CXL® Attached Flash Memory Economics

Mahinder Saluja
Director of Strategy, SSD BU
KIOXIA America, Inc.

Agenda

- **Why Memory Expansion?**
- **CXL® Attached Flash Memory**
- **Performance and Cost**
- **Challenges and Opportunities**

Why Memory Expansion?

Today’s data center workloads require increased memory capacity and bandwidth.

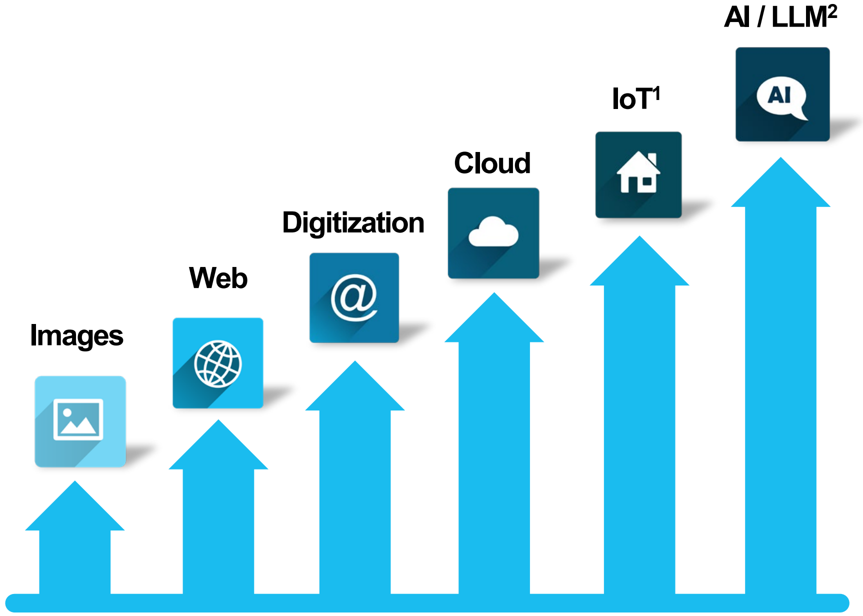

[1] that Internet of Things (IoT) [2] Artificial Intelligence (AI) / Large Language Models (LLM)

Central processing unit (CPU) / graphics processing unit (GPU) core idles due to gap in density growth compared to memory capacity and bandwidth

- Memory ~40%-50%* of total server cost

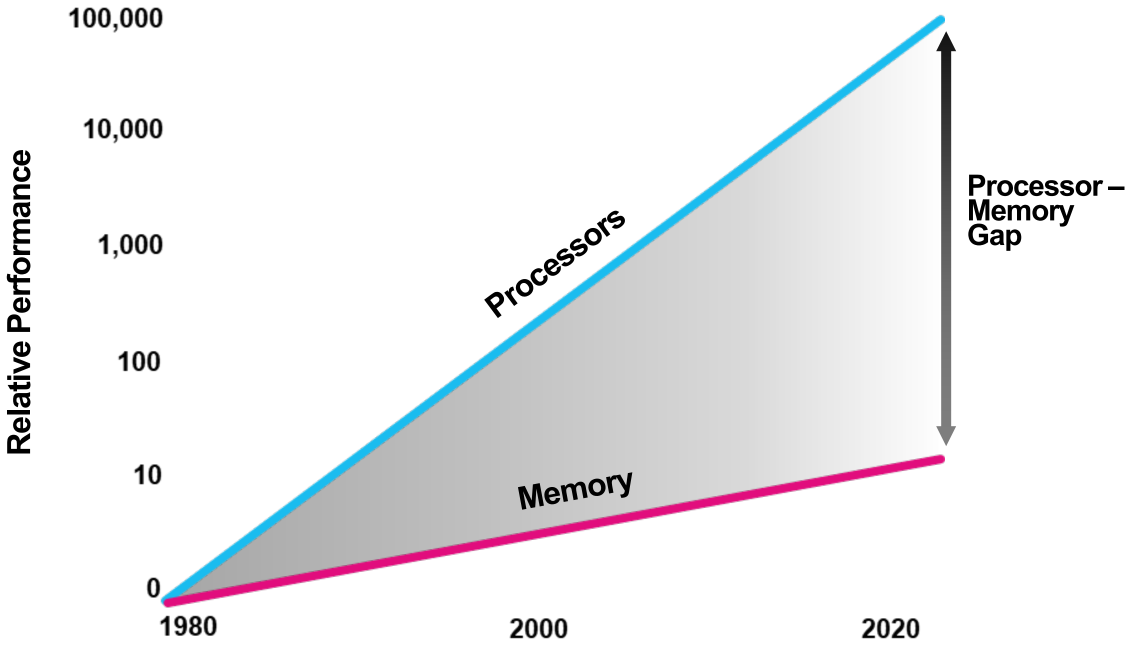

Image source: created by KIOXIA

* Based on KIOXIA’s engineering team’s perspective and cost of memory solutions today

Memory Expansion with Flash

- **CXL® Benefits**
  - Cost effective memory capacity and bandwidth expansion
  - Enables memory pooling and sharing with DRAM
  - Abstracts memory media interface

- **CXL® Technology Creates the Perfect Opportunity to…**
  - Explore alternative to costly DRAM
  - Flash media to jump over the semantic wall

But can flash media jump over the sematic wall and reduce cost?

