Modularity and disaggregation

Jean Laurent PHILIPPE
EMEA CTO

intel.

Forum TERATEC 23

Building Blocks for largest system

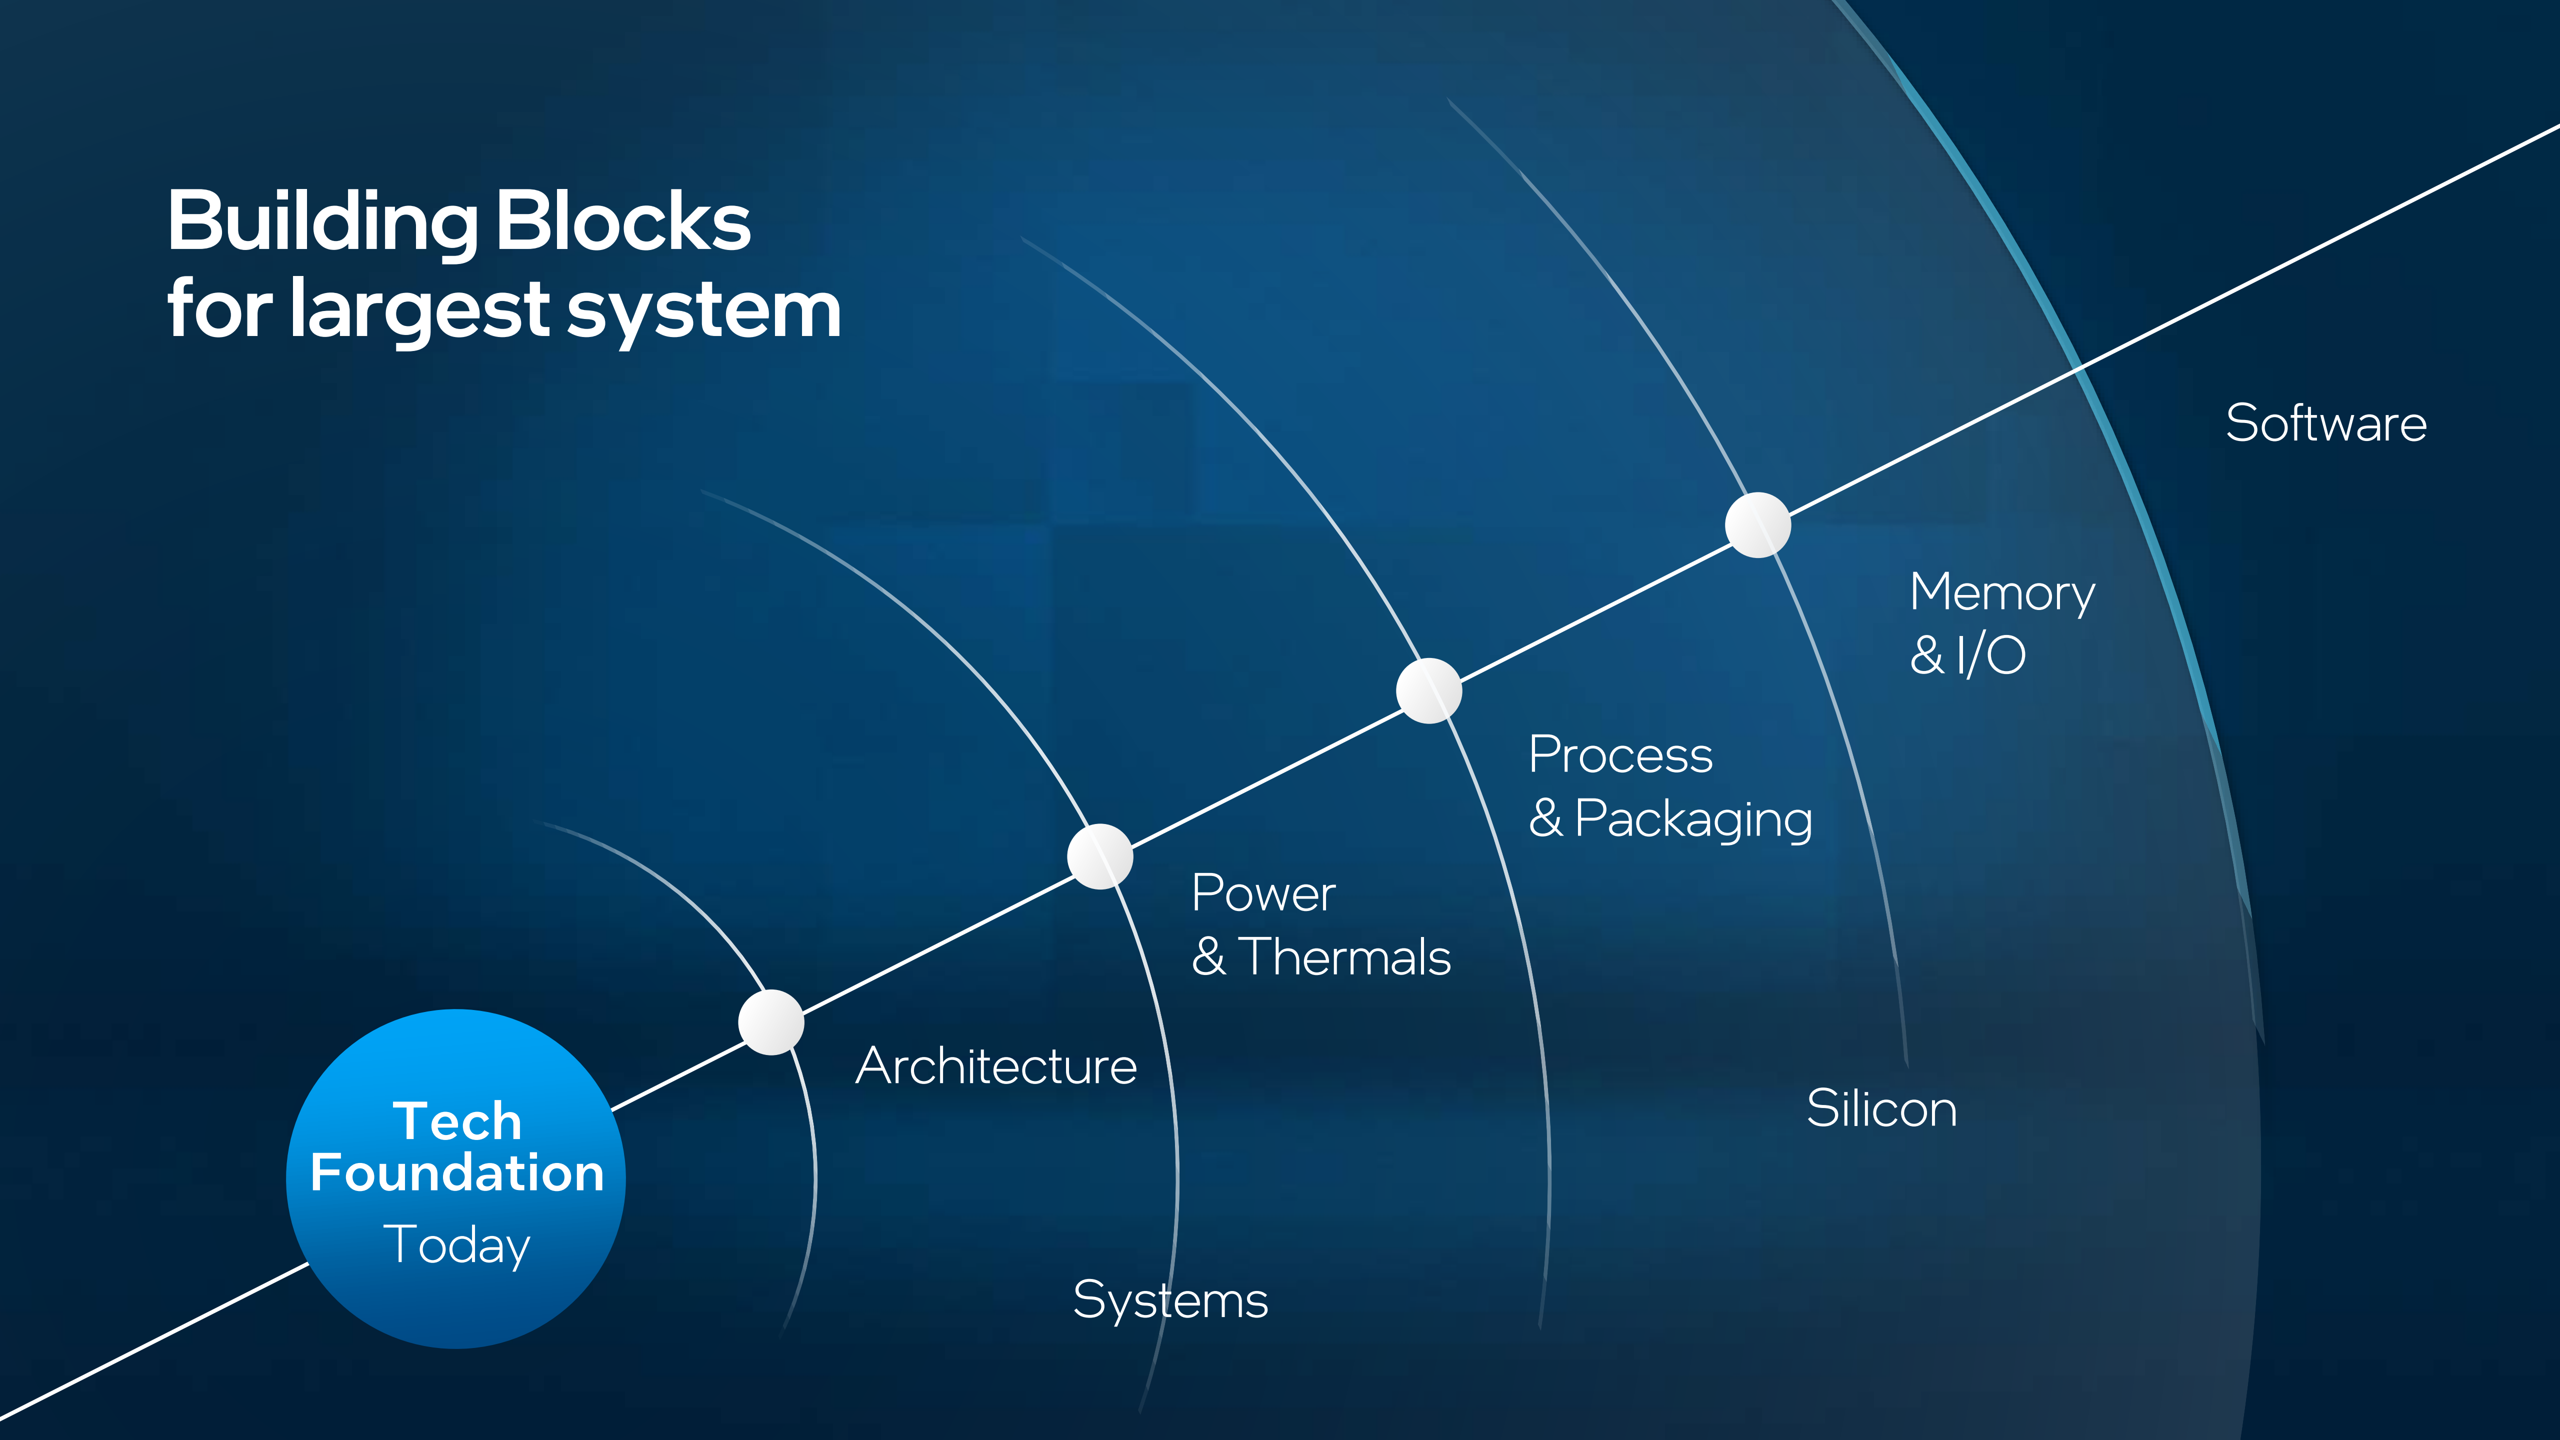

Silicon

Systems

Software

Silicon

Systems

Software

Workload Heterogeneity

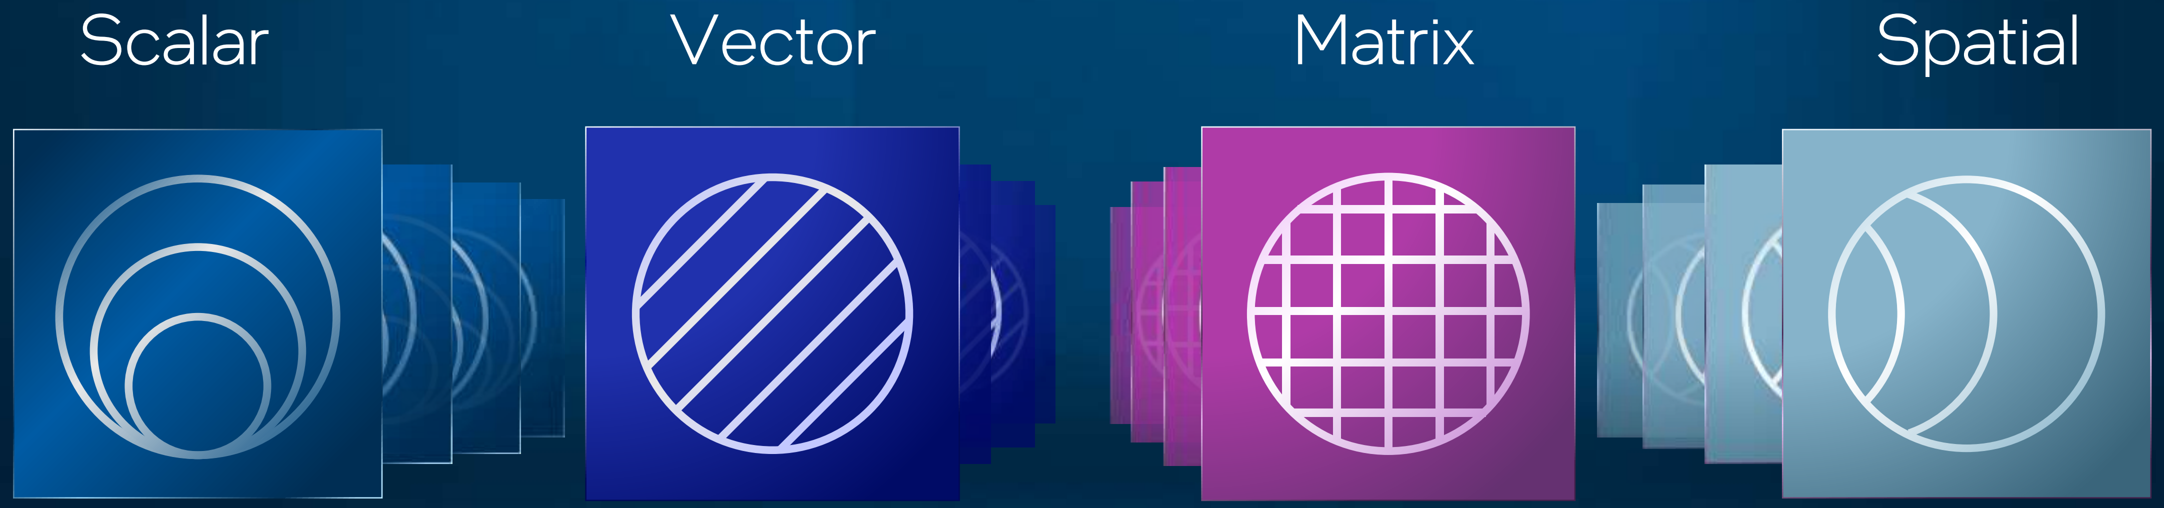

Scalar

Vector

Matrix

Spatial

Hybrid Compute Cluster

in a Package

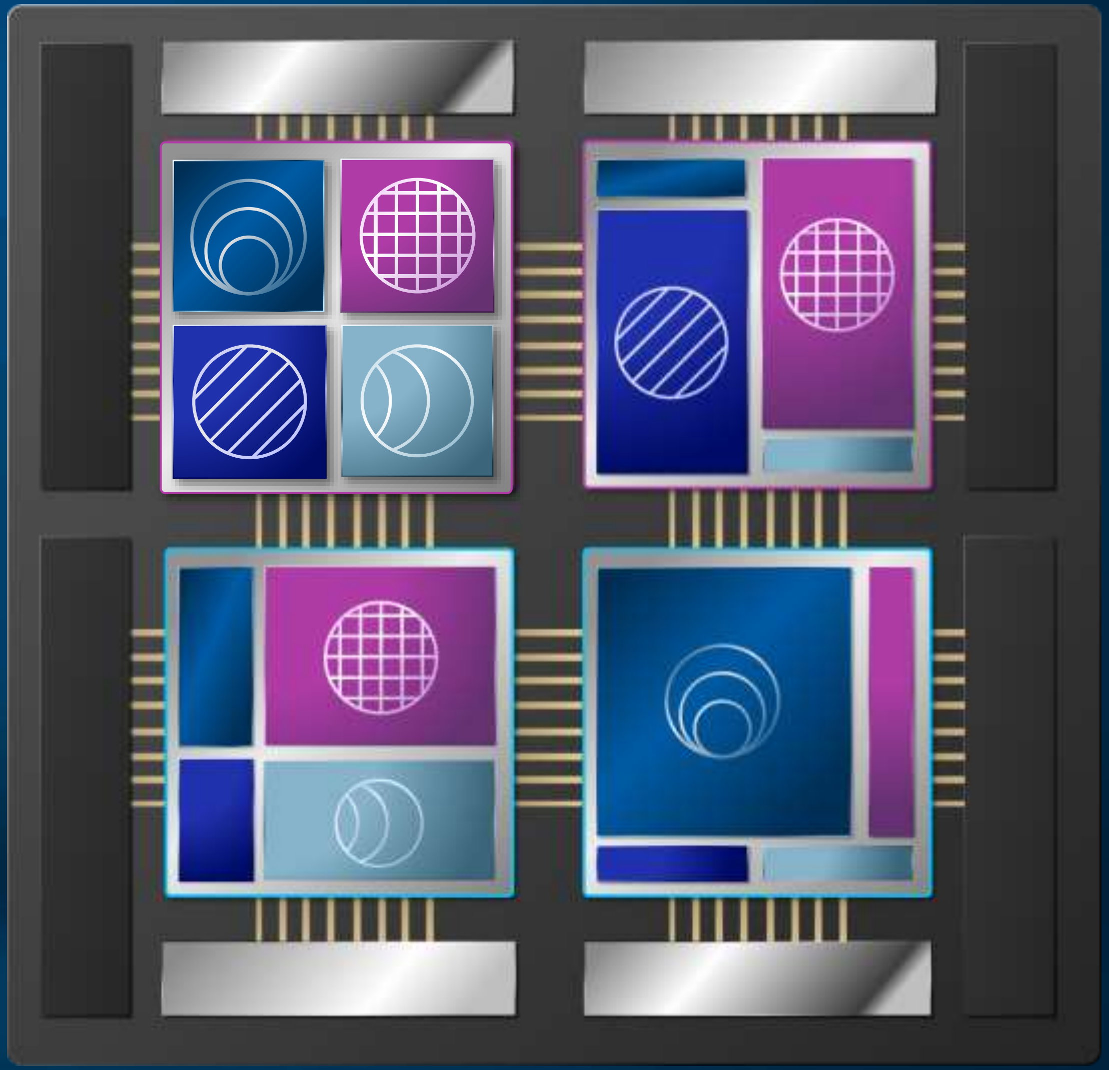

