Advanced Modeling & Simulation (AMS) Seminar Series
NASA Ames Research Center, February 28th, 2023

Vision of oneAPI:

Innovation through Abstraction

James Reinders and Vasanth Tovinkere
Intel



oneAPI Empowers Developers to Accelerate Performance Across Architectures



Open

Enable developer productivity on high-performance open platforms



Choice

Foster choice and interoperability of software platforms and ecosystems



Trust

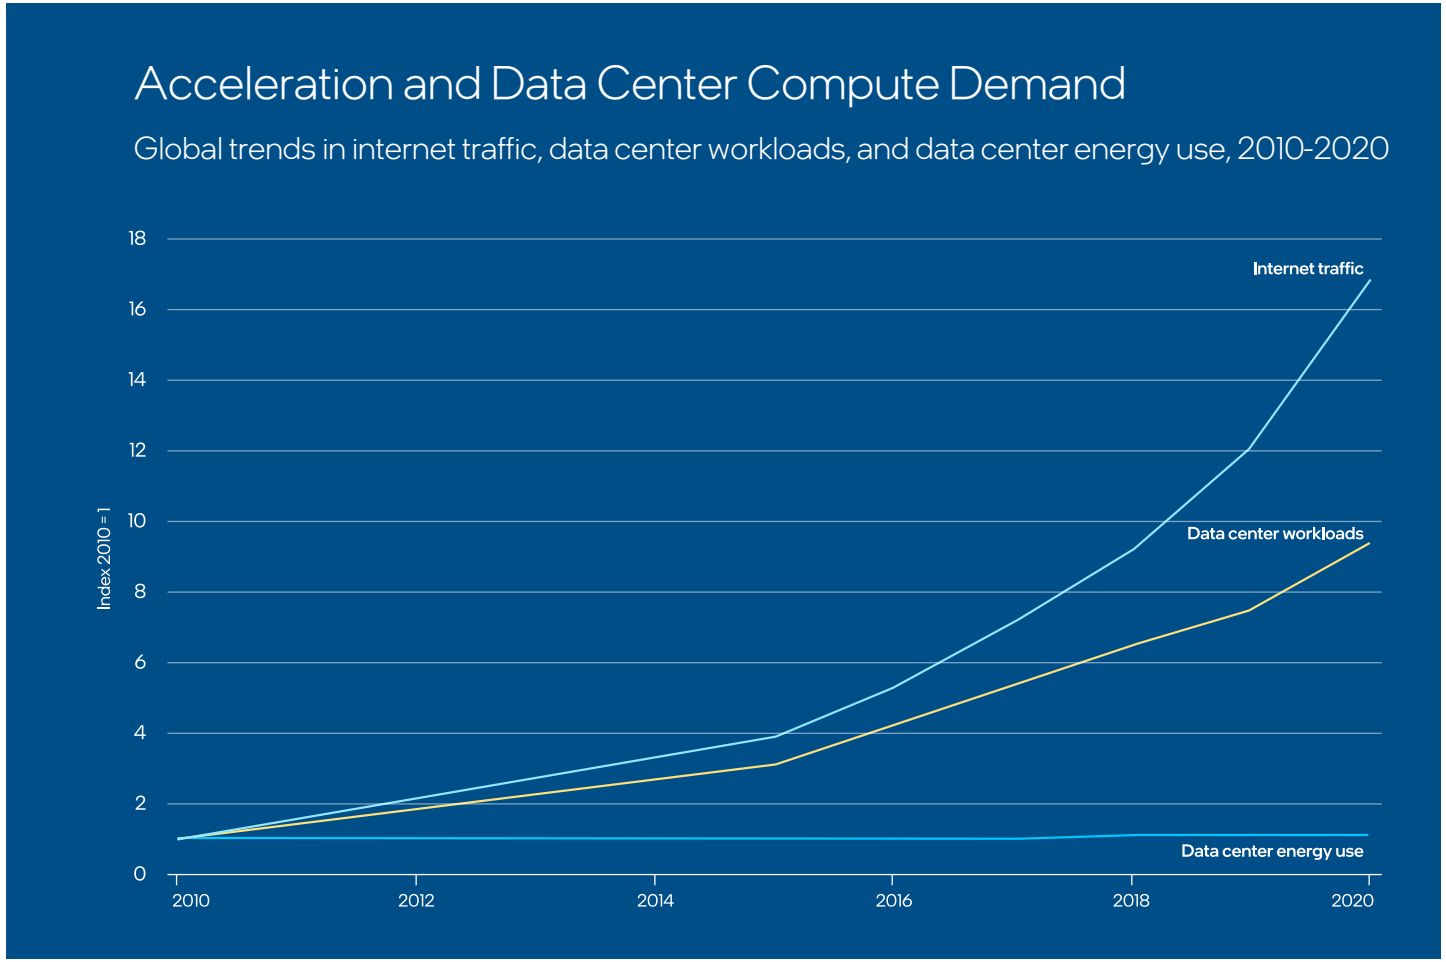
Built on a confidential computing platform you can trust

Three reasons oneAPI/open accelerated computing make sense for IHVs and SW developers:

1. Technologies need standards to scale from niche to volume
2. Developers, customers, and industry are looking for an open alternative
3. The oneAPI initiative provides developers the open, trustworthy path to hardware choice