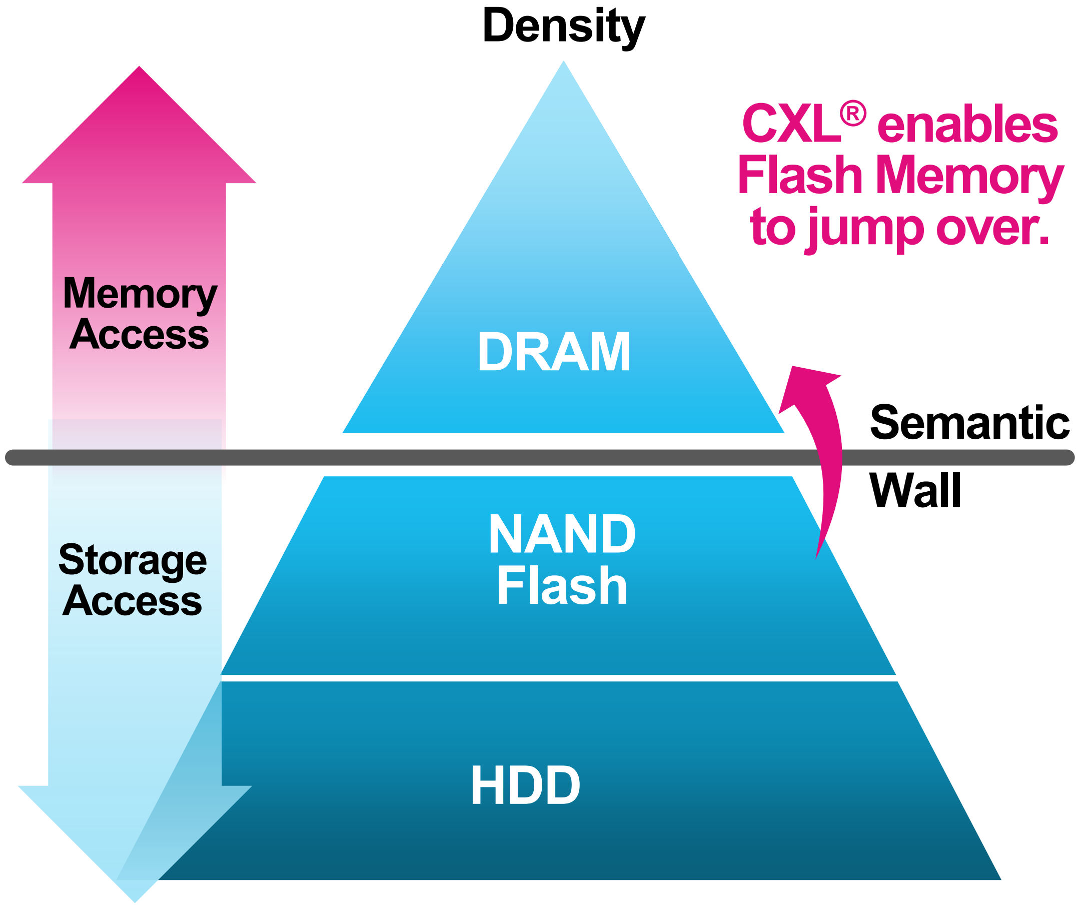

Memory Hierarchy & Density

Image source: created by KIOXIA

CXL and Compute Express Link are registered trademarks of the Compute Express Link Consortium, Inc.

KIOXIA CXL® Flash Memory Exploration

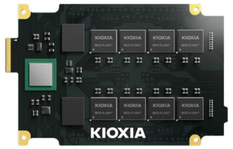

Combination of CXL® and BiCS FLASH™ technologies

Read intensive high capacity and high bandwidth memory

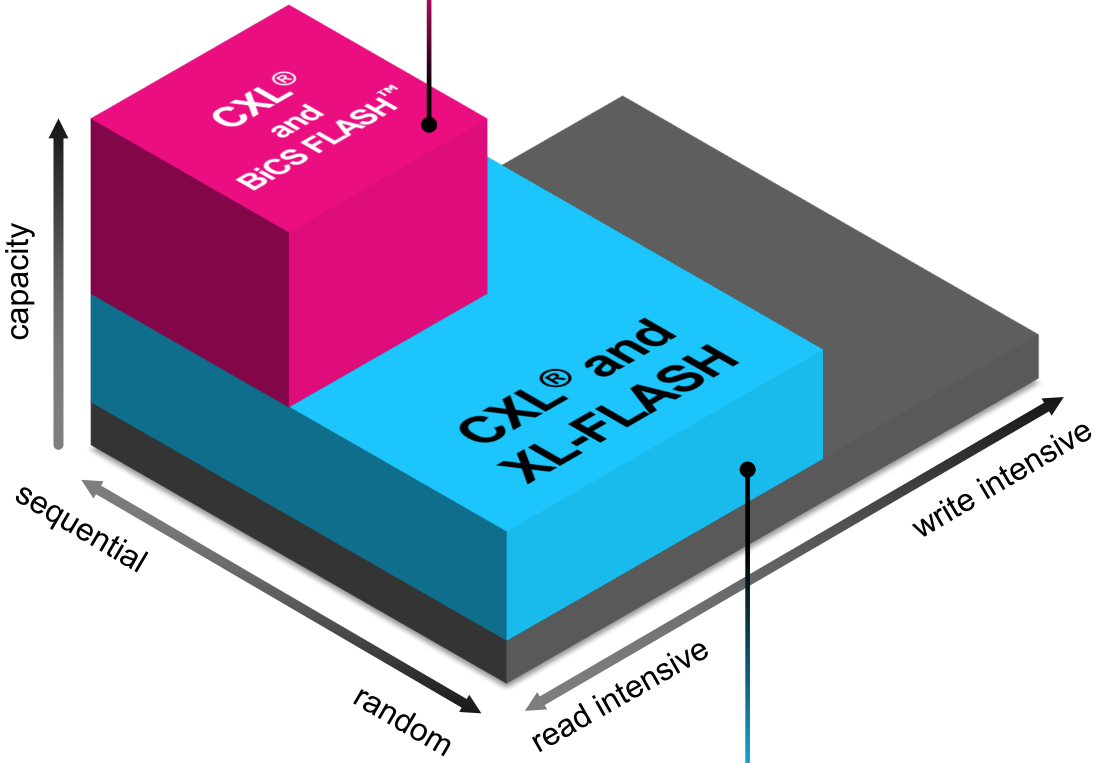

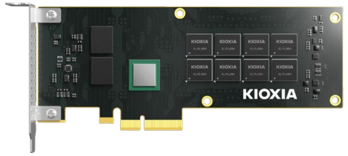

Combination of CXL® and XL-FLASH technologies

Random read/write access memory expansion

	CXL® and XL-FLASH Technologies	CXL® and BiCS FLASH™ Technologies
Media	BiCS FLASH™ (XL-FLASH)	BiCS FLASH™
Value Pillar	Low latency (<5us (single-level cell), <10us (multi-level cell); DRAM cache tier)	High bandwidth and High capacity
CXL Access	CXL.mem, CXL.io	CXL.mem,CXL.io
Capacity	>256 gigabytes (GB)	> 1 terabytes (TB)
Suitable Applications	Artificial intelligence (AI) / machine learning (ML) inference, In-memory data bases (DB), graph processing, cache, tiering	AI/ML training & inference, big data analytics

Image and table source: *created by KIOXIA*

CXL and Compute Express Link are registered trademarks of the Compute .The product images shown are a representation of the design models and not an accurate product depiction. Definition of capacity: KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes. A computer operating system, however, reports storage capacity using powers of 2 for the definition of 1GB = 230 = 1,073,741,824 bytes and therefore shows less storage capacity. Available storage capacity (including examples of various media files) will vary based on file size, formatting, settings, software and operating system, and/or pre-installed software applications, or media content. Actual formatted capacity may vary.

