Open Emerging Technology Standards

Presentation to Australian Bureau of Meteorology
Annual R&D Workshop 2024
9th September 2024

Peter Kerney – Senior Solution Architect Manager
Intel Australia
peter.kerney@intel.com

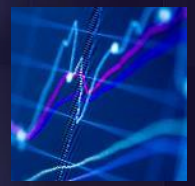

Agenda

01 Intel today

02 OneAPI Industry Momentum

03 Unified Acceleration Foundation (UXL)

04 Compute Express Link (CXL)

Our Priorities

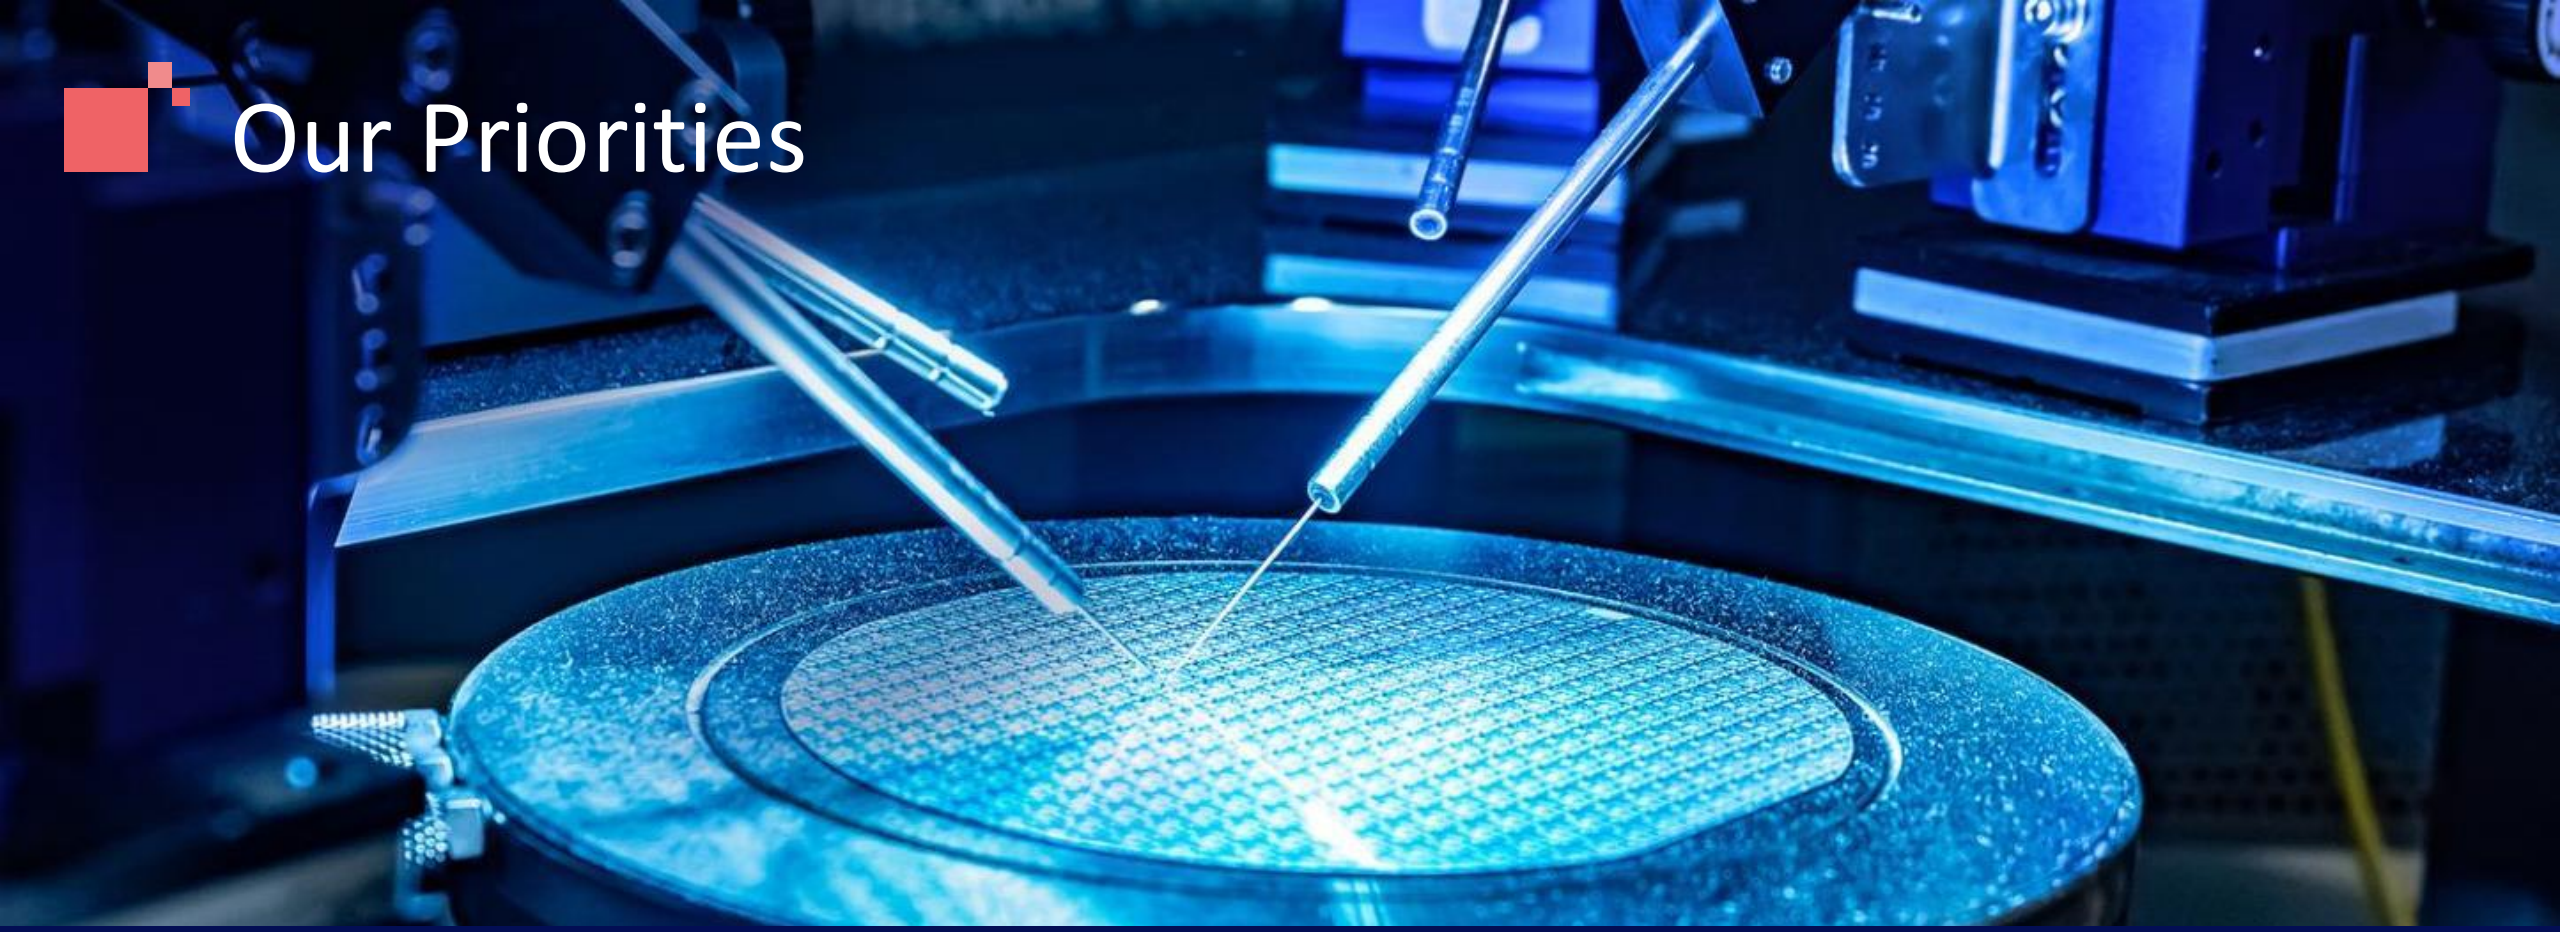

Product Leadership

We lead and democratize compute with Intel x86 and xPU.

