses to 2nd tier memory or memory pool.
- Understanding behaviors of OS, control plane, QoS are important to find optimized data distributions.
- More results on tiered memory performance evaluation coming up; followed by more studies in RAS & security.

Motivations for Investigation:

Many rooms for improvement in data efficiency (i.e. large portions as cold & seldom accesses)

Target Workloads Experimented:

In-house Cloud, Redis, Spark, TPC-H, SPEC, PARSEC, Web Server, Data warehouse (batch query), etc.

Changes Made in Research:

Building prediction models, changes to swap/compression engine, kernel modification in allocation, changes to NUMA policies with optimized page table migrations, fabric manager model, etc.

Pros & Cons:

Pros: Improved data efficiency -> TCO reduction on total fleet & less congestion in local memory
Cons: Performance drop that can be contained with learned management in page table & events



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Creating More PoC Evaluation Opportunities



- Alongside research, PoC development must be accompanied to create more market.
- Beginning with CXL memory expanders & computational memory solutions.
- Coming up: pooling appliance PoC with industry partners.
- Aim to build PoC such that customers can evaluate with ease:
 - ✓ Add options in the PoC sample to support legacy slots
 - ✓ Accompany SW tools & guidelines to run the PoC more effectively



SK hynix 96GB E3.S CXL Memory Module PoC (8/1/22)



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Conclusion (#1)



- CXL shows promises in enhancing the scaling of memory capacity and data efficiency of the memory system.
- With many value-added use cases introduced, however, CXL introduces a new paradigm to memory industry and the development process of CXL memory will face uncharted territories in executing the process.
- With new requirements and features as memory, enabling CXL memory will be challenging as the ways to evaluate the requirements and processes to validate the memory product may include many foreign aspects.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Conclusion (#2)



- SK hynix is working hard to achieve bandwidth and capacity scaling of CXL memory, by going through extensive performance evaluation and concept engineering to define robust product definitions.
- Industry-wise, many researches in finding optimal use cases look promising and more PoC developments from SK hynix will accompany the journey to achieve successful adoption of CXL memory.
- Looking forward to close collaboration with all ecosystem partners to build value-added CXL memory products!



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA

EMPOWERING OPEN.

Thank you!



EMPOWERING OPEN.



OCP
GLOBAL
SUMMIT

OCTOBER 18-20, 2022
SAN JOSE, CA