CXL® and BiCS FLASH™ Application

Generative Artificial Intelligence (AI) Inference & Training

Large Language Models (LLMs) require large memory to refill HBM DRAM on GPU

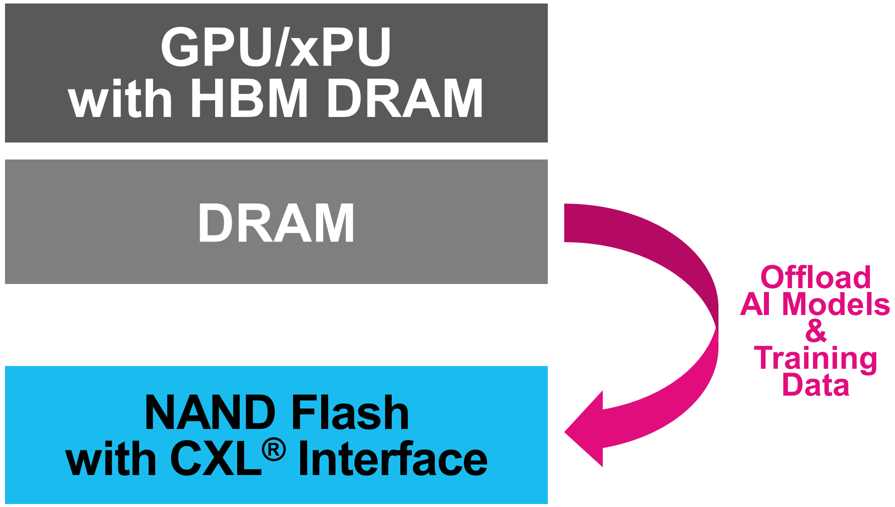

GPU/xPU with HBM DRAM
DRAM
NAND Flash with CXL® Interface
Offload AI Models & Training Data