Manufacturing at Scale

Our IDM 2.0 strategy combines three capabilities; our internal factory network, strategic use of foundry capacity, and Intel Foundry.

Open Platforms

We deliver open software and hardware platforms with industry-leading standards.

Our People

Our world-class talent is at the heart of everything we do. Together, we strive to have a positive effect on business, society and the planet.

oneAPI Industry Momentum

End Users

Accrad Accelerated Radiology
@rchanan
FOUNDRY MODO
BioDataAnalysis GmbH GeoEast
BENTLEY
CINESITE
LAIKA
GE
kt
AUTODESK ARNOLD
中国石油集团东方地球物理勘探有限责任公司 BGP INC.,CHINA NATIONAL PETROLEUM CORPORATION
iOmniscient Autonomous
EURECOM
PHILIPS
Brightskies
大势智慧 DASPATIAL
WeBank
SAMSUNG MEDISON
TANGENTANIMATION
allegro.ai
ILLUMINATION MACGUFF
Verizon
MediaKind

National Labs

Argonne NATIONAL LABORATORY
UT-BATTELLE Oak Ridge National Laboratory
CERN openlab
CINECA
Peraton Labs
Laboratório Nacional de Computação Científica
Leibniz Supercomputing Centre of the Bavarian Academy of Sciences and Humanities
SANKHYA SUTRA

ISVs & OSVs

codeplay
AI SINGAPORE
CANONICAL
CERTIFACE
ANACONDA.
Ansys
MAX PLANCK COMPUTING & DATA FACILITY
MPCDF
sas
Red Hat OpenShift Data Science
KATANA GRAPH
CGG
AIBLE
CHAOSGROUP
MAXON
Ospirent
FOUNDRY
mercenaries engineering Guerilla
Hisense Medical 海信医疗
SENAI CIMATEC FIEB
SAP
E4 COMPUTER ENGINEERING
YUAN
GIGASPACES
MEGH COMPUTING
KFBIO
Tech Mahindra
AsiaInfo 亚信科技
SUSE
vmware
UNITED IMAGING

OEMs & SIs

Bittware a molex company
Hewlett Packard Enterprise
DELL Technologies
Atos
Lenovo
HCL
RENIAC
MEGWARE

Universities & Research Institutes

LOBACHEVSKY UNIVERSITY
TECHNION Israel Institute of Technology
UNIVERSITY OF CAMBRIDGE
ICT 中国科学院计算技术研究所 INSTITUTE OF COMPUTING TECHNOLOGY, CHINESE ACADEMY OF SCIENCES
Ben-Gurion University of the Negev
Berkeley
UCDAVIS
THE UNIVERSITY OF TENNESSEE KNOXVILLE
UNIVERSIDAD COMPLUTENSE MADRID
OLD DOMINION UNIVERSITY
ctg
北京大学软件与微电子学院 School of Software & Microelectronics
THE UNIVERSITY OF UTAH
ILLINOIS
PURDUE UNIVERSITY Elmore Family School of Electrical and Computer Engineering
UNIVERSITY OF OREGON
TÉCNICO LISBOA
inesc id lisboa
ZIB
FACULTY OF MATHEMATICS AND PHYSICS Charles University
University of Stuttgart Germany
IIT GOA
Indian Institutes of Technology Delhi / Kharagpur / Roorkee
University College London
Durham University
SDSC SAN DIEGO SUPERCOMPUTER CENTER
CDAC
UKRI Science and Technology Facilities Council Scientific Computing
ATLAS EXPERIMENT
TACC
Indian Institute of Science Bangalore
IISER PUNE
Indian Institute of Science Education & Research Pune
HEIDELBERG UNIVERSITY COMPUTING CENTRE
Northern Illinois University
University of BRISTOL
Stockholm University

CSPs & Frameworks

Microsoft Azure
Alibaba Cloud
Google Cloud
TensorFlow
Taboola
Tencent 腾讯
DataRobot
Baidu 百度
飞桨 PaddlePaddle
QCT
NAVER CLOVA
PyTorch

intel

These organizations support the oneAPI initiative for a single, unified programming model for cross-architecture development.
It does not indicate any agreement to purchase or use of Intel's products. *Other names and brands may be claimed as the property of others.

Commitment to Open, Scalable Acceleration

Freeing the Developer Ecosystem from the Chains of Proprietary Software

UXL
Unified Acceleration Foundation

Linux Foundation governed open industry foundation driving a vendor-neutral software ecosystem for multiarchitecture accelerated computing

Founding Members: ARM, Fujitsu, Google Cloud, Imagination Tech, Intel, Qualcomm, Samsung, VMware

oneAPI Specification (oneAPI.io) is the starting point

oneAPI

Open, Standards-based, Multiarchitecture Programming

Performance | Productivity | Freedom from Vendor Lock-In

CPU

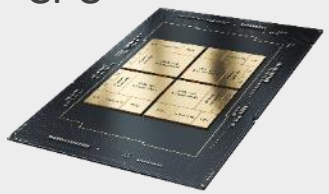

GPU

FPGA

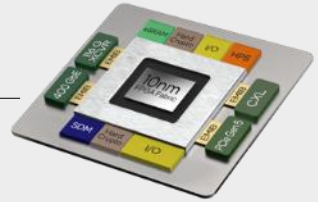

Other Accelerators

Unified Acceleration Foundation (UXL)

Building a multiarchitecture, multivendor accelerator software ecosystem

UXL
Unified Acceleration Foundation

- Governance: Linux Foundation's Joint Development Foundation
- Mission: Unify the heterogeneous compute ecosystem around open standards
- Starting point: oneAPI Specification (oneAPI.io)
- Goal: broad-based industry participation and contributions
- SIGs: AI, Hardware, Language, Math, Safety Critical
- Join Us: Participate in SIGs
  - www.UXLFoundation.org

Steering Members

arm FUJITSU Google Cloud Imagination
intel Qualcomm SAMSUNG vmware

UXL Foundation Structure

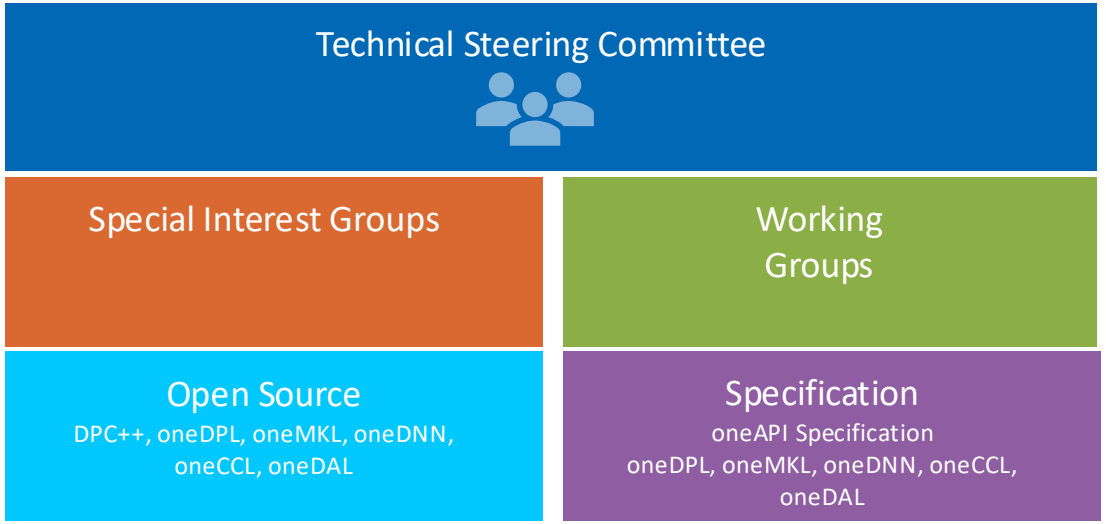

Compute Express Link

A high-speed interconnect offering coherency and memory semantics, CXL uses high-bandwidth, low-latency connectivity between the host processor and devices such as accelerators, memory buffers, and smart I/O devices.

