Memory Centric 시대

- SK Hynix Solution 방향 (Storage 중심)

2024.10

SKHY / Solution AT
Junghyun Joh

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

AI Pipeline and Storage Workload

AI Data Pipeline

Data Ingestion	Data Preparation	Training	Checkpoint & Restore	Inference
				
Sequential Write	Mostly Sequential	Mostly Random	Mostly Sequential Read/Write	Mostly Sequential Read

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Challenges to AI Deployment

AI Datacenter

Ingest | Prep | Training | Checkpoint | Inference

STORAGE

Limited rack power & space, growing dataset

COMPUTE

Power optimized SSD as cache device

On-device AI

Inference

USER EXPERIENCE

AI model loading time and user experience

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Memory Centric

Memory is system competitive

Past System
Calculator : Data Simple Calculation

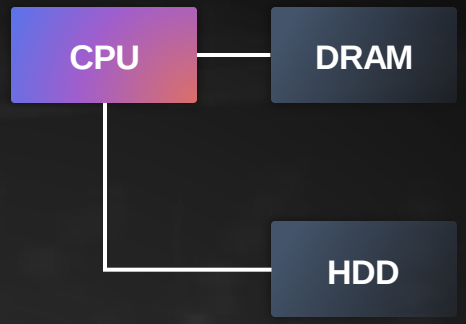

Present System
Workable : Data/information Reading

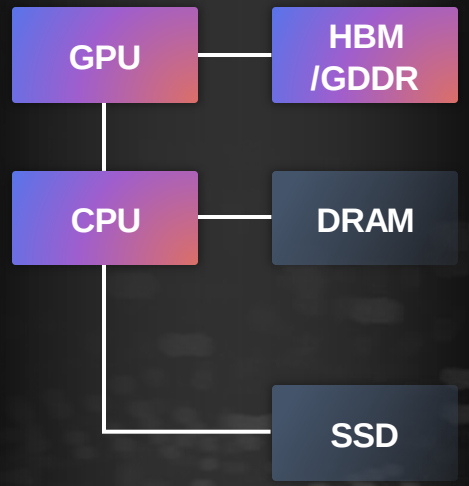

Future System
Thinkable AI : Data/information Creation

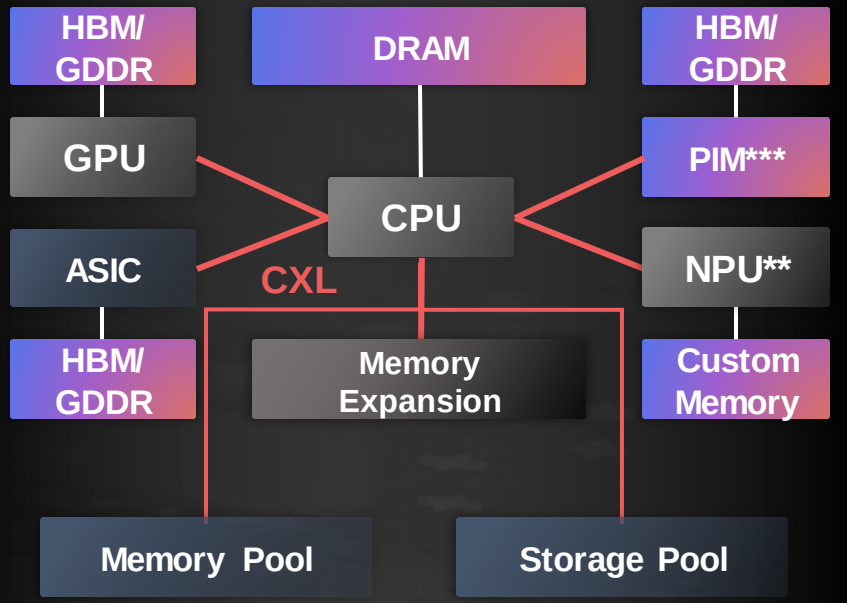

* CXL: Compute eXpress Link
** NPU: Neural Processing Unit
***PIM: Processing in Memory

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

SK hynix Solution for the AI Era

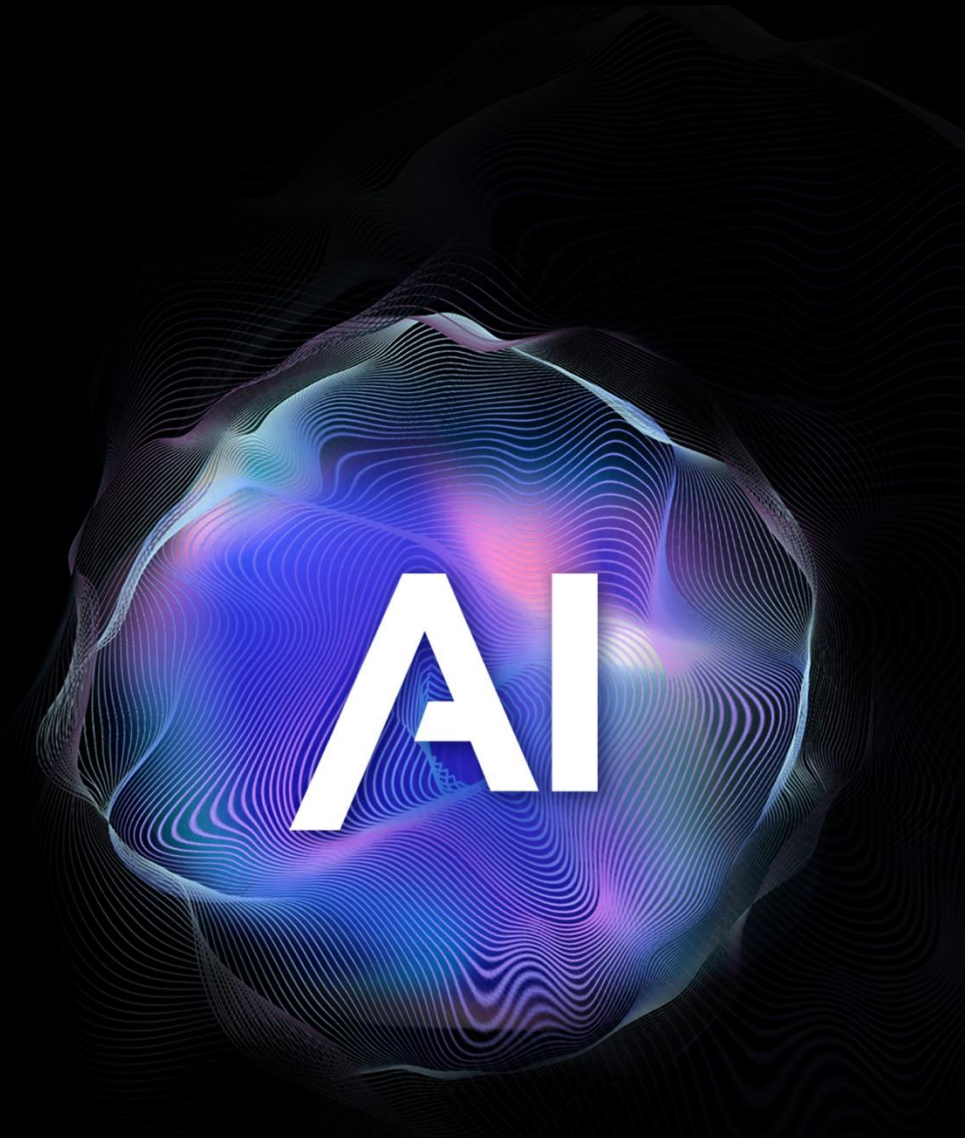

Junghyun Joh
SK Hynix / Solution AT

SK hynix

FMS 2024
SK하이닉스 미래포럼
AI HW Summit

Contents

SK hynix

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Facing Problems

Risk increasing due to Energy, Cost, Environmental issue and etc.

Cloud Cover: Cloud Prices Rise as the Era of Generative AI Dawns

NOVEMBER 29, 2023

By Jordan Galhardo-Burnett, Pankaj Sherawat, Peter Dankert, Engelhar Scognam

BCG

Microsoft's and Google's AI plans clouded by concerns of rising costs

Tech giants tout new tools that will need significant investment as the technology takes hold

FINANCIAL TIMES

Meta's Costs Rise Rapidly as Zuckerberg Vows to Keep Spending on AI Arms Race

Shares in the social-media company fell more than 12% after it revealed AI investment plans while reporting record revenue

By Salvador Rodriguez Follow

Updated April 24, 2024 6:01 pm ET

THE WALL STREET JOURNAL.

AI Is Pushing The World Toward An Energy Crisis

Ariel Cohen Contributor

I cover energy, security, Europe, Russia/Eurasia & the Middle East

Sam Altman Invests in Energy Startup Focused on AI Data Centers

Investment by OpenAI CEO highlights artificial intelligence's electricity appetite

By Amrith Ramkumar Follow

April 22, 2024 5:00 am ET

THE WALL STREET JOURNAL.

