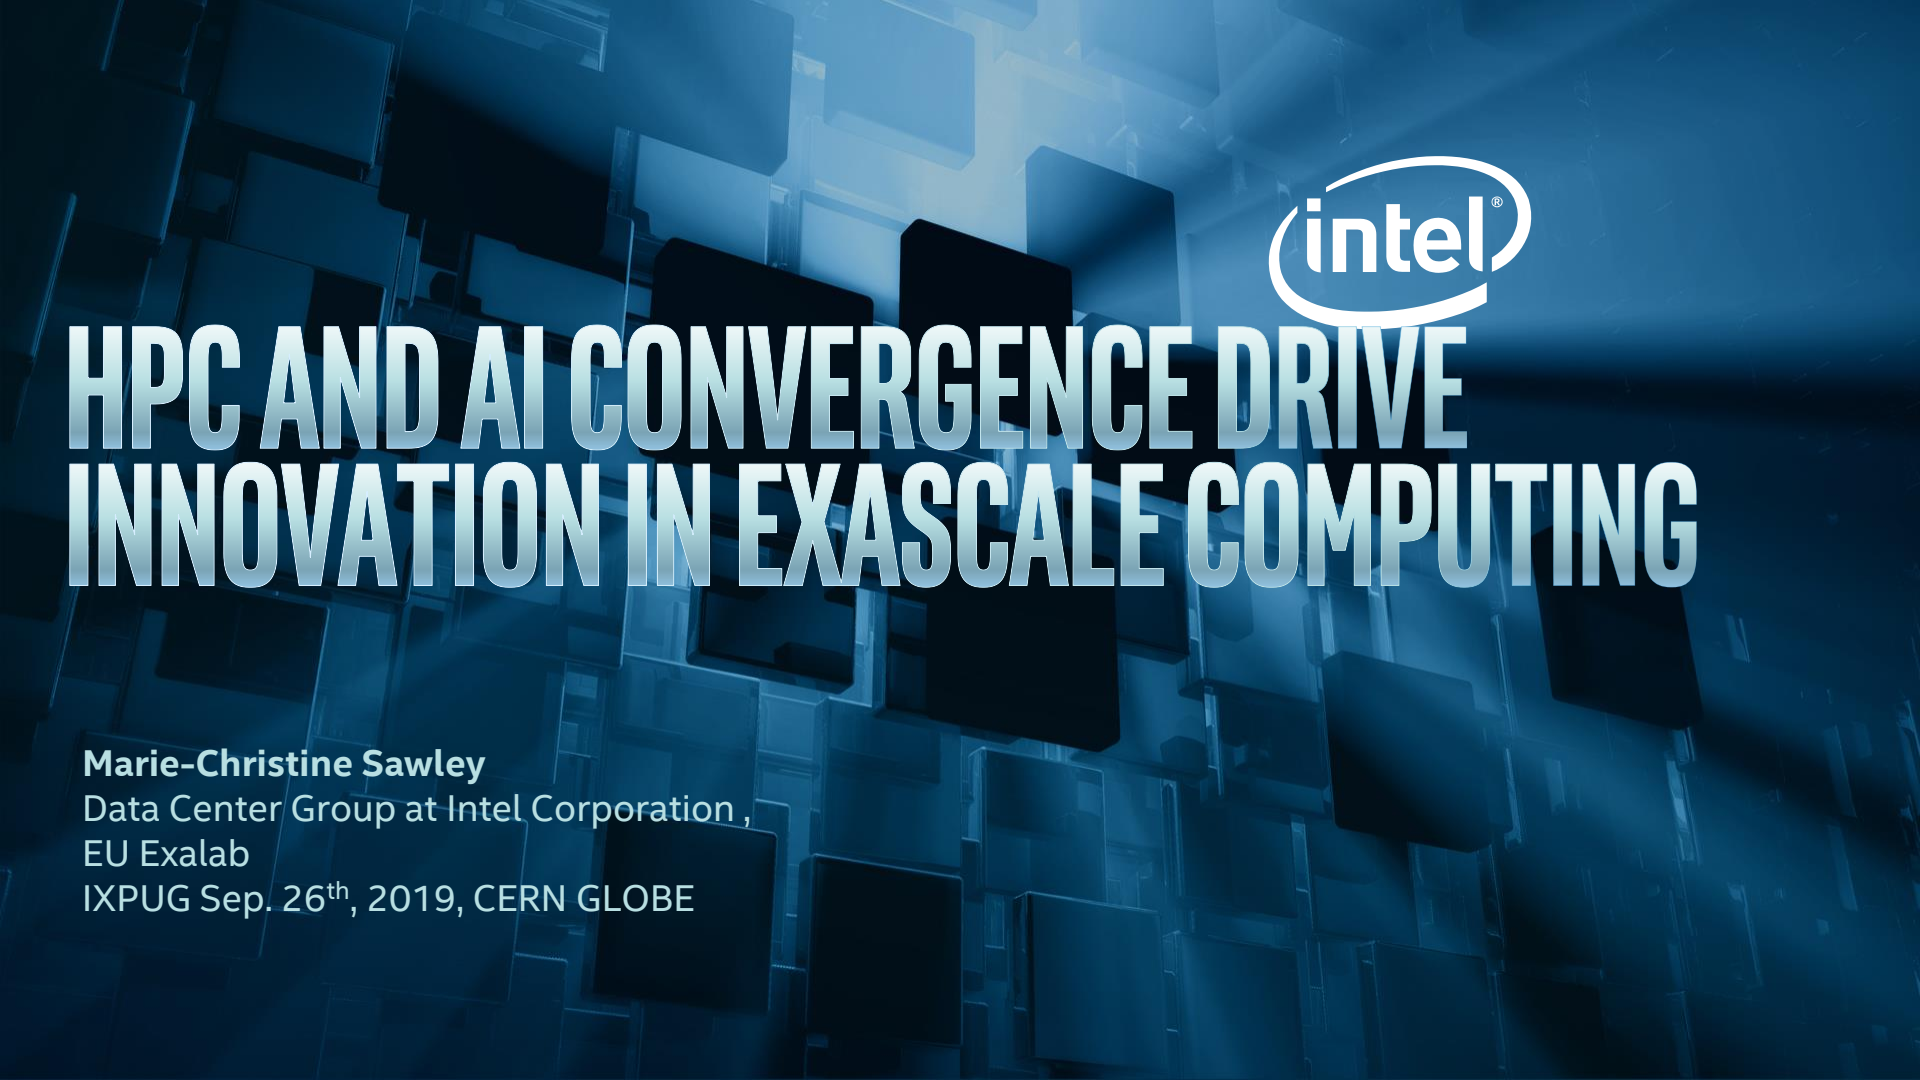

HPC AND AI CONVERGENCE DRIVE INNOVATION IN EXASCALE COMPUTING

Marie-Christine Sawley
Data Center Group at Intel Corporation ,
EU Exalab
IXPUG Sep. 26th, 2019, CERN GLOBE

TODAY'S OBSERVATIONS

- Convergence of AI and HPC: is it really taking place?
- 3 exemples of Extraordinary WL needing Exascale
- Technologies that will get us there

INTERTWINED FUTURES: OPPORTUNITIES AND CHALLENGES

AI DATASETS GROWING TO EXASCALE

- Added Speed and Operability

HIGH-PERFORMANCE DATA ANALYTICS

- Added Flexibility and Efficiency

BREAKTHROUGH COMBINATION

- Unveiling New Technology Opportunities
- Exploring New Technology Challenges

MARKET TRENDS ARE CREATING AN INFLECTION POINT FOR ARTIFICIAL INTELLIGENCE

BREAKTHROUGHS IN DATA SCIENCE
Graph Search -> Neural Nets -> Machine Learning -> Deep Learning

EXPONENTIAL GROWTH OF TRAINING DATA
Tabular -> Structured -> Unstructured | 50x Growth by 2020[1] (Google)