Memory Density and Costs are not Scaling to Meet Data Center Workloads and Infrastructure Cost Requirements

Compute Performance Growth Is Accelerating[1]

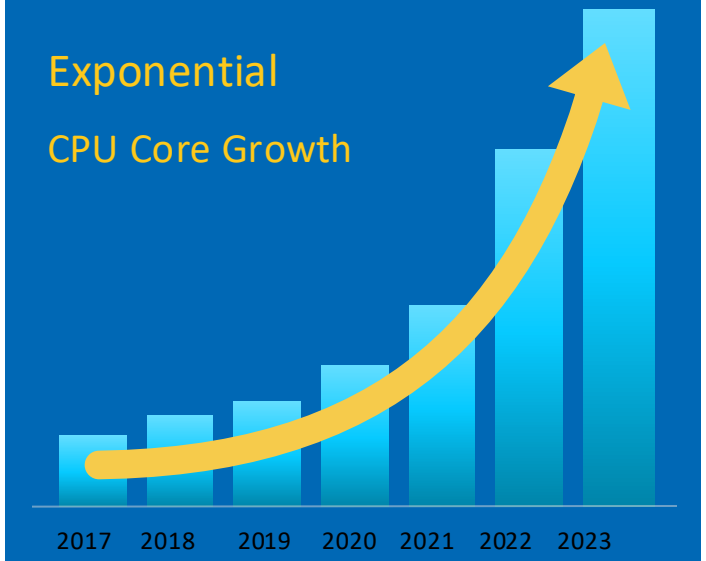

CPU Core Growth Projection Over Time[1]

Scaling of DRAM Density Is Slowing[2]

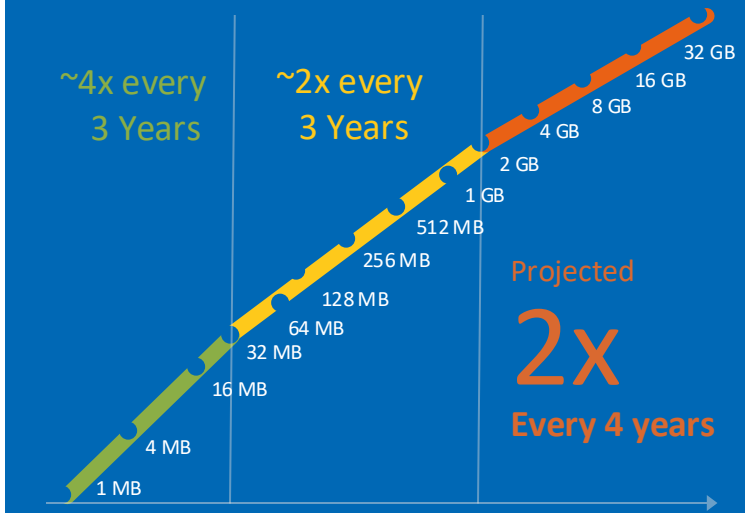

DRAM Density Over Time[2]

Memory Is a Growing Percent of Server Costs[3]

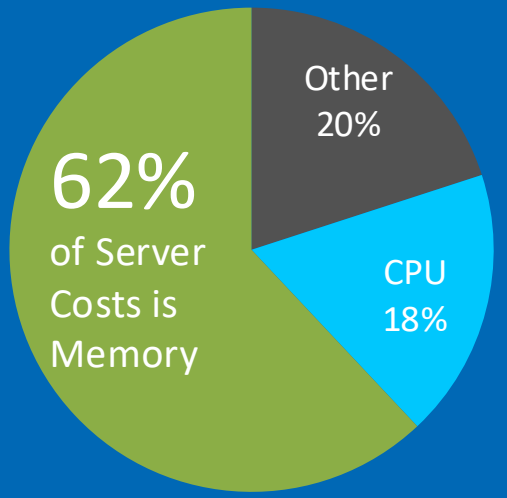

Current cost percentage of server memory compared to other components

1 Source: Intel. Results may vary.
2 Source: https://flashmemorysummit.com/English/Collaterals/Proceedings/2018/20180809_NEWM-301A-1_Gervasi.pdf
3 Source: Intel Internal. Estimates – based on large scale deployments and does not include software costs.