Memory Requirements by AI Model Size

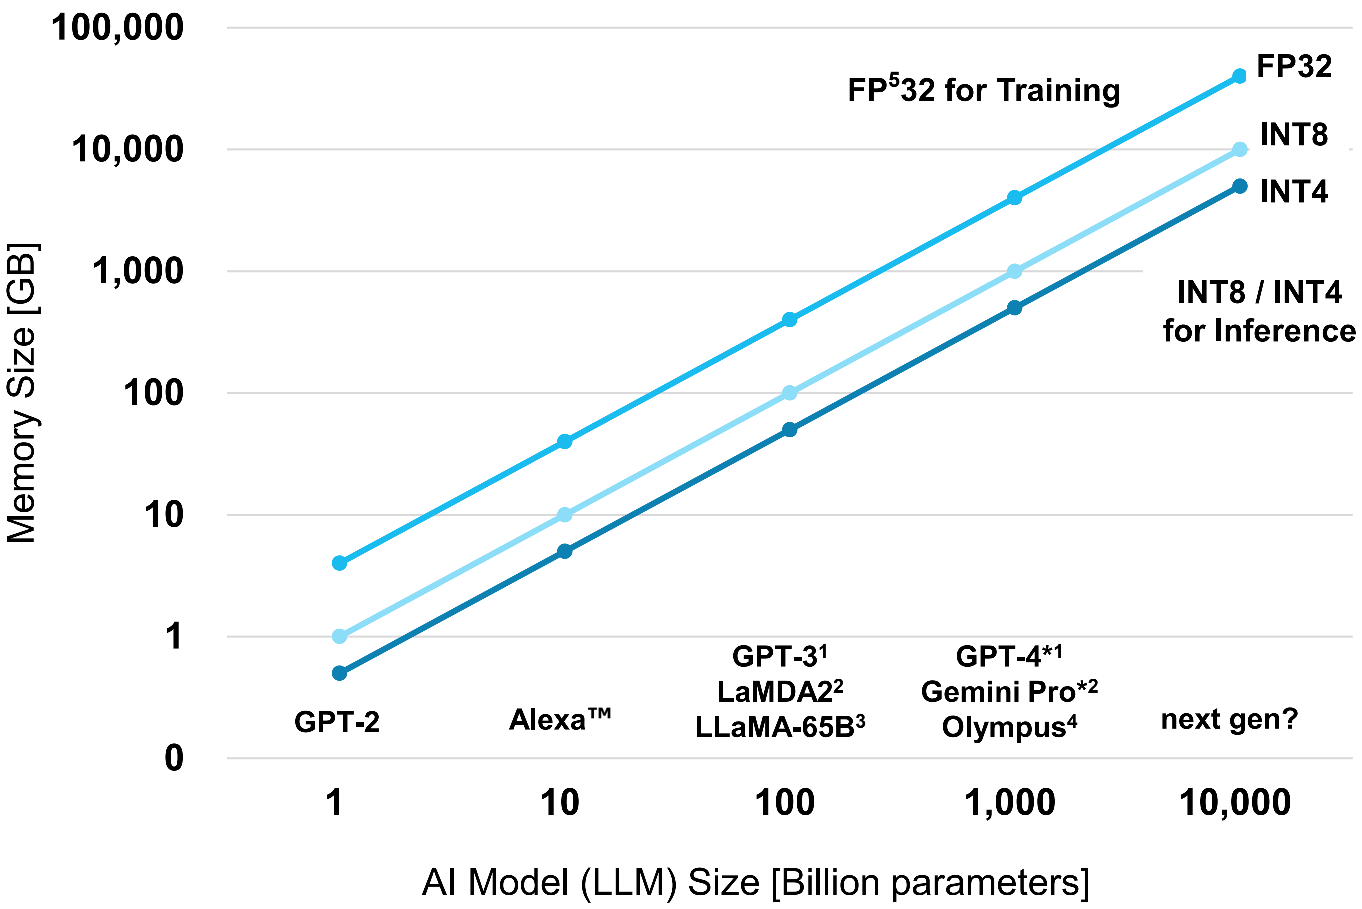

Image and graph source: created by KIOXIA

* Some model sizes are estimated by KIOXIA

CXL and Compute Express Link are registered trademarks of the Compute Express Link Consortium, Inc. Alexa and all related logos are trademarks of Amazon.com, Inc. or its affiliates. [1] Generative Pre-trained Transformer (GPT-3 and GPT-4) are a LLMs released by OpenAI. [2] LaMDA2 and Gemini Pro are a language models from Google. [3] LLaMA-65B is an LLM released by Meta. [4] Olympus is an AI model codenamed by Amazon. [5] Floating point (FP) number represents real numbers. FP32 = 32bit Definition of capacity: KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes.. Actual formatted capacity may vary.

GPU Graph Processing with Low Latency Flash

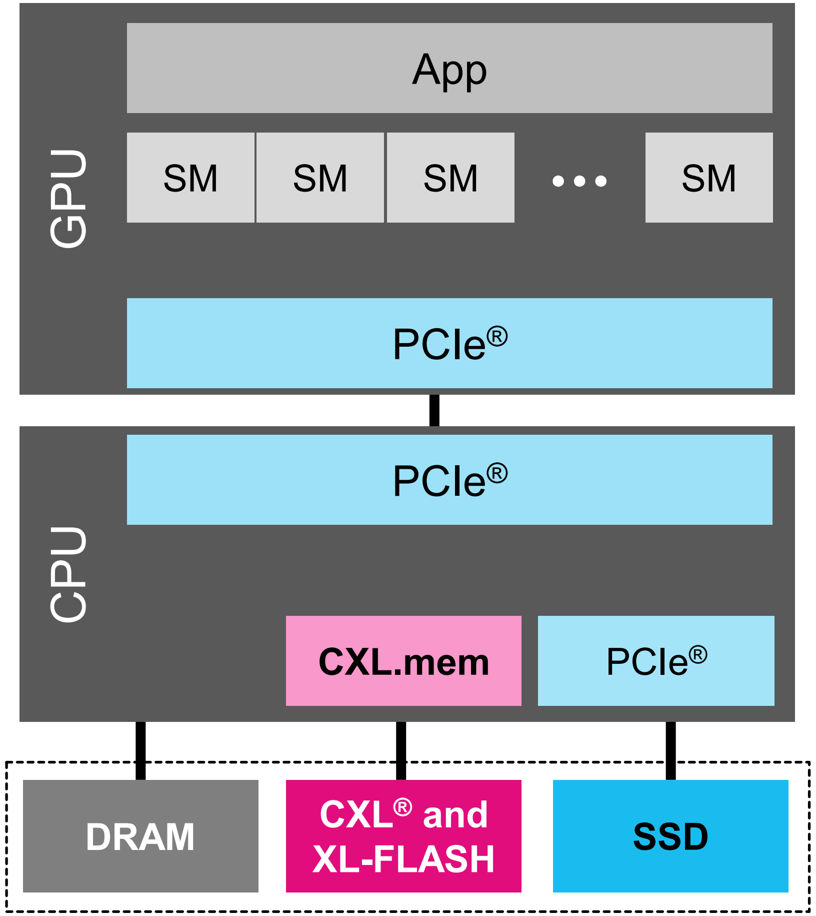

External Memory for GPU

BFS (Breadth First Search)

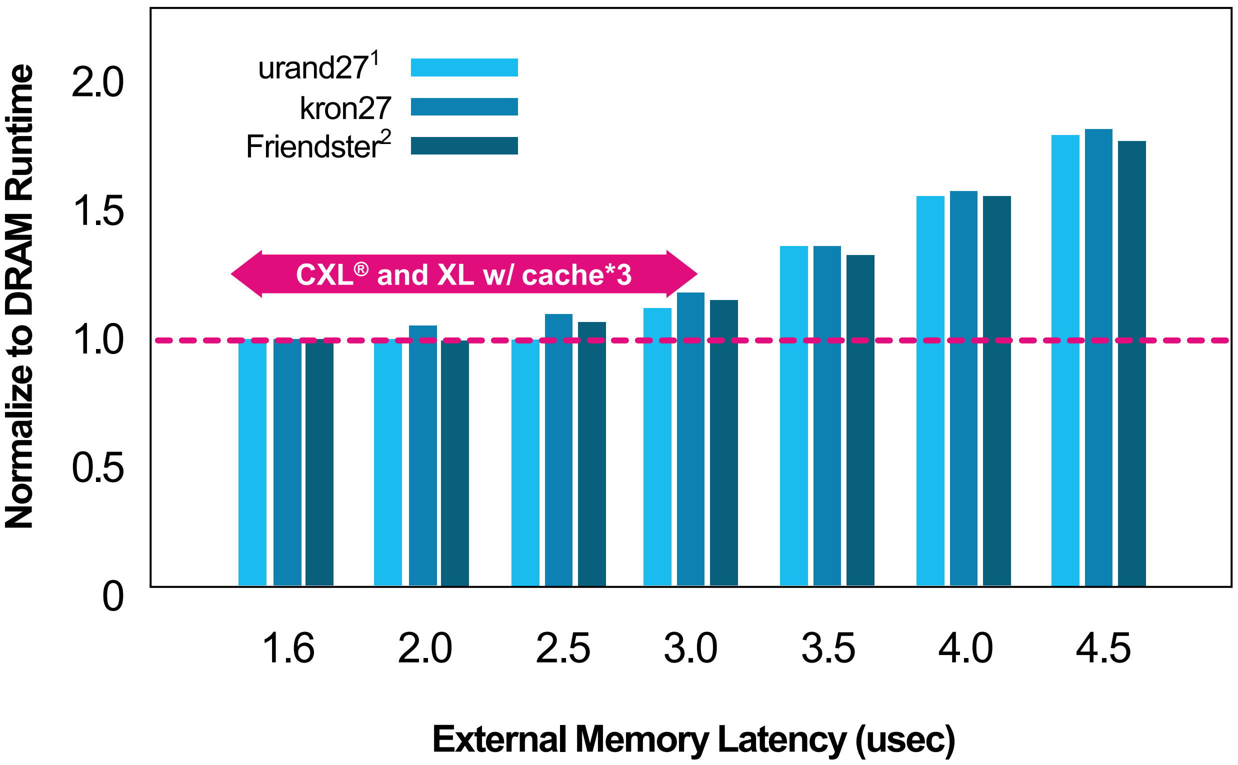

GPU Graph Processing on CXL®-Based Microsecond-Latency External Memory (SC23 MSTA)

Low latency XL-FLASH with cache can deliver DRAM-like application performance.