INNOVATION IN COMPUTING
Integrated Circuits -> GPUs -> FPGAs -> NNP

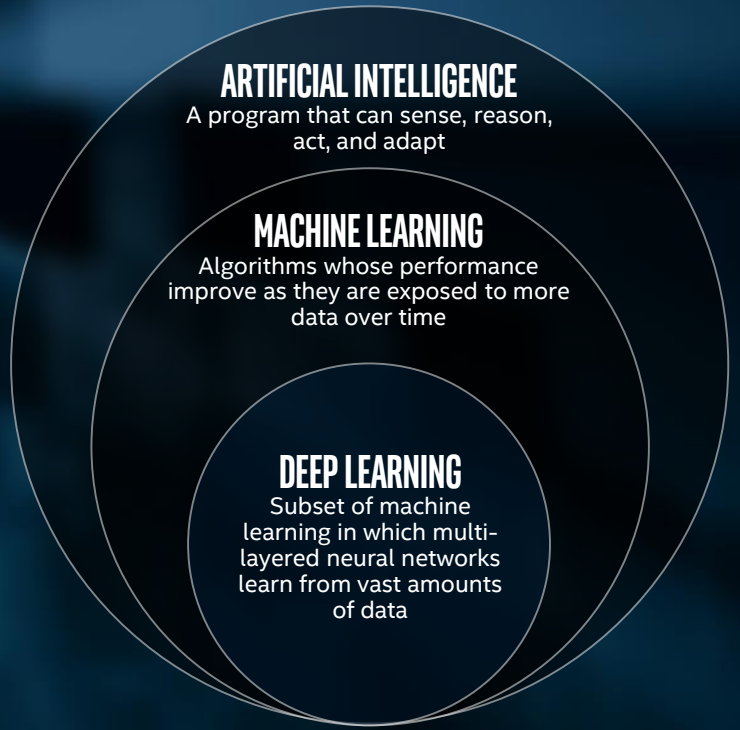

1. Eric Schmidt, Former CEO at Google. https://www.theatlantic.com/technology/archive/2010/07/quote-of-the-day-google-ceo-compares-data-across-millennia/344989/

3 EXEMPLES OF CHALLENGES WAITING FOR EXA

A Roadmap to the Implementation of 1km Earth System Model Ensembles

«...Within the ESiWACE project, we have demonstrated that 1 km simulations with global atmosphere models are technically feasible (see ESiWACE Deliverable D2.1). However, we are still far from being able to run a 1 km coupled ensemble in an operational mode even if some of the fastest sup ercomputers that are currently available are used (Neumann et al., 2019, Fuhrer et al., 2018, Schulth ess et al. 2019, Dueben et al. 2019). "

www.esiwace.eu

CHEESE

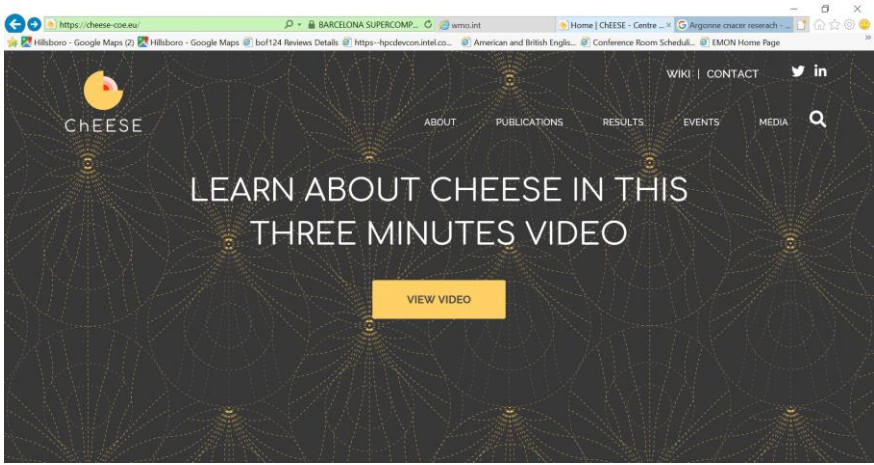

Exascale Pilot demonstrators

“The objectives …. are to provide geophysical simulations with a potential usefulness for management of geological resources, civil protection or insurance, among others. Reduce the gap between HPC algorithms capabilities and end-users needs (response time, throughput, etc.), provide an easy access to HPC codes for SE Applications and build PD as an end-to-end solution that can be easily compared to the state of the art in the fields of hazard assessment, civil protection or fundamental R&D”