Intel's History of Developing I/O Ecosystems

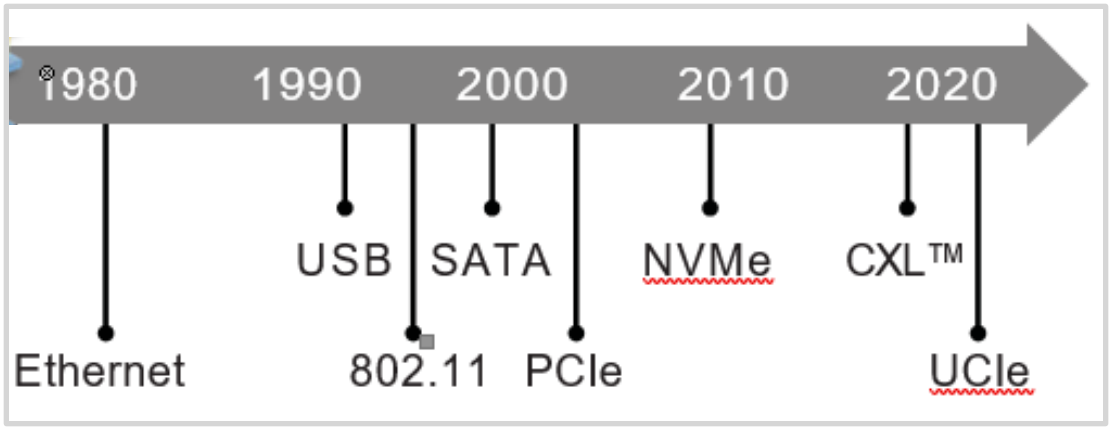

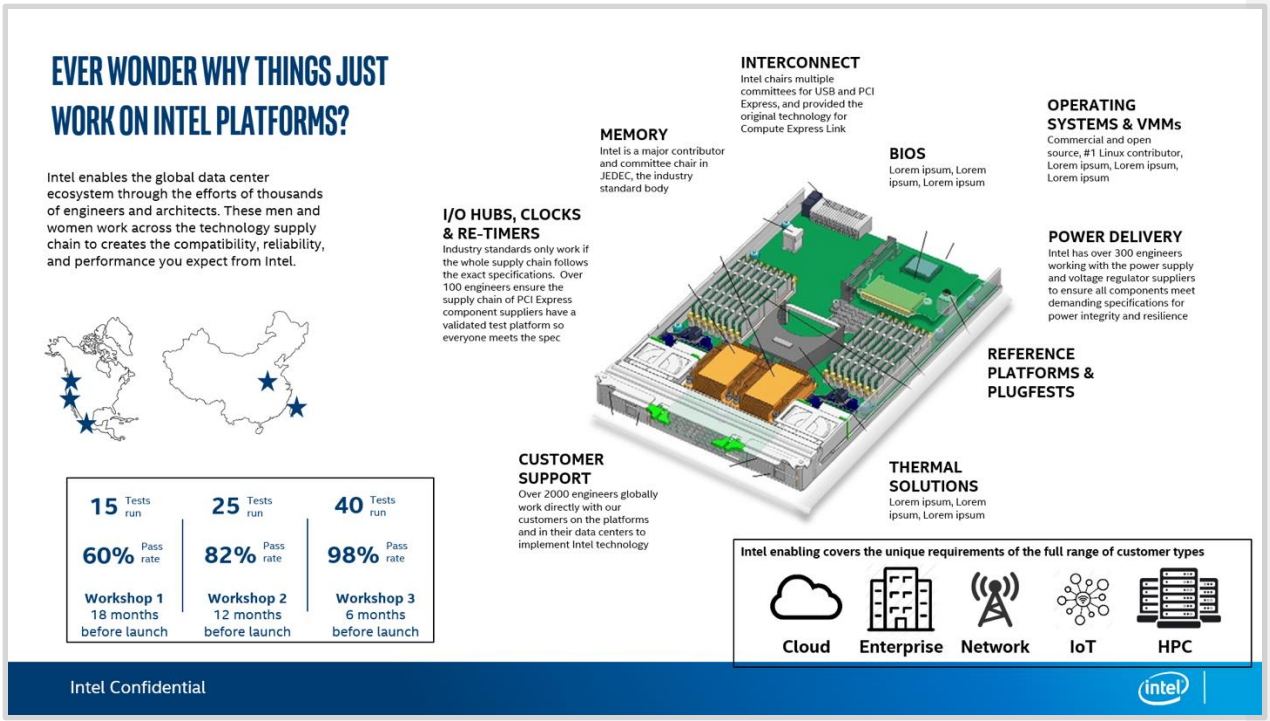

CXL is no different

Industry Focal Point

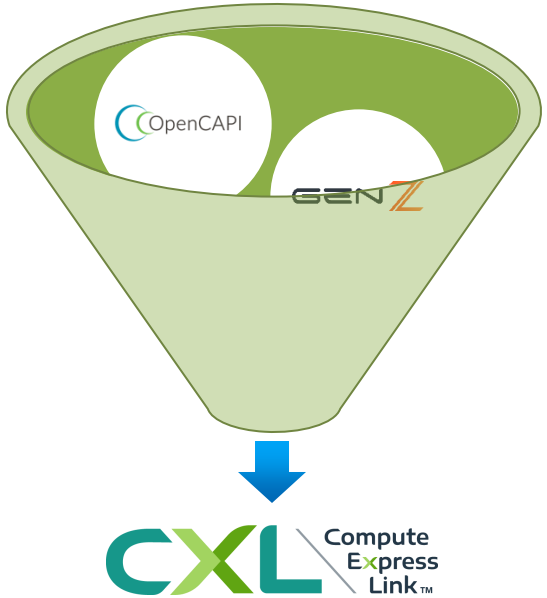

Alibaba Group · AMD · arm · cisco · DELL EMC · Google · Hewlett Packard Enterprise · HUAWEI · IBM · intel · Meta · Micron · Microsoft · NVIDIA · SAMSUNG

CXL Compute Express Link™

Consortium and OpenCAPI sign letter of intent
to transfer OpenCAPI specification and assets to the CXL Consortium

OpenCAPI CXL Compute Express Link™

August 2, 2022, Flash Memory Summit
CXL Consortium and OpenCAPI Consortium Sign Letter of Intent to Transfer OpenCAPI Assets to CXL

- In February 2022, CXL Consortium and Gen-Z Consortium signed agreement to transfer Gen-Z specification and assets to CXL Consortium

Industry Open Standard for High Speed Communications

250+
Member Companies

CXL on Motherboards: One slot for PCIe OR CXL

- **Motherboards featuring 'FlexBus' slots**
  - Starting with SPR generation CPUs
  - Flexible port configured for PCIe or CXL during link-up
  - Up to x16 lanes @32 Gb/s => 64 GB/s full duplex BW

