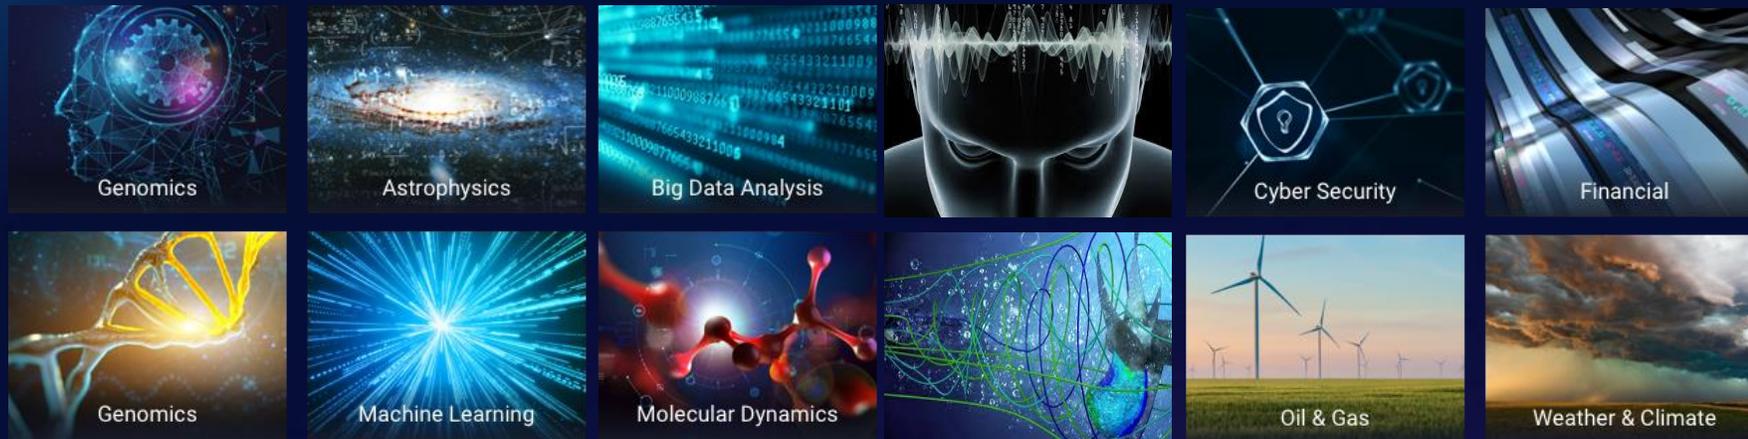




Solving the world's complex computational challenges with Intel's data center GPU, codenamed Ponte Vecchio
AMIL011