©Intel Corporation

Data Center/Network Demand and Power

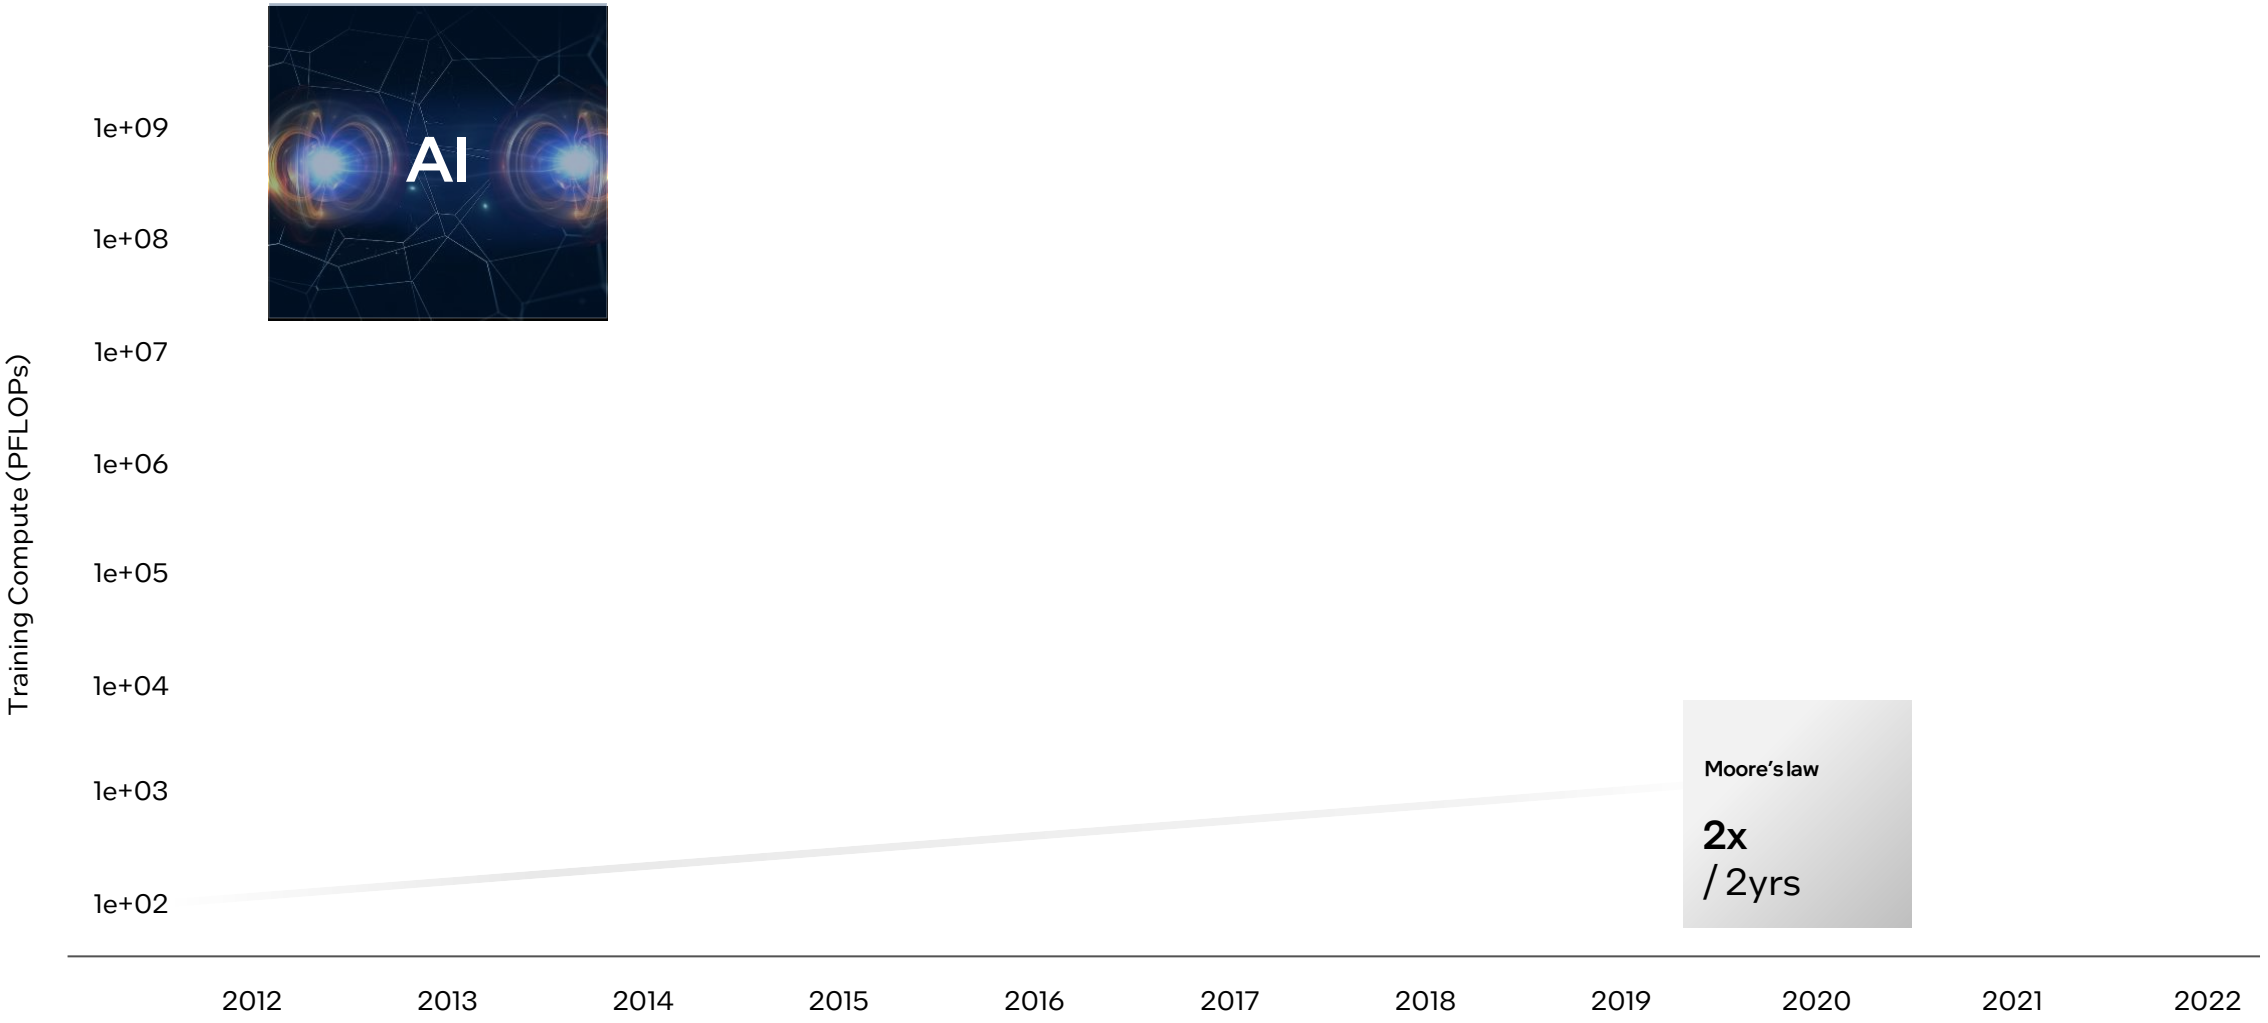


Source: International Energy Agency, November 2021. Data Centers and Data Transmission Networks, Lead Author George Kamiya. https://www.iea.org/reports/data-centres-and-data-transmission-networks