AI boom sparks concern over Big Tech's water consumption

Microsoft, Google and Meta are using more water to cool down data centres that power artificial intelligence products

FINANCIAL TIMES

AI Is Accelerating the Loss of Our Scarcest Natural Resource: Water

Cindy Gordon Contributor

CEO, Innovation Leader Passionate about Modernizing via AI

Forbes

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

AI Datacenter Challenges: Storage

In the AI datacenter:

Avg. processor active power[1)]

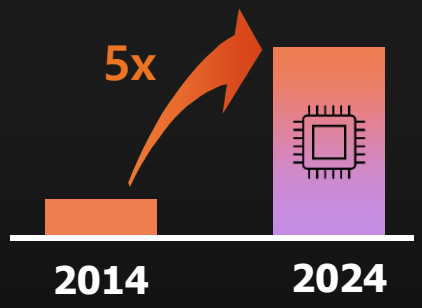

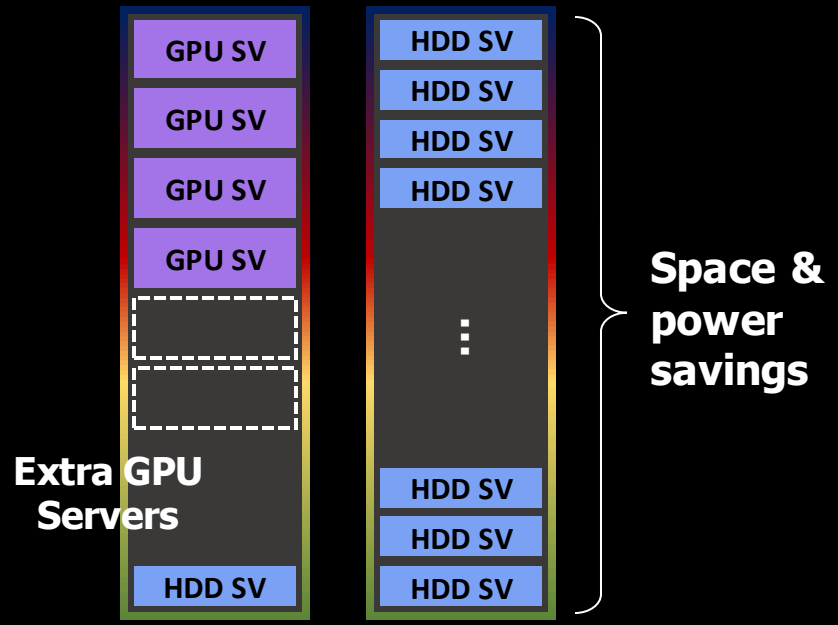

Avg. rack power[1)]

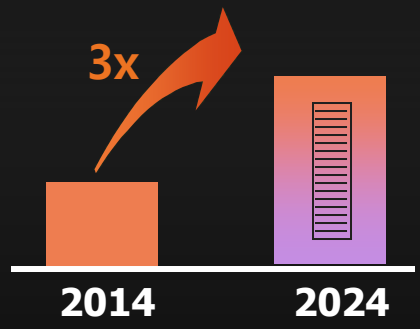

Limited Datacenter Power and Space

1) Source: SDxCentral

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

AI Datacenter Challenges: Compute

General-purpose server for Data Prep (ETL)

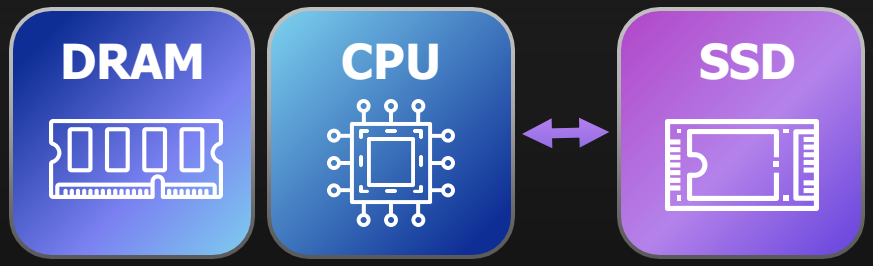

GPU Server for Training/ Inference

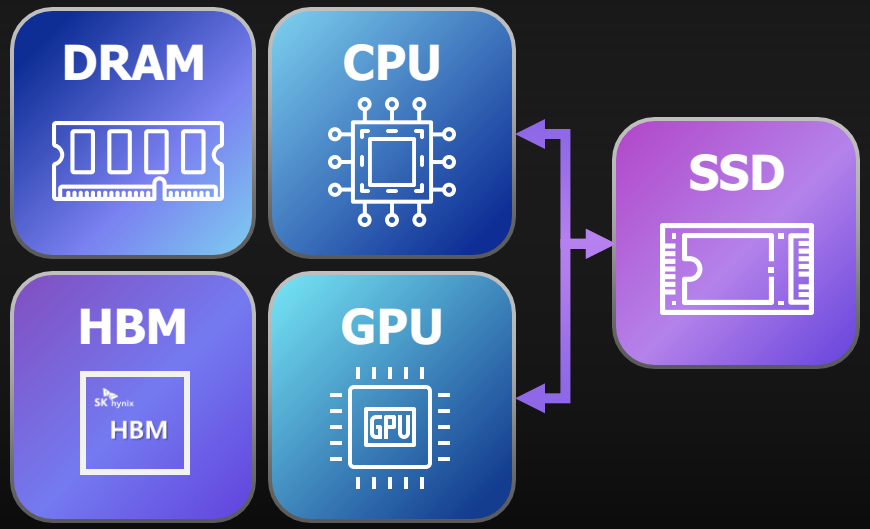

Need Something with optimized power & high performance

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

On-device AI Challenges: PC/Mobile

Computing Power

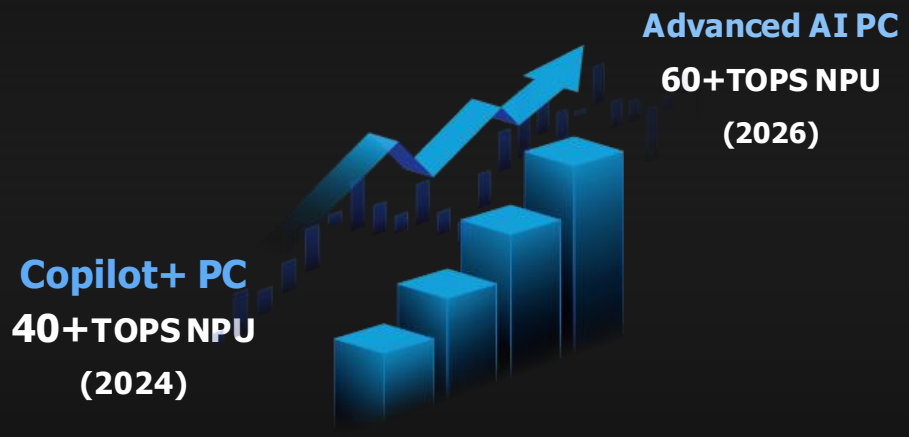

High-performance SSD required

User Needs

Minimal AI Model load time (<1s)
from Storage to all memory states

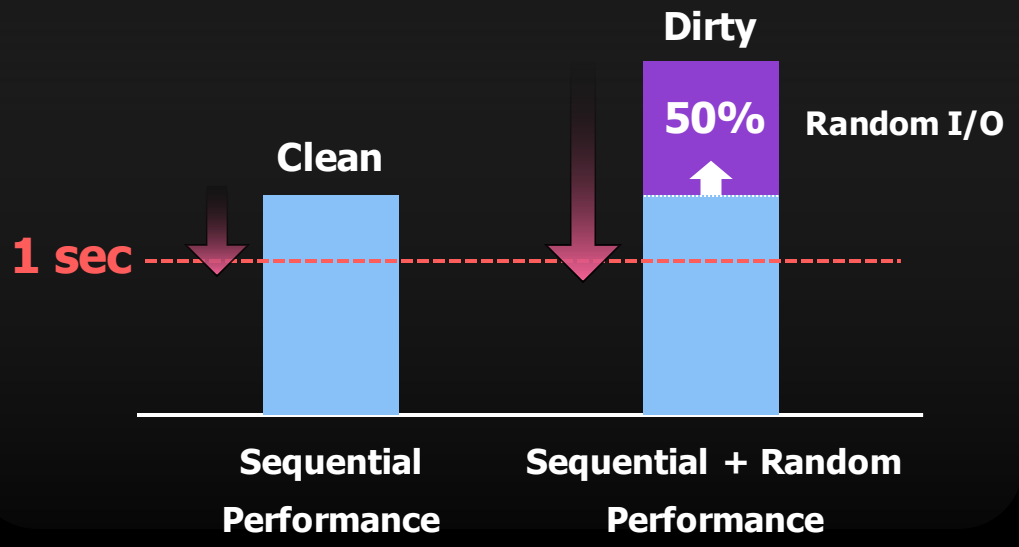

(based on SK hynix mainstream Gen4 SSD)

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Contents

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

GenAI-ready Solutions

For AI Datacenter:

PCIe Gen5 Enterprise SSD
Best-in-class Performance, IOPS/W

61TB QLC Enterprise SSD
World's highest-capacity PCIe SSD

For On-device AI:

PCIe Gen5 Client SSD
World-first for mainstream client

Zoned UFS
World-first, vertical optimized mobile storage

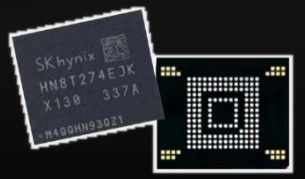

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

High-density QLC eSSD: vs. HDD

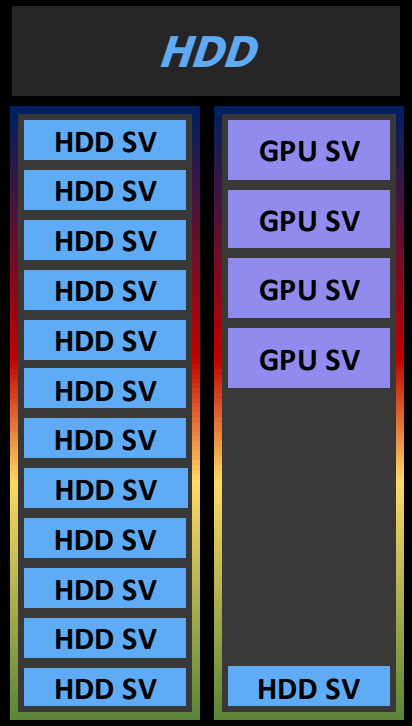

30TB HDD		122TB QLC SSD
4	# GPU servers configurable	5 (+1 vs. HDD rack)
4U x 13ea	# Storage servers configurable	2U x 7ea
39.78kW Compute 65% Storage 35%	Power Consumed	36.7kW Compute 89% Storage 11%
2 Racks	# Racks required	1 Rack