Ogi Brkic , GM & Vice President
Radhika Rao, Sr. Director

Meeting the compute needs of the future



SOLVE THE WORLD MOST CHALLENGING PROBLEMS BY PROVIDING ZETTA-SCALE COMPUTING CAPABILITIES BEFORE THE END OF THE DECADE

HPC & Data Center GPU Segments

Super Compute

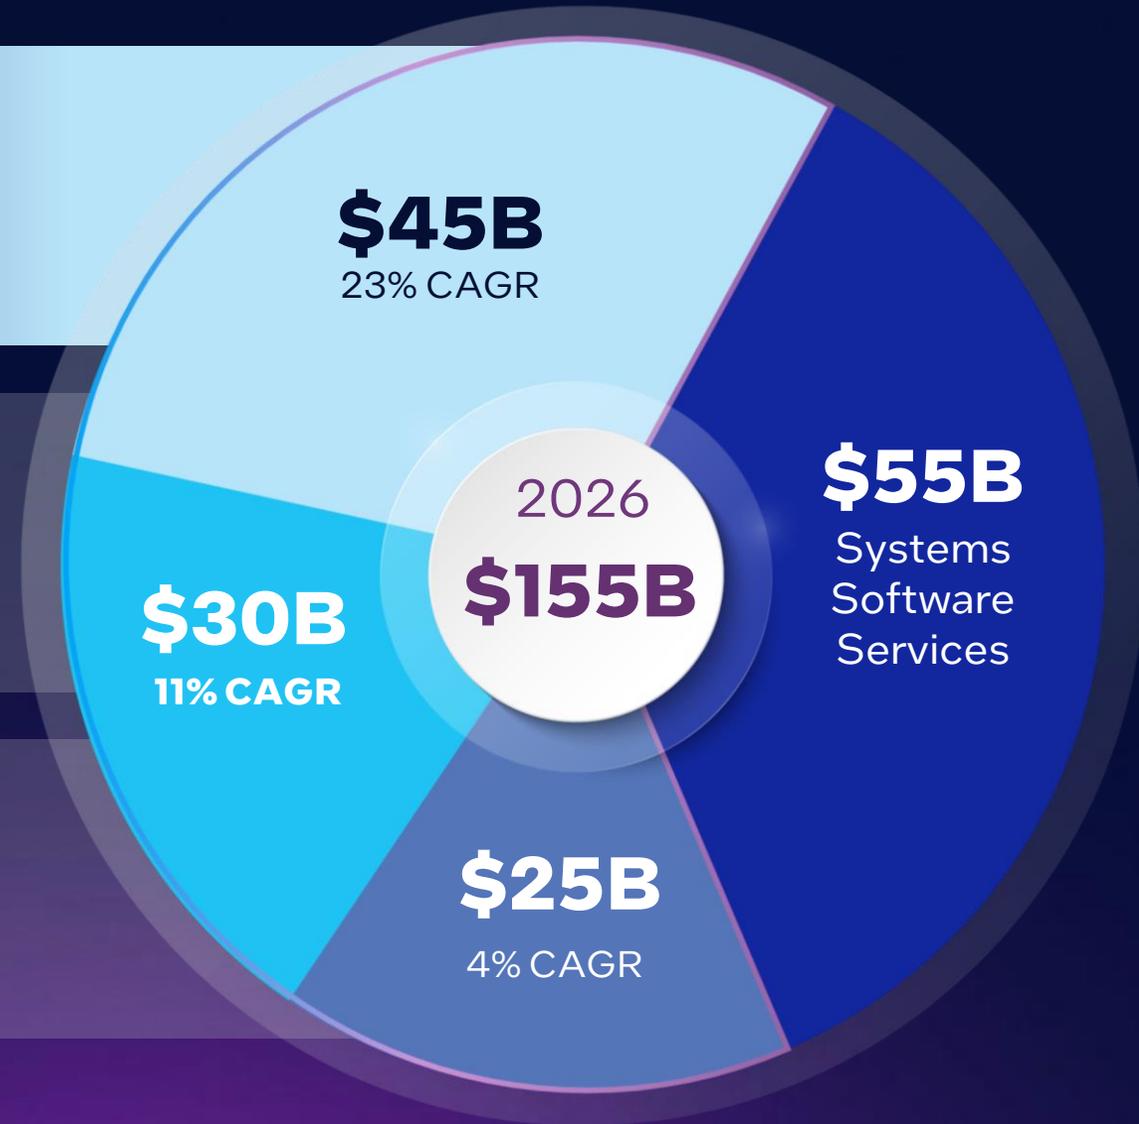
HPC - AI
Media & Visual Cloud

Custom Compute