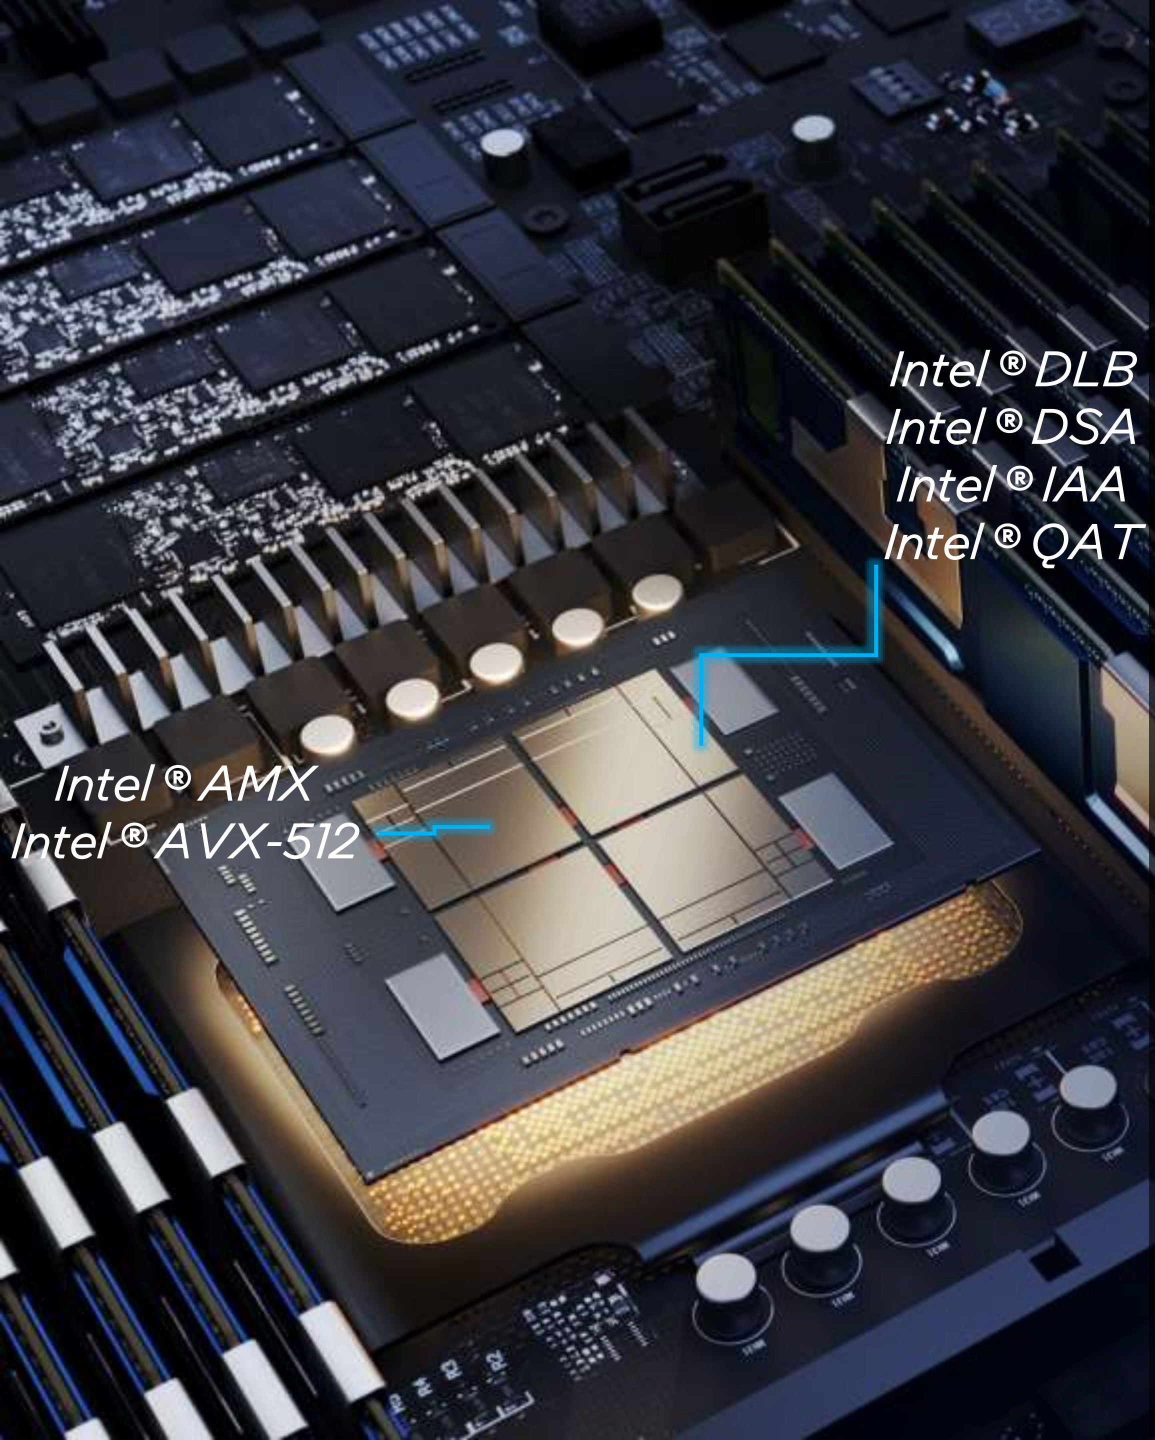

Intel® Xeon® CPU Max Series

codenamed

Sapphire Rapids HBM

AMX
Advanced Matrix Extensions

Integrated Acceleration Engines

64GB
HBM2e

intel
DATA CENTER
GPU
MAX SERIES

Falcon Shores GPU

Next Gen GPU
for AI & HPC

- Converged Habana & Xe IP
- Modular Tile-Based Architecture
- HBM 3 & I/O designed to scale
- Standard Ethernet Switching
- Flexible CPU-GPU ratio