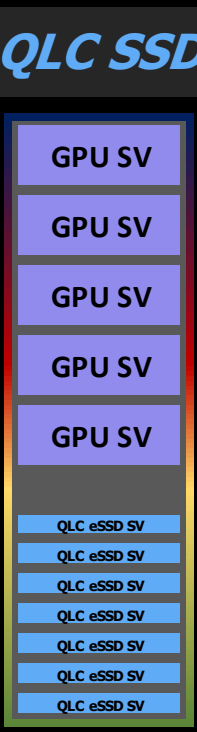

QLC SSD = +1 GPU server (more compute) per rack;
reduced power consumption & rack footprint

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

PCIe Gen5 cSSD: Gen-on-Gen Advanced

Performance

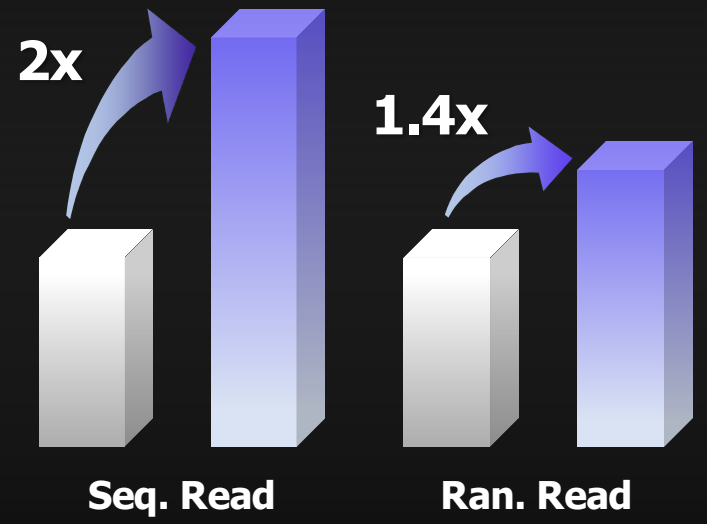

Gen4 vs. Gen5

Power Efficiency

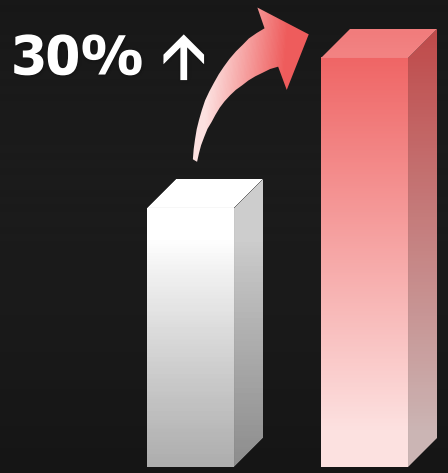

Gen4 vs. Gen5

AI Model[1)] Launch

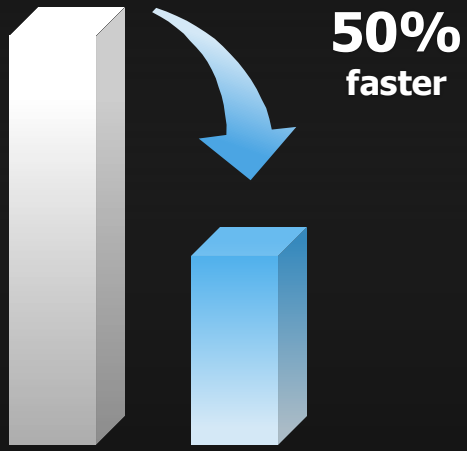

LOADING 100% @ Clean

Gen4 vs. Gen5

1) AI Model: Llama2 7BQ8

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

PCIe Gen6 & UFS 5.0 Solutions

PCIe Gen6 SSD

Performance[1)]

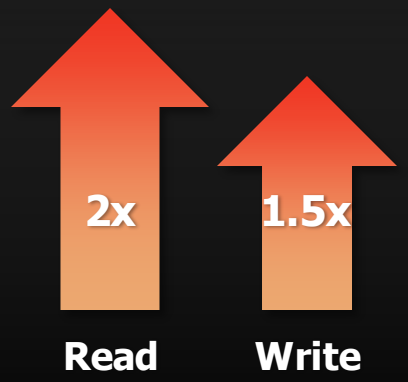

Power Efficiency

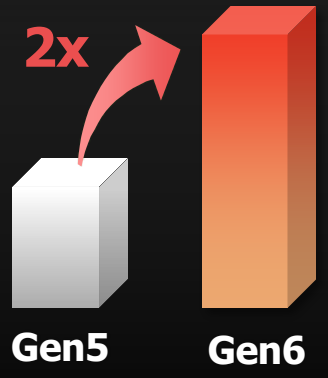

UFS 5.0[2)]

Random Performance

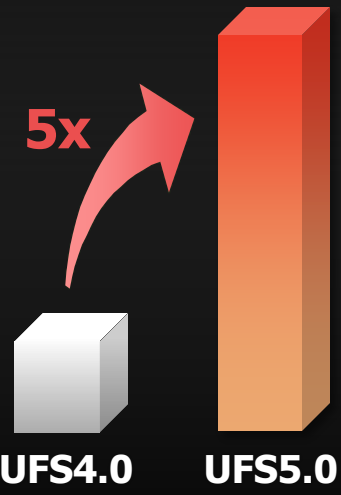

Power Efficiency

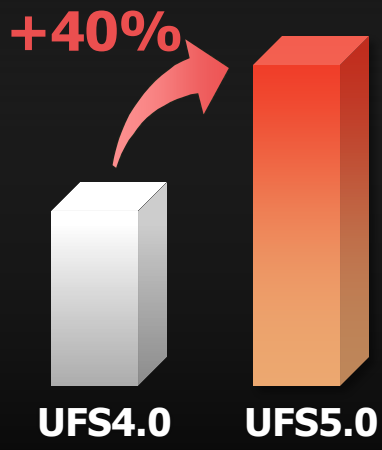

Using SK hynix-proprietary, H-TPU™ architecture

1) Performance bound by max. power: 25W on Gen5 and Gen6 (estimate)

2) Estimates

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] High-density QLC eSSD: Leading Capacities

Specification

- PCIe 4.0, NVMe 1.4c, OCP 2.0

Capacity

- 7.68TB, ~ 61.44TB
- 122TB upcoming, early ’25

Performance

- 7GB/s, 1M IOPs

“D5-P5336” QLC eSSD

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] PCIe Gen5 eSSD: BIC Performance

Specification

- PCIe 5.0, NVMe 2.0, OCP 2.0

Capacity

- 1.92TB ~ 15.36TB (Read-Intensive)
- 1.6TB ~ 12.8TB (Mixed-Use)

Performance (Max.)

- Seq. R 14.5GB/s , Seq. W 9.3GB/s
- Ran. R 3,200K IOPS , Ran. W 400K IOPS
- Up to 12% better sequential write of vSAN HCI bench vs. competition

"PS10x0" Compute eSSD

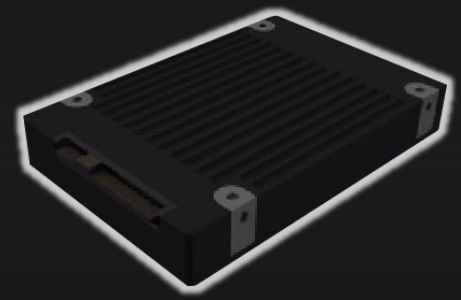

U.2/U.3

E3.S

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] PCIe Gen5 cSSD: World's First Mainstream

Specification

- PCIe 5.0, NVMe 2.0c
- HYPERWRITE™ Cache Technology

Capacity

- 512GB, 1TB, 2TB

Performance

- Seq. R 14GB/s (@10W)
- Seq. W 12GB/s (@10W)

"PCB01" mainstream cSSD

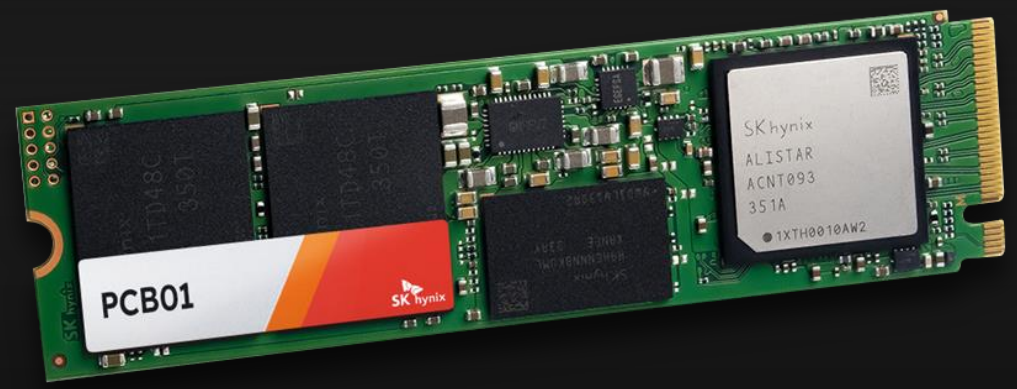

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Zoned UFS: World's First for Advanced Mobile