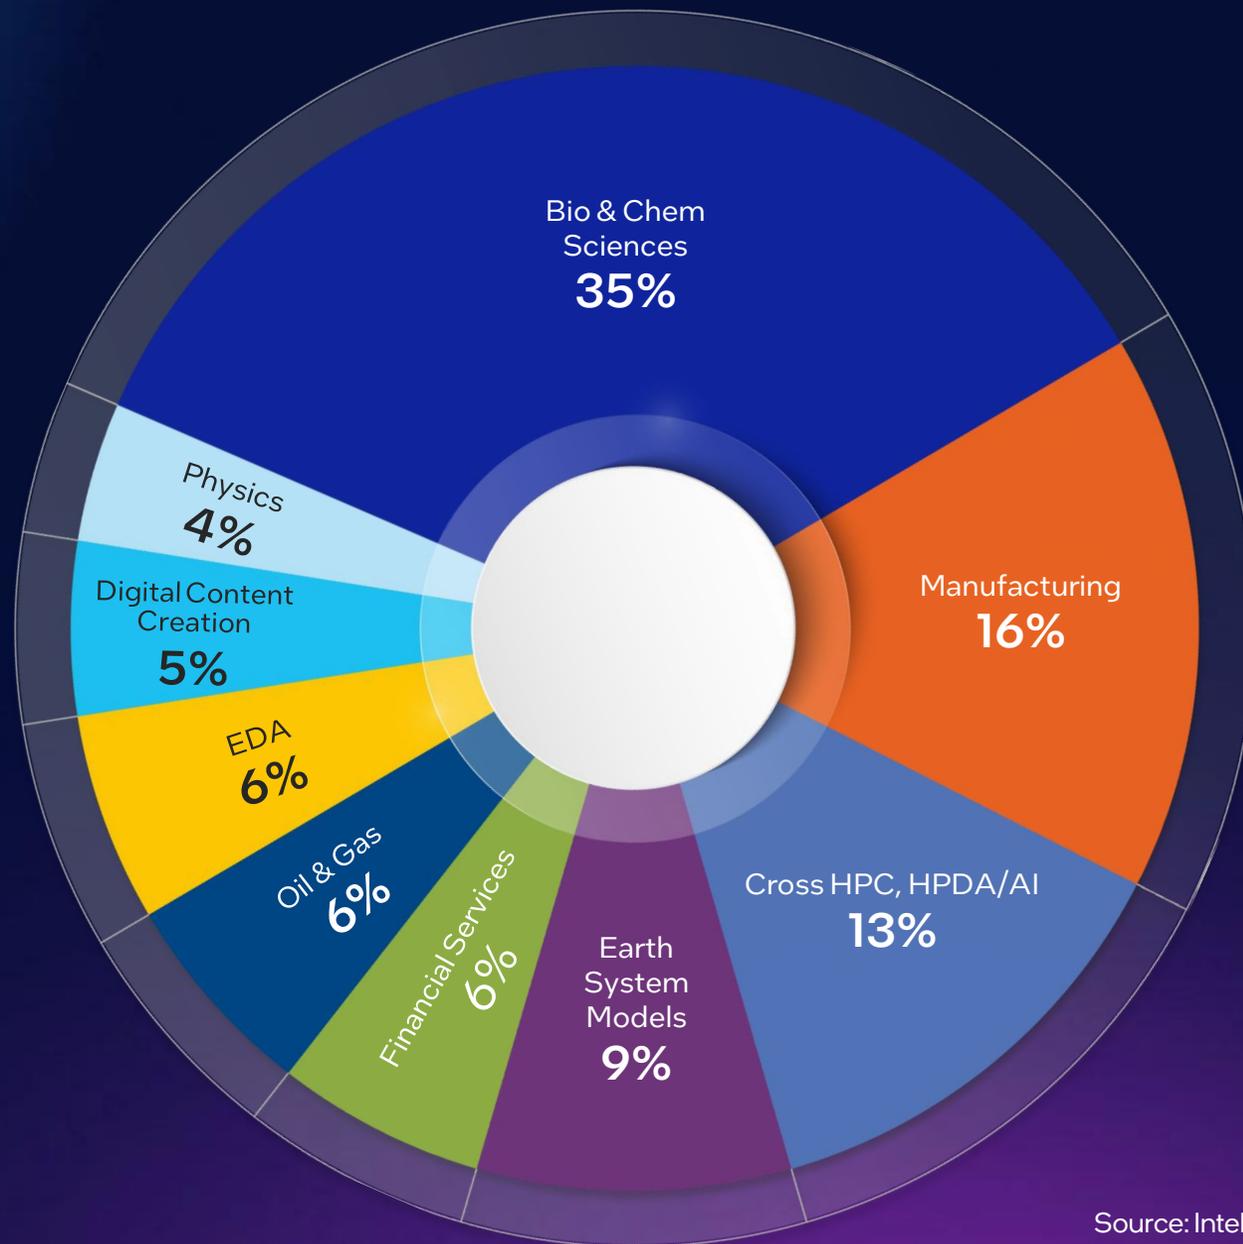
Blockchain
Supercomputing at the
Edge

Visual Compute

Gaming
Content Creation
Metaverse

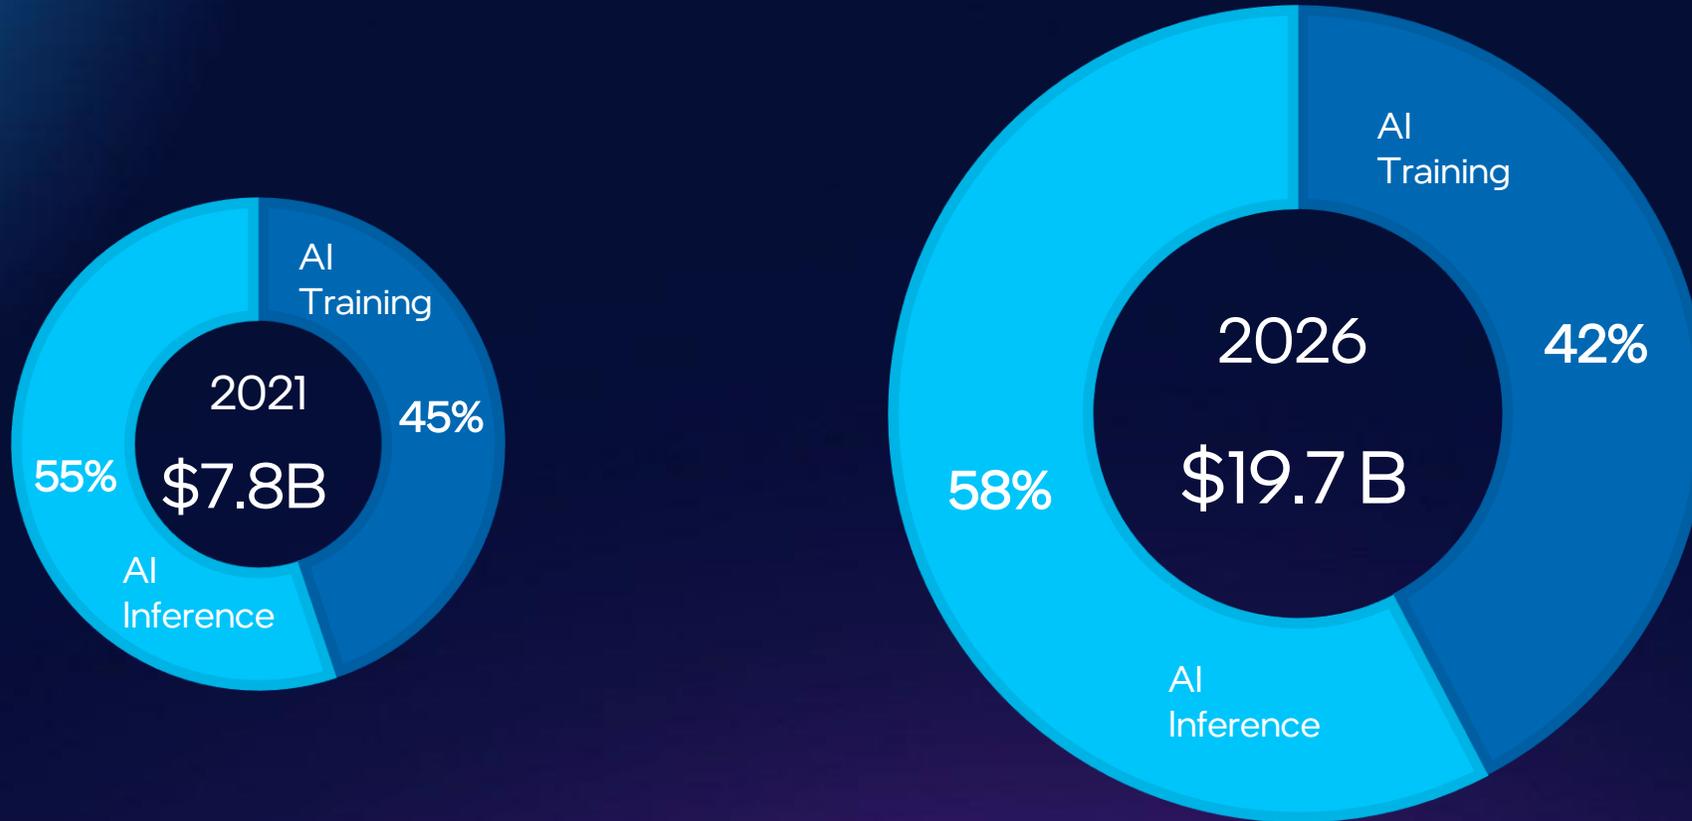


Top HPC Applications by Vertical

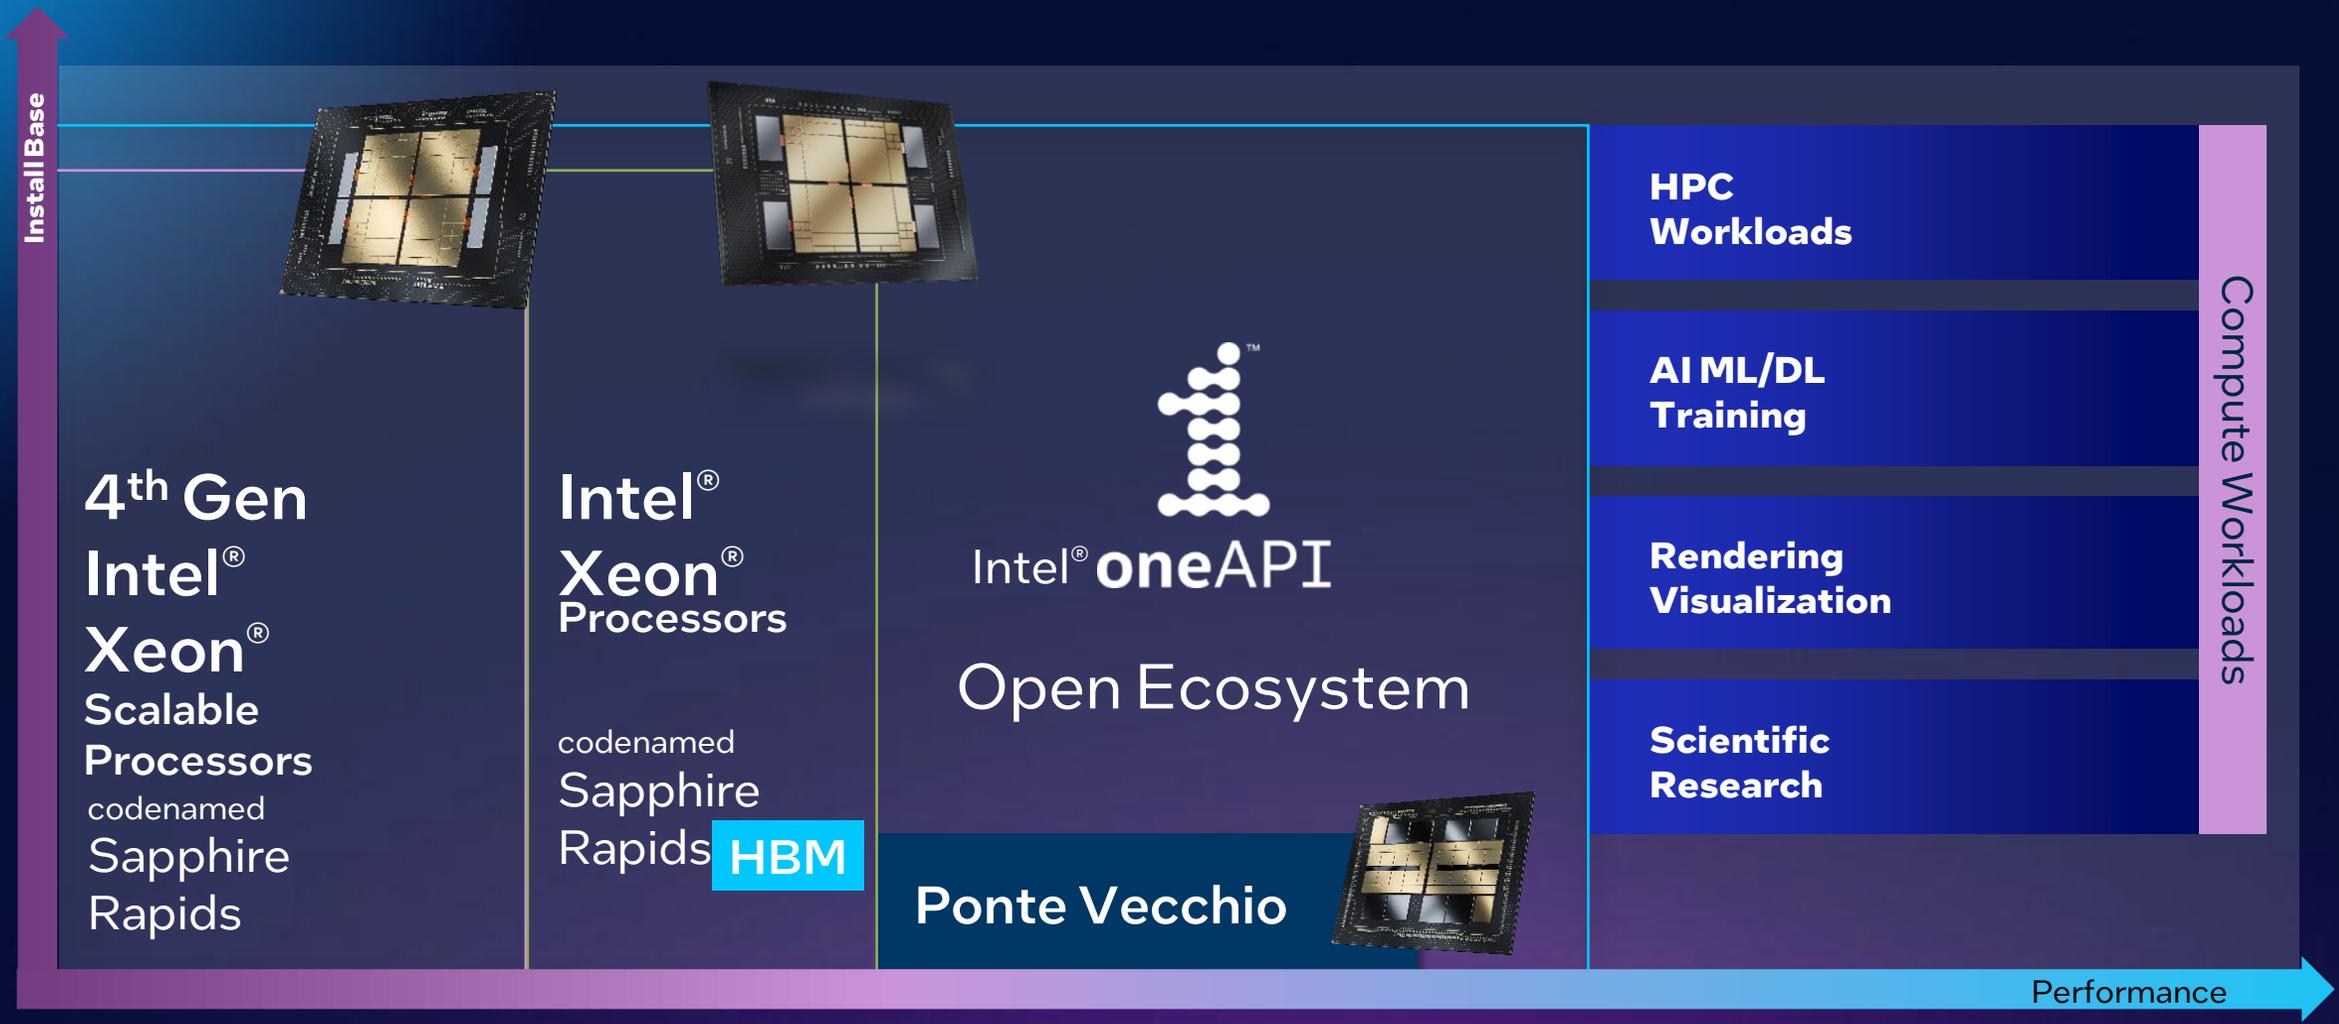


AI Training and Inference Opportunity

CPU and GPU



HPC - AI Super Compute Strategy



Annual Refresh Cycle



Ponte Vecchio

4th Gen Xeon HBM

Arctic Sound-M

4th Gen Intel Xeon[®] processors

2022



Ponte Vecchio Next

Xeon HBM Next

Arctic Sound Next

Xeon Emerald Rapids

2023

Falcon Shores

XPU

>5x

Memory Capacity & B/W
Compute density in x86 socket
Performance/Watt

New Tile Based Flexible & Scalable Architecture

Scalable Architecture for all Super Computing Workloads



+



2024

Compute

Up to
128 Ray
Tracing Units

Highest
Compute Density
socket & node

Up to
128 Xe
Cores



Memory

Up to
64MB
L1 cache

Up to
408MB
L2 Cache

Up to
128GB
HBM2e



I/O

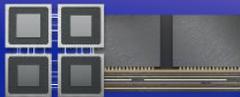
Up to **8**
Fully Connected
GPUs

PCIe
Gen 5

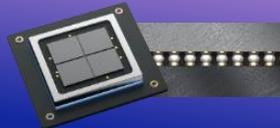
Xe Link
High-Speed
Coherent
Unified GPU Fabric