https://cheese-coe.eu/

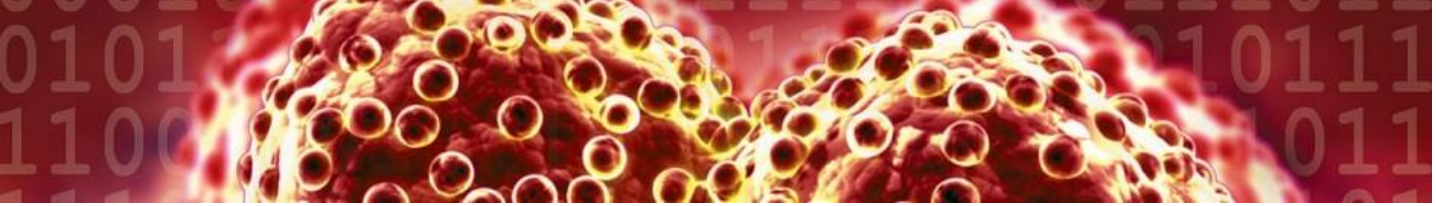

CANDLE
Exascale Deep Learning and Simulation Enabled Precision Medicine for Cancer

.....In the drug response problem, we use supervised machine e learning methods to capture the complex, non-linear relationships between the properties of drugs and the properties of the tumors to predict response to treatment and therefore develop a model that can provide treatment recommendations for a given tumor. The scale in this problem derives from the number of relevant parameters to describe properties of a drug or compound (O(106)), number of measurements of important tumor molecular characteristics (O(107), and the number of drug/tumor screening results (O(107)).....

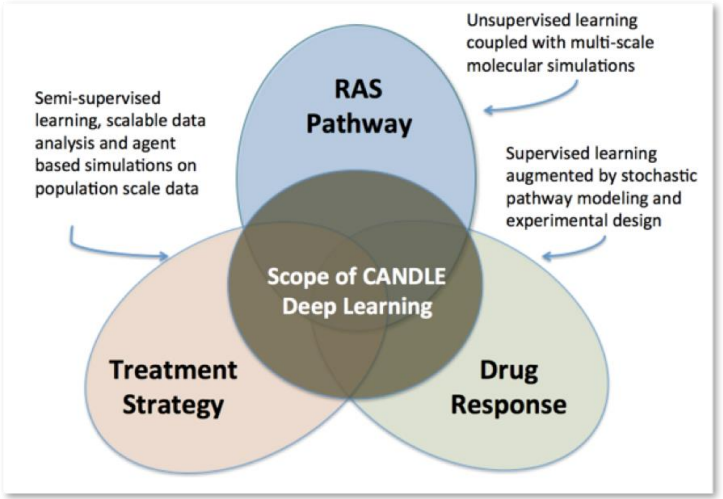

https://candle.cels.anl.gov/

DRIVING FORCES

CONVERGENCE: CHANGING THE SYSTEM PARADIGM

ARCHITECTURE FLEXIBILITY

HIERARCHICAL, DYNAMICALLY RE-CONFIGURABLE, COMPUTE DIVERSITY

EXPANDED MEMORY STORAGE HIERARCHY

PERFORMANCE, CAPACITY, CAPABILITY

SOFTWARE ABSTRACTION

ARCHITECTURE AGNOSTIC, READABLE, MAINTAINABLE, PORTABLE

HETERO. WORKLOADS: SIMULATION, DATA ANALYTICS AND AI

Exascale machines must treat all as first class citizens

- This is both a hardware and a software requirement

AI has become a major consumer of computing cycles and it is expected to grow

- Compute deployment both at edge and in large cloud
- Will drive economies in fabric, compute with a large focus on power and perf/W

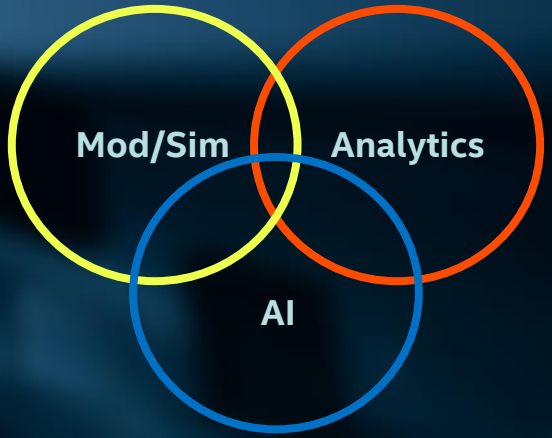

NAVIGATING THE ARCHITECTURE SPACE

Multiple offerings covering similar problem domains

HPC is truly the birthplace of many architectures… and the testing ground

The dimensions that we are exploring now to get exceptional performance

- Data level parallel (from fine grain to coarse grain)
- Energy efficient accelerators. (compute density and energy efficiency often are correlated)
- Exploiting predictable execution at all levels. (cache to coarse grain)
- Integrated fixed function data flow accelerators.
- General purpose data flow accelerators

