



SAMSUNG

Memory-Semantic SSD™

CXL-Era Memory Expansion with DRAM and NAND Flash Media


International Conference on Massive Storage Systems and Technology

Dr. Rekha Pitchumani

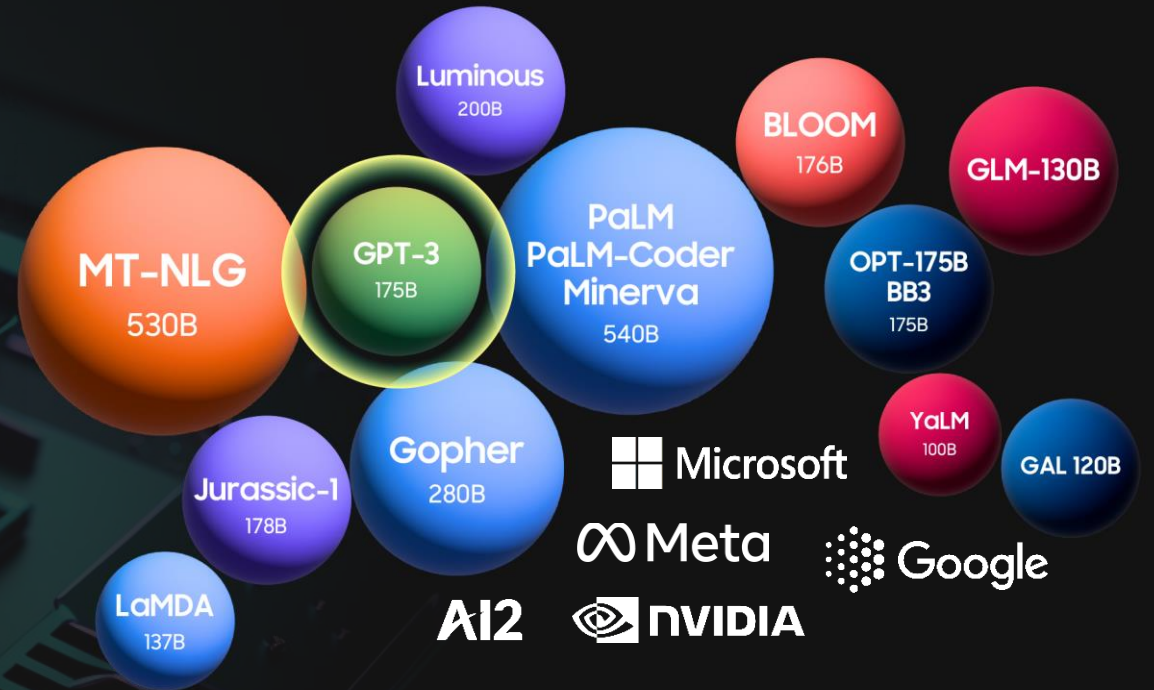
PhD, Senior Manager of Memory Solutions Lab.
Samsung Semiconductor Inc.

In the Era of AI & ML

Swift increase in demand for memory capacity and performance





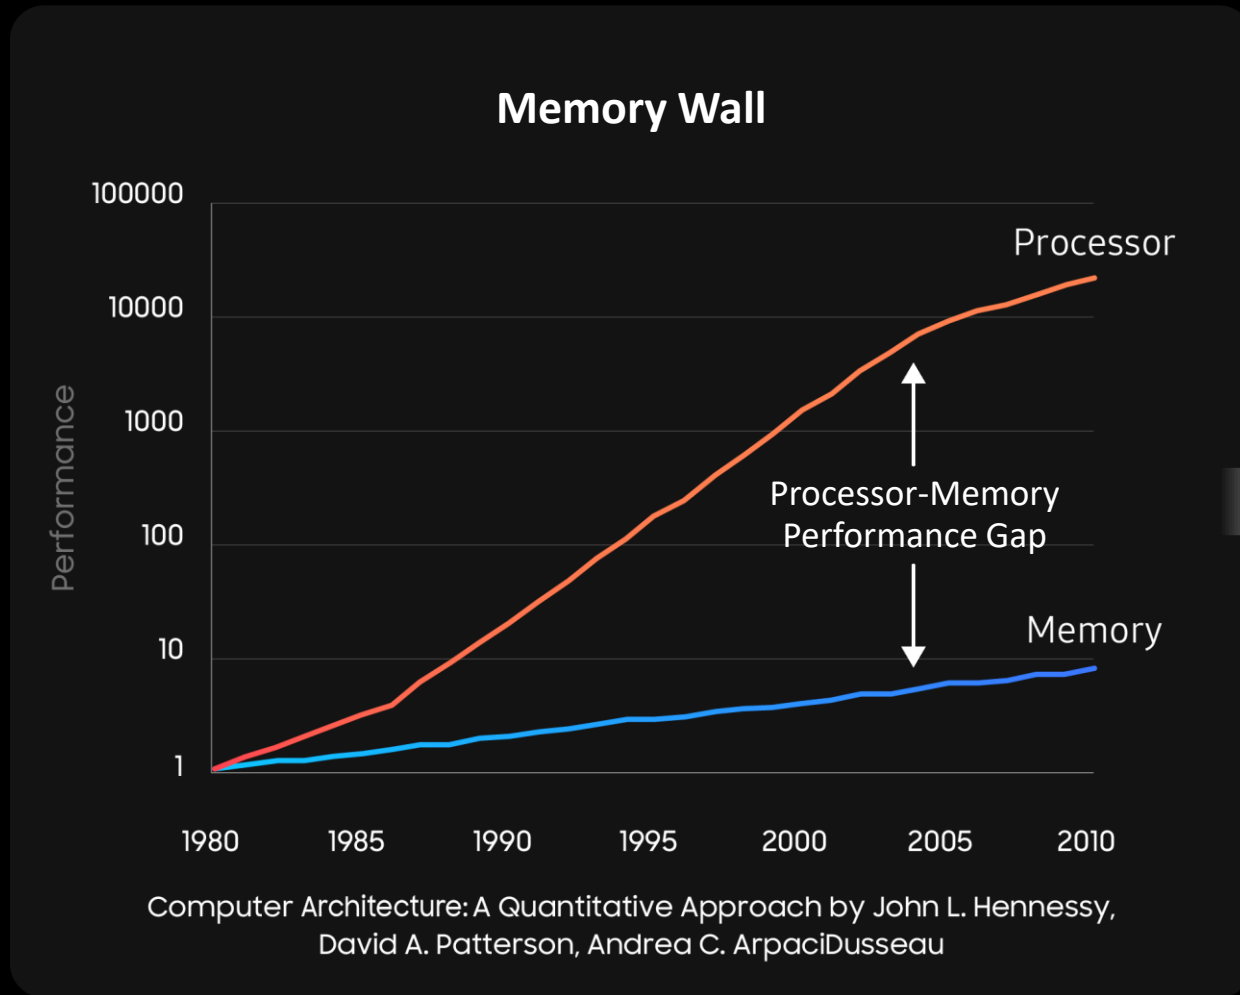
Language Model Sizes to Dec 2022



*Number of parameters/ Source: <https://lifaearchitect.ai/models>

Memory Wall

Challenge in memory to keep up with CPU performance



Source: <https://bitbashing.io/memory-performance.html>

CXL Memory Device Types

Memory Expander

CXL Type 3 device


CXL device with high bandwidth and low latency without a long tail



Tiered Memory Solution (Memory-Semantic SSD™)