Intel Archer City reference design w/ FlexBus slots

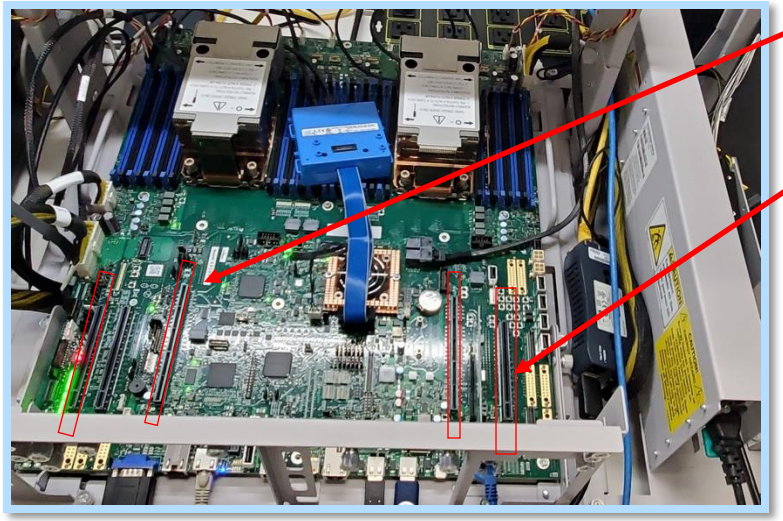

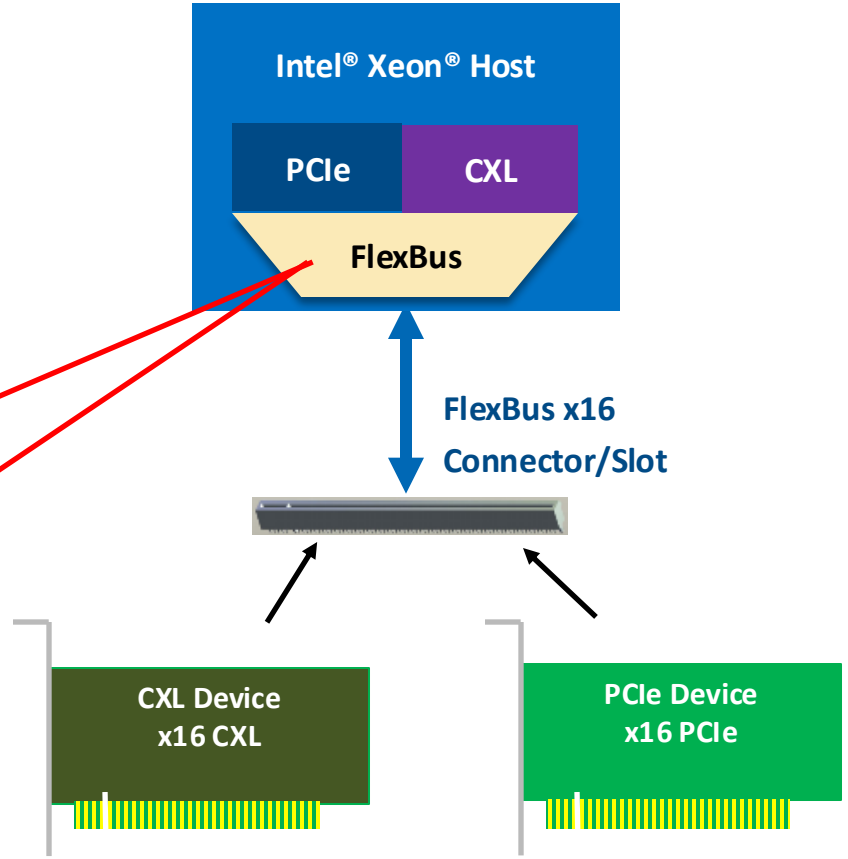

Representative CXL Usages

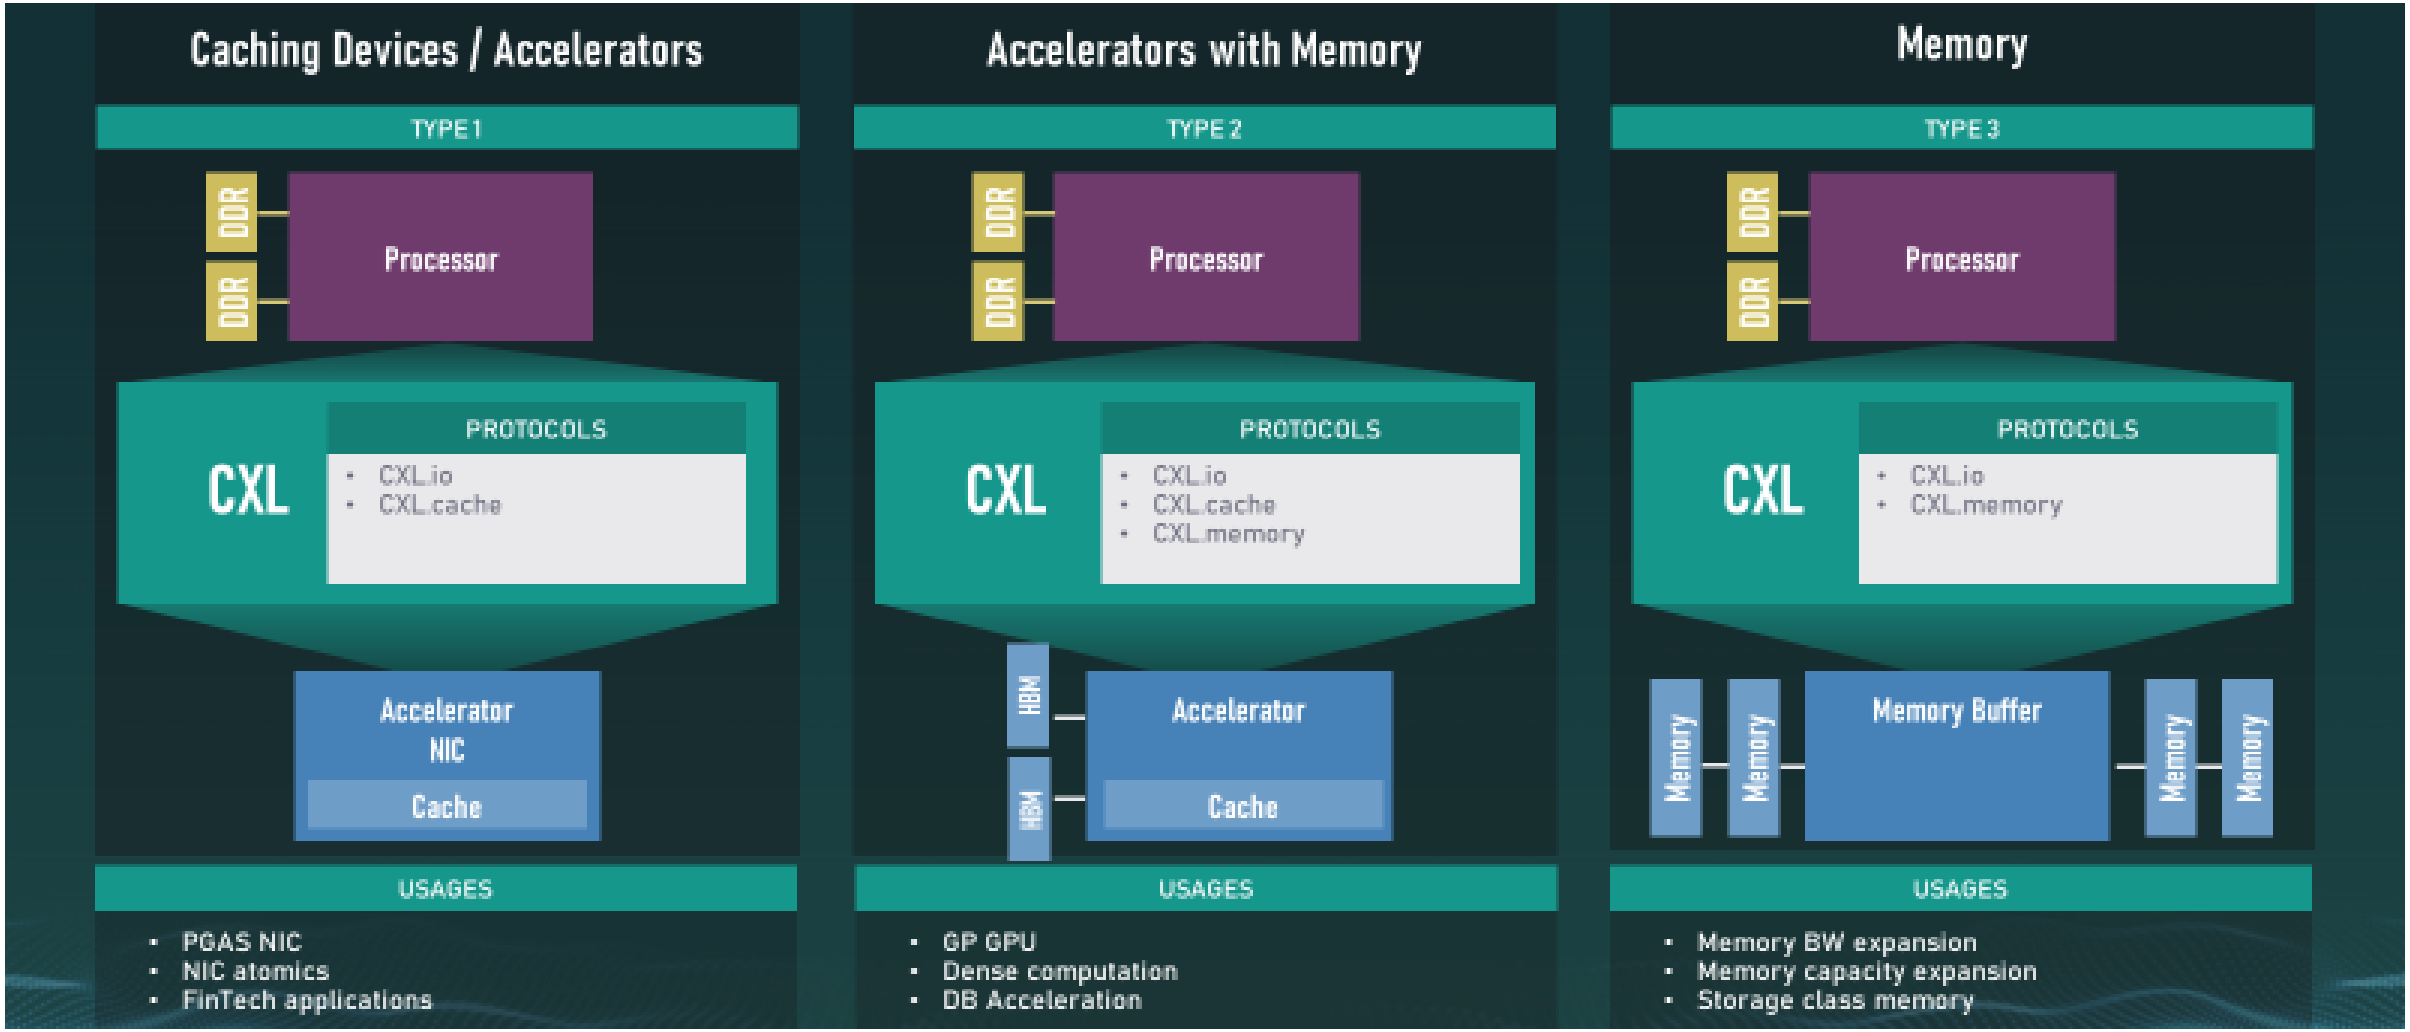

Evolution of CXL

CXL 3.X — "Scale out"

- Composable Fabric growth for disaggregation / pooling / accelerator

CXL 2.0 — "Scale up"

- Multiple nodes inside a Rack/ Chassis supporting pooling of resources
- Memory/accelerator pooling with single logical devices (SLD)
- Memory pooling with multiple logical devices (MLD)

CXL 1.1 — "Single server"