Specification

- UFS 4.0

Capacity

- 512GB, 1TB

ZUFS Spec

- Zone type SWR[1)]
- # Max Open Zones: 6

Vertical Optimized Mobile Storage

Zoned UFS 4.0

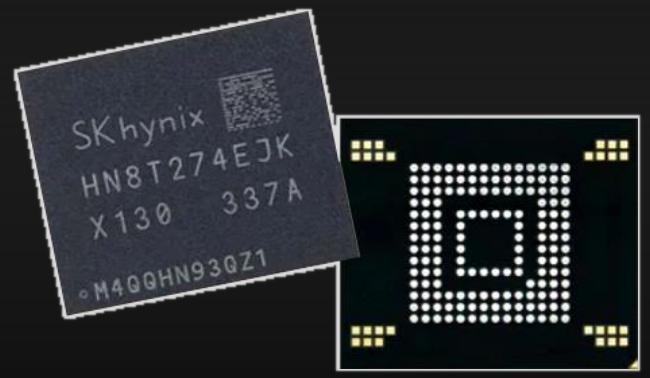

1) SWR : Sequential write required zone type

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Zoned UFS: Optimized for Advanced Mobile

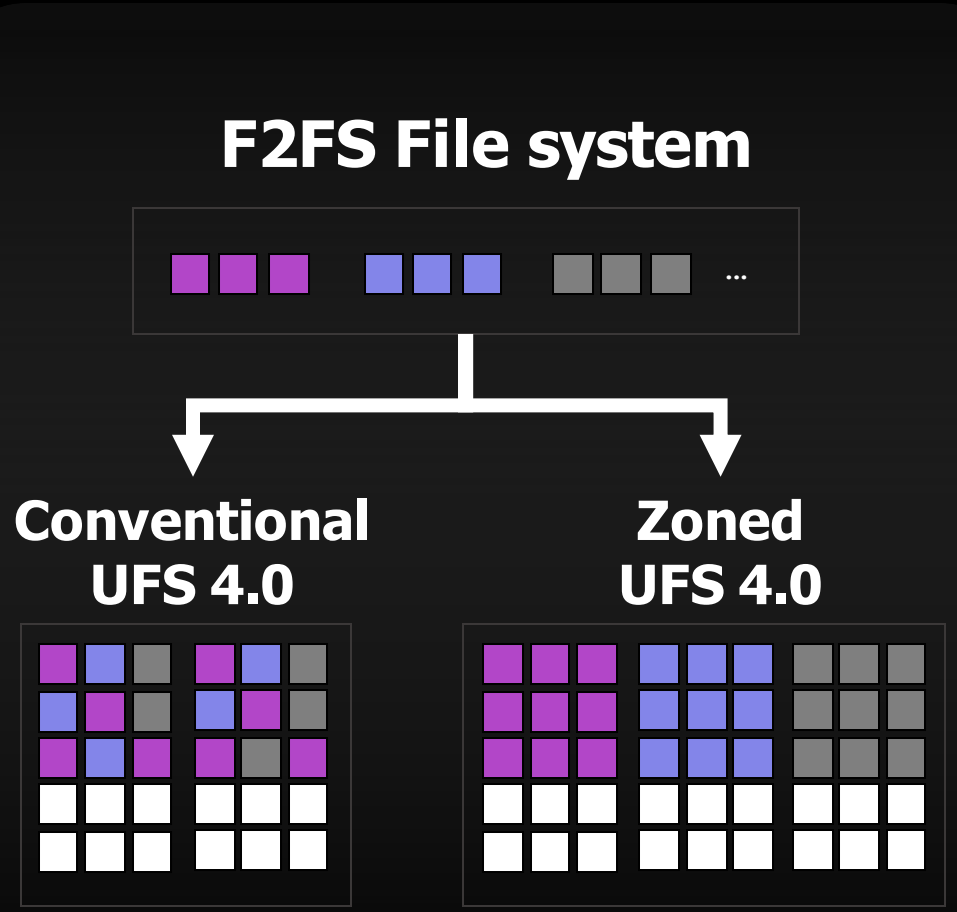

Application Launch Time[1)]

0.45 x

Memory Allocation Time[2)]

0.67 x

Product Lifetime

40% longer

1) Condition: long-hours use

2) Avg. of memory & storage backup/restore time

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Contents

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Complete AI Memory Solutions Provider