"Strong growth in demand for data center services continues to be mostly offset by ongoing efficiency improvements…"

©Intel Corporation

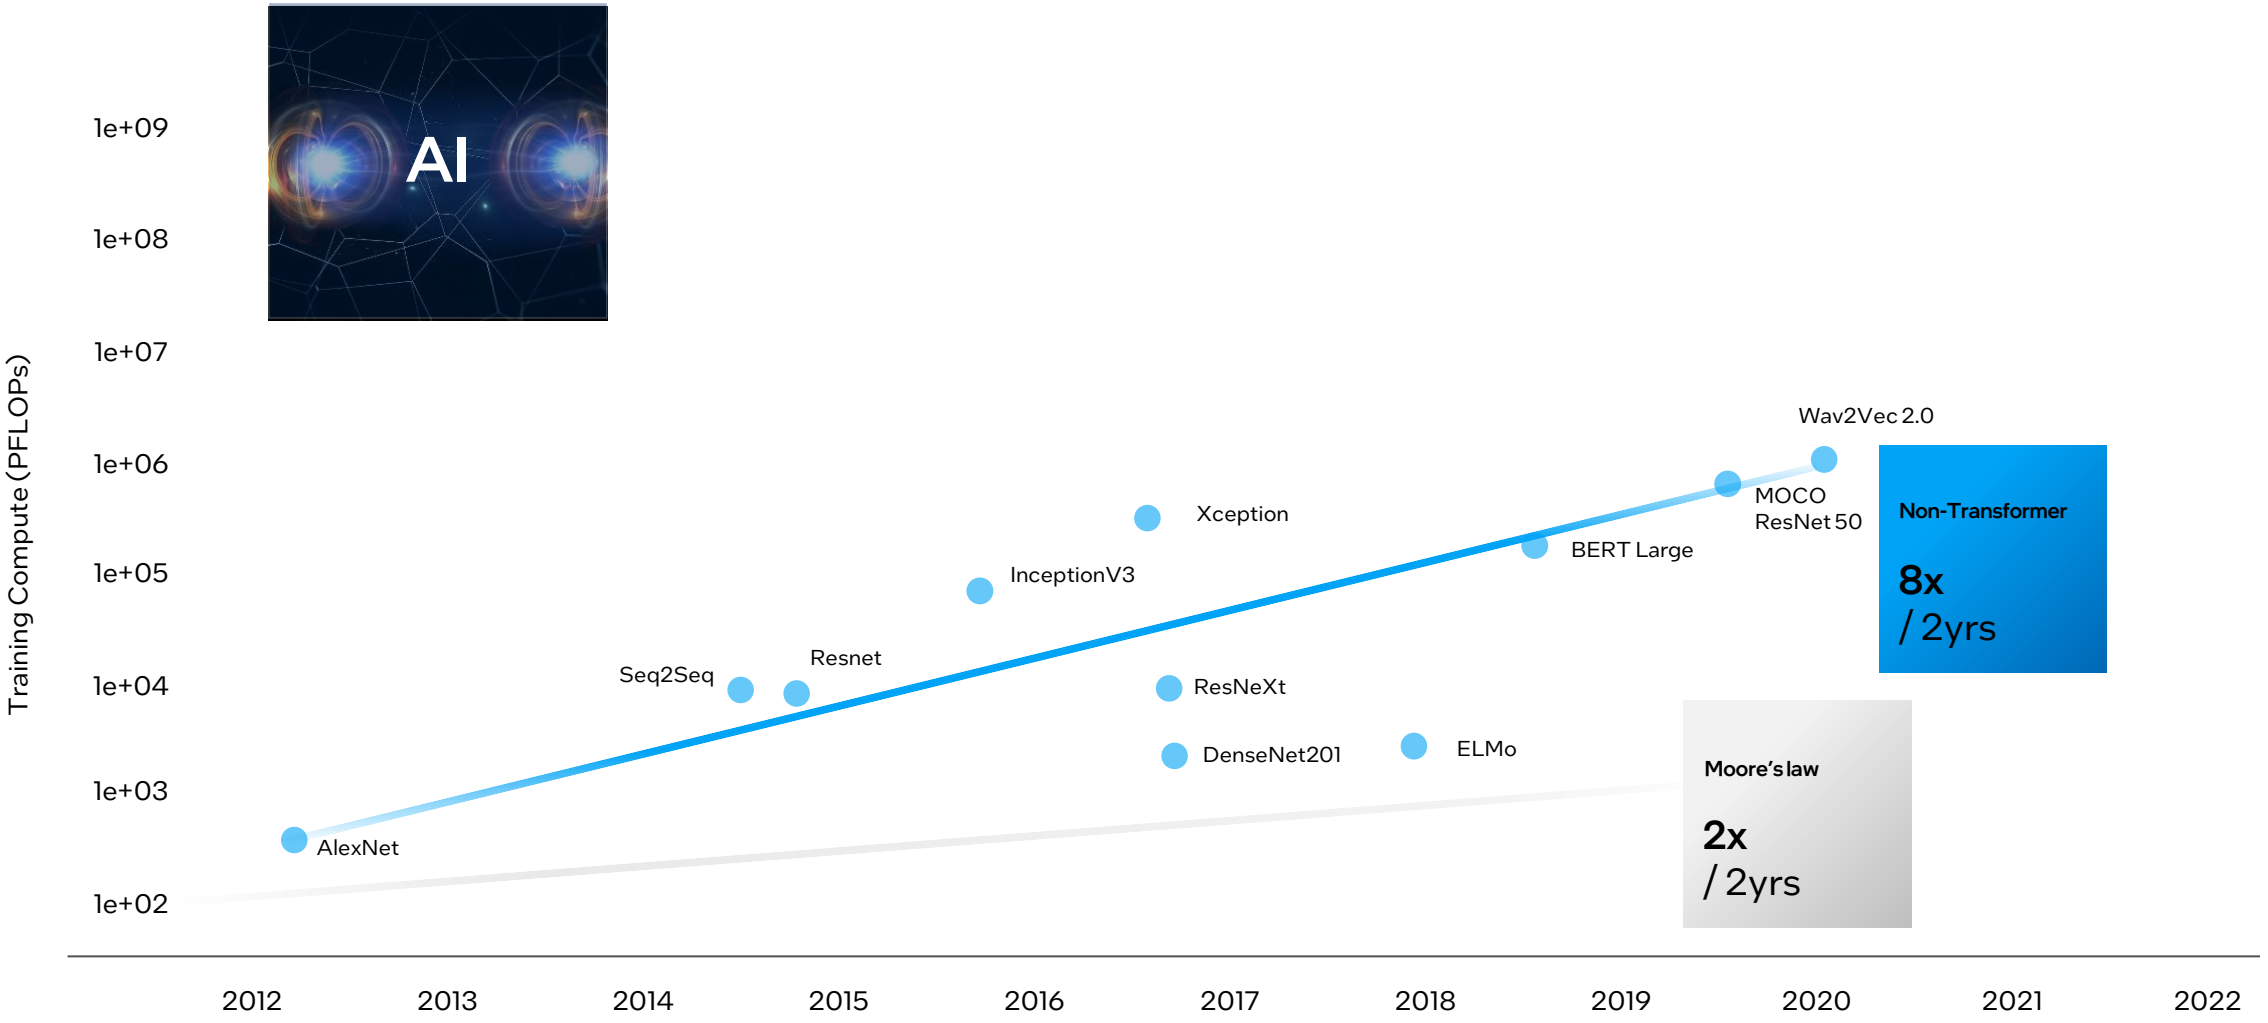
Exponential Compute Demand Growth



Source: Moore, S. (2022), "Nvidia's Next GPU Shows That Transformers Are Transforming AI", IEEE Spectrum.