CXL Type 3 device

CXL device with .mem and .io as active data path



Accelerator Attached Solution

CXL Type 2/3 device

Accelerator with CXL interface



Memory-Semantic SSD™

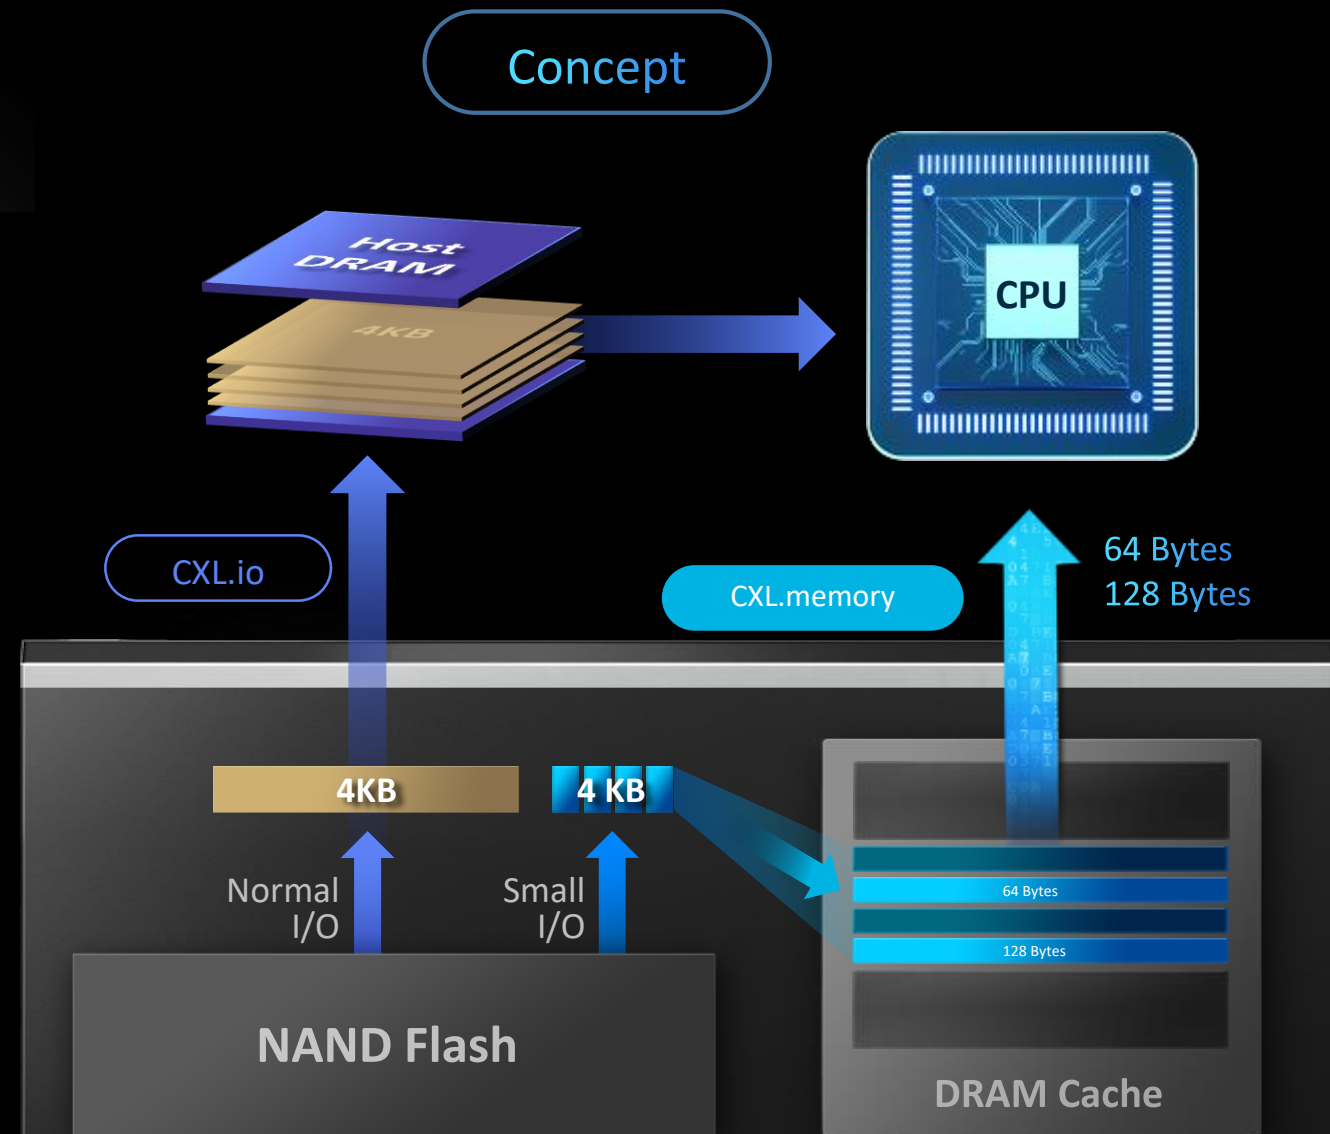
CXL-based SSD with built-in DRAM

Built-in DRAM

- Processing AI and ML applications, usually need relatively small-sized data chunks
- Applications can write data to the DRAM cache at DRAM speed