Technology



EMIB



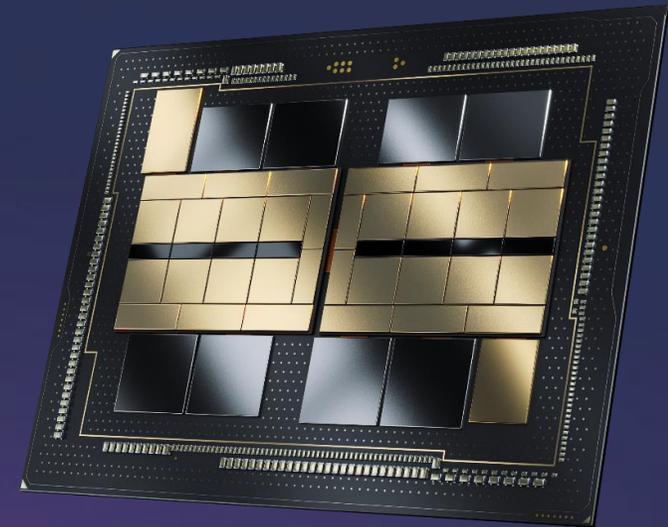
Foveros

Intel 7
TSMC N5
TSMC N7

Ponte Vecchio

Xe HPC based GPU

Up to 2.6x¹ Perf
over best in market today



On Track
for Aurora 2 Exaflop Supercomputer²

¹Based on pre-production measurements vs A100.
Learn more at www.intel.com/PerformanceIndex. Results may vary.
²>2 exaflop peak performance

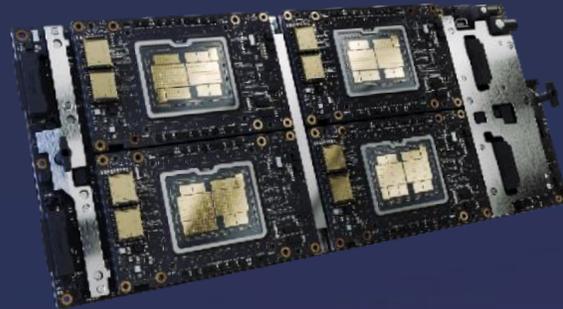
Ponte Vecchio GPU Boards & Systems

GPU BOARDS Intel Branded



PCIe Add In Cards

OAM SUBSYSTEMS



[x4 GPU Subsystem](#)

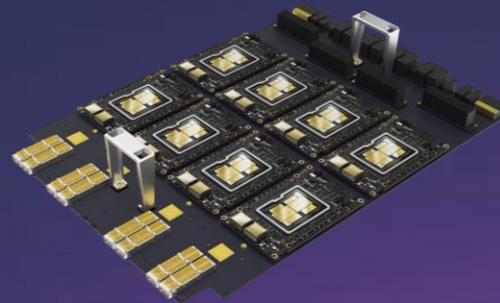
FULL SYSTEMS



[1U 4 GPU Server](#)



OAM Modules



[x8 GPU Subsystem](#)



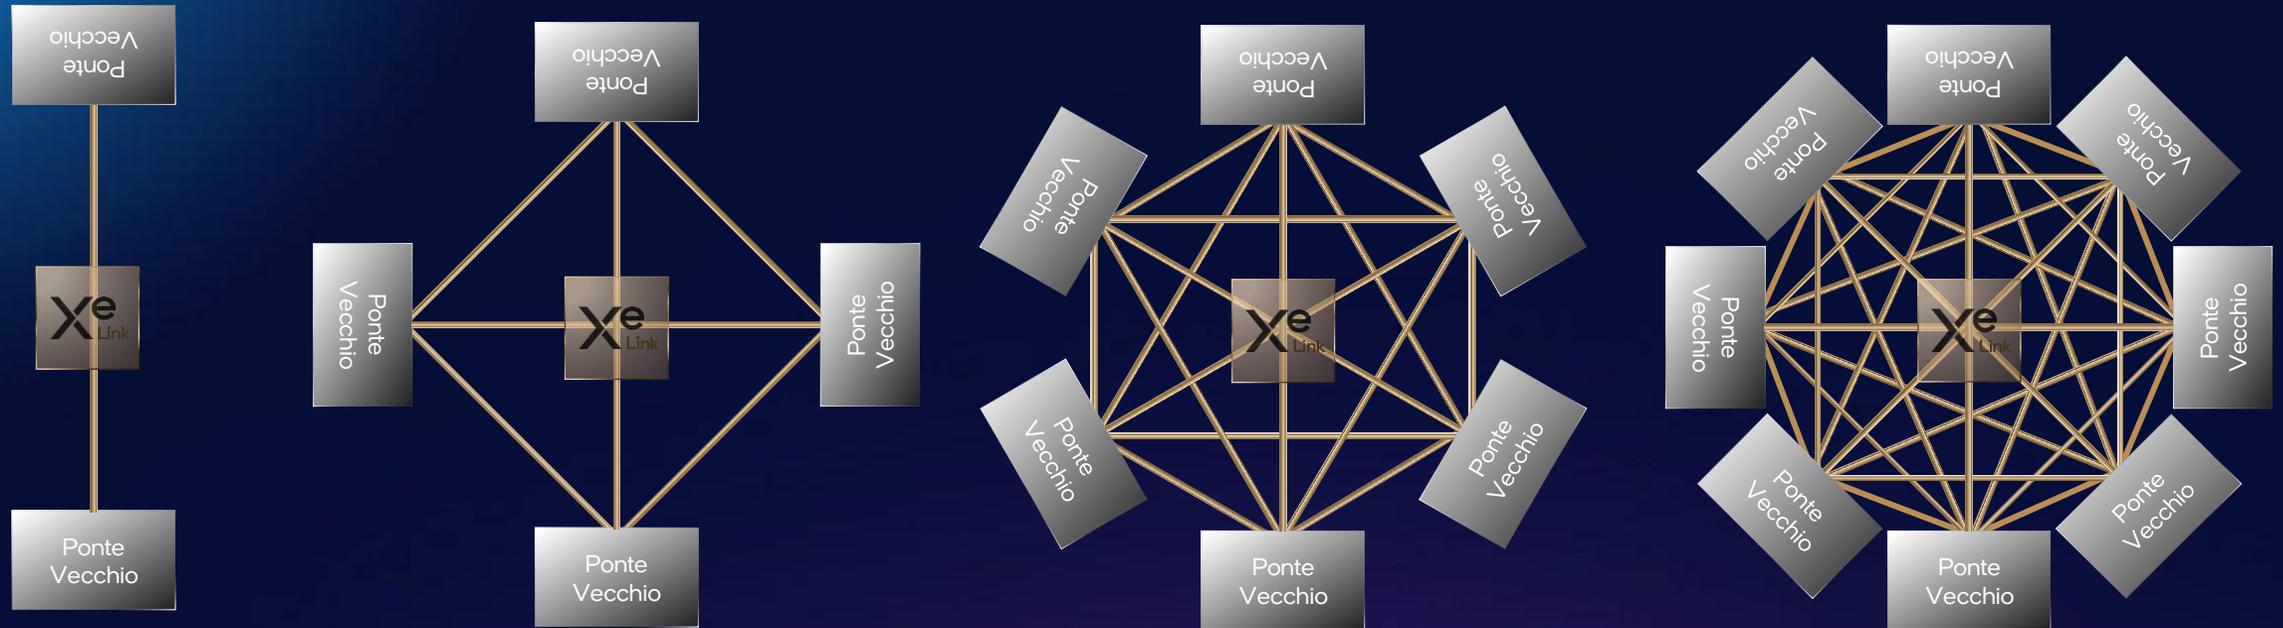
[4U-7U GPU Server](#)