- A single GPU programming Interface
- CXL Programming Model

UCIe

Universal Chiplet Interconnect Express

Reduced Time-to-Market

Reduced IP Porting Costs

Smaller, Higher Yielding Components

Optimal, per-Component Tech

Leaders in semiconductors, packaging, IP, cloud service providers joining forces

Alibaba Group | AMD | arm | ASE GROUP

Google Cloud | intel | Meta | Microsoft

NVIDIA | Qualcomm | SAMSUNG | tsmc

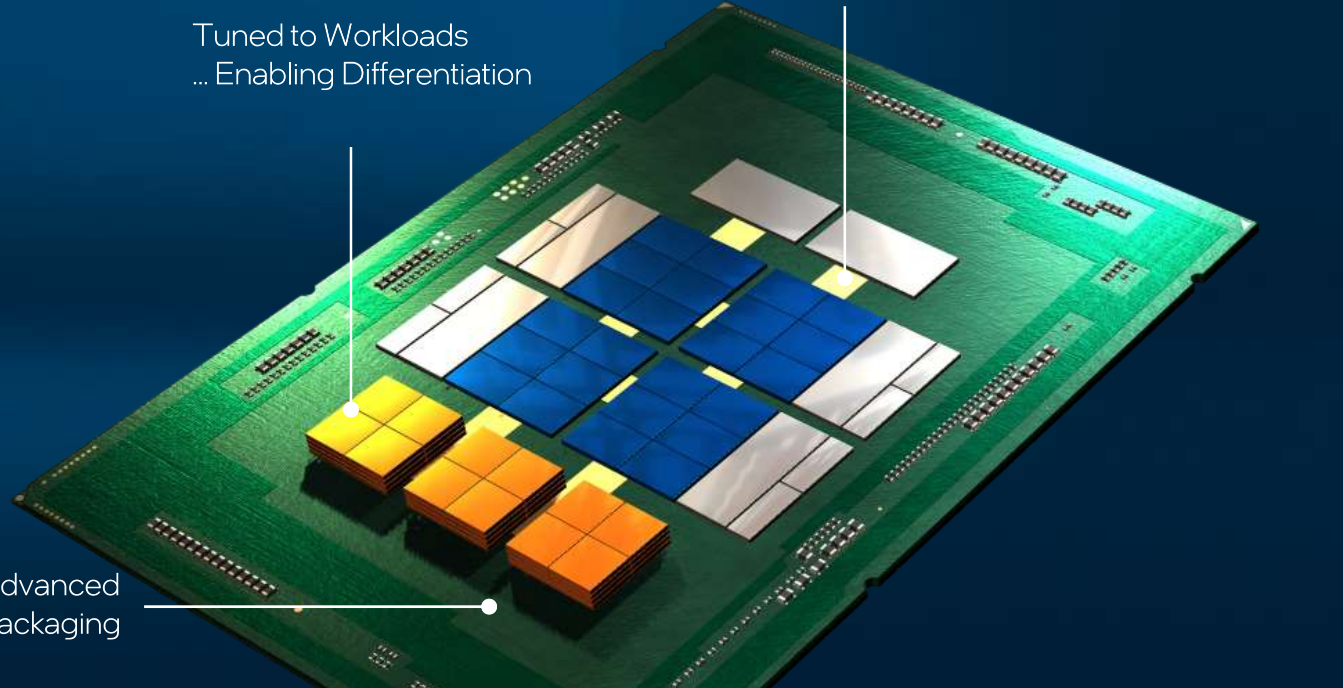

High-Speed Standardized Chip-to-Chip Interface

Tuned to Workloads
… Enabling Differentiation

Advanced Packaging

Integrated Optical I/O

Key innovation for growing data rates, energy efficiency and channel loss minimization needs

2023 target

Bandwidth	Reach	Shoreline Density	Energy Efficiency	Latency
~1 Tbps per fiber	>100m	>4x PCIe6	Trending 3pJ/b	<10ns + TOF