CXL Technology

- Low latency enabled by CXL.memory protocol

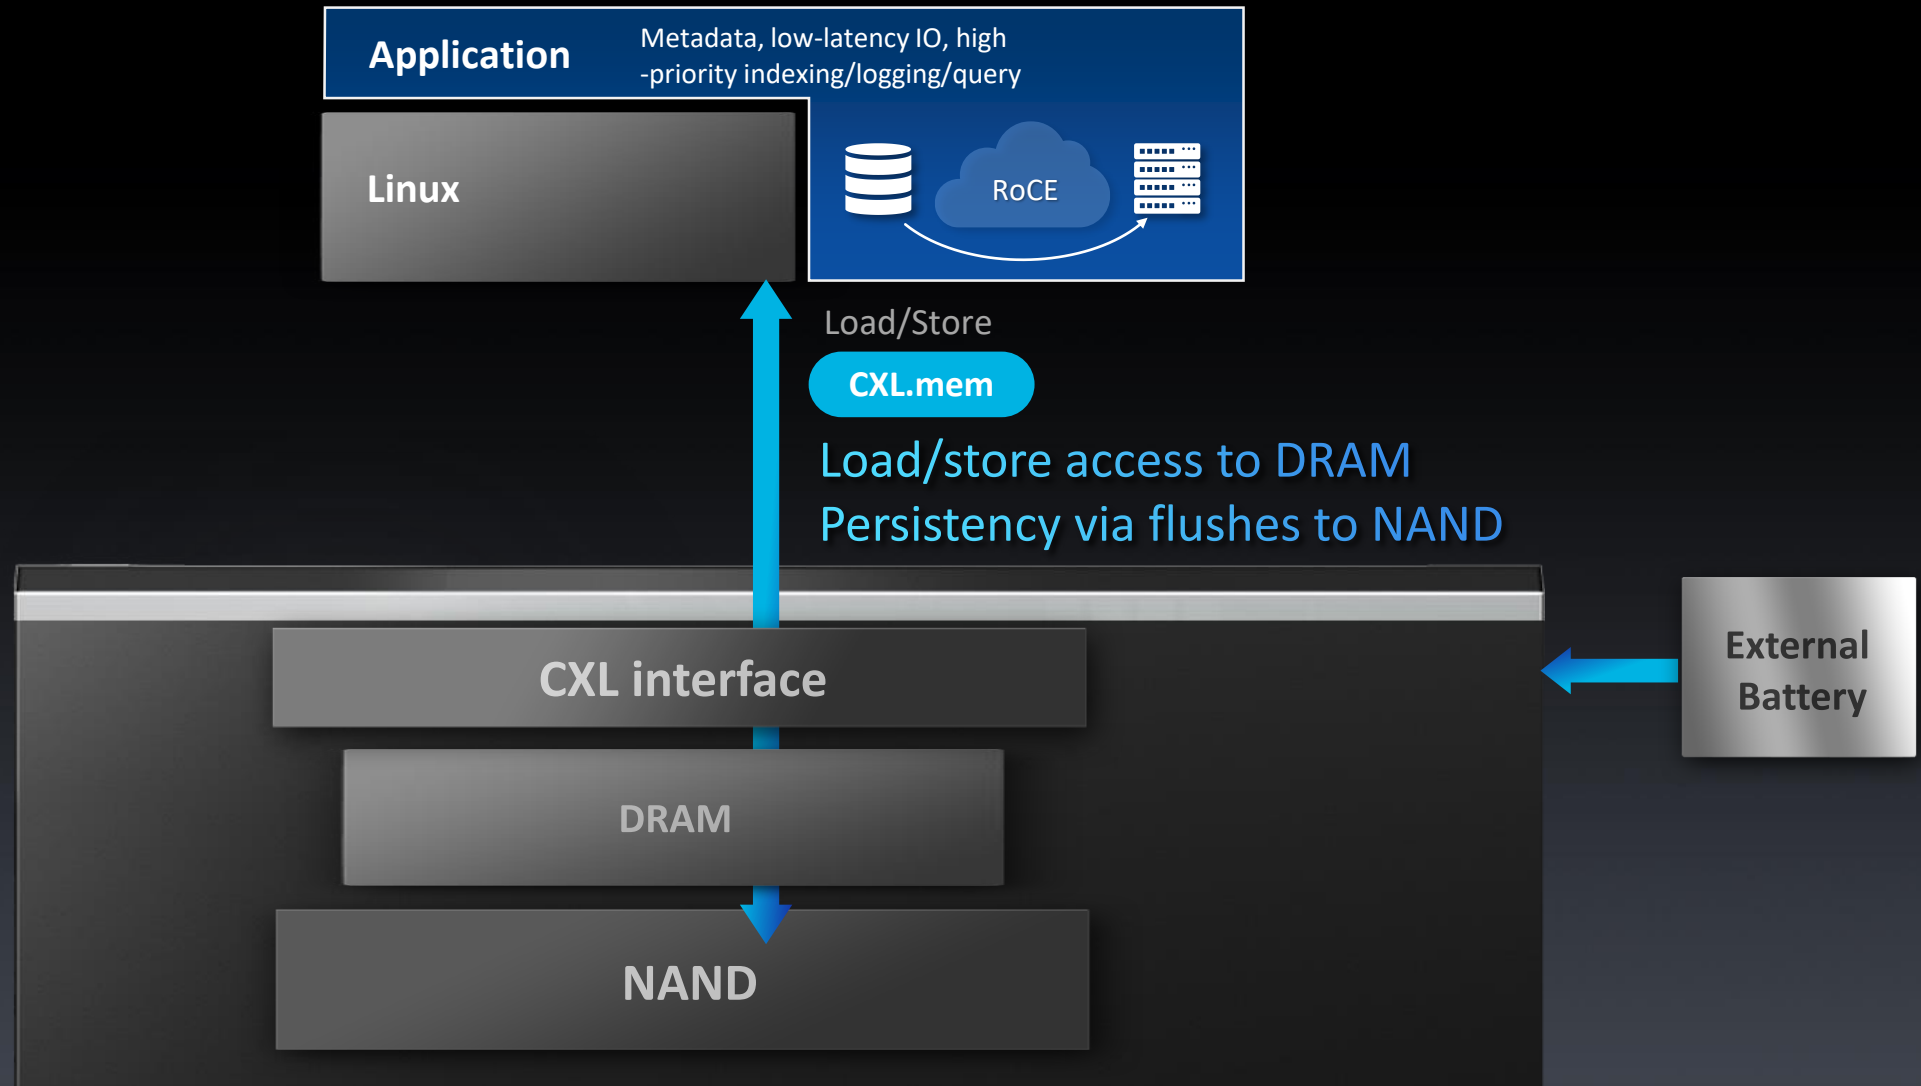


Storage Class Memory Options

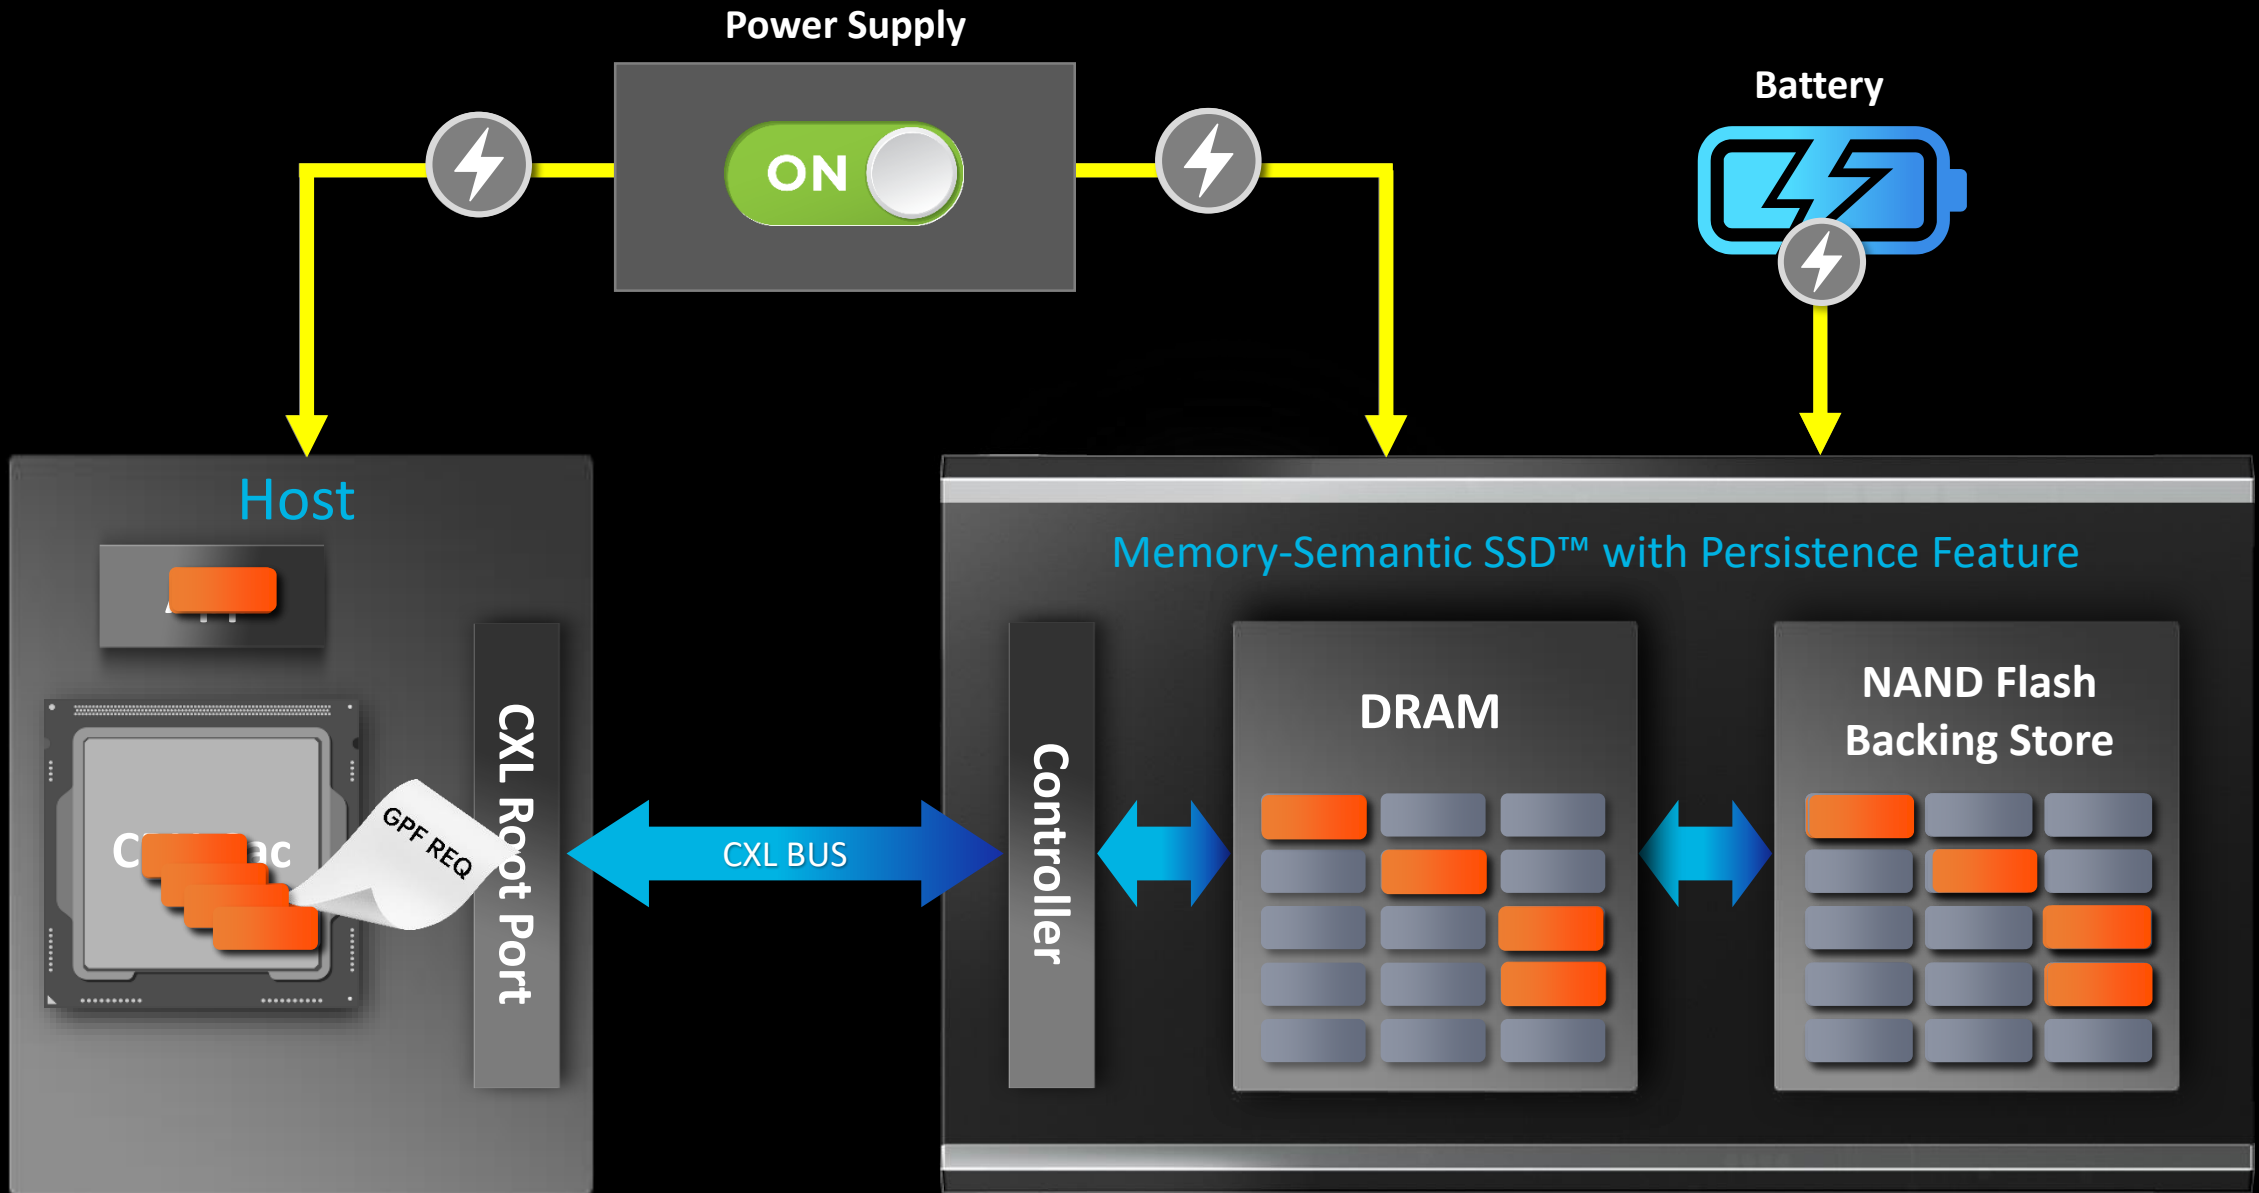
	Battery-Backed DRAM	3D-XPoint	NAND Flash
High Write Bandwidth	Best	Poor	Good
High Small IO Bandwidth	Best	Poor	Poor
Low Latency	Best	Better	Poor
Low Transient Writes Endurance Impact	Best	Good	Poor
Performance/\$	Better	Poor	Best