Compute Accelerator Market Segmentation by GPU

	Segment	Number of GPUs
	Exascale	> 10K
	Exascale Follow On	Up to 10K
	Hyperscalers CSP	Up to 10K
	Large Enterprise and Next Wave CSPs	Up to 1K
	Enterprise AI and HPC	Up to 100

Xe Link for Scalability

Enabling a high number of coherent and unified accelerators



Flexibility for Scale Up and Scale Out across GPUs and Nodes

Note:
Xe Links connect gluelessly between PVCs
The Xe Link Logo above is not representing an additional System Device



Open, Standards-Based Unified Software Stack

Freedom from proprietary programming
models

Full performance from the hardware

Piece of mind for developers

CPU & XPU - Optimized Stack

Applications & Services

Middleware, Frameworks & Runtimes

TensorFlow PyTorch mxnet learn NumPy dmlc XGBoost openVIN

Low-level Libraries

oneMKL oneDNN oneDAL oneVPL
 oneTBB oneCCL oneDPL Other Libraries

Languages

DPC++
 Other Languages

Hardware Abstraction Layer Level Zero

Compute Hardware



CPU



GPU

Monte
Carlo

Black
Scholes

Binomial

Leadership Performance on X^e-HPC

intel[®]

Ponte Vecchio Performance

Relative Perf. Higher is better



Research
Scientists

Students

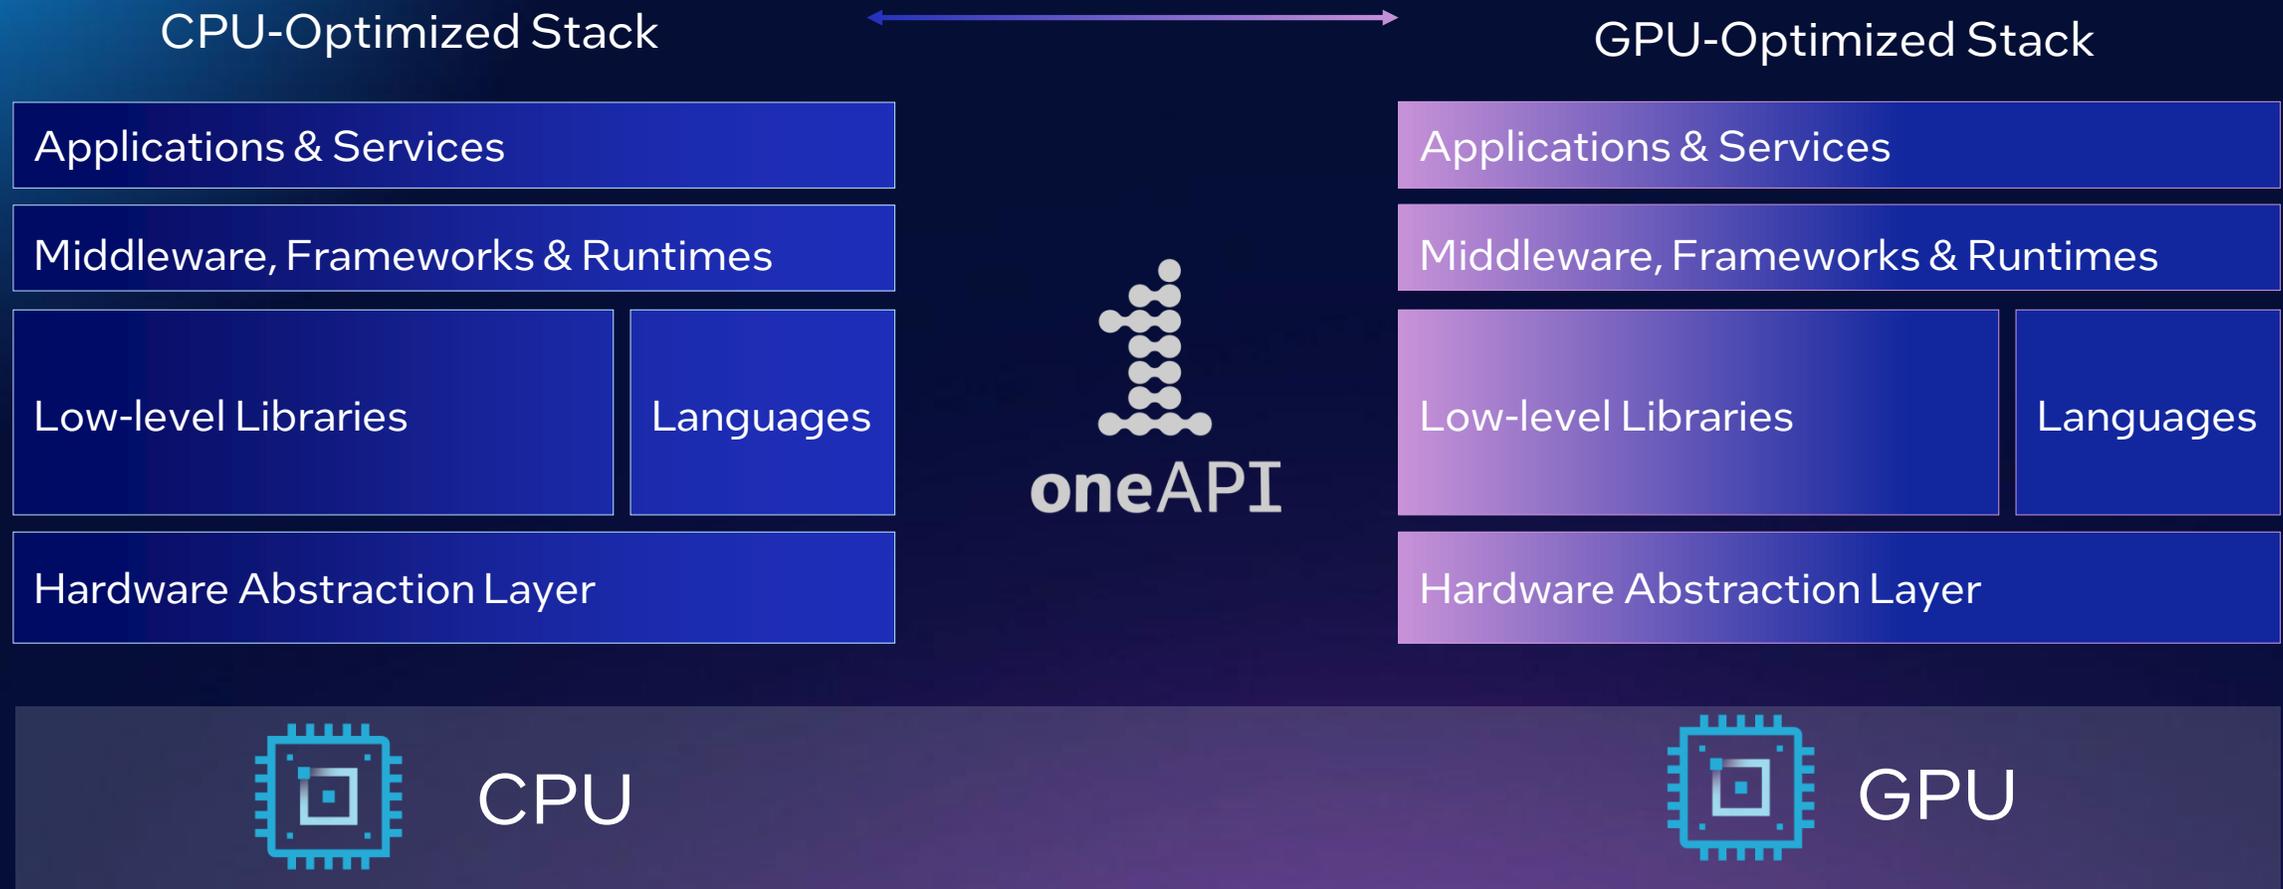
SW Developers

Data Engineers

Win Developers for HPC/AI Everywhere

 Agriculture	 Energy	 Education	 Government
 Finance	 Telecom	 Transport	 Smart Home
 Media	 Health	 Industrial	 Retail
 Academia	 Cloud/DC	 Edge	 Client

Overcoming Separate CPU and GPU Software Stacks with oneAPI





intel.
VISION
venture on

A decorative graphic in the bottom-left corner consisting of a large blue square partially overlapping a smaller white square.

Thank you