Image and graph source: created by KIOXIA

CXL and Compute Express Link are registered trademarks of the Compute Express Link Consortium, Inc. PCIe is a registered trademark of PCI SIG. [1] Uniform random graph datasets (urand27 and kron27). [2] Network dataset consists of ground-truth communities (Friendster)

SPEC CPU® Benchmark with Low Latency Flash

SPEC CPU: SPEC CPU is a benchmark suite designed to measure and compare the performance of CPUs, memory subsystems, and compilers through a series of compute-intensive tests.

Test Setup: 32 copies x 4 hours,10%,15% & 20% offload

10%
DRAM offload*

95.5% of DRAM performance

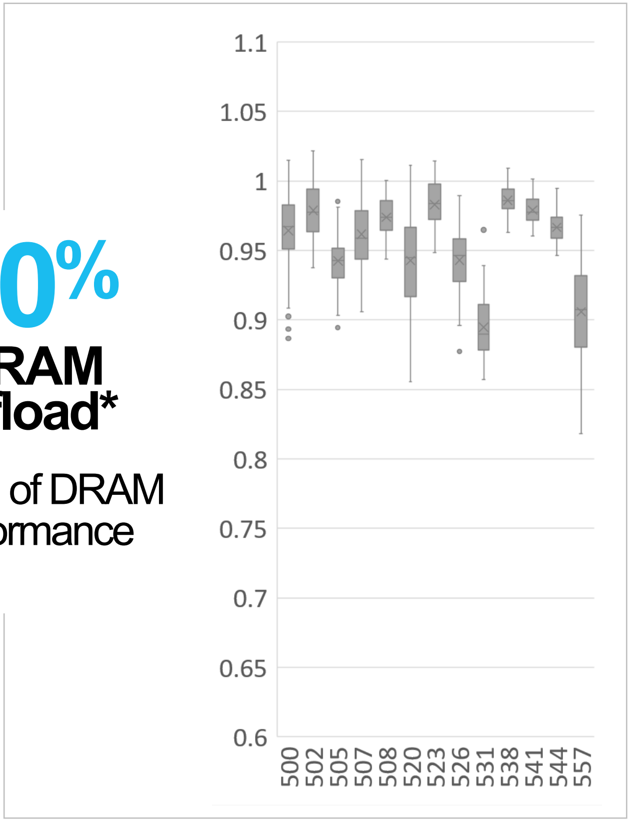

15%
DRAM offload*

94.5% of DRAM performance

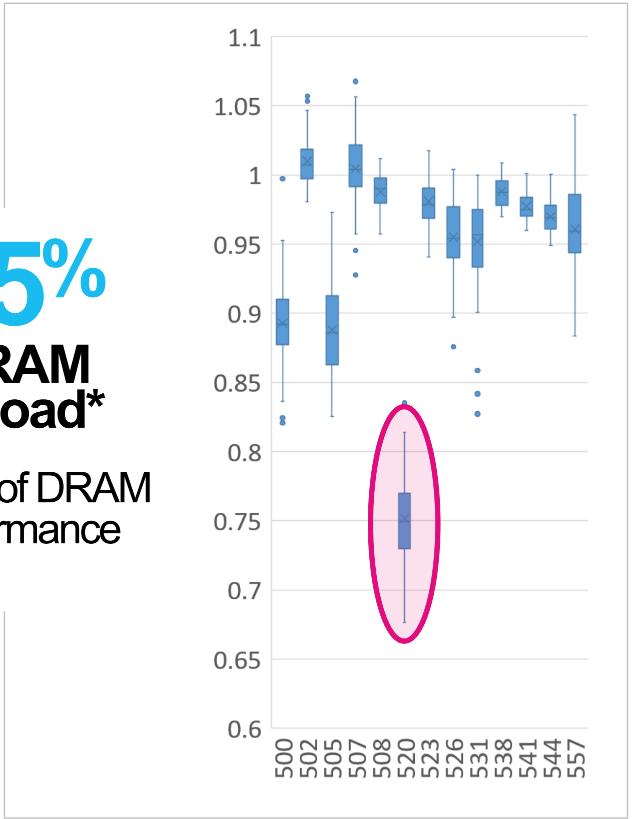

20%
DRAM offload*

90.9% of DRAM performance

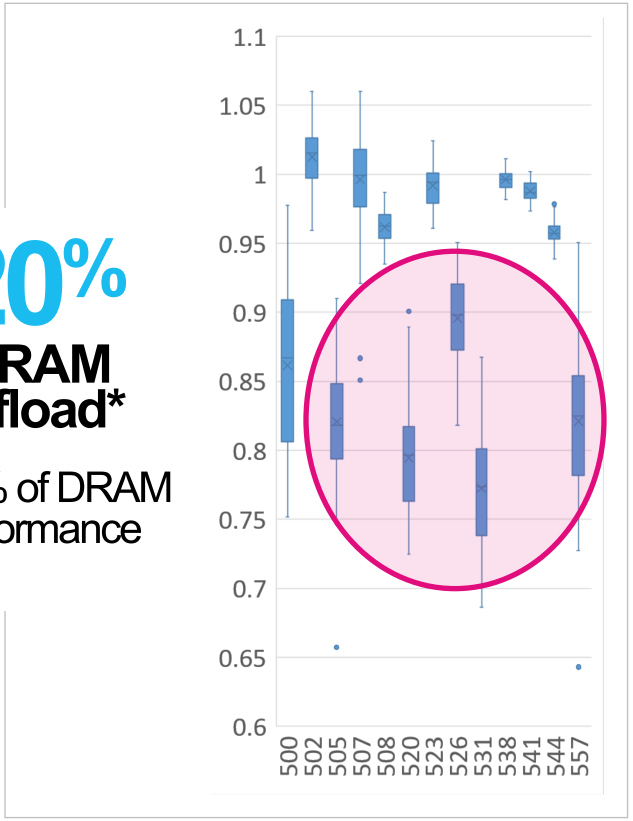

*FPGA emu. for memory offload

Image and graph source: created by KIOXIA

Low Latency XL-FLASH can offload memory with nominal performance degradation.