- Single Node Coherent interconnect

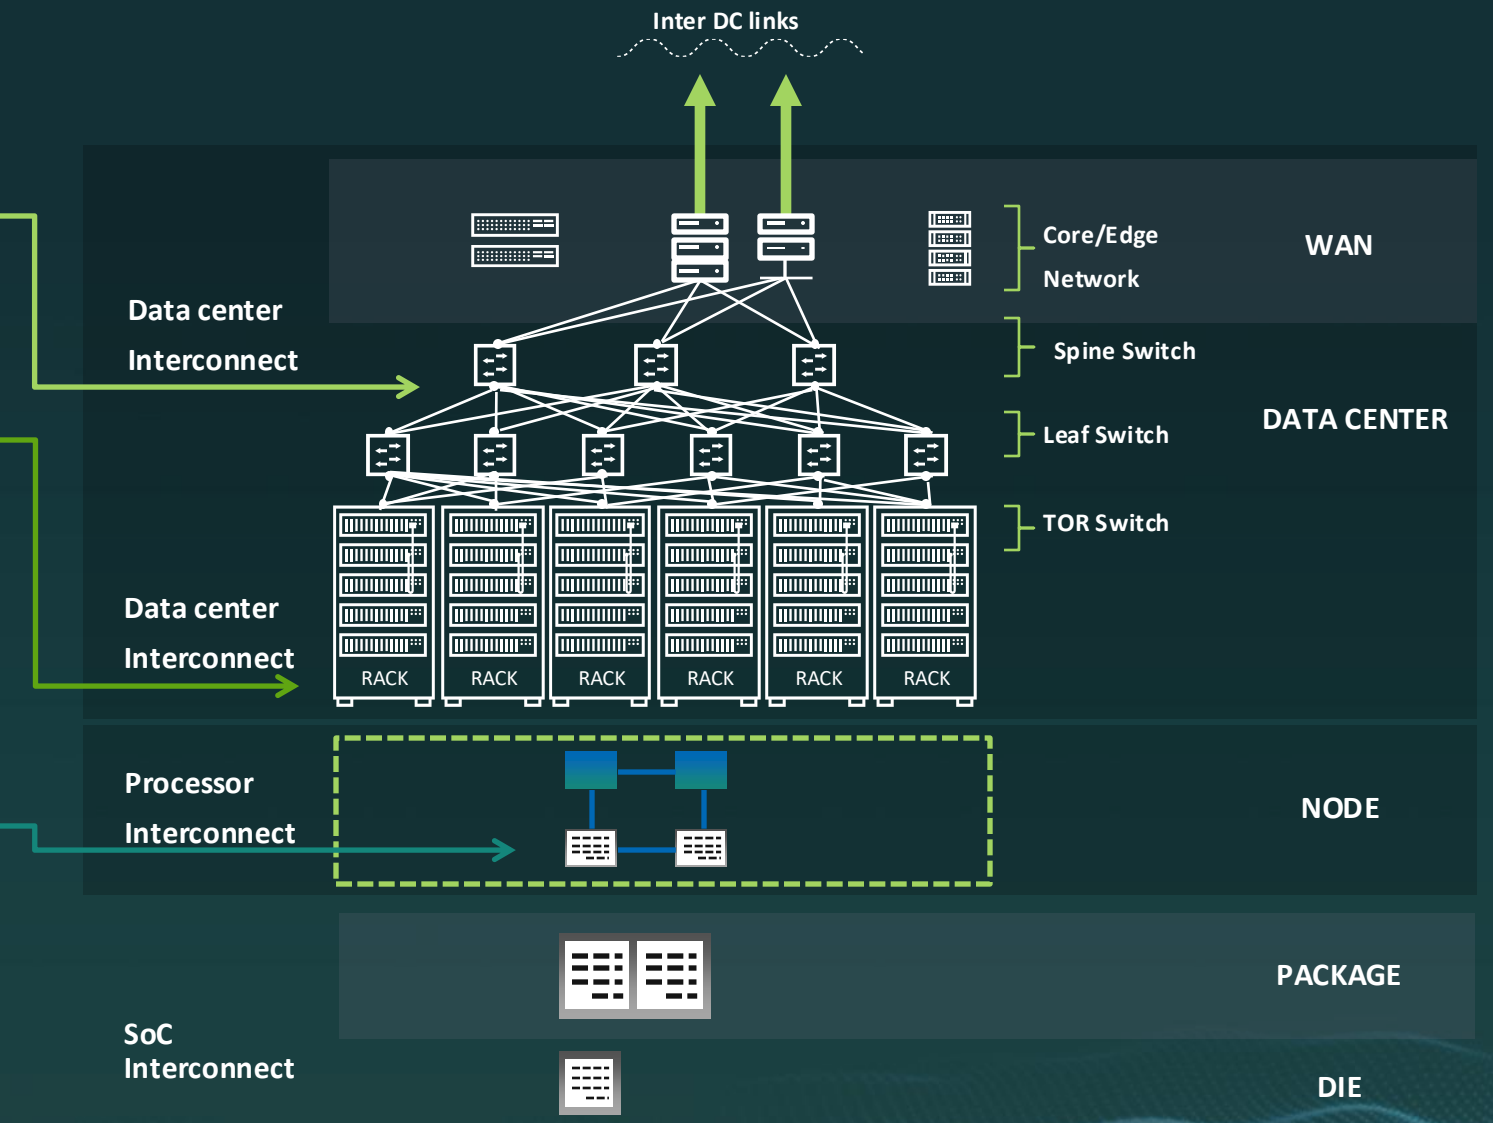

Memory Pooling vs Sharing – roadmap feature

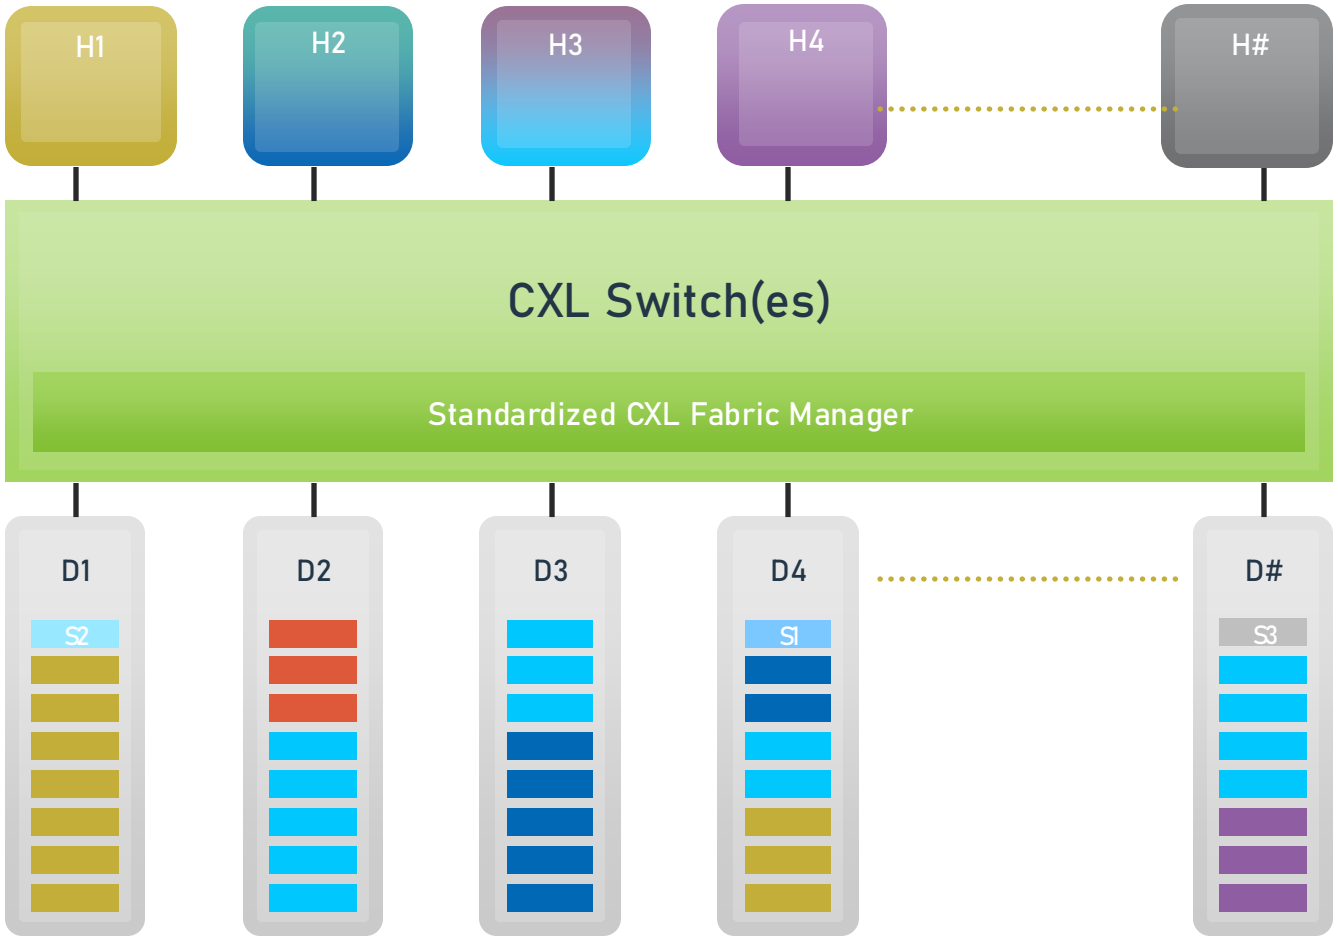

- Memory Pooling (CXL 2.0) allows a host or multiple hosts to expand their memory capacity and bandwidth by using unique areas in a memory device with allocation of Memory segments tied to a unique host at any given time
- Memory Sharing (CXL 3.0) allows multiple hosts to coherently share a common memory area in one or multiple memory devices. Multiple hosts can concurrently access and operate on that shared memory.

Growing CXL™ Ecosystem — Public Examples

Vendor	Type	Info/website
Astera Labs	LEO CXL™ Memory Accelerators	https://www.asteralabs.com/products/leo/
Montage	CXL® Memory eXpander Controller (MXC)	https://www.montage-tech.com/MXC
Micron	Memory expansion module using CXL	https://www.micron.com/products/memory/cxl-memory
Samsung	CXL DDR5 Module	https://news.samsung.com/global/samsung-electronics-introduces-industrys-first-512gb-cxl-memory-module
Samsung	CXL Memory-Semantic SSD	Next-Generation Memory Solutions at Flash Memory Summit 2022