©Intel Corporation

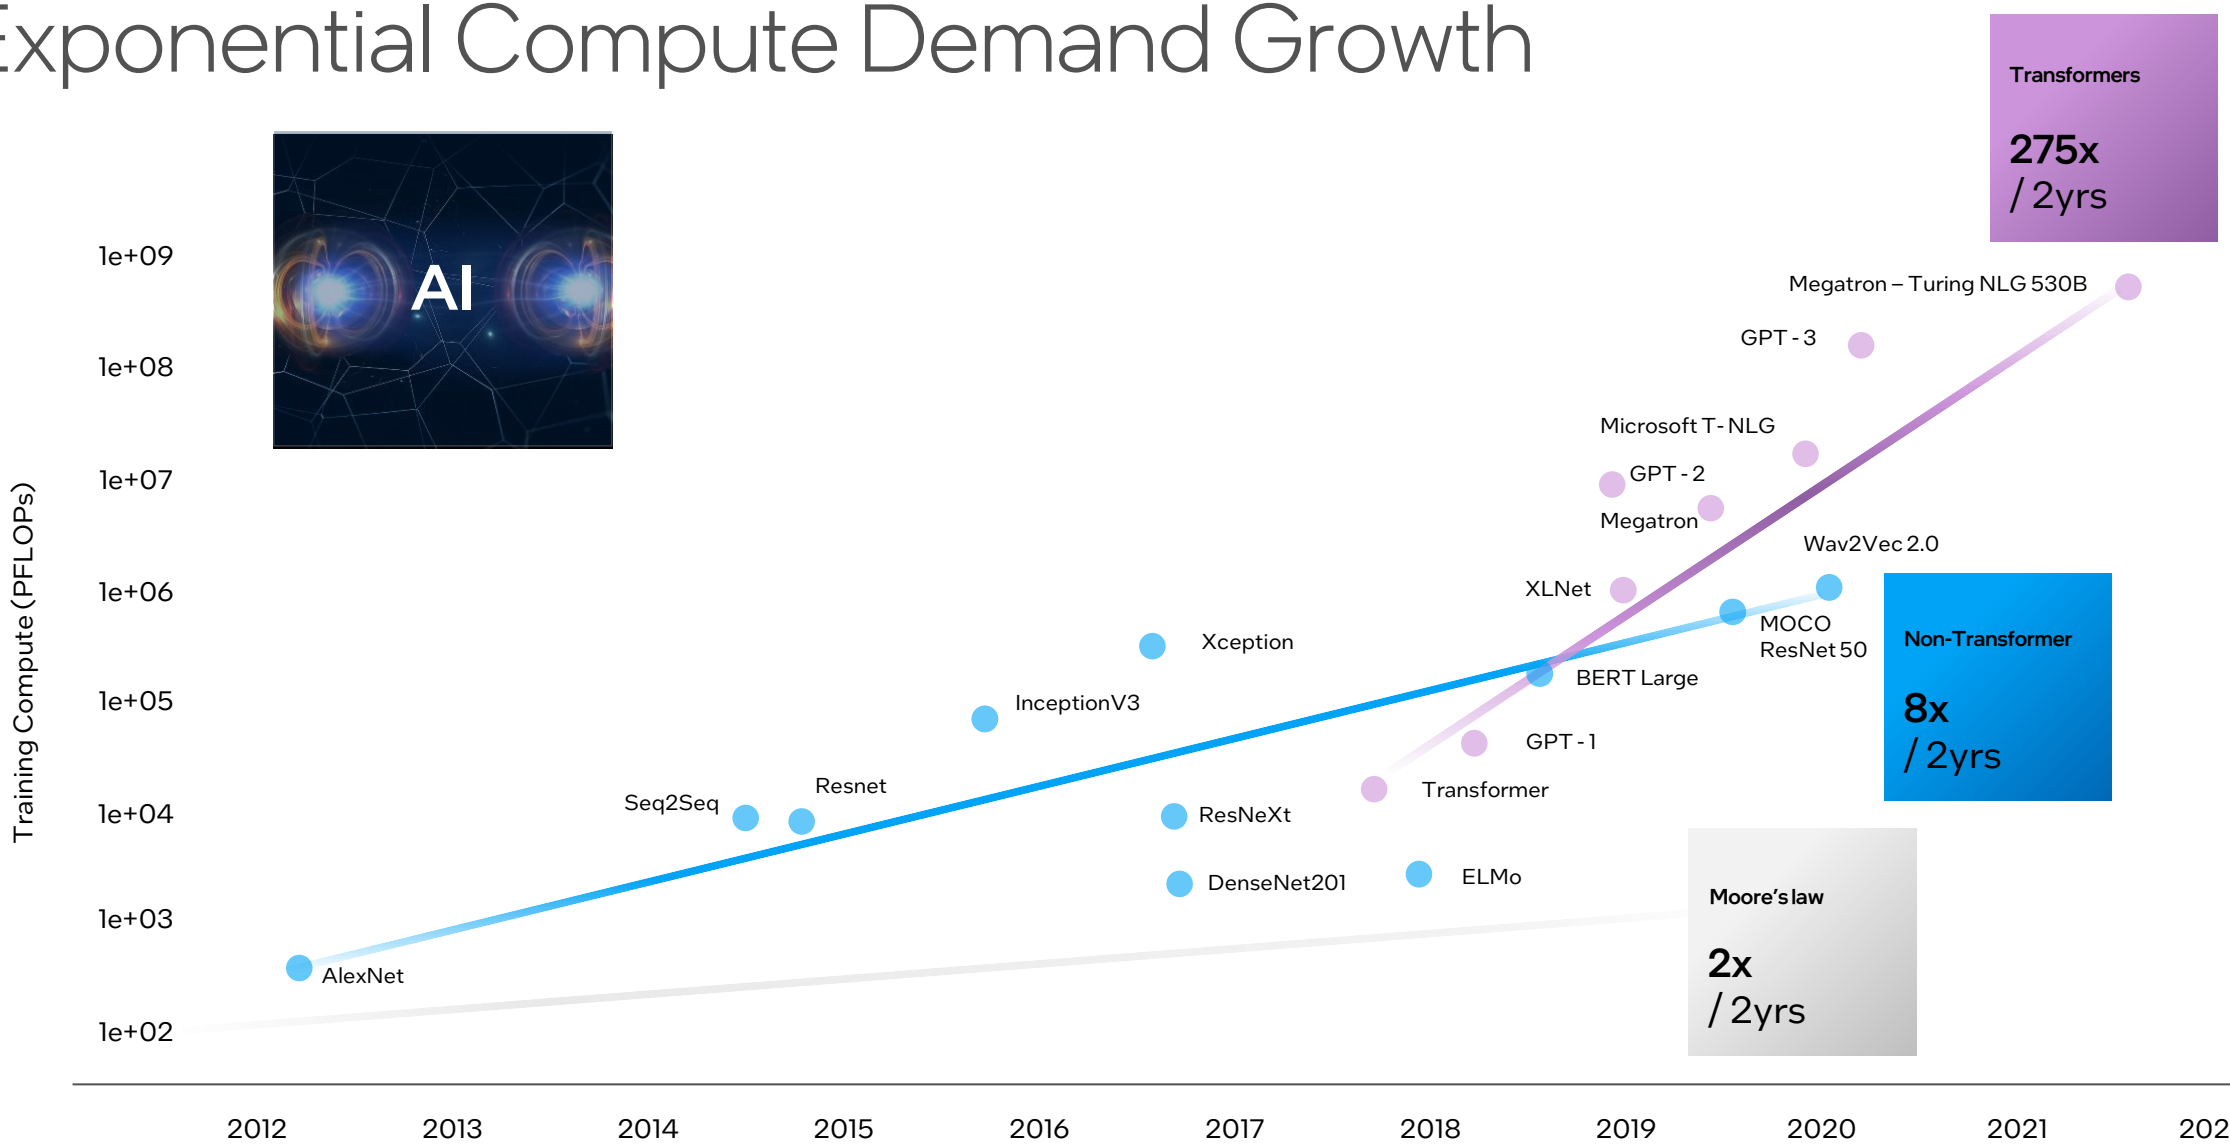
Exponential Compute Demand Growth



Source: Moore, S. (2022), "Nvidia's Next GPU Shows That Transformers Are Transforming AI", IEEE Spectrum.

©Intel Corporation

Exponential Compute Demand Growth



Source: Moore, S. (2022), "Nvidia's Next GPU Shows That Transformers Are Transforming AI", IEEE Spectrum.

©Intel Corporation

Cambrian Explosion of Computer Architecture

High-level, domain-specific languages and architectures, freeing architects from the chains of proprietary instruction sets, along with demand from the public for improved security, will usher in a new golden age for computer architecture.



Diverse and evolving workloads enable hardware innovation

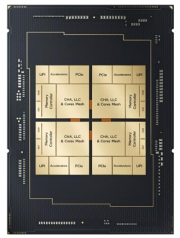