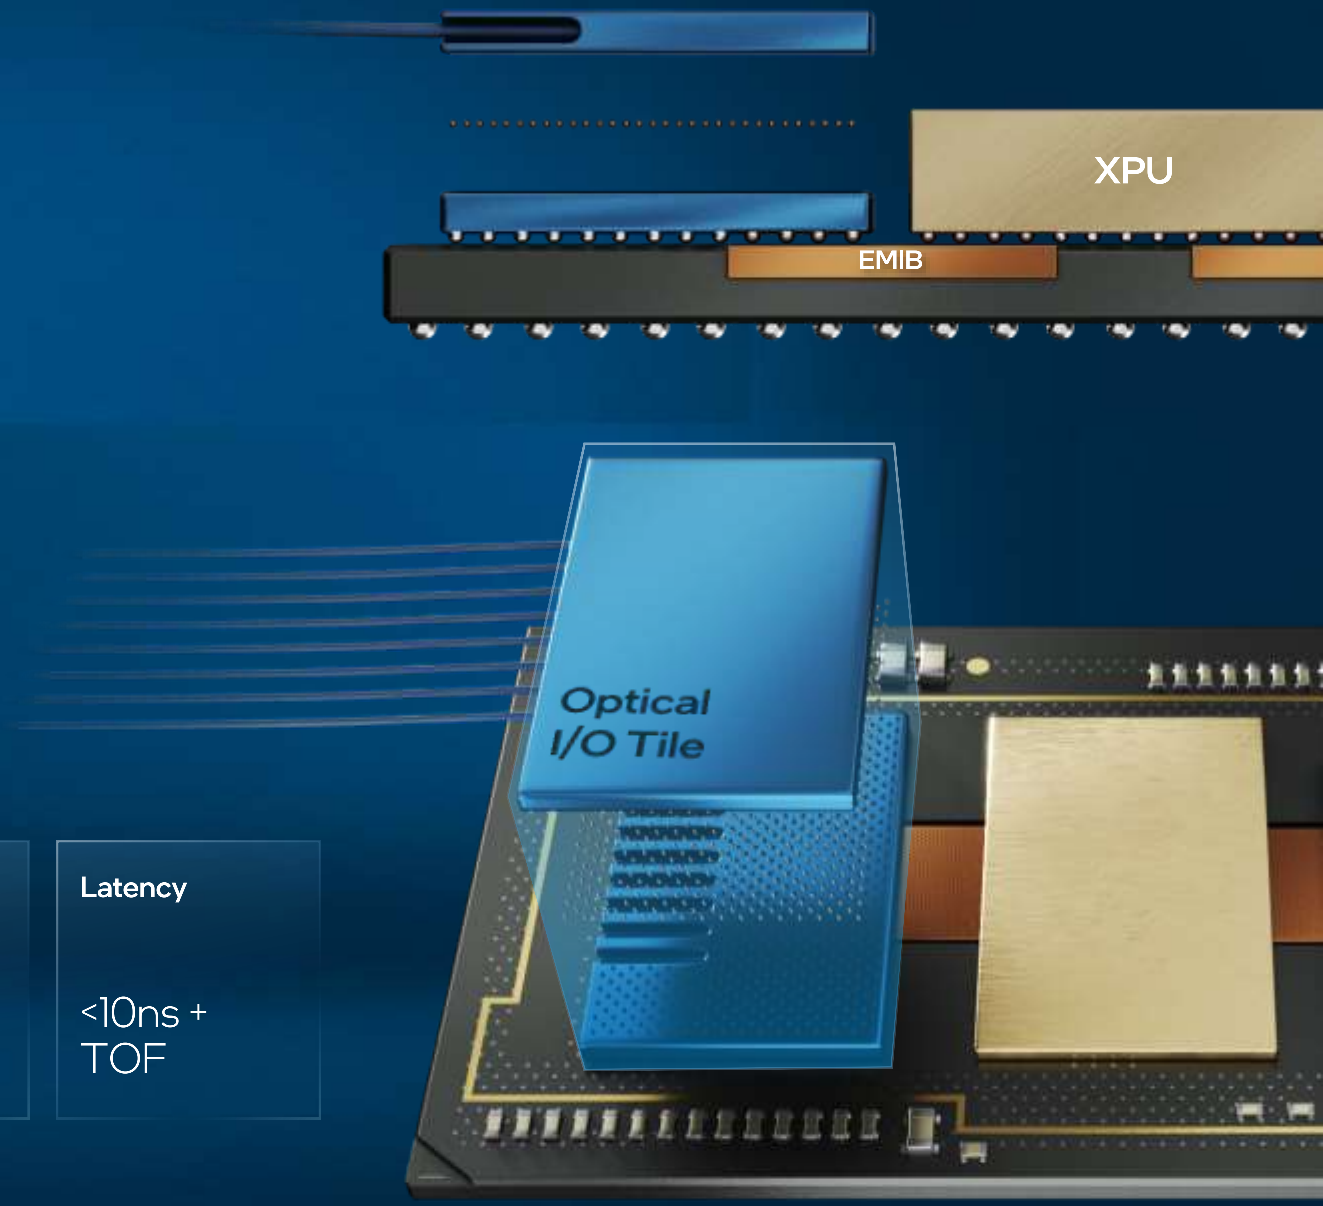

Silicon

Systems

Software

The Open Standards—CXL™ is The Future

Solve these challenges with CXL

- Scaling challenges: latency, bandwidth, capacity
- Increase of heterogenous computing + massive datasets + demanding workloads = need for more accessible data, faster
- These computing demands require a more efficient interconnect between CPU and traditional I/O interface

CXL Compute Express Link

Go Faster

- Improves data handling and reduces I/O bottlenecks

Do More

- Memory bandwidth and capacity expansion, efficient access across shared memory
- Enables CPU and accelerators to share memory resources for higher performance

Save More

- Improves TCO

Sources: Introduction to Compute Express Link White Paper, March 2019 and https://www.computeexpresslink.org/

CXL™ Consortium – Scope and Feature

Coherent Interface

Leverage PCIe with three multiplexed protocols

Built on top of PCIe® infrastructure

Low Latency

CXL.Cache/CXL.Memory targets near CPU cache coherent latency (<200ns load to use)

Asymmetric

Complexity

Eases burden of cache coherence interface designs for devices

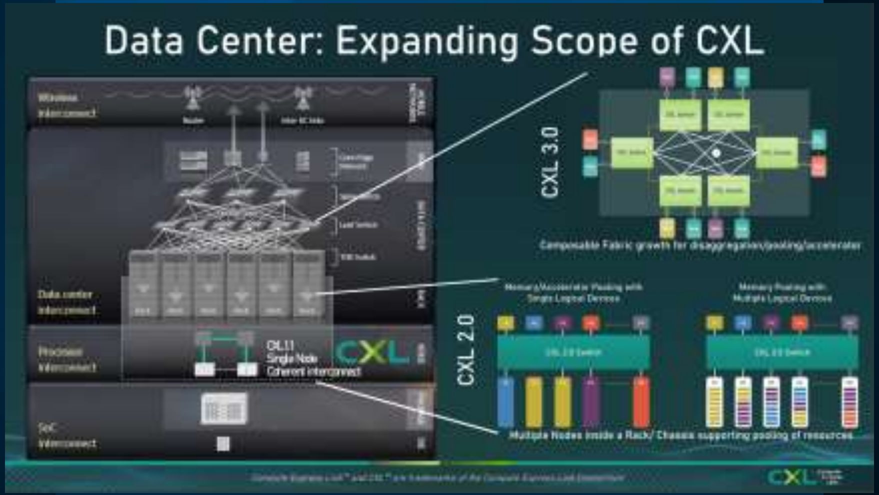
Data Center: Expanding Scope of CXL
Wireless interconnect
Data center interconnect
Processor interconnect
SoC interconnect
CXL1.1
Single Node
Coherent interconnect
CXL
CXL 3.0
Composable Fabric growth for disaggregation/pooling/accelerator
CXL 2.0
Memory/Accelerator Pooling with Single Logical Devices
Memory Pooling with Multiple Logical Devices
CXL 2.0 Switch
CXL 2.0 Switch
Multiple Nodes inside a Rack/ Chassis supporting pooling of resources

Heterogeneity

Enable 3 type of devices based on 3 protocols:
memory, cache, IO

Modularity & disaggregation

Allocate and deallocate resources on demand,

Enhance memory and cache coherency

Peer to peer memory access

Scalable

High speed low latency fabric

Composable fabric

enable efficient resources sharing and modularity

Possible Usage Model Transition Examples

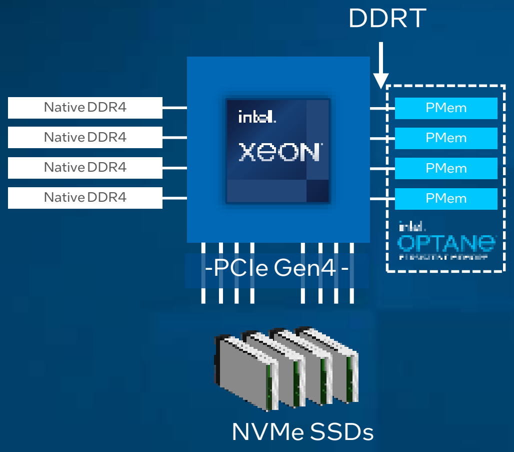

Today: Intel® Optane™ Technology

Intel® Xeon® CPU + DDR4 +
Intel® Optane™ Technology on DDRT

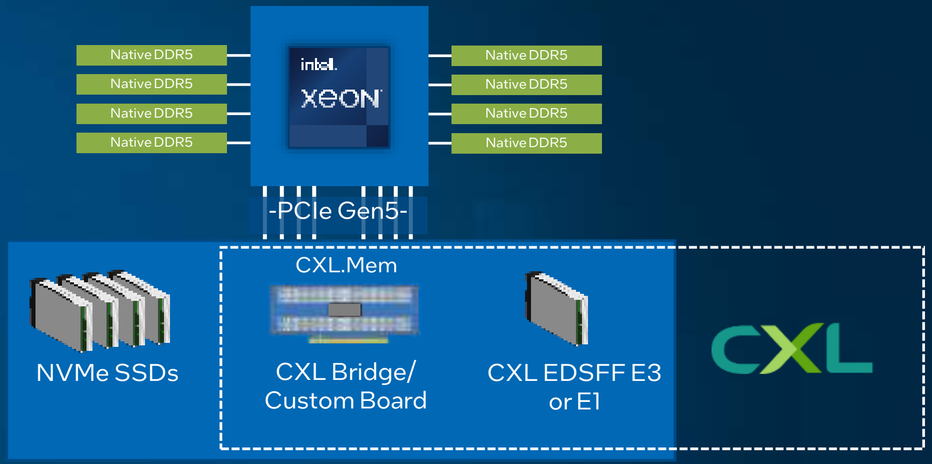

Future: CXL 2.0

Intel® Xeon® CPU + DDR5 +
Third-Party CXL Memory Products

Memory Mode	Cost-effective memory to meet workload needs and expanded memory capacity for workload scale-up.	Adds industry standard protocol for software to cache or tier memory.
Memory Expansion, Augmentation	Path for higher capacity than DRAM for workloads to scale-up.	Adds additional memory capacity and memory bandwidth to existing DDR attached DRAM.
Persistent Memory	Fast storage for meta data storage, fast write logs, and caching/tiering acceleration of other storage.	Adds persistence support (non-volatile memory over CXL).
High-Endurance, Performant SSDs	Very fast storage "replacement" for meta data storage, fast write logs, caching/tiering acceleration.	Adds persistence support (memory-semantic storage over CXL).