intel AI HARDWARE

MOVE FASTER

INTEL® SILICON PHOTONICS

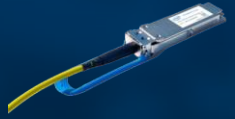

INTEL® ETHERNET

INTEL® OMNI-PATH FABRIC

STORE MORE

intel OPTANE™ DC SOLID STATE DRIVE

intel OPTANE™ DC PERSISTENT MEMORY

PROCESS EVERYTHING

CPU

intel ATOM inside

intel XEON PLATINUM inside

intel CORE i7 8th Gen

AI ACCELERATORS

2H'19 intel NERVANA inside

intel MOVIDIUS inside

FPGA, GPU

intel STRATIX 10 inside

FUTURE GPU (Integrated & Discrete)

POWERING THE FUTURE OF COMPUTE AND COMMUNICATIONS

EXASCALE AND SYSTEMS THINKING

NEW STORAGE STACK

EXASCALE STORAGE STACK

- Designed from the ground up for NVM storage
  - Built over **new userspace** NVMe/pmem software stack
  - High **throughput/IOPS** @**arbitrary** alignment/size
  - **Ultra-fine grained** I/O
- Manage **massively** distributed NVM storage
  - **Scalable** communications & I/Os over homogenous, shared-nothing servers
  - **Software**-managed **redundancy**
    - declustered replication & erasure code with self healing

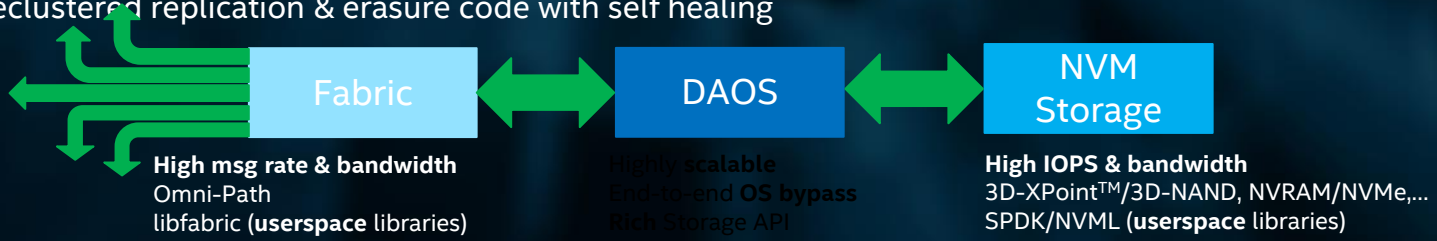

FUTURE ARCHITECTURAL DIRECTION

KEY ATTRIBUTE OF SYSTEMS IN THE FUTURE WILL BE HIGHLY CONFIGURABLE AND NEED TO BE ABLE TO SUPPORT UPGRADES TO FUNDAMENTALLY NEW TECHNOLOGIES.

- Many new technologies on the horizon. Want to be able to accommodate when they are available.
- Already see need for larger degree of configurability in systems based on workload focus.
- Same architecture should cover HPC, ML and Data Analytics through configuration.
- Need a consistent software story across these different workloads.

Node Architecture View:

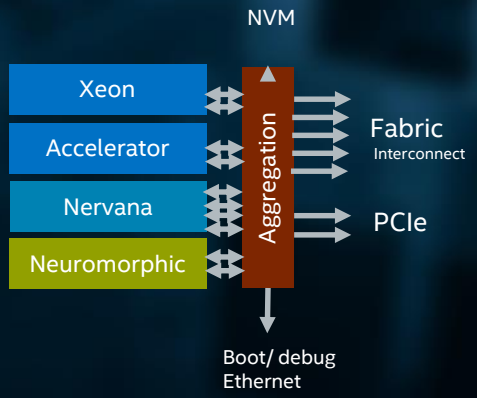

Should allow for tight integration of…

Intel® Xeon®
Intel® Nervana™
Neuromorphic
FPGA
3DXPT
3DNAND
custom

....MORE FEATURES TO WATCH AND TUNE

Change of workloads

- Simulation, when models exist
- ML, when models are lacking; or with large datasets; or to speed up DA of experimental data
- handling different numerical precisions