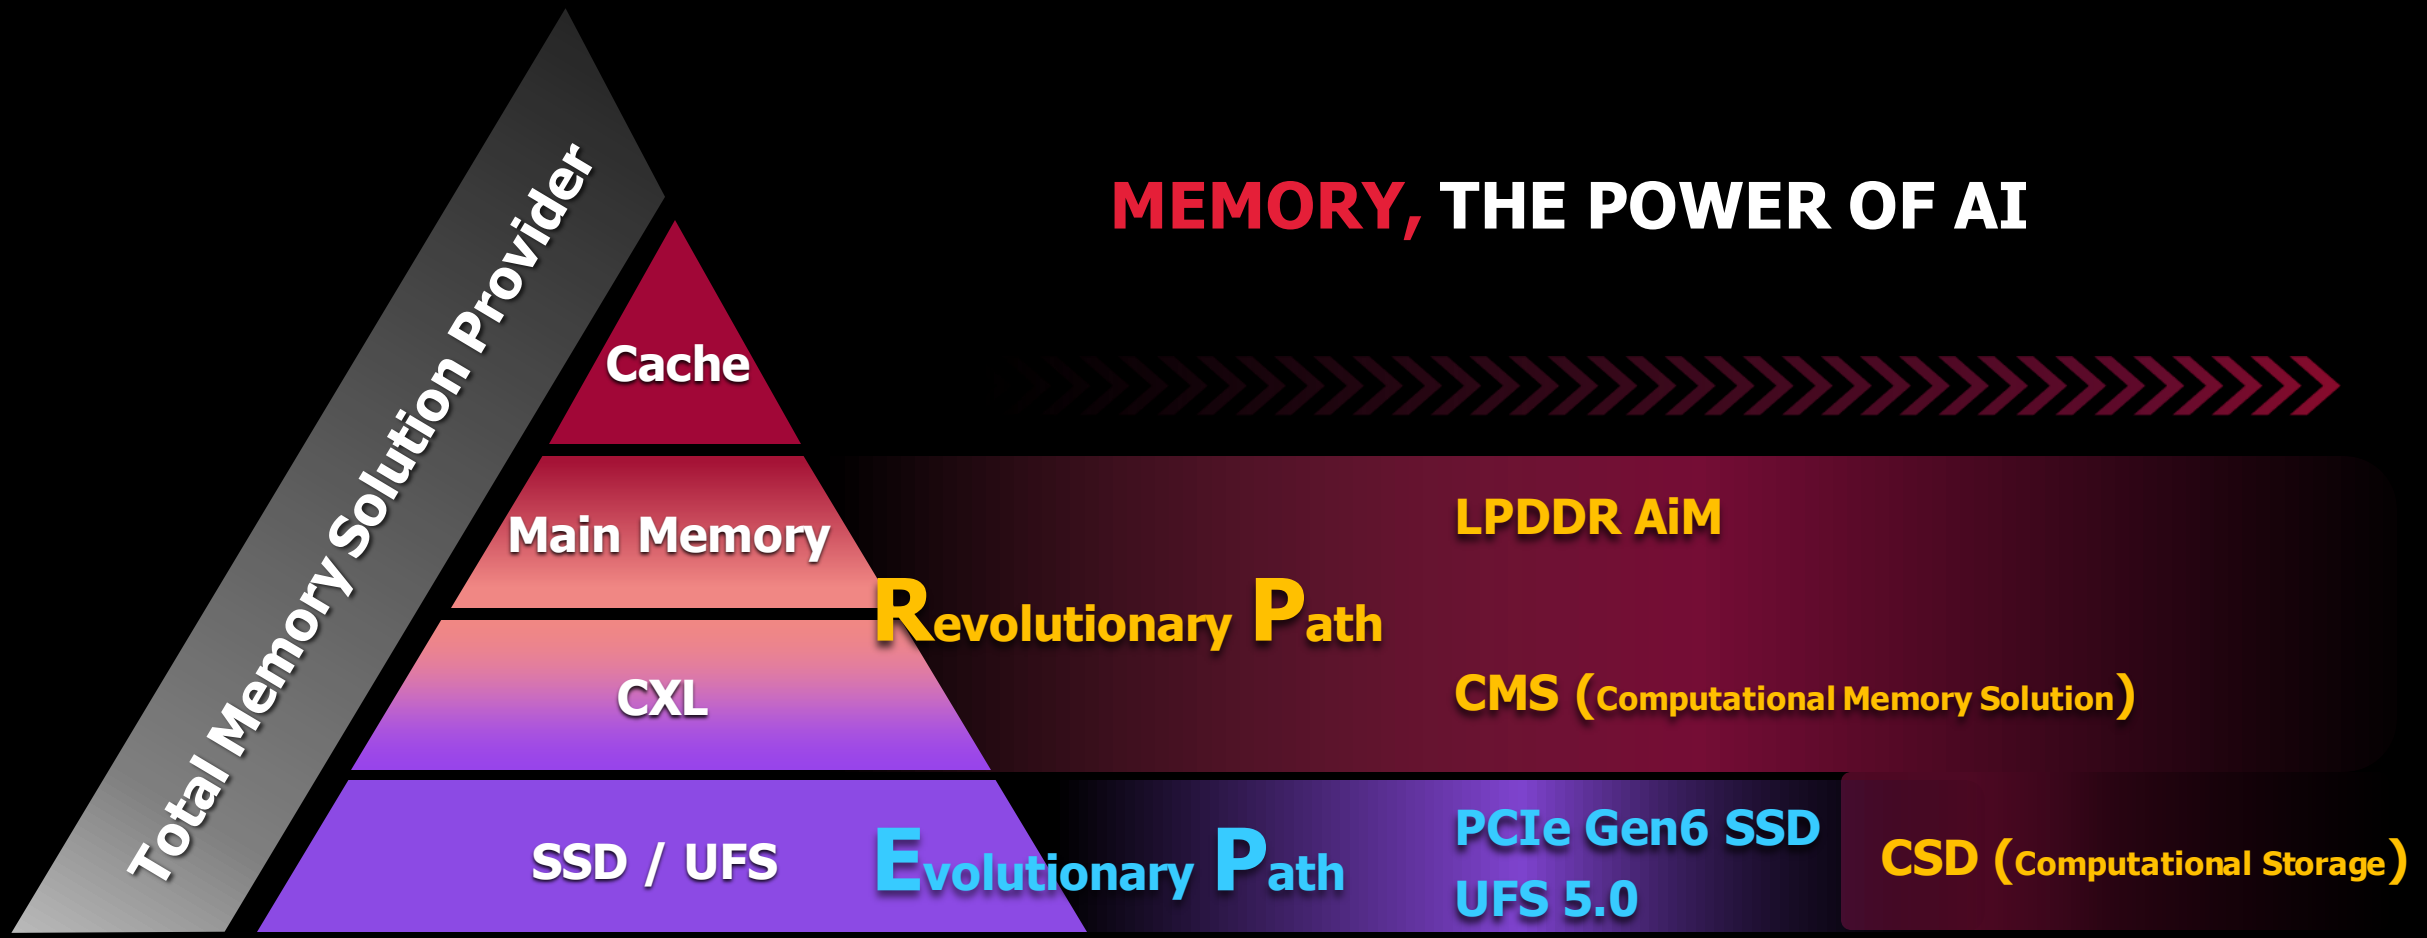

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Memory Centric

Fusion in All Memory Layers

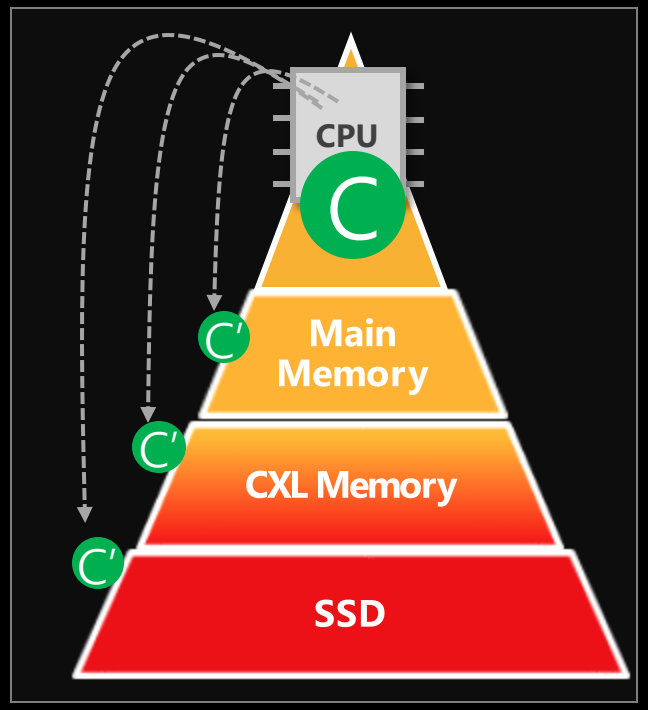

C' : computing unit

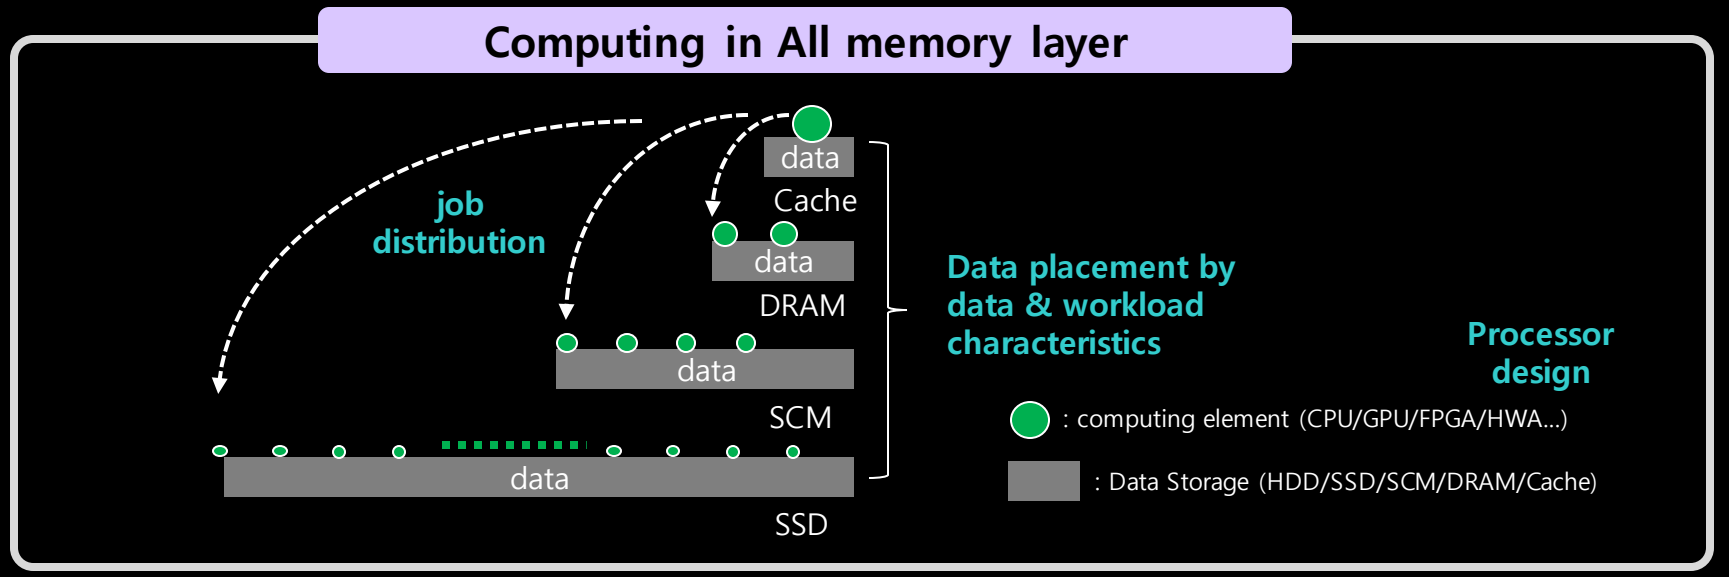

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Memory & Computing Fusion

Various Memory + Computing Solution from AiM to CSD

		Workload density	Workload
Memory Die level (AiM) 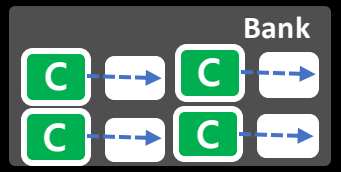	• Computing in memory die • High density workload	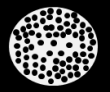	Transformer
Memory Card level (CMS) 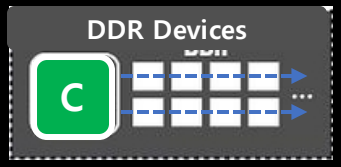	• Computing in memory card • Low density workload	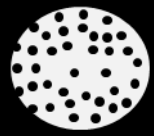	Embedding Data Analytics
Storage level (CSD) 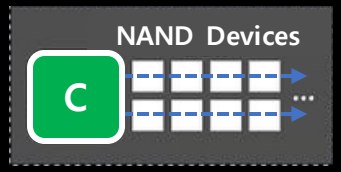	• Computing in SSD • Low density & high capacity workload		Data Analytics

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Processing in Memory

- Even with multi batch processing, due to the attention processing for long context, there is huge portion of memory intensive function

LLM block architecture

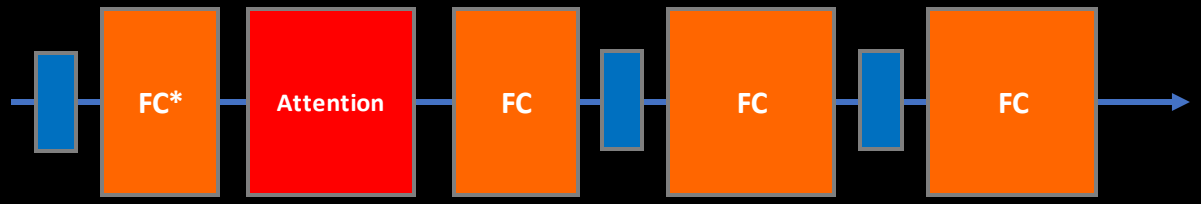

* FC : Fully connected layer

: FC weight GEMV
- In Multi batch, weight data can be reused → computing intensive & high GPU utilization