Persistent Memory Mode

NVDIMM like operation with external battery

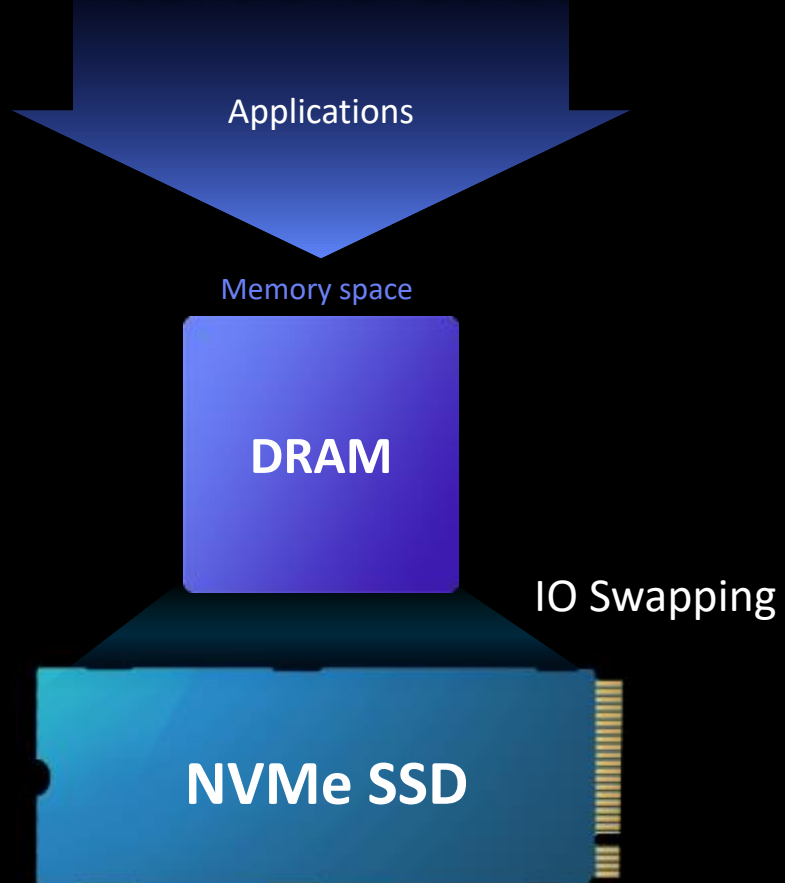


Persistent Memory Mode Illustration

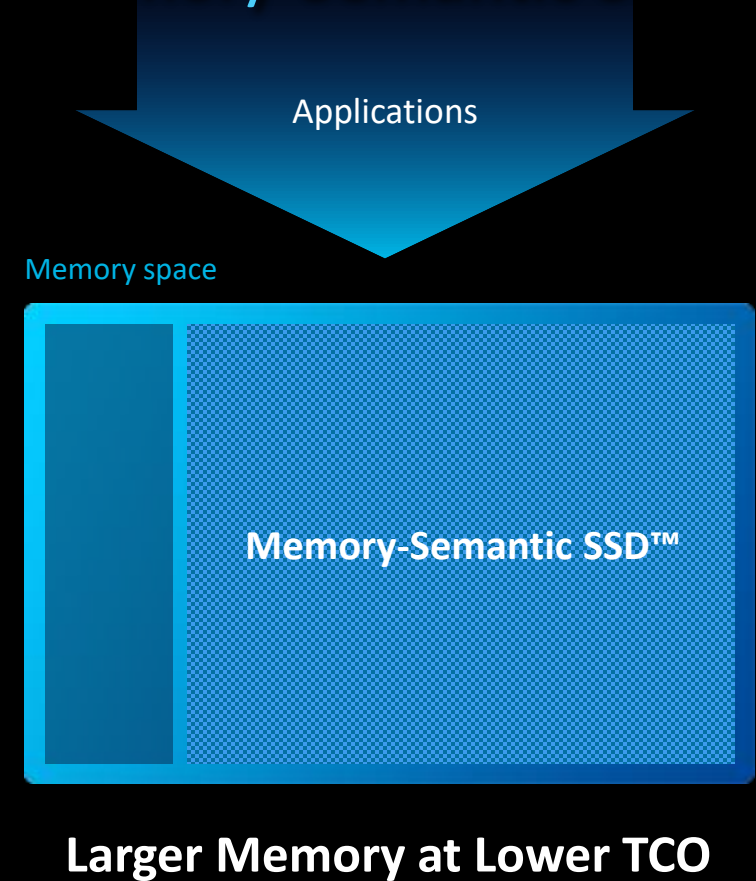


Lower System TCO Desired with Memory-Semantic SSD™

Conventional Solution

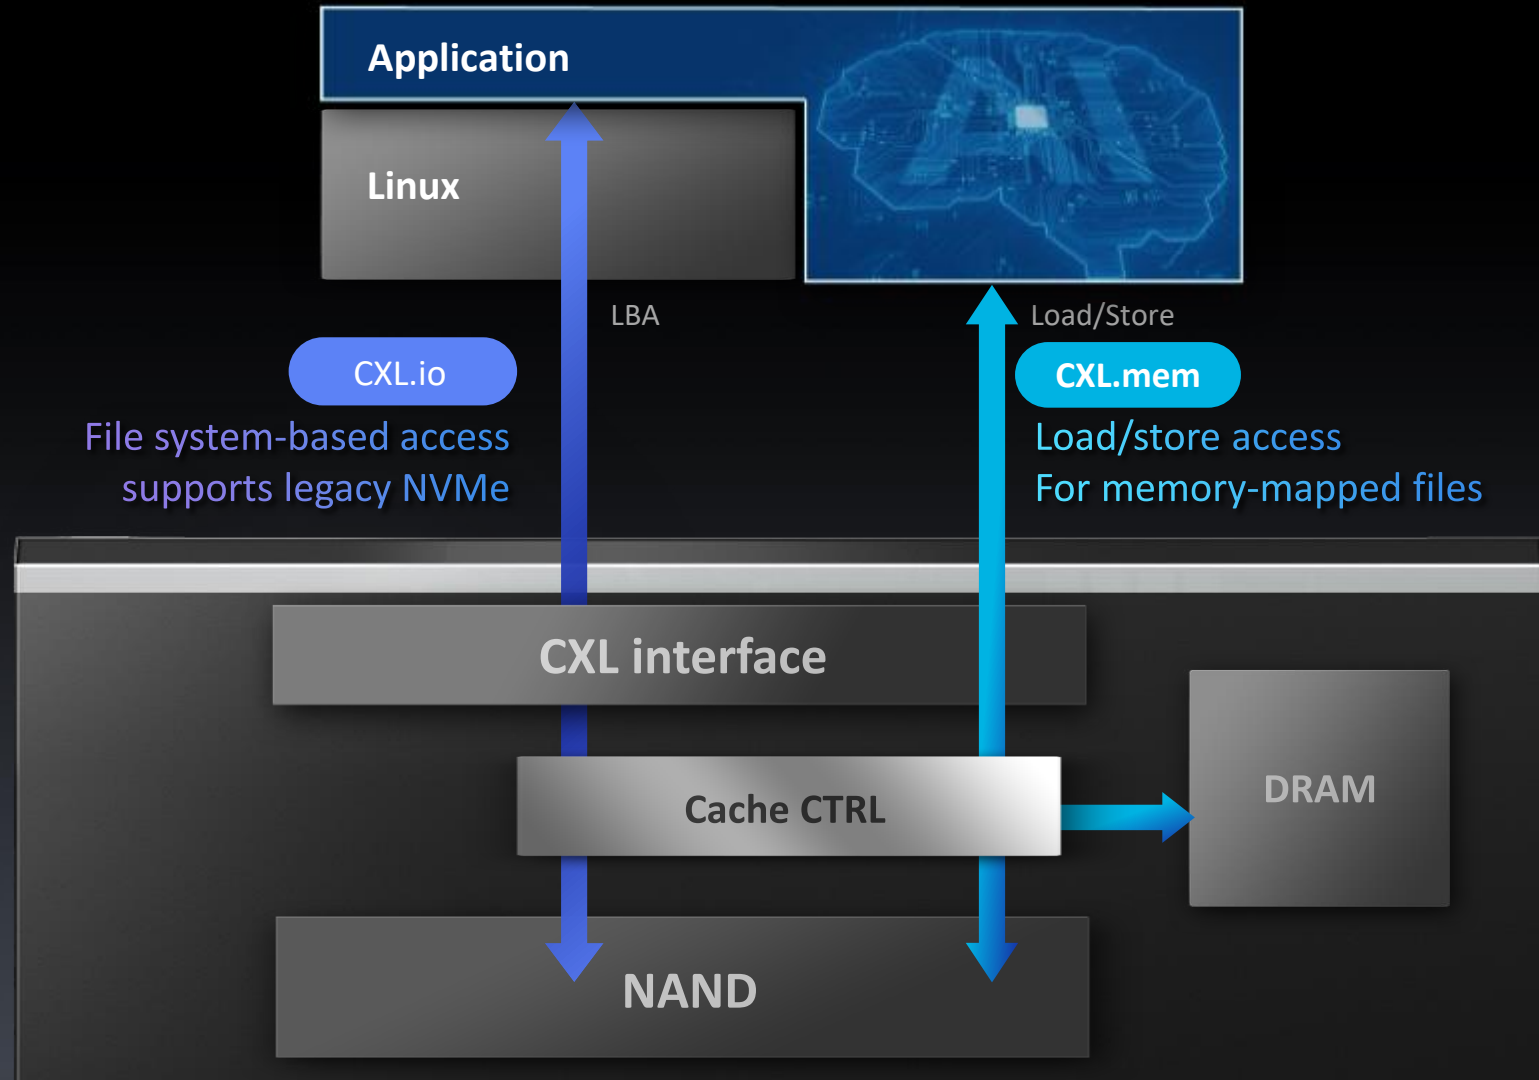