Focus on

- computing (CPU, accelerators) & stacked memory
- data management, integrated storage
- complex worflows

VISION: UNIFYING THE "3 PILLARS"

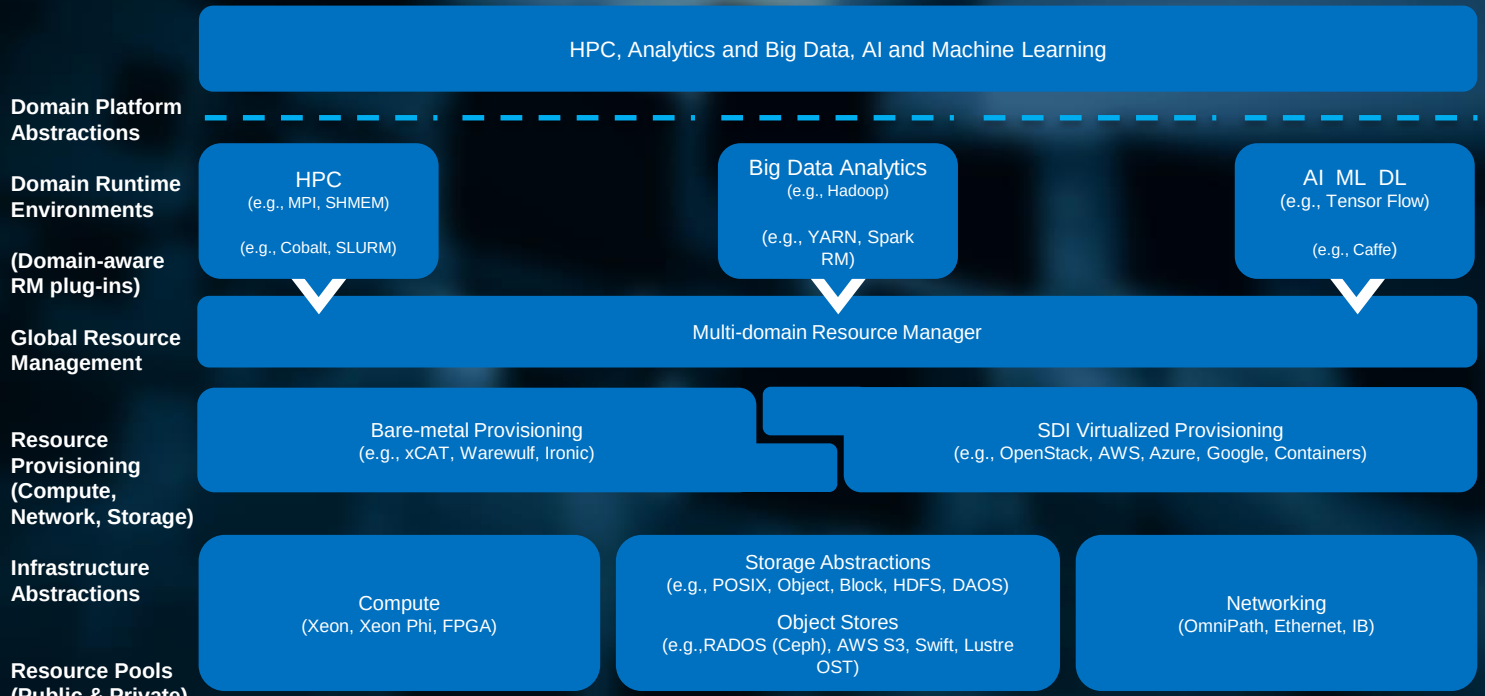

FOCUSED INVESTMENTS TO ACCELERATE HPC & AI

ADVANCED ARCHITECTURES

Intel Xeon Platinum inside | Xe | Intel Nervana inside | Intel Movidius inside | Intel Stratix inside

SCALAR | VECTOR | MATRIX | SPATIAL

+ Quantum & Neuromorphic

COMPUTE ARCHITECTURES
FOR ALL YOUR WORKLOADS

MANUFACTURING EXCELLENCE

PROCESS

FPGA

Intel AGILEX inside

10 NM in 2019

10th Gen Intel® Core™ Processor

10 NM in 2019

PACKAGING

EMIB + "FOVEROS"