SPEC CPU is a registered trademark of the Standard Performance Evaluation Corporation.

Redis™ In-Memory Database with Low Latency Flash

Tested with Yahoo!™ Cloud Serving Benchmark (**YCSB**) **tool**
Setup: 10M records(14 GB), 32 client threads

Data Type: 100B*10 fields/record
Offload with Linux® TPP (Transparent Page Placement)

Test C: Get 100%

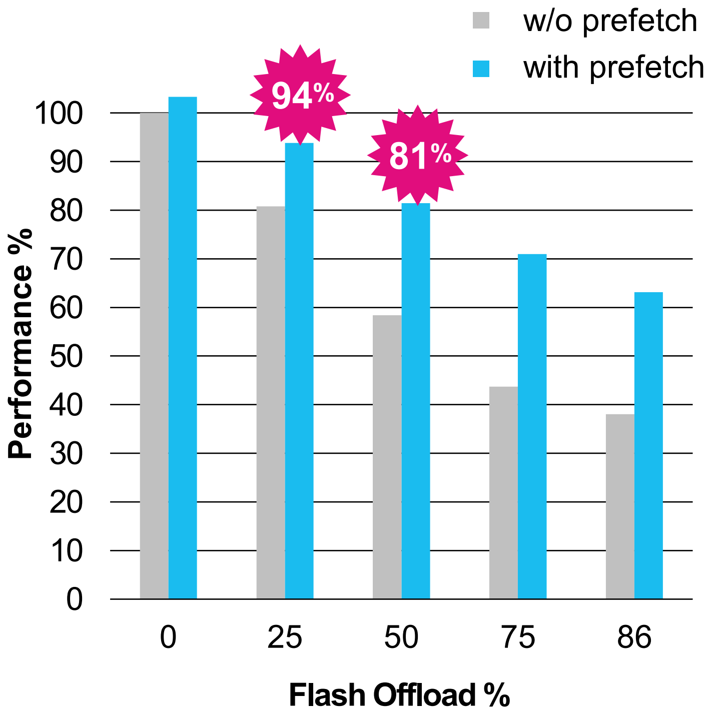

Test A: Get 50%, Put 50%

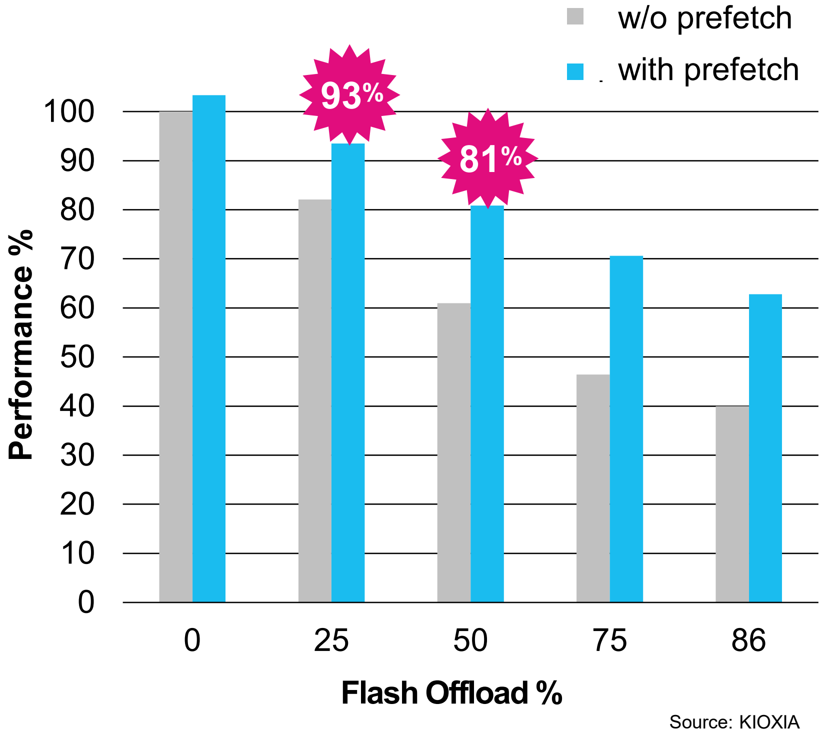

Source: KIOXIA

Zipf Distribution Workload A,C

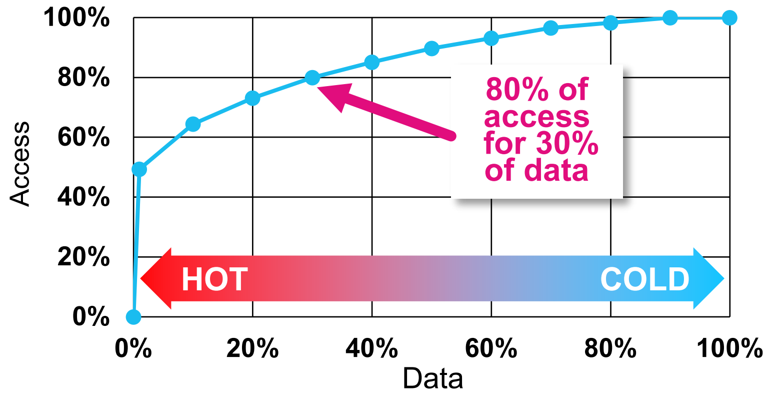

Core Scaling for Redis

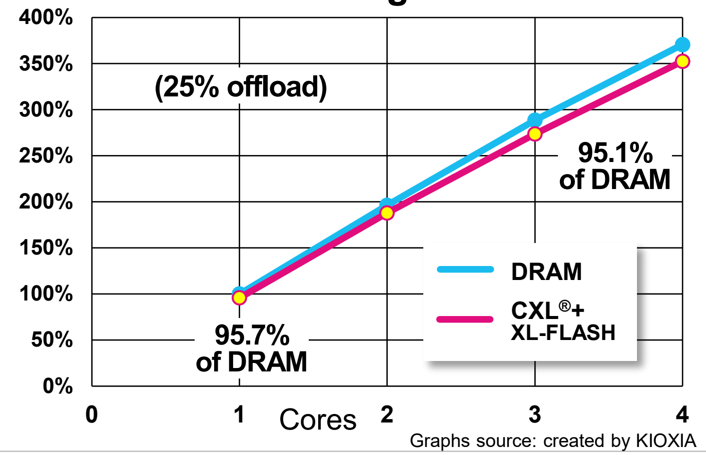

Graphs source: created by KIOXIA

YCSB demonstrates CXL® and XL-FLASH technologies can offload 25% of memory with ~5% of performance degradation.