Coherent Discrete Accelerator for 4th Gen. Intel® Xeon® Scalable Processors

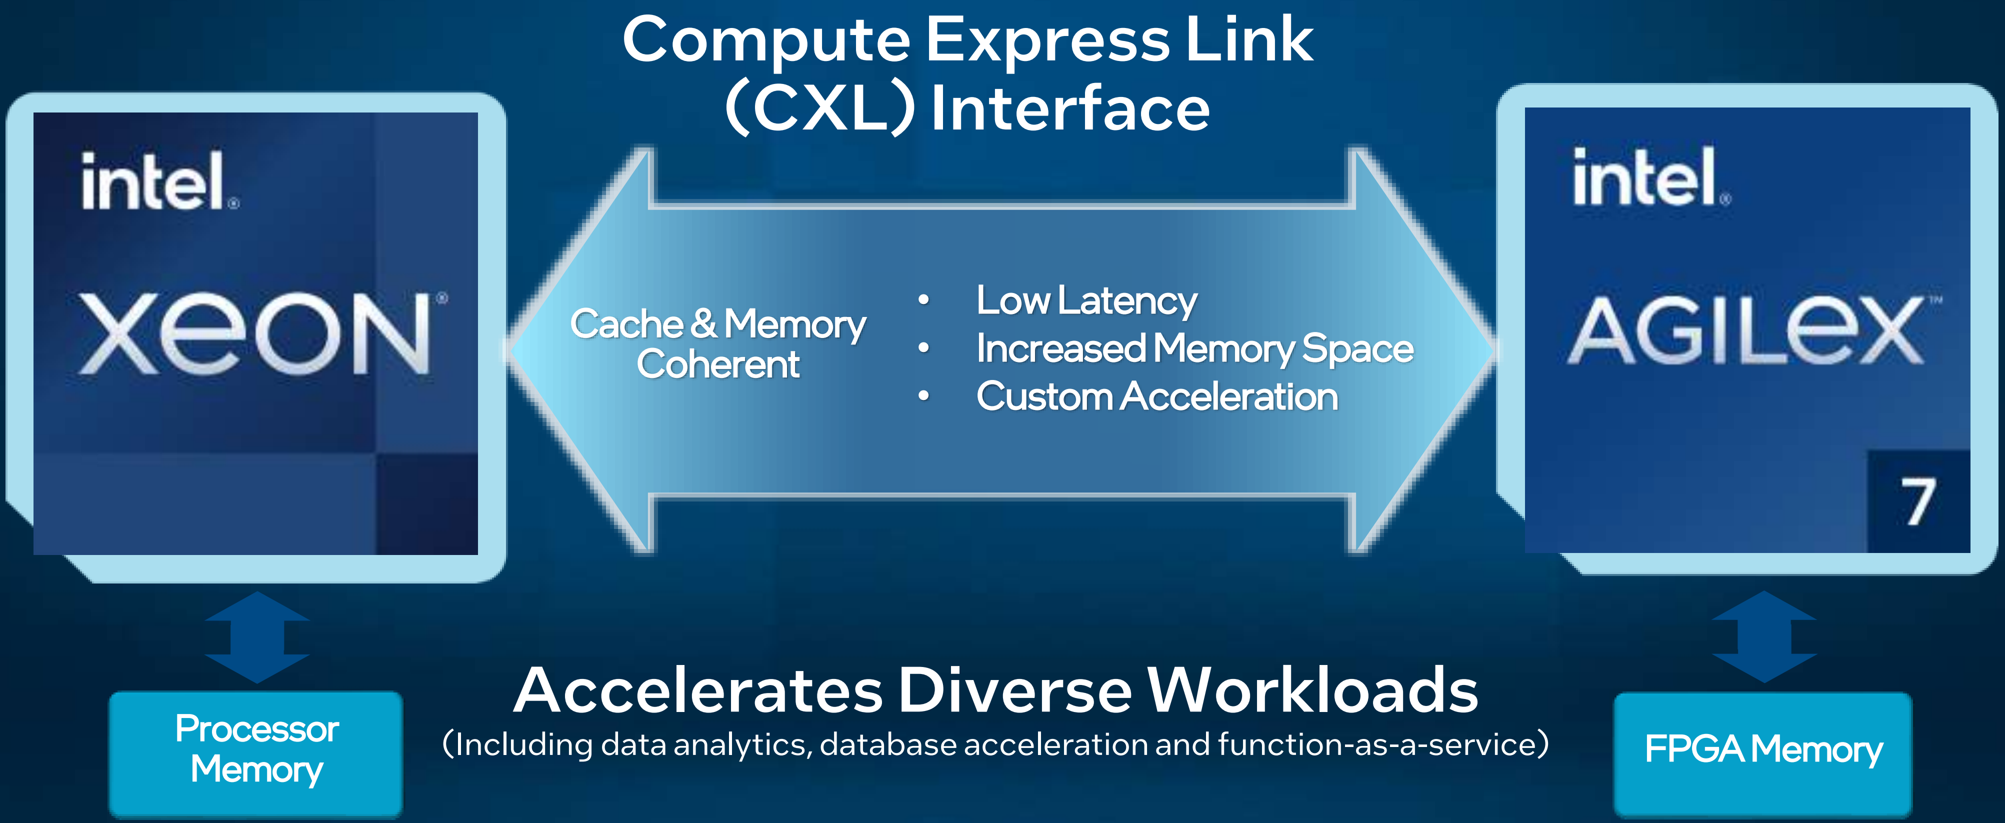

- Low Latency
- Increased Memory Space
- Custom Acceleration

Accelerates Diverse Workloads

(Including data analytics, database acceleration and function-as-a-service)

*Available on selected Intel Agilex 7 I-series and M-series FPGAs which contain at least one R-Tile.

https://www.computeexpresslink.org/

Silicon

Systems

Software

Unifying Software Stacks Critical....