ADVANCED PACKAGING
FOR HETEROGENEOUS INTEGRATION

SIMPLIFIED PROGRAMMING

OPEN STANDARDS

oneAPI
Services and Tools

openHPC

UNIFIED SINGLE SOFTWARE ABSTRACTION
AND DOMAIN-SPECIFIC LIBRARIES

TRANSFORMING MEMORY & STORAGE

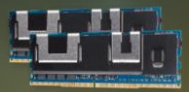

OPEN STANDARDS

DAOS

intel OPTANE DC PERSISTENT MEMORY

PMDK

RE-ARCHITECTING THE MEMORY
HIERARCHY AND FILE SYSTEMS

GROWING THE ECOSYSTEM AND ADVANCING HPC & AI THROUGH OPEN STANDARDS

intel

EXASCALE : WHAT THIS MEANS FOR USERS

Single unified stack with resource allocation and scheduling across all pillars and ability for frameworks and libraries to seamlessly compose

- Cohesive software stack, Unified control system

Minimize data movement: keep permanent data in the machine via distributed persistent memory and maintain availability requirements

- Enhance with data-aware scheduling

Support standard file-based IO with high performance and provide incremental path to a memory coupling model for highly efficient HPC, BDA, and AI interaction

- DAOS provides high performance file-based and object-based capabilities

Continue driving scalable high performance HPC simulation capability

- Example areas: MPICH, OFI, mOS, optimized libraries

CONCLUSION

- New and emerging technologies are enabling new building blocks from which we will architect future systems
- Future systems will enable both excellent legacy performance and a transitional path to exascale performance.
- User focus on computation per memory will pay big dividends across architectures.
- We have an exciting path to exascale which will mean multi-PF racks for broad HPC usage
- System Software and User applications will continue to have opportunity to stay on the exponential performance growth curve through exascale and beyond.
- Power of systems is an increasing challenge. More focus on TCO optimization.
- The convergence of AI, data analytics and traditional simulation will result in systems with broader capabilities and configurability as well as cross pollination.