SKhynix	CXL2.0 DDR5 Module	SK hynix Develops DDR5 DRAM CXLTM Memory | SK hynix Newsroom
Smart Modular	XMM (E3.S CXL2.0 DDR5 module)	https://www.smartm.com/media/press-releases/SMART_Modular_Technologies_Launches_its_First_Compute_Express_Link_Memory_Module
Marvell	Structera™ Memory expansion & memory pooling SoC	https://www.marvell.com/company/newsroom/marvell-introduces-breakthrough-structera-cxl-product-line-to-address-server-memory-bandwidth-and-capacity-challenges-in-cloud-data-centers.html

Compute Express Link® 2.0 (CXL® 2.0)

Enhancements in Intel® Xeon® 6 processors

Benefits:

- CXL 2.0 support for all device types and backward compatible with CXL 1.1
- Managed hot plug support for device addition/removal
- Link bifurcation for multiple devices per host port
- Advanced memory interleaving for improved bandwidth
- QoS enhancements to prevent CXL memory congestion

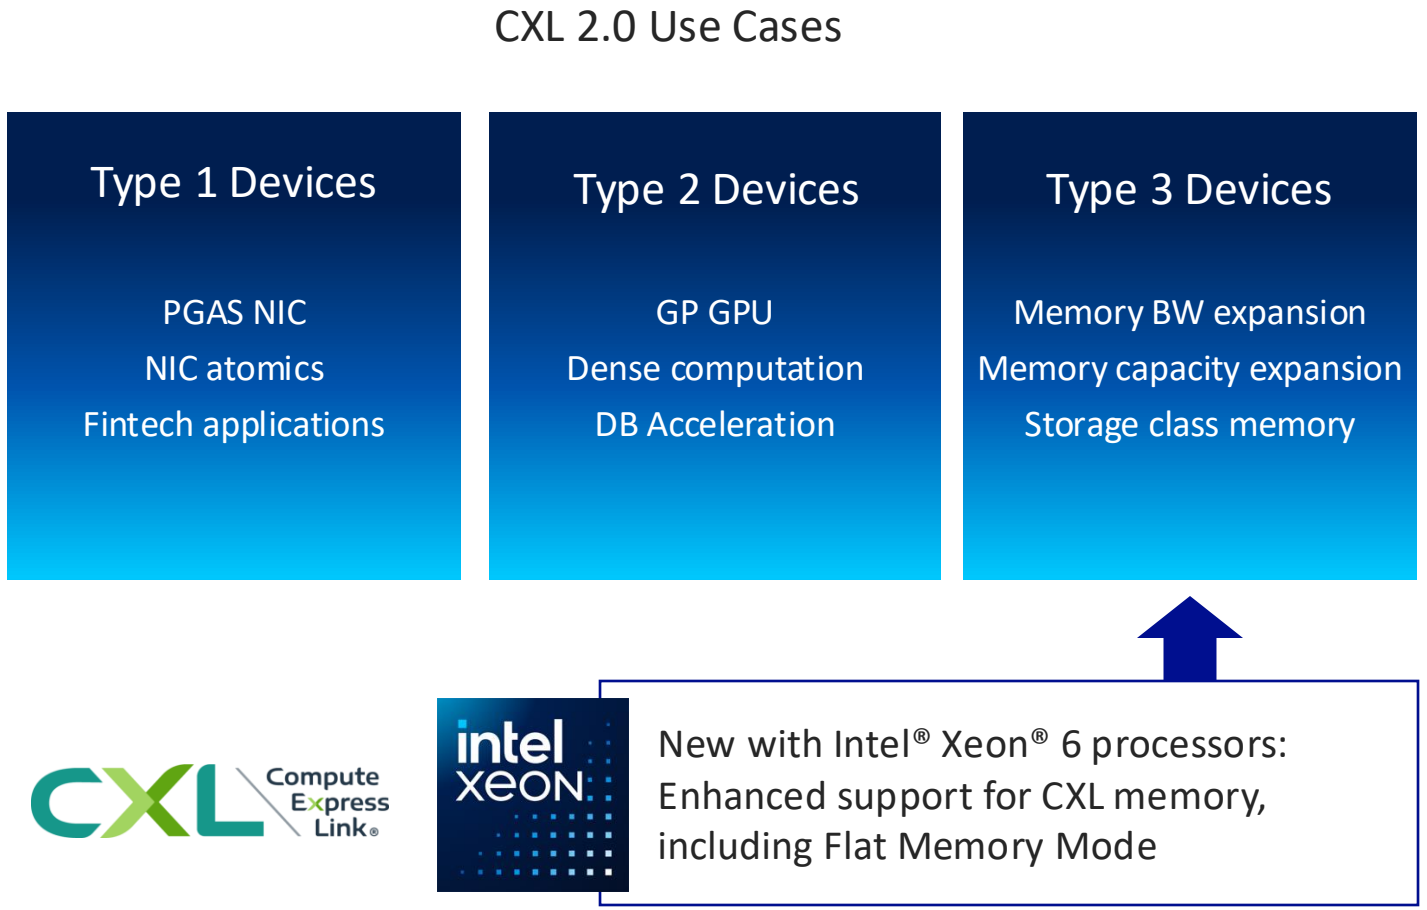

Intel® Xeon® 6 Processors with CXL® Attached Memory

CXL Memory is a cost efficient, flexible, first-class solution

CXL Numa Node

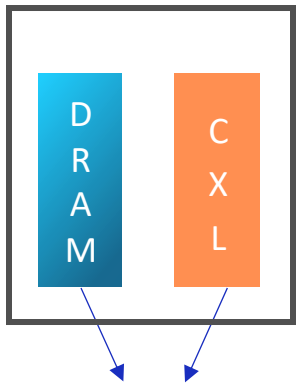

2 numa nodes

- SW (Hypervisor/OS/App)-assisted tiering
- 4KB page transfer, with movement overheads