Source: John L. Hennessy, David A. Patterson, Communications of the ACM
https://cacm.acm.org/magazines/2019/2/234352-a-new-golden-age-for-computer-architecture/fulltext

© INTEL CORPORATION

©Intel Corporation

A New Golden Age for Computer Architecture

Products



Novel On-Die Accelerators



GPU/Data Parallel



Spatial/ Dataflow



Deep Learning Optimized

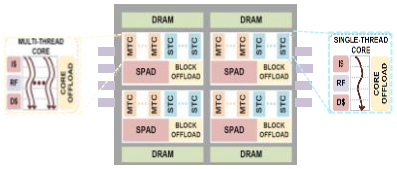


Blockchain

Research



Neuromorphic



Graph Analytics

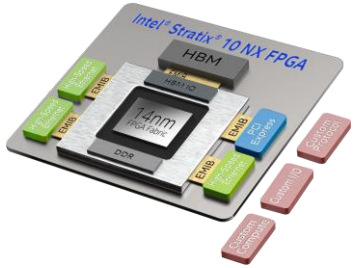
Source: International Energy Agency, November 2021. Data Centers and Data Transmission Networks, Lead Author George Kamiya. https://www.iea.org/reports/data-centres-and-data-transmission-networks

"Strong growth in demand for data center services continues to be mostly offset by ongoing efficiency improvements..."