KIOXIA

Redis is a trademark of Redis Labs Ltd. Any rights therein are reserved to Redis Labs Ltd. Yahoo! is trademark of Yahoo INC. Linux is the registered trademark of Linus Torvalds in the U.S. and other countries. CXL and Compute Express Link are registered trademarks of the Compute Express Link Consortium, Inc. Definition of capacity: KIOXIA Corporation defines a megabyte (MB) as 1,000,000 bytes, a gigabyte (GB) as 1,000,000,000 bytes and a terabyte (TB) as 1,000,000,000,000 bytes.

Challenges and Opportunities with CXL® Attached Flash Memory

All Applications Are Not The Same

- It is not suitable for latency/bandwidth sensitive applications
- Application not tuned for leveraging memory hierarchy optimally

Leverage Industry Efforts

- Transparent Page Placement - Automatically manages large memory pages
- Transparent memory tiering solutions optimizes data placement across different memory types
- Application specific libraries can further increase the efficiency and reduce cost

Applications

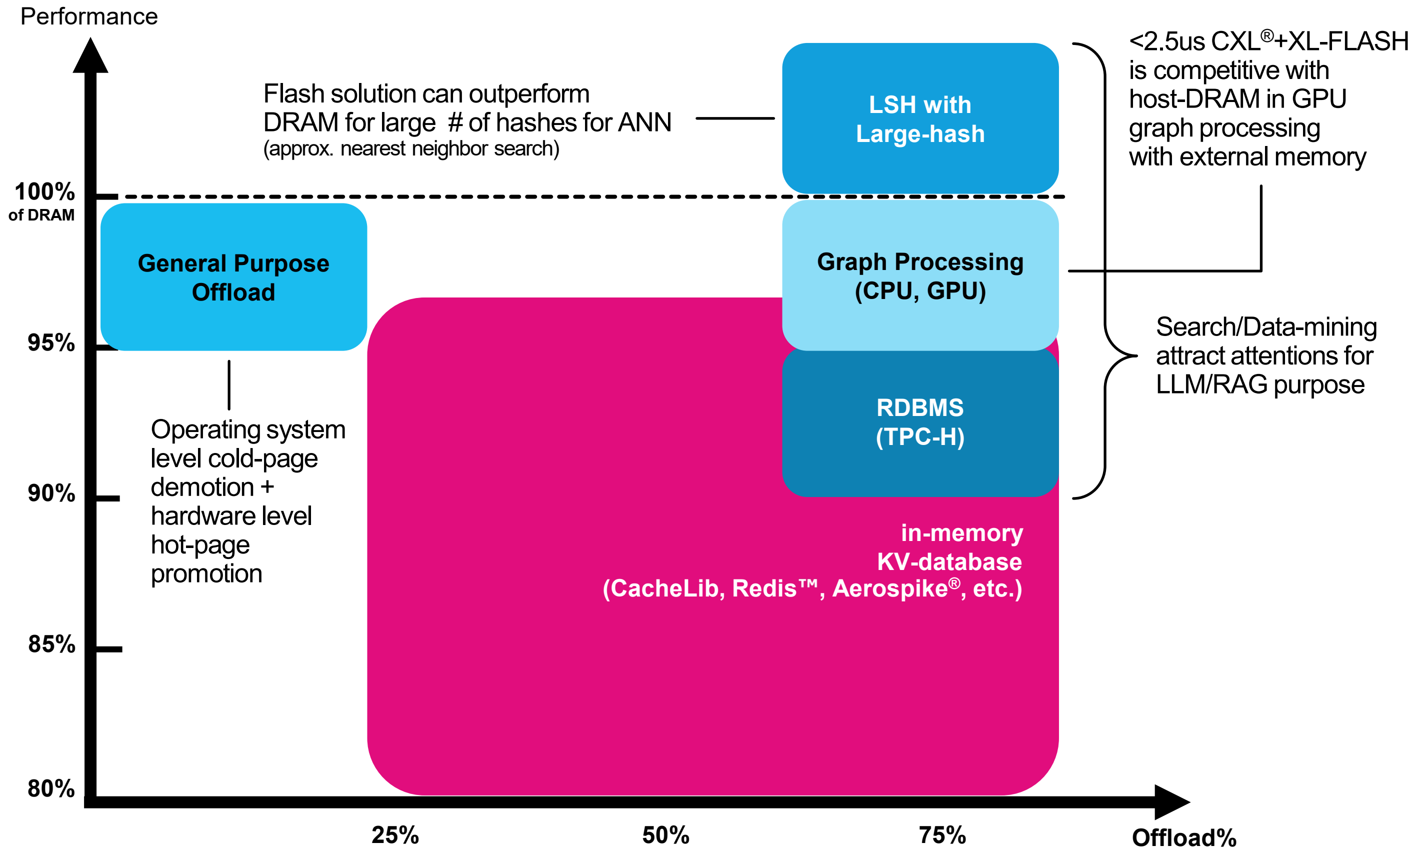

Source: KIOXIA

KIOXIA

CXL and Compute Express Link are registered trademarks of the Compute Express Link Consortium, Inc.
Redis is a trademark of Redis Labs Ltd. Any rights therein are reserved to Redis Labs, Ltd. Aerospike is a registered trademark of Aerospike, Inc.