CXL Hetero Interleaved

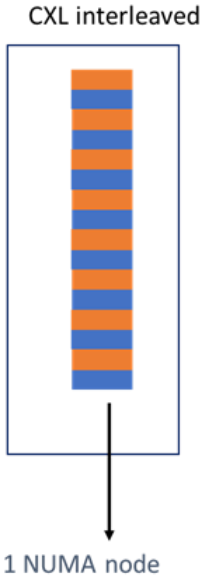

- HW-based memory interleaving
- Interleaving DRAM and CXL memory address space
- Only on 6900/6700 with P-cores, not on 6900/6700 with E-cores

Flat Memory Mode
Only available on Intel Xeon 6 CPUs

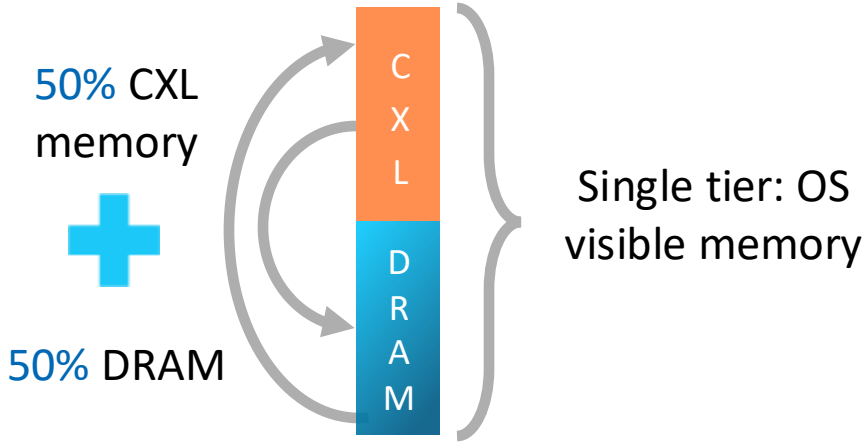

- HW-assisted tiering, fully SW transparent
- Single native DDR + CXL memory region
- 1:1 Near/Far memory ratio

intel

Compute Express Link Performance with Intel® Xeon® 6

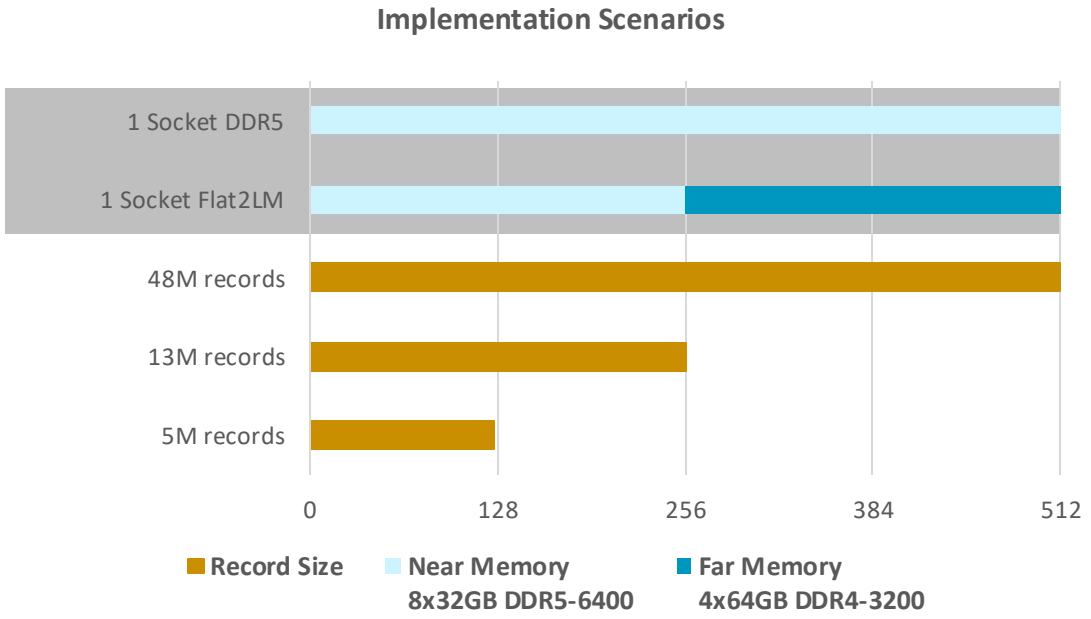

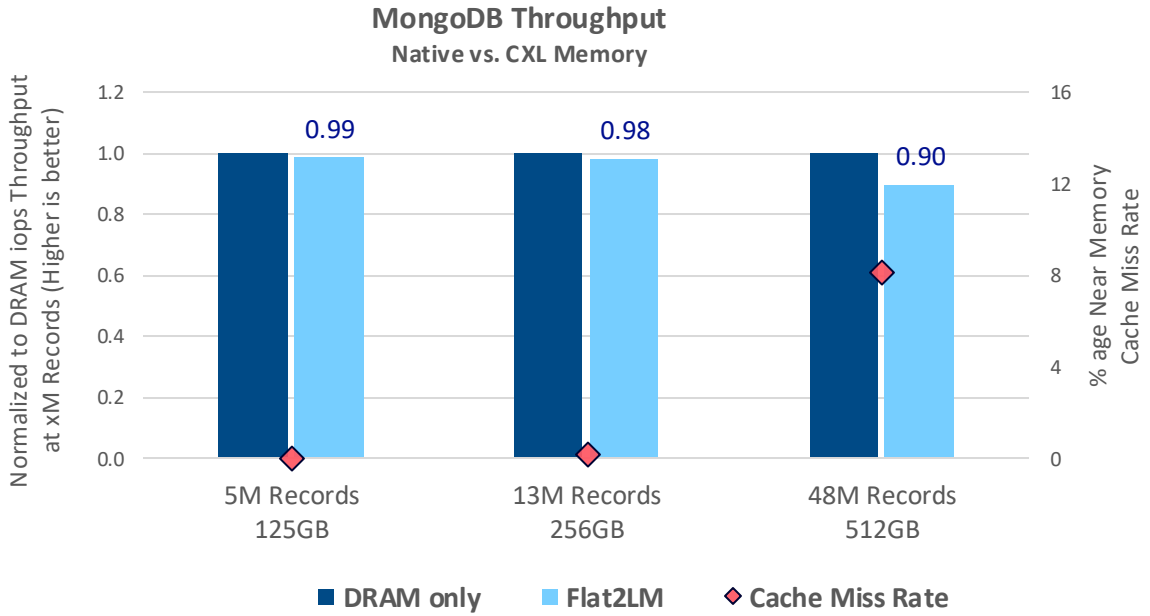

- Customers can reuse existing DDR4 modules (no additional expense) resulting in TCO savings
- Slight degradation in performance in exchange for repurposing the DRAM
- No dependence on OS data-tiering capability and hence, no OS-related overhead
- Lower latency compared to OS-based tiering (64B vs. 4KB page movement)

intel

See backup for workloads and configurations. Results may vary.

CXL Resources

[Resource Library | Compute Express Link]

- Webinars
- White Papers
- Videos
- Presentations

Educational blogs: https://www.computeexpresslink.org/blog

Technical Training

- CXL 1.1 Technical Training | Compute Express Link
- CXL 2.0 Technical Training | Compute Express Link

Gen-Z Specification

- Gen-Z specification archive


intel