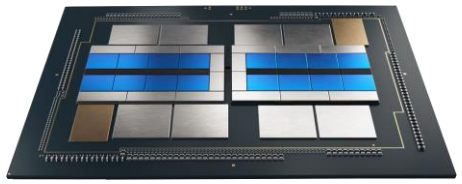
©Intel Corporation

Time to Market and Scale through Integration

Chiplet interconnects accelerated by standards

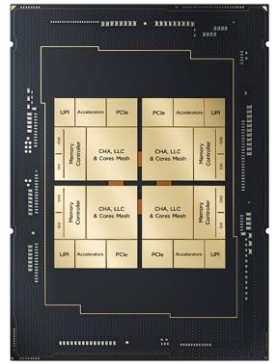


UCIe
Universal Chiplet
Interconnect Express



CXL
Compute
Express Link

On-die/on-package accelerators



"It may prove to be more economical to build large systems out of smaller functions, which are separately packaged and interconnected."*

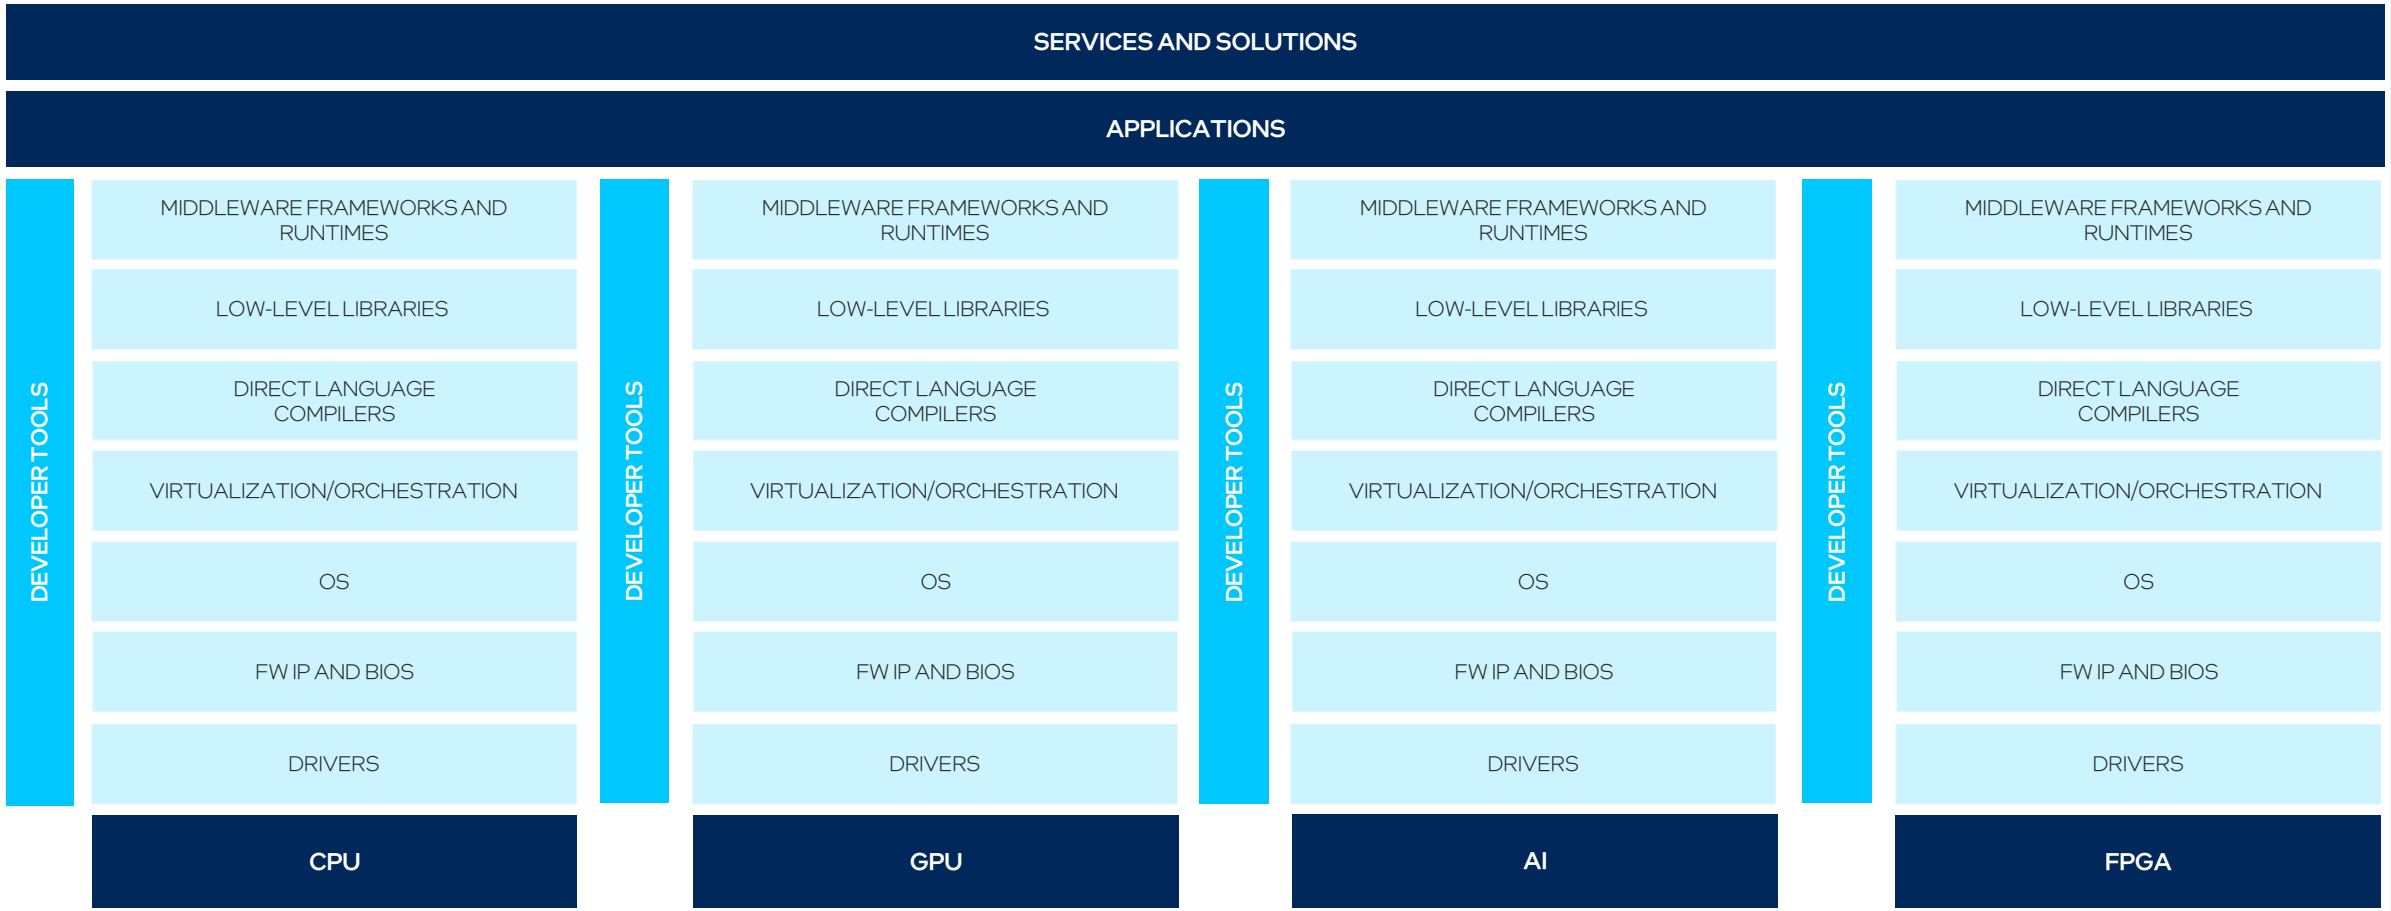
- Gordon Moore

*Source: "Cramming More Components onto Integrated Circuits," Electronics, Volume 38, Number 8, April 19, 1965