: Attention token GEMV
- Even in Multi batch, token data cannot be reused → still memory intensive & low GPU utilization

: Misc. functions
- Regardless batch → computing intensive & high GPU utilization

Processing time portion

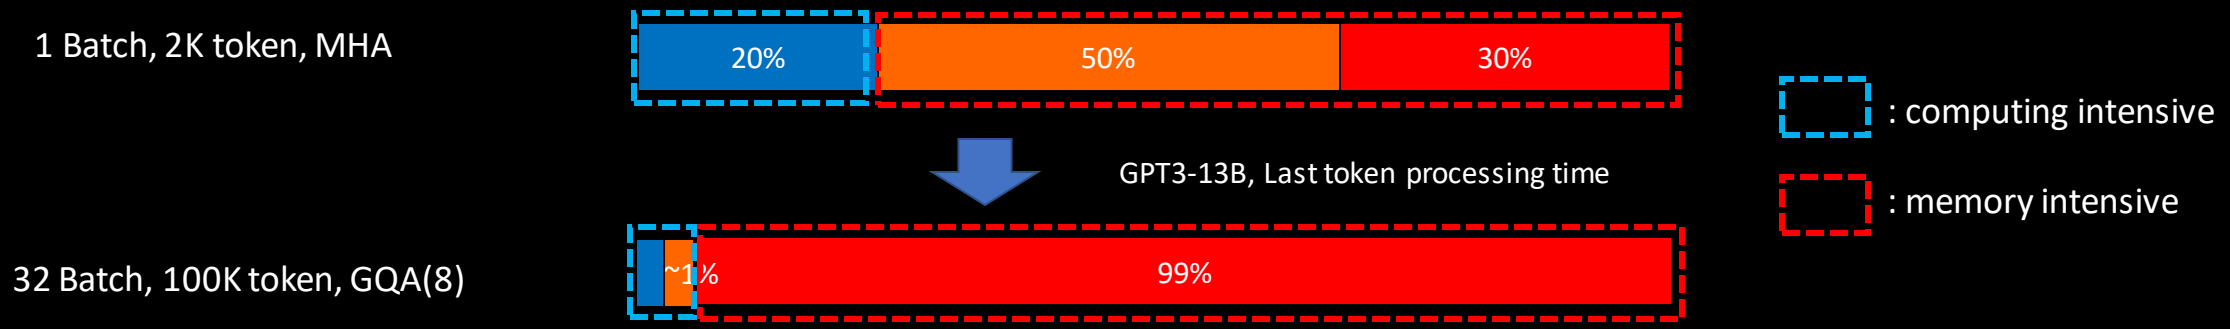

LLM service trend : longer context

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] AiMX system from DC to Edge device LLM

- To provide different level of customer experience and save your operating cost

Datacenter LLM		Edge device LLM
Larger LLM size Multi-batch Longer context	99%	75%	Smaller LLM size Single-batch Long context

Both system has memory intensive portion as a majority

Boosting your Datacenter LLM
by adding AiMX as an attention accelerator

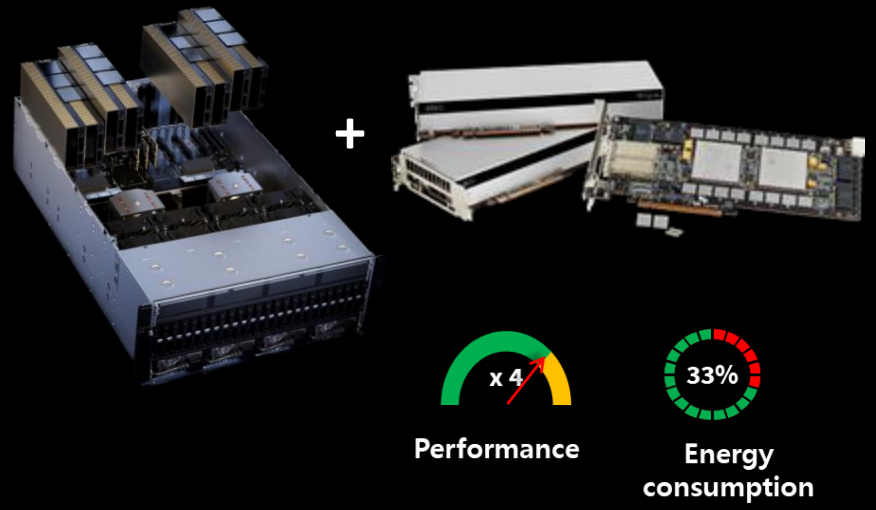

Boosting your Edge device LLM
by replacing memory with AiM

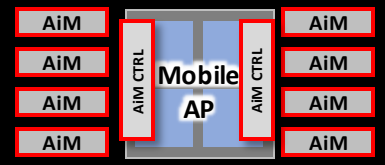

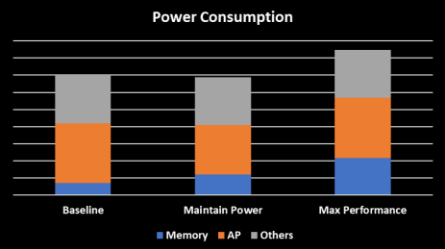

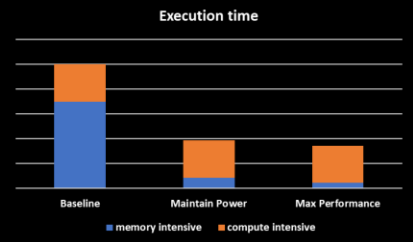

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Computational Memory Solution

CMS

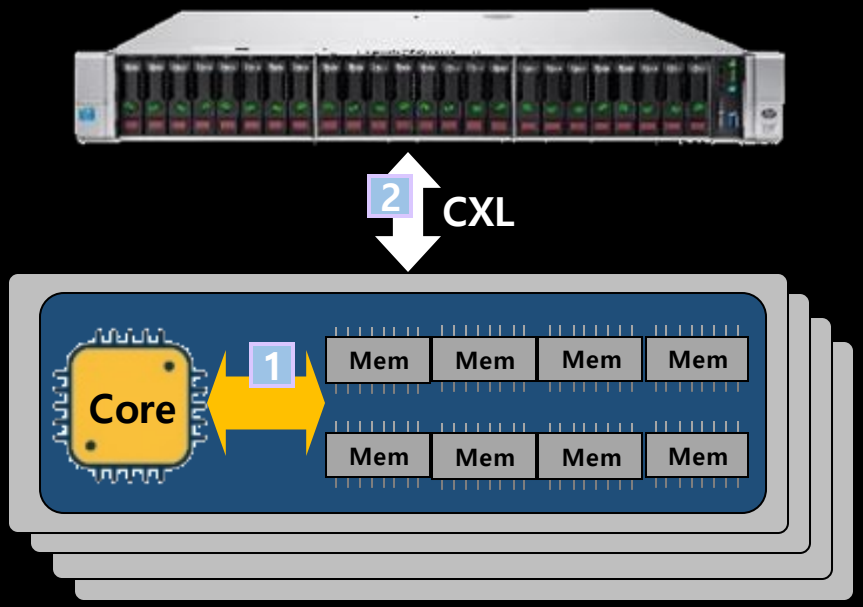

1	**Performance Improvement** By utilizing the higher Bandwidth inside the CXL module
2	**Energy Efficiency Improvement** By minimizing data movement between host and CXL module