Application and TCO

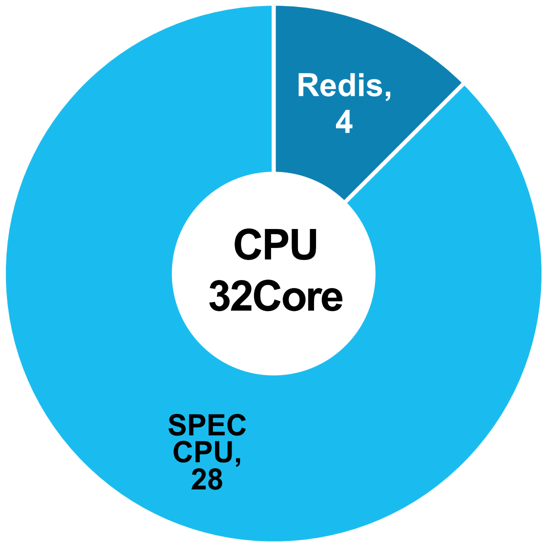

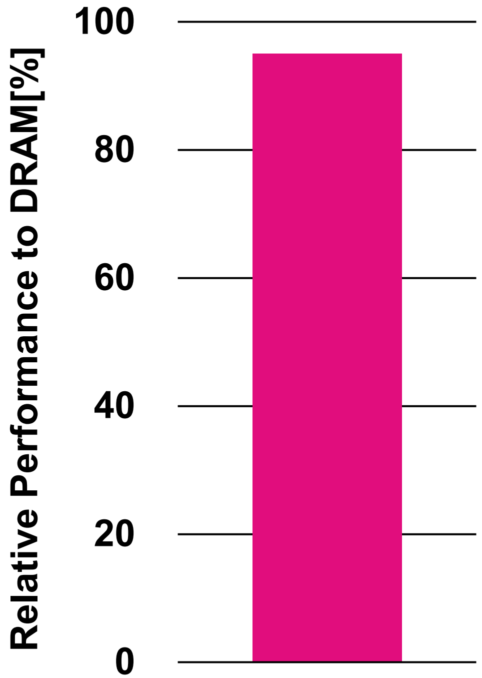

Case B	Offload	Performance
Redis™ x 4 cores	25%	94%
SPEC CPU® x 28 cores	15%	95%
total	**23%**	**95%**

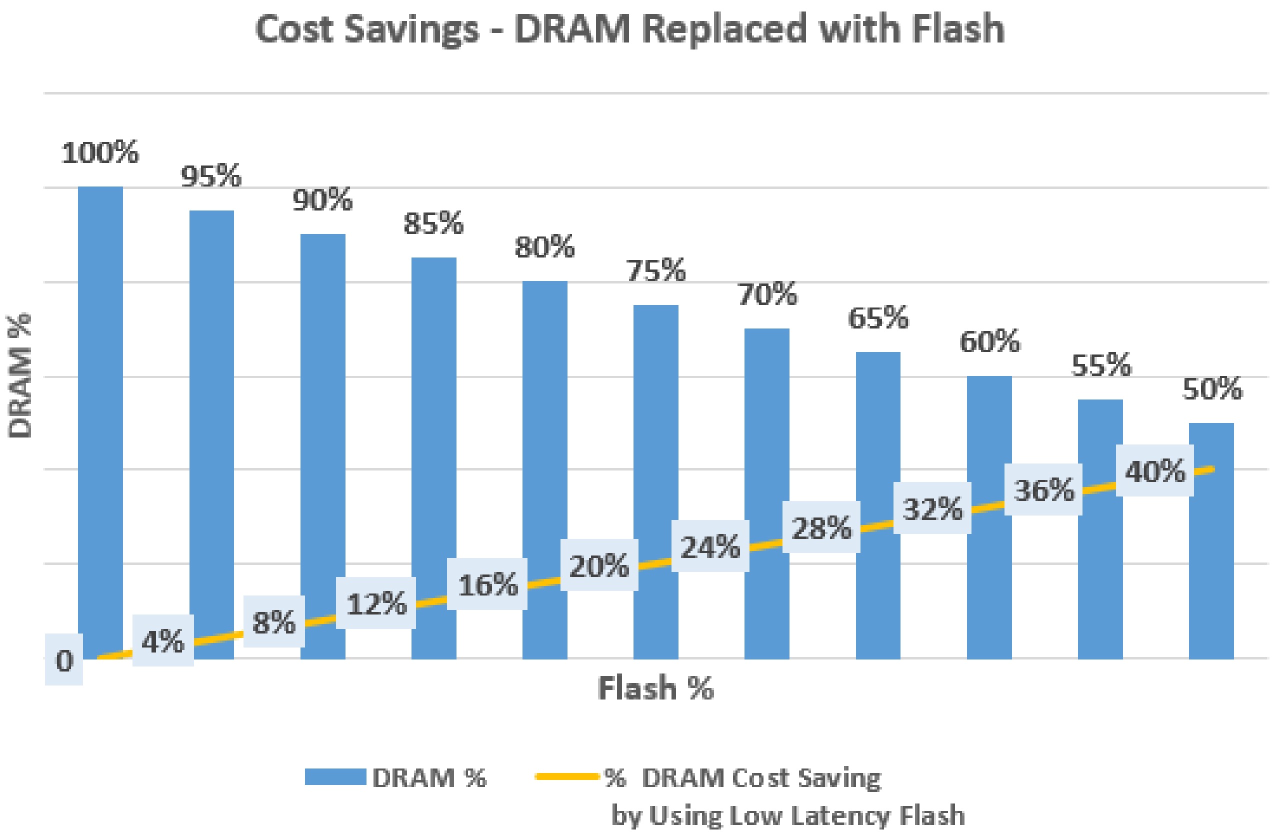

KIOXIA assumption: Price of Low Latency Flash is 1/5 DDR Price.

Image and graph source: created by KIOXIA

Redis is a trademark of Redis Labs Ltd. Any rights therein are reserved to Redis Labs Ltd.
SPEC CPU is a registered trademark of the Standard Performance Evaluation Corporation

Flash memory can jump the semantic wall.

Flash memory is proven and reliable media.

Flash memory lowers the system TCO.

Flash memory can further perform and reduce cost with software.

If you are working on large memory intensive applications like **Data Mining, Artificial Intelligence (AI), Machine Learning (ML), Analytics, High Performance Computing (HPC), Graph Processing Applications,** Please visit **KIOXIA Booth #307** for collaboration opportunities.

Images and icons source: licensed to KIOXIA from Shutterstock

The product images shown are a representation of the design models and not an accurate product depiction.

KIOXIA