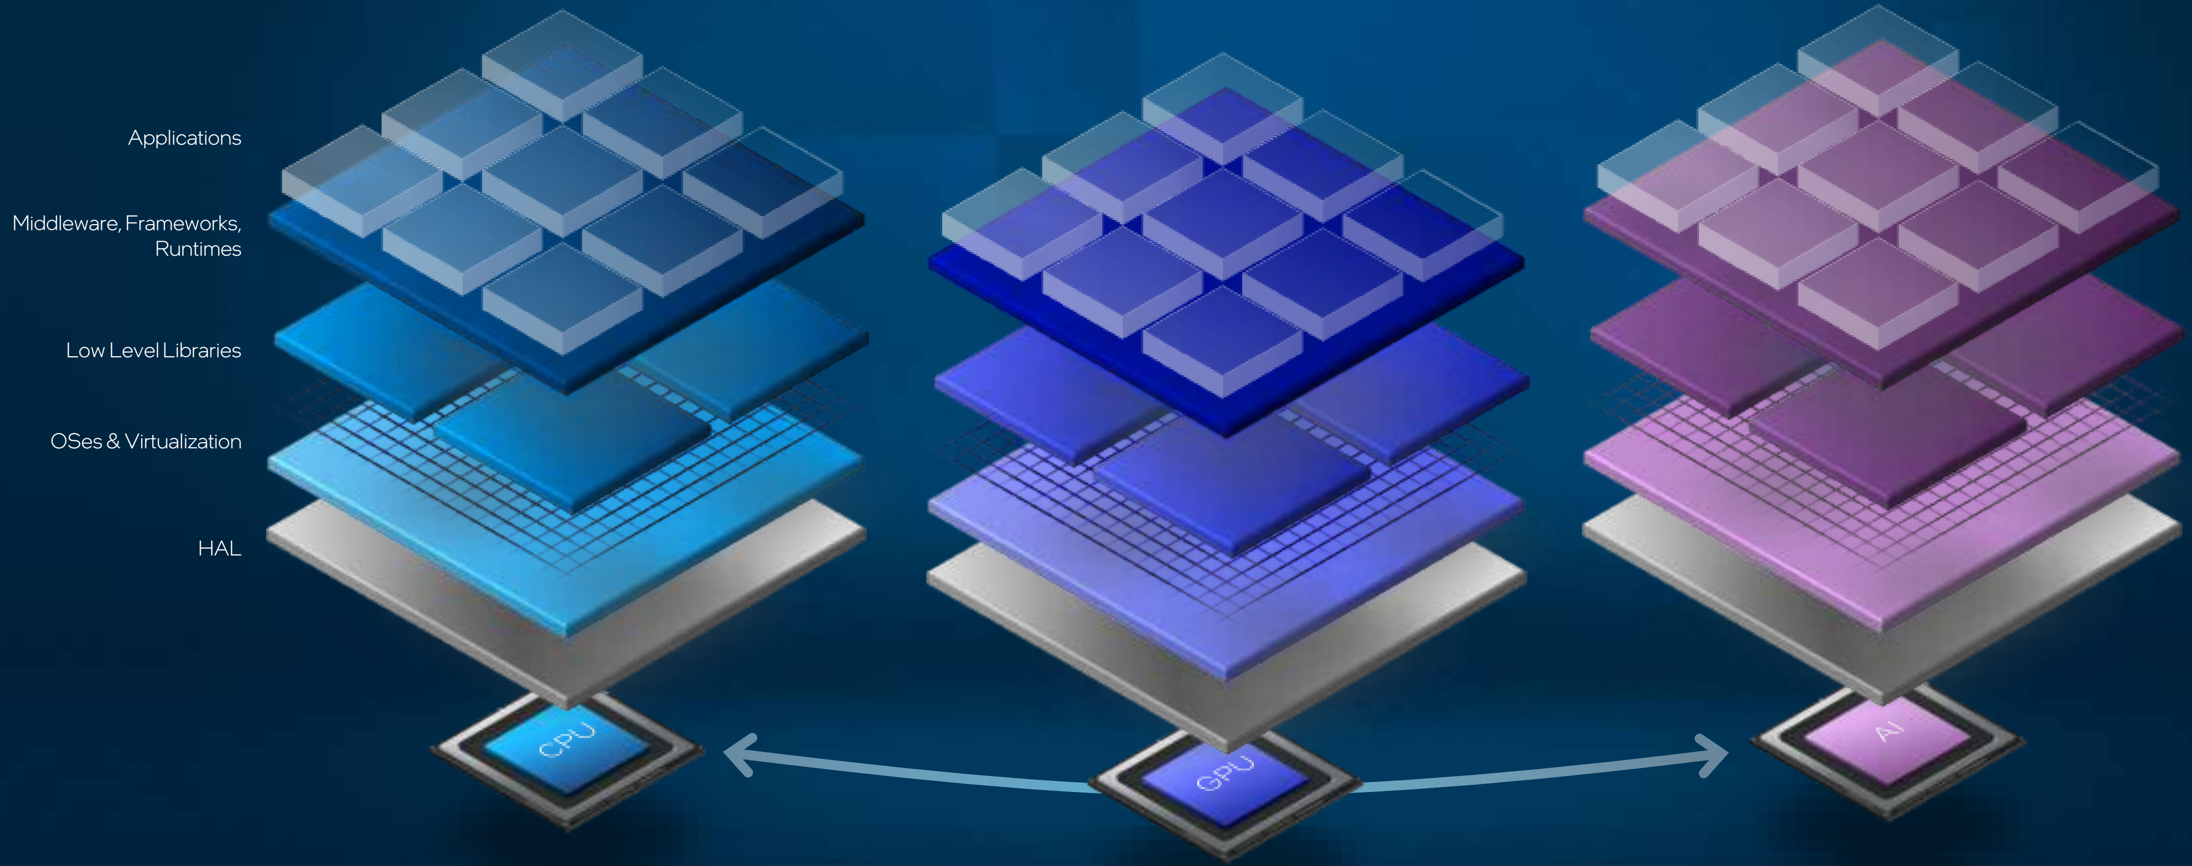

Maximizing Impact,
Minimizing “energy to solution” through...

Open
Ecosystem

Choice
Compatibility

Trust
Workloads

Scale
Delivery & Deployment

oneAPI

open, multiarchitecture, multivendor programming

- Open industry specification
- Freedom in hardware choice
- Performance, productivity & portability
- Standards-based, community-driven