Memory-Semantic SSD™

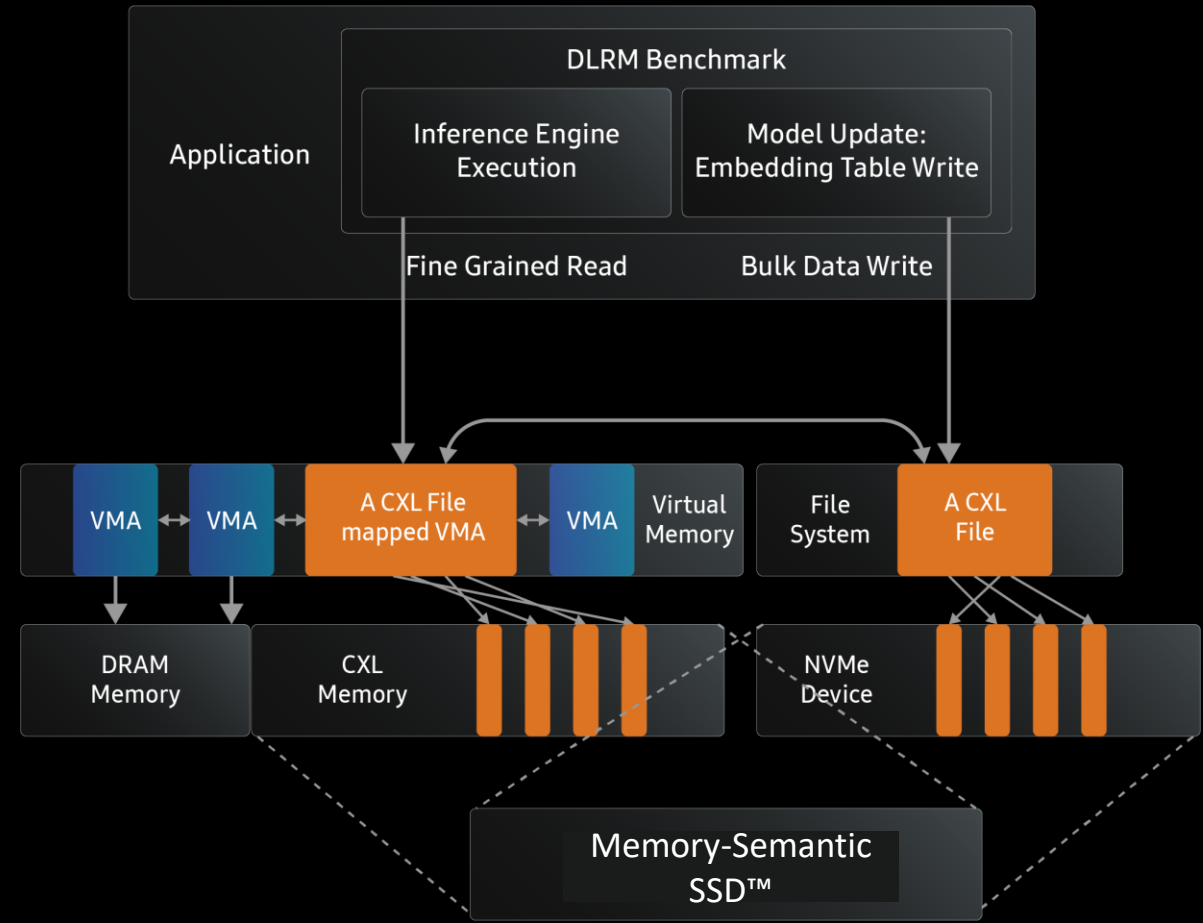
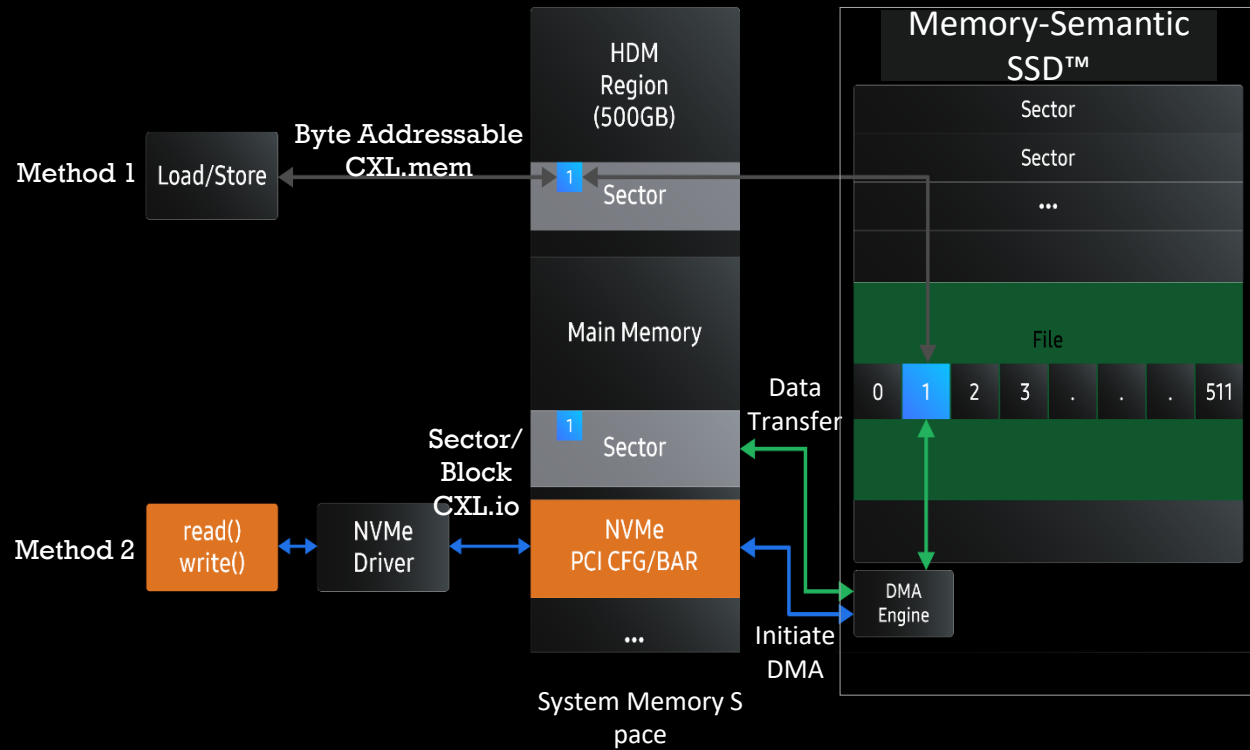


High Performance Dual Mode

NVMe IO mode and CXL memory mode



Dual Mode Device Access and Operation



Rack Level Battery Backup Options

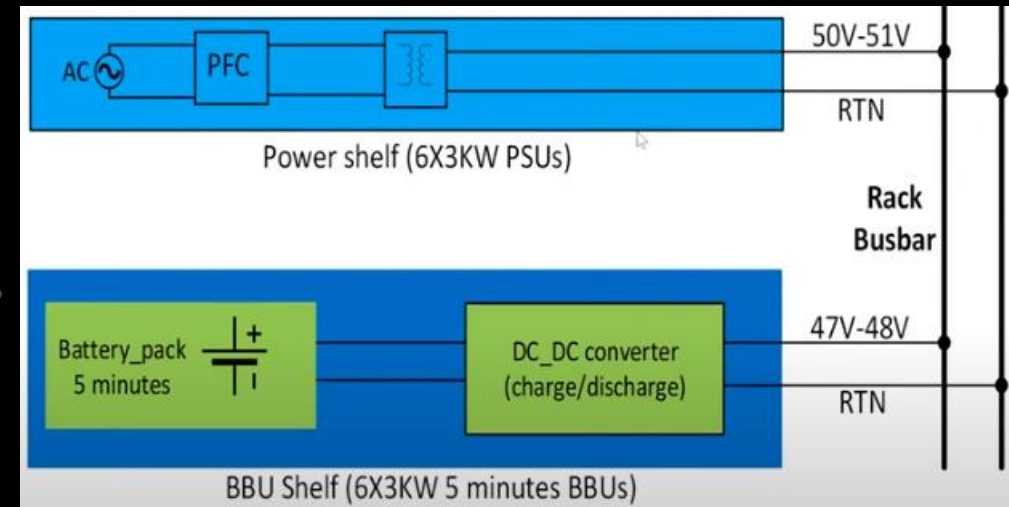

OCP Open Rack V3 Specs

Battery backup shelf for the rack

- BBU shelf shall be able to provide backup power with 6X 3KW BBU modules for a 5 mins backup time period

Server platform and CPU support

- CPU cache flush and CXL GPF support
- Supply backup power to device until flush completes



Source: [Open Rack/SpecsAndDesigns - OpenCompute](#)