Use case

Recommendation Model

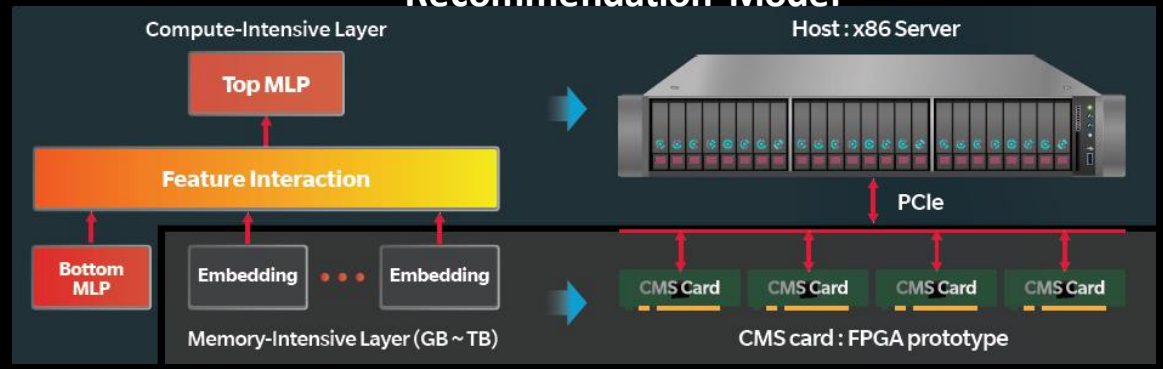

Lightning DB for LLM

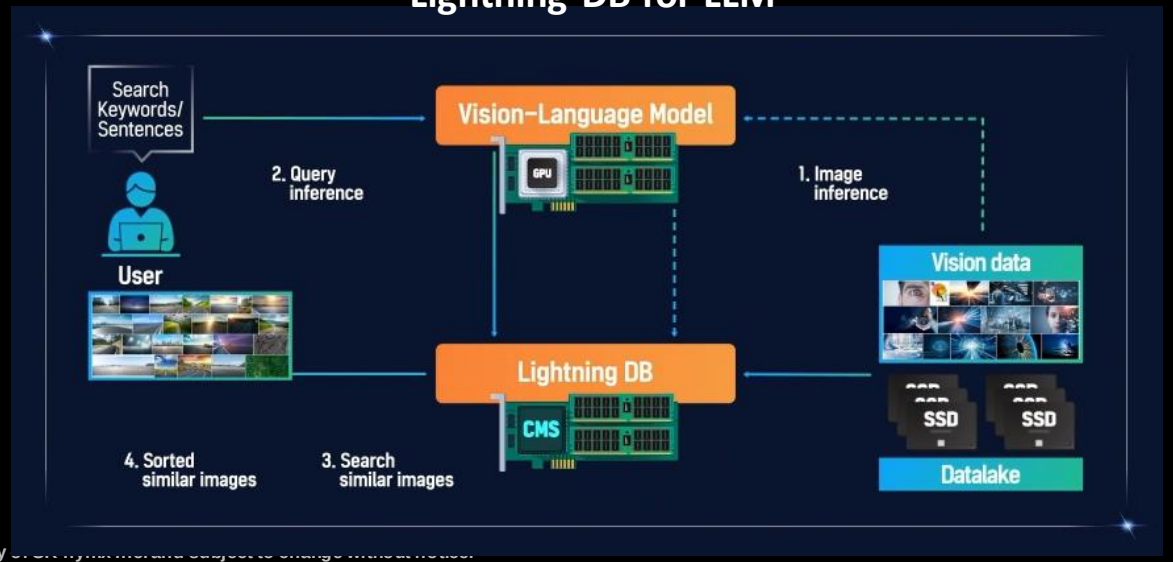

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] CXL™ Computational Memory Solution (CMS)

CMS 2.0

- CXL Memory Expansion
- Near-memory Processing
- Reduce data movement between CPU & CXL
- Free up CPUs to do other useful work

Data Analytics Performance

Perf. & CPU Utilization

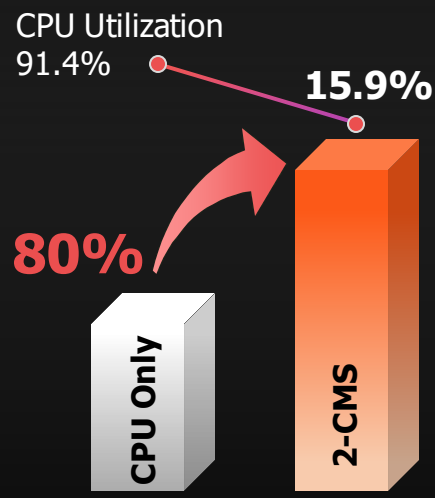

Energy Efficiency

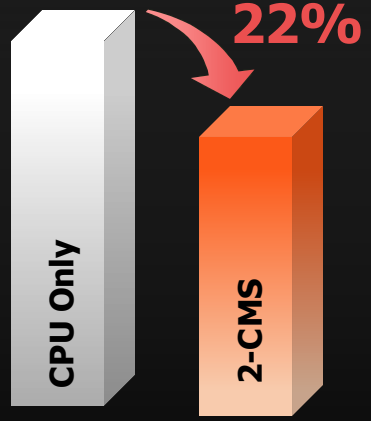

Benchmark : TPC-DS Q28 (Dataset : 38.67GB)

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Computational Storage Device

CSD

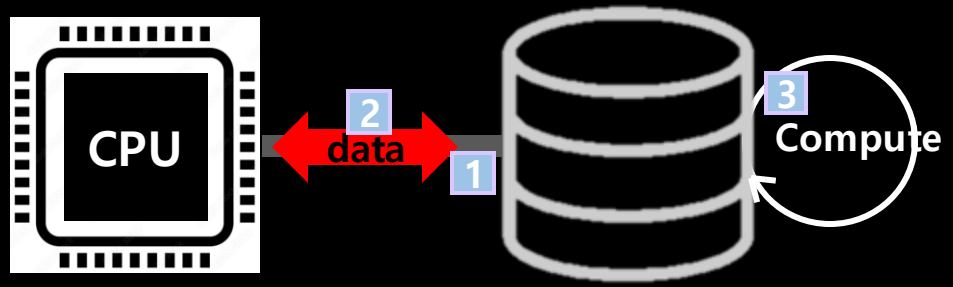

1	**Performance Improvement** By avoiding interface bottleneck
2	**Energy Efficiency Improvement** By minimizing data movement
3	**Reduce CPU overhead** By offloading tasks from Host CPU

Use case

Object based Computational Storage for Big Data Analysis

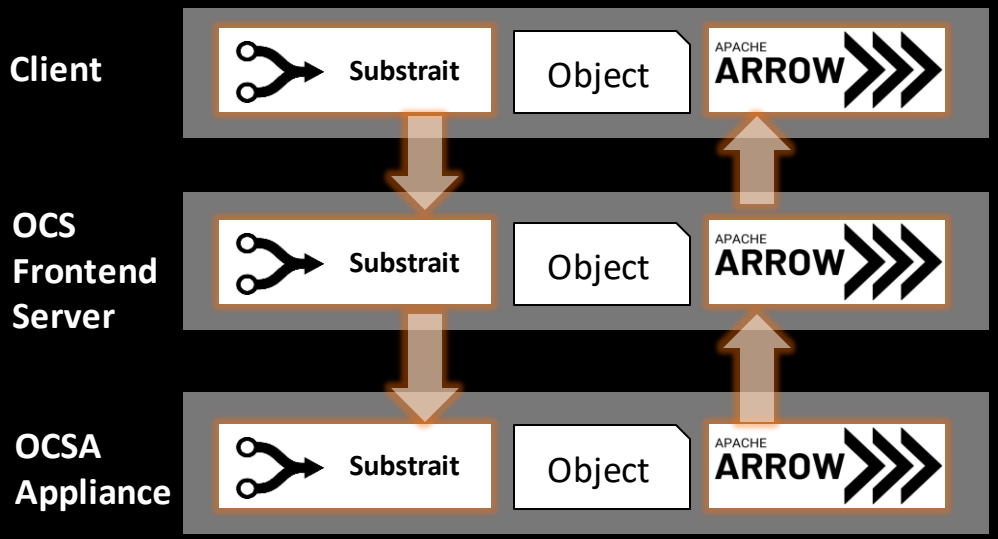

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Appendix] Object-based Computational Storage (OCS)

Speed and efficiency from data-awareness

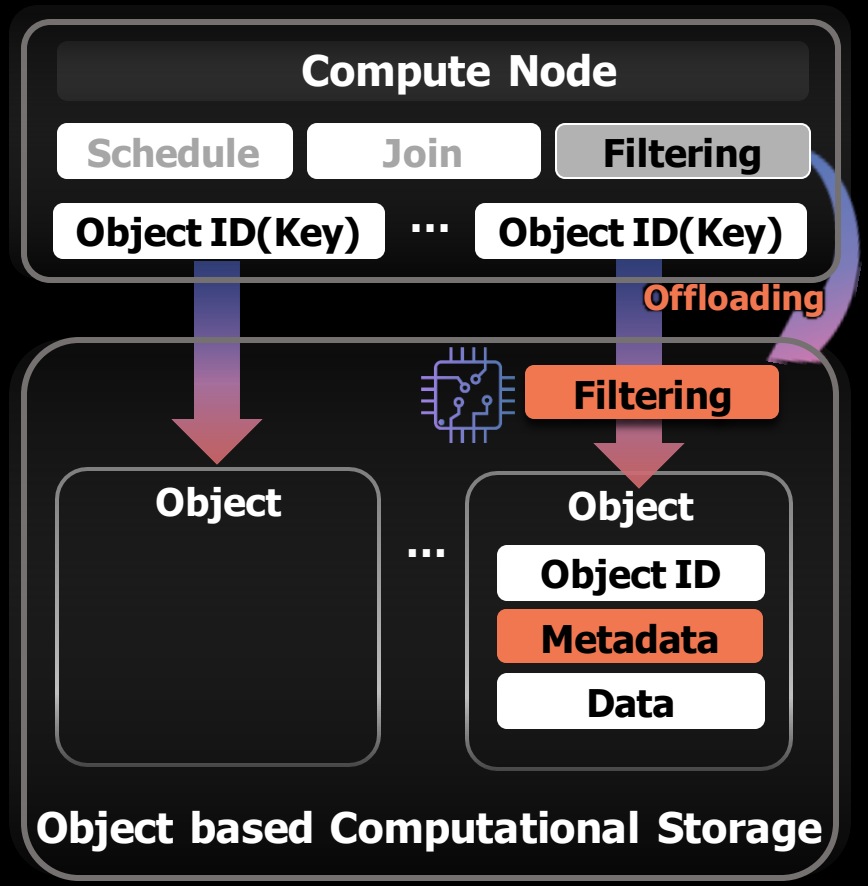

Enhanced Data Analytics

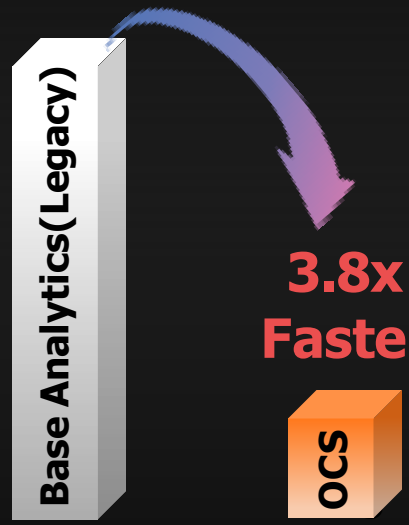

Execution Time [sec]