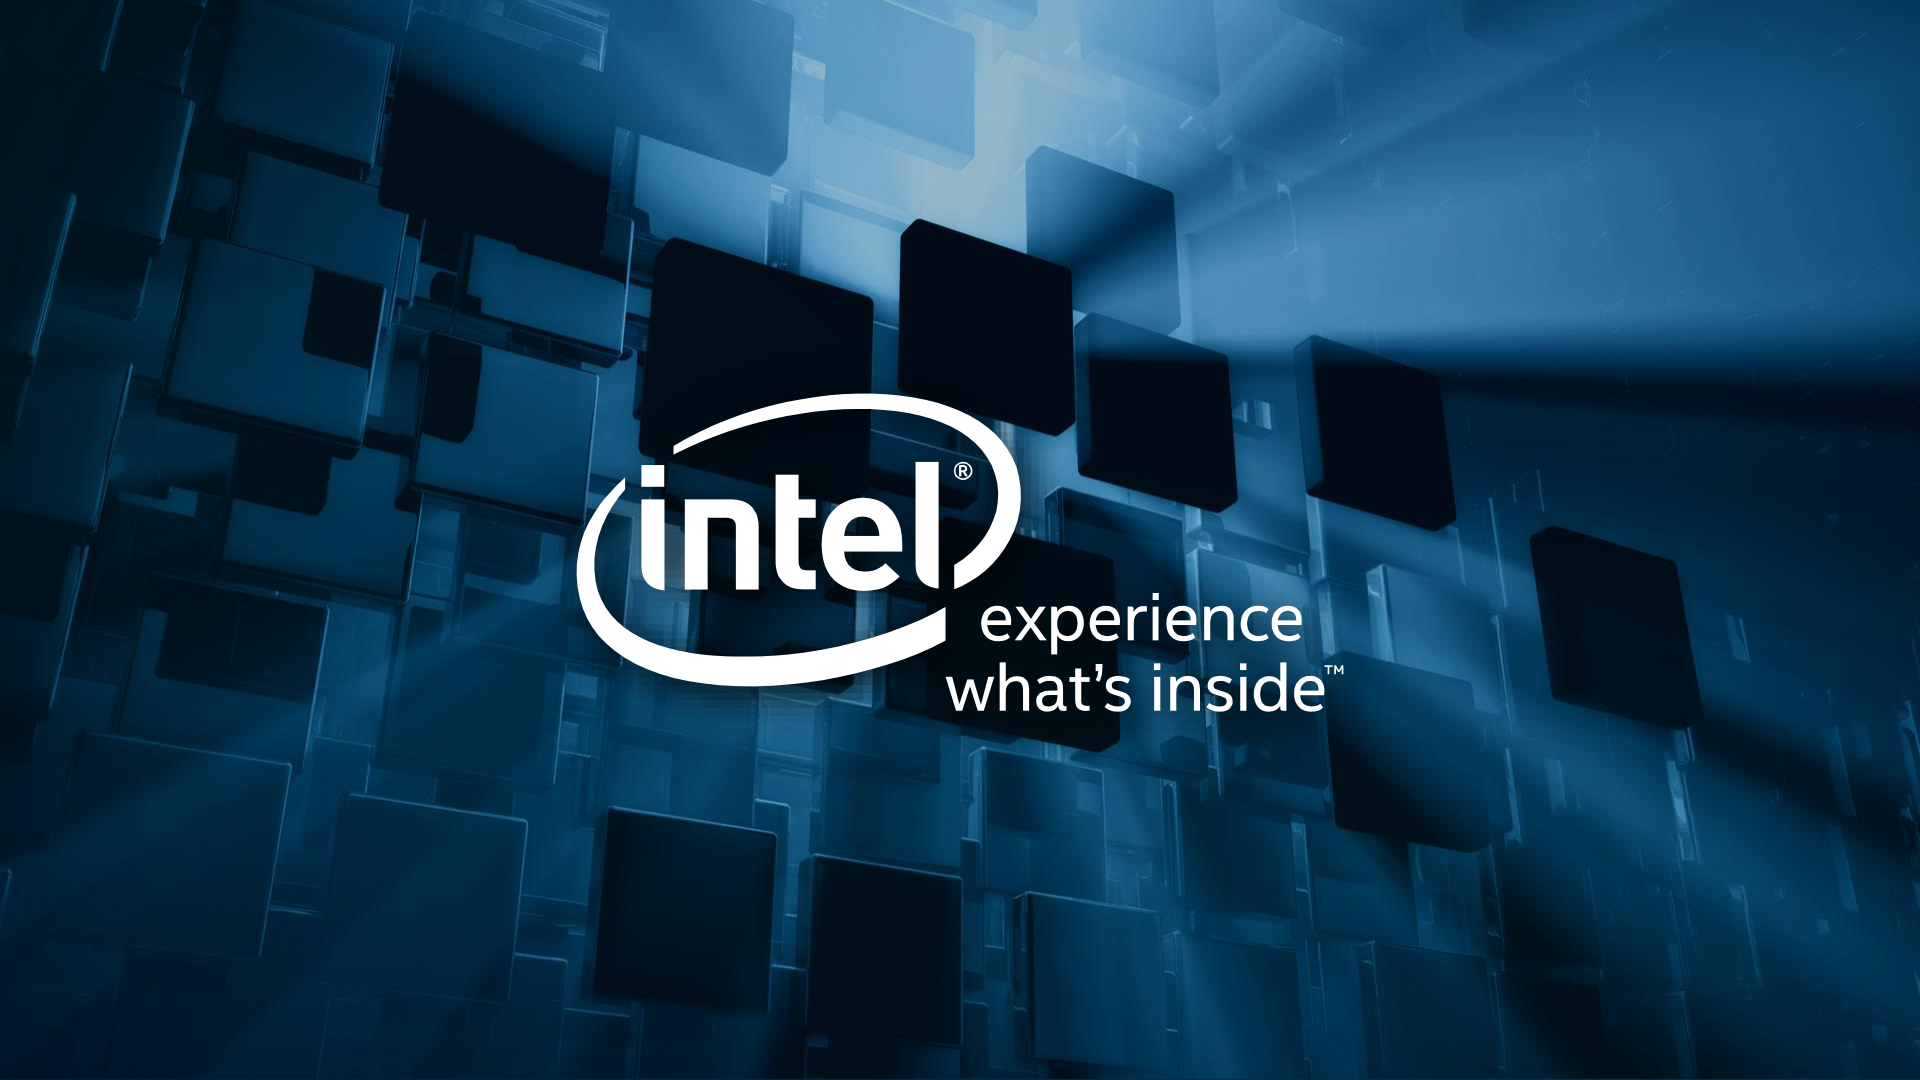
intel®
experience
what's inside™