Reducing Battery Backup Power Requirement


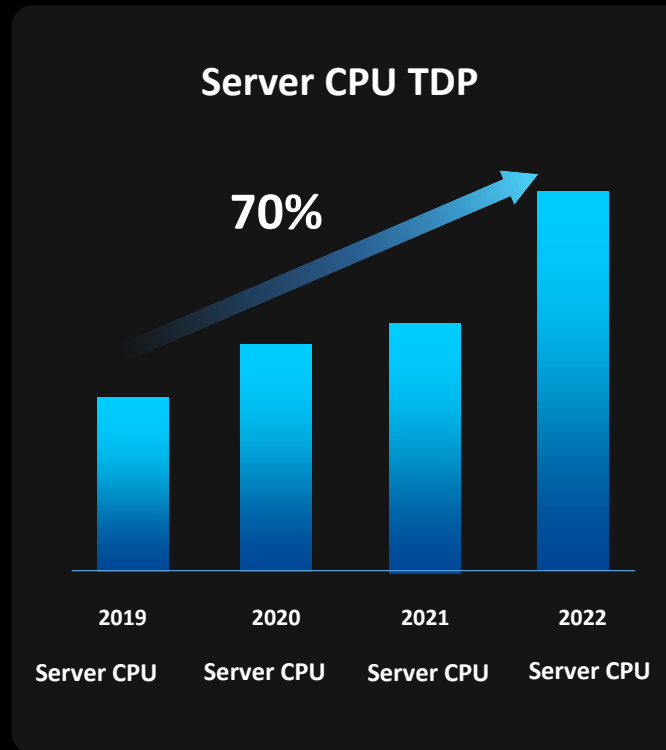
Focus on power backup to only CXL/ PCIe devices for extended time

CPU and Memory DIMMs

- Higher power requirements
- Short duration backup should be sufficient for compute cache flushes and GPF issue

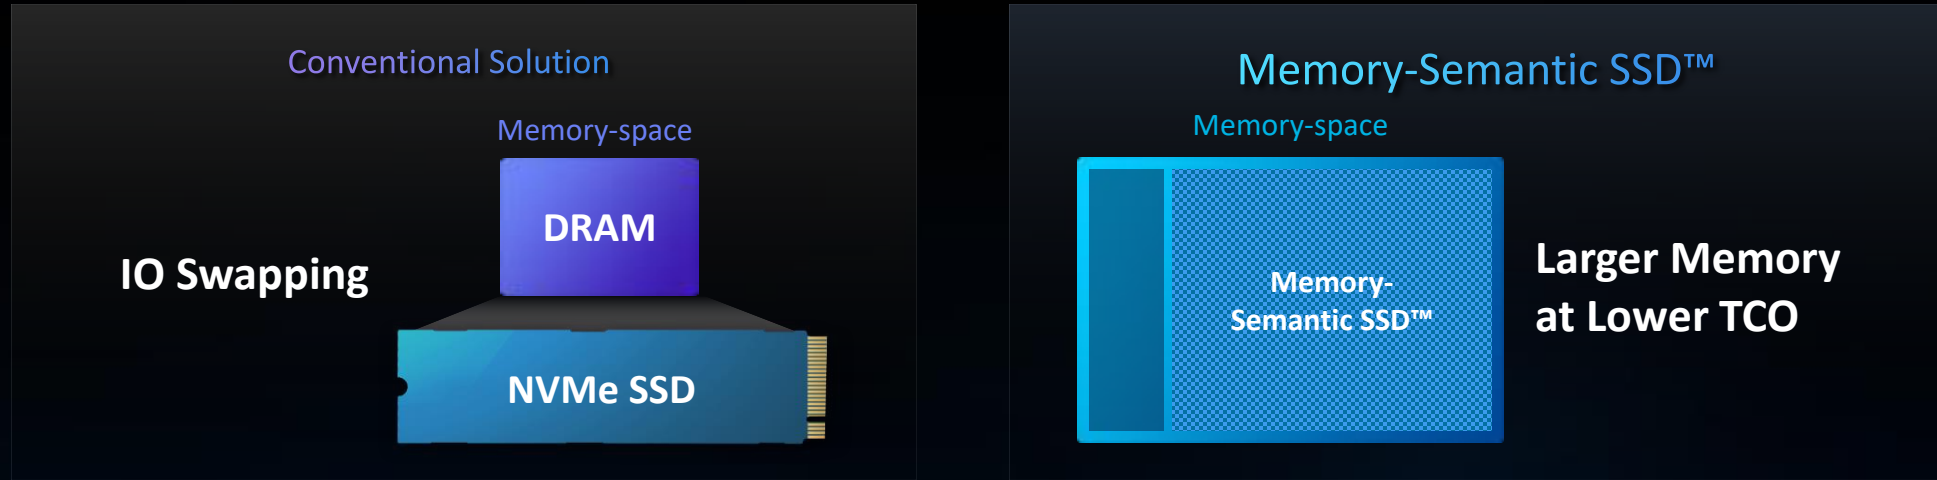
CXL/PCIe devices

- Lower power requirement for achieving persistency
- Since memory and storage devices occupy same slot, can cover both needs



Source: <https://m.samsungsem.com/global/newsroom/news/view.do?id=4893>

Lower System TCO with Memory-Semantic SSD™



- ✓ Tiered Memory Solution using DRAM and NAND with Small Cacheline Granularity Access
- ↑↑ Overall Better TCO with Larger Capacity Memory Expansion
- 🔄 Persistency Provided with NAND Flash Backup & External Battery Backup
- 📦 With CXL Standard, Memory-Storage Convergence Becomes a Reality
- 🔌 Paves Way for New Class of Higher Performance Devices with External Power Backup

Thank You

SAMSUNG

Disclaimer: Caution for Document License

THIS DOCUMENT AND ALL INFORMATION PROVIDED HEREIN (COLLECTIVELY, "INFORMATION") IS PROVIDED ON AN "AS IS" BASIS AND REMAINS THE SOLE AND EXCLUSIVE PROPERTY OF SAMSUNG ELECTRONICS CO., LTD. CUSTOMER MUST KEEP ALL INFORMATION IN STRICT CONFIDENCE AND TRUST, AND MUST NOT, DIRECTLY OR INDIRECTLY, IN ANY WAY, DISCLOSE, MAKE ACCESSIBALE, POST ON A WEBSITE, REVEAL, REPORT, PUBLISH, DISSEMINATE OR TRANSFER ANY INFORMATION TO ANY THIRD PARTY. CUSTOMER MUST NOT REPRODUCE OR COPY INFORMATION, WITHOUT SPECIFIC WRITTEN CONSENT FROM SAMSUNG. CUSTOMER MUST NOT USE, OR ALLOW USE OF, ANY INFORMATION IN ANY MANNER WHATSOEVER, EXCEPT FOR CUSTOMER'S INTERNAL EVALUATION PURPOSE. CUSTOMER MUST RESTRICT ACCESS TO INFORMATION TO THOSE OF ITS EMPLOYEES WHO HAVE A BONA FIDE NEED-TO-KNOW FOR SUCH PURPOSE AND ARE BOUND BY OBLIGATIONS AT LEAST AS RESTRICTIVE AS THIS CLAUSE. BY RECEIVING THIS DOCUMENT, IT IS UNDERSTOOD THAT CUSTOMER AGREES TO THE FOREGOING AND TO INDEMNIFY SAMSUNG FOR ANY FAILURE TO STRICTLY COMPLY THEREWITH. IF YOU DO NOT AGREE TO ANY PORTION OF THIS CLAUSE, PLEASE RETURN ALL INFORMATION AND ALL COPIES (IF ANY) WITHIN 24 HOURS OF RECEIPT THEREOF.