Saving Host Resource

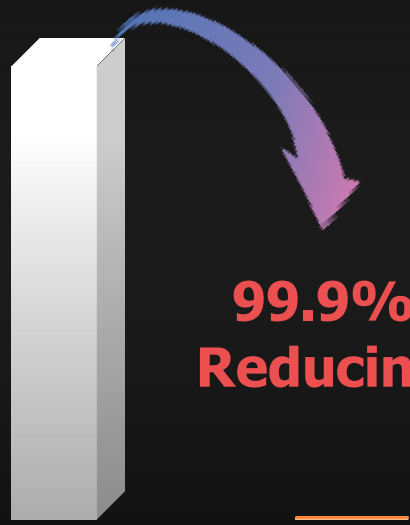

Data Movement [MB]

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Memory Centric

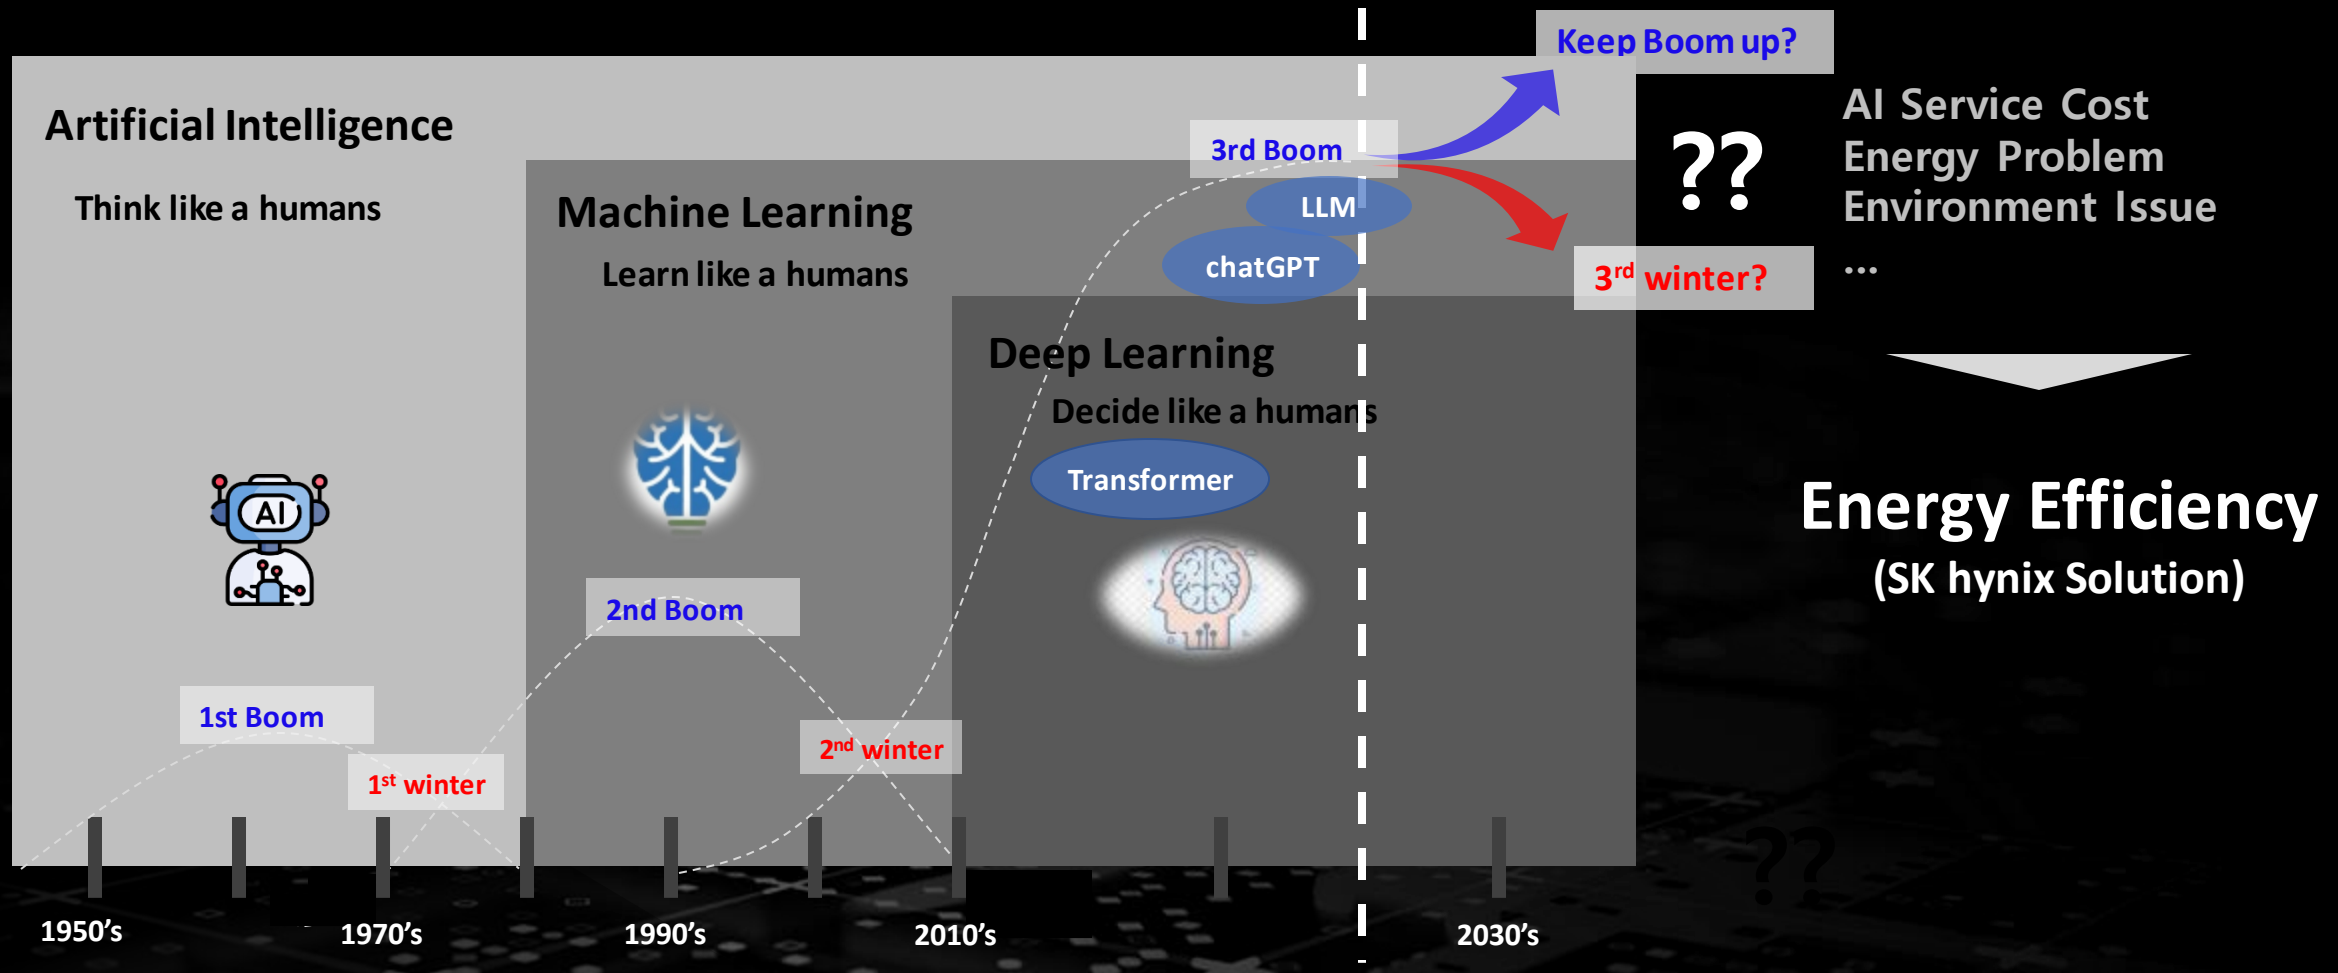

Energy Efficiency
(SK hynix Solution)

© SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

Total Memory & Storage Portfolio

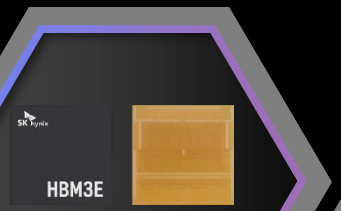

HBM3E

AiM

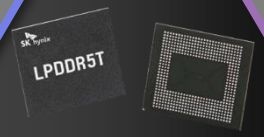

LPDDR5T

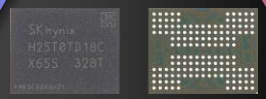

NAND

Mainstream cSSD

PCIe Gen5 16ch eSSD

Boot eSSD

CMM DDR5 CMM MDS

DDR5 RDIMM

Value TLC cSSD

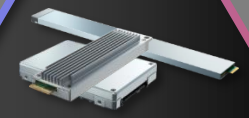

Value PCIe eSSD

AiMx

UFS / ZUFS / uMCP

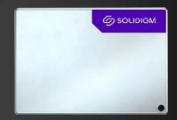

High Capacity QLC eSSD

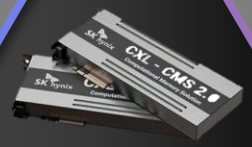

CXL CMS

MCRDIMM

Value QLC cSSD

SATA eSSD

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.

End of Document

SK hynix | © SK hynix Inc. This material is proprietary of SK hynix Inc. and subject to change without notice.