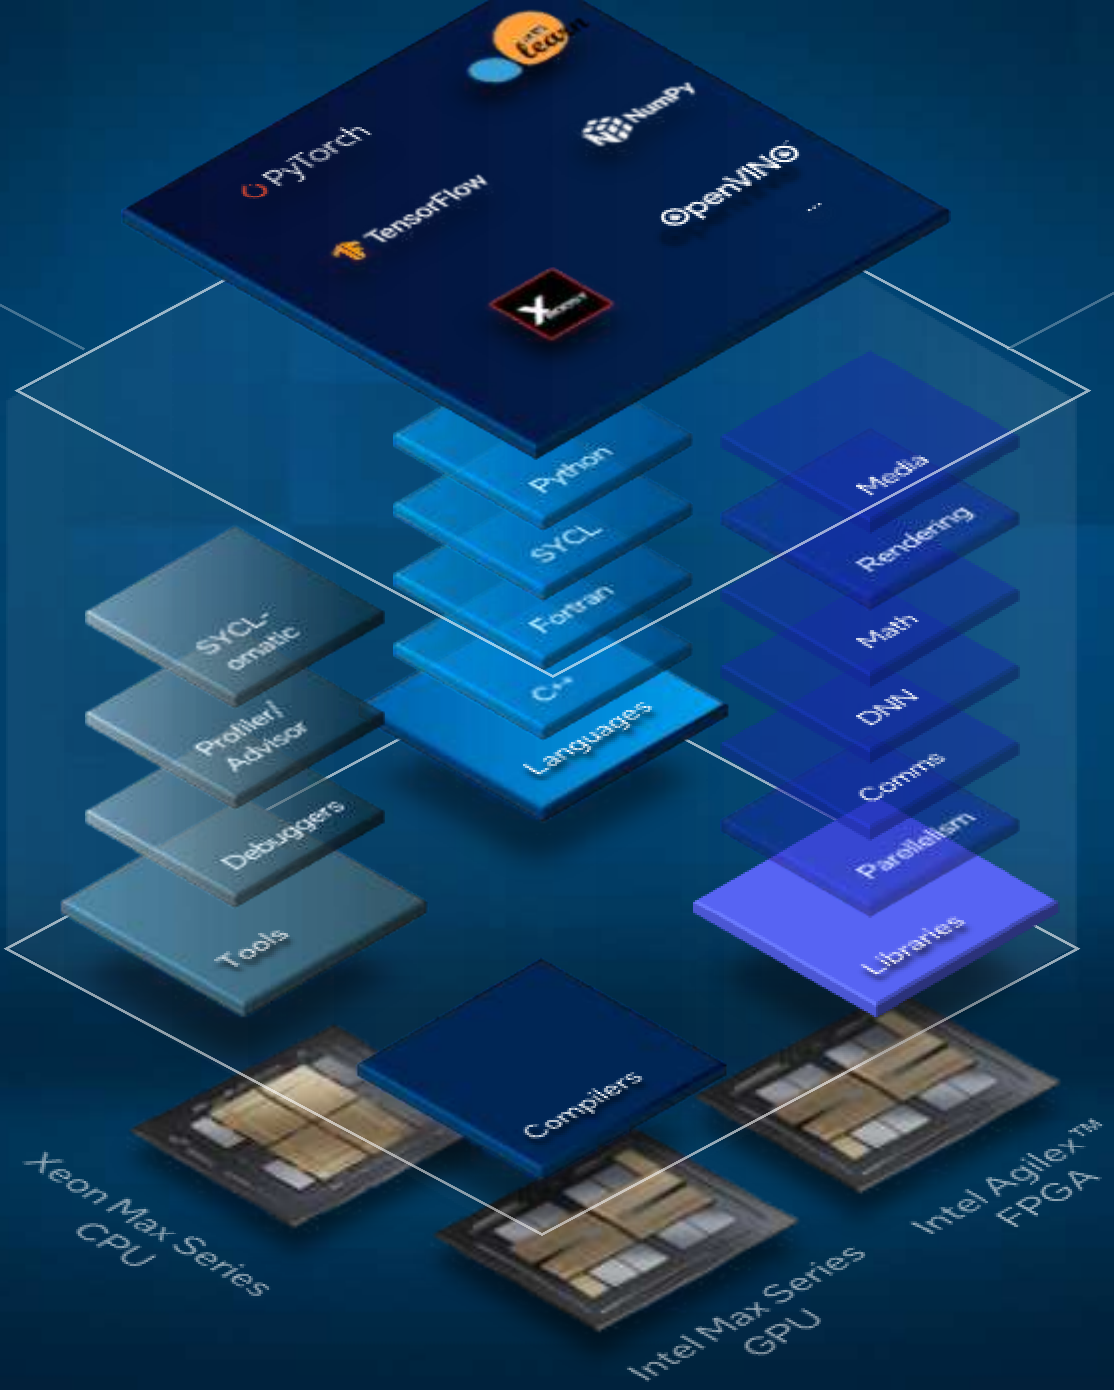

oneAPI

Intel's implementation with a set of tools

- Optimized for Intel hardware
- Proven performance, best-in-class capabilities
- Supports SYCL, Fortran, Python, OpenMP, MPI...
- Enhanced CUDA-to-SYCL code migration capabilities
- Optimizations for TensorFlow & PyTorch
- Download free, commercial support available

lrz Leibniz-Rechenzentrum
der Bayerischen Akademie der Wissenschaften

Write Once, Run Source Code

Across Architectures & Across Vendors

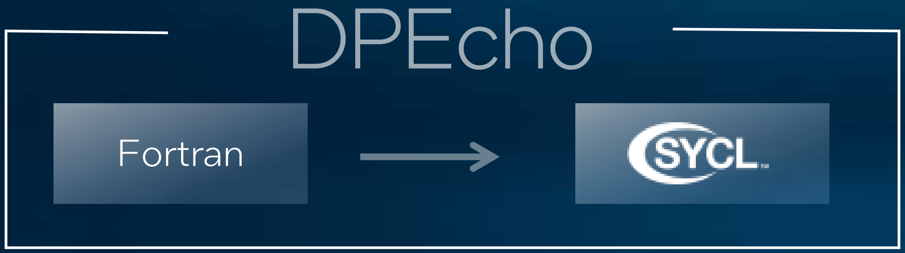

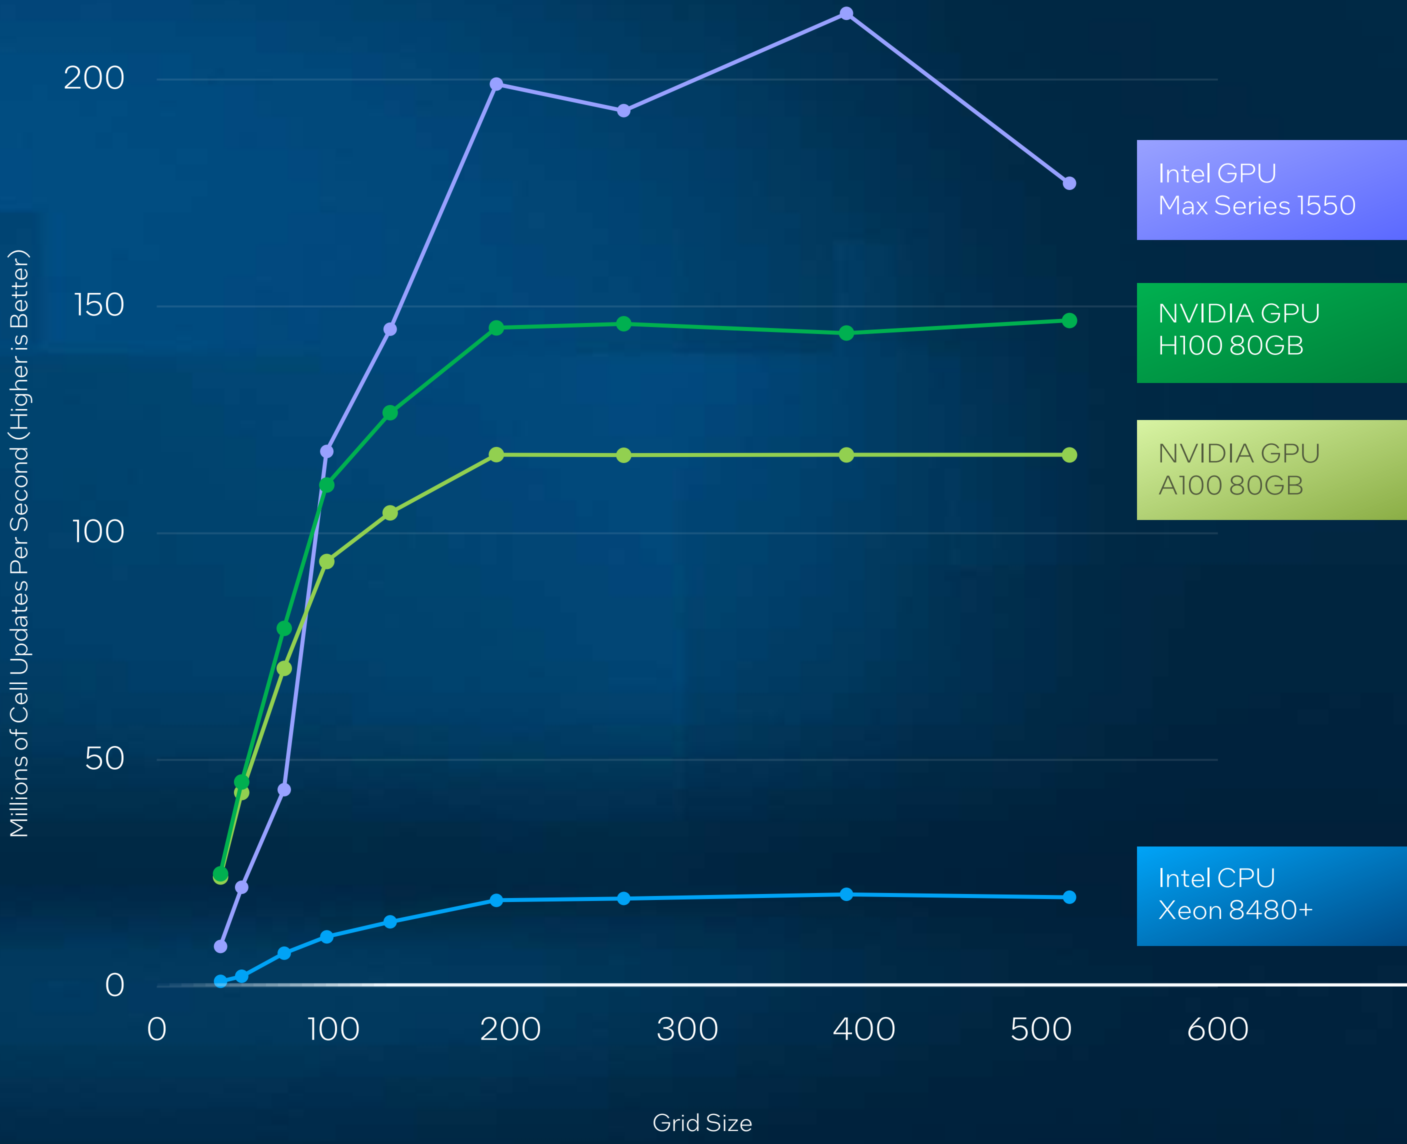

See backup for workloads and configurations. Results may vary.

Drive innovation with Open Standards

UCIe
Universal Chiplet
Interconnect Express

CXL

Integrated
Optical I/O

Intel® AMX

intel.
xeon

Silicon

intel.
AGILEX
7

Systems

oneAPI

oneAPI

Software

intel®