© INTEL CORPORATION

©Intel Corporation

Software Enabling Costs Increase with Hardware Support

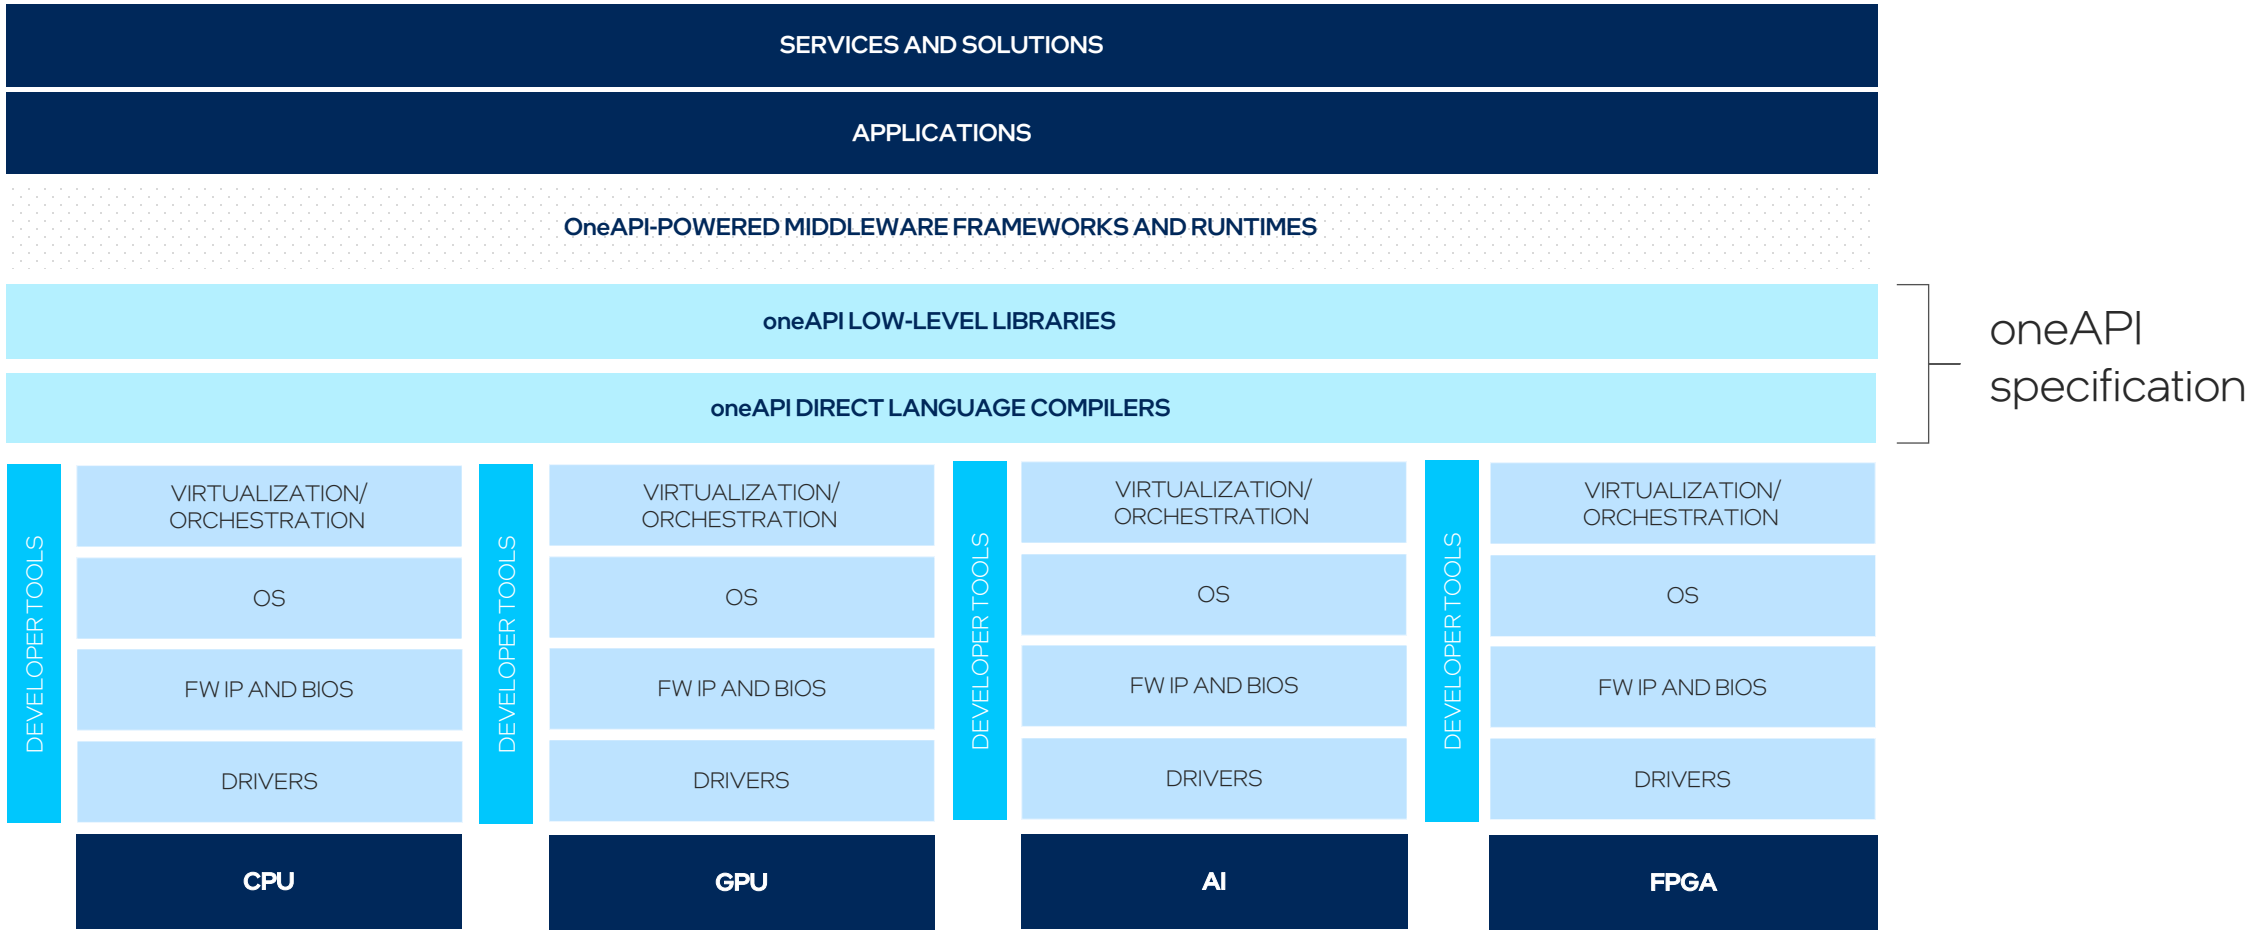


IHV and SW developers incur cost, time to market, and complexity impact of per device developer stacks.

© INTEL CORPORATION

©Intel Corporation

Standards Drive Scale for Accelerated Computing



© INTEL CORPORATION

©Intel Corporation

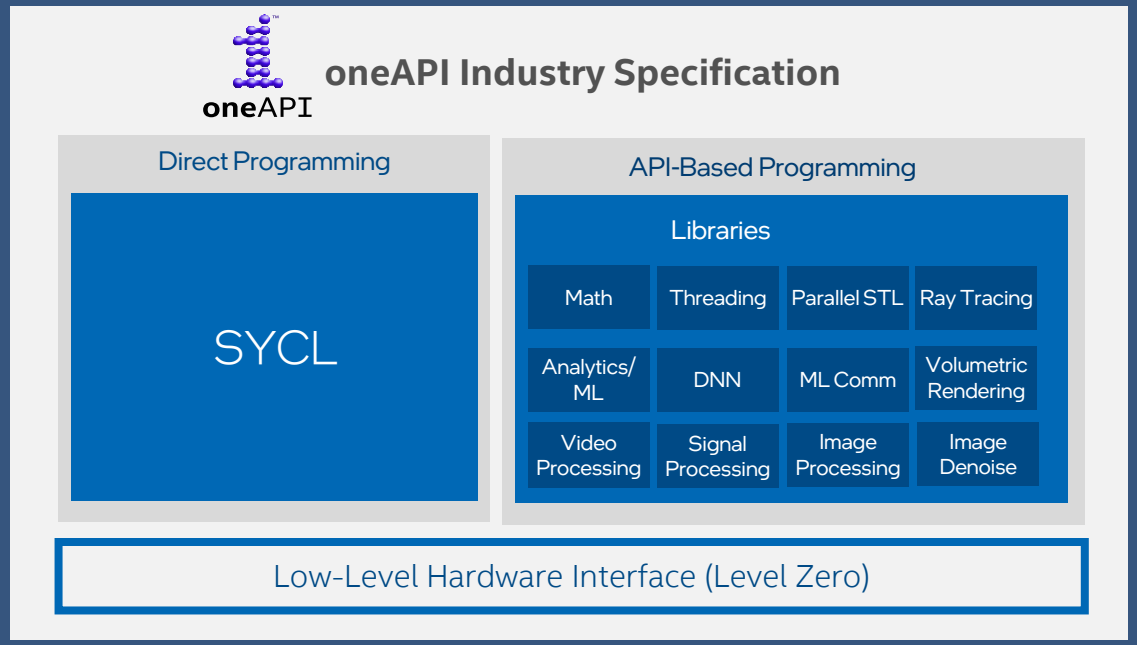
oneAPI: One Name, Two Distinct Objectives

oneAPI

- Open industry specification
- Open-source repo and development
- Community driven
- Supports multi-vendor implementation

oneAPI

- Intel's implementation
- Toolkits optimized for Intel HW
- Available for free download



- Open, standard language: SYCL*
- Standardized interfaces for common libraries
- Standardized hardware interface

*SYCL and the SYCL logo are trademarks of the Khronos Group Inc. in the U.S. and/or other countries

©Intel Corporation

Why C++ with SYCL?

```
sycl::queue q(cpu_selector{});

auto A = sycl::malloc_shared<float>(n, q);
auto B = sycl::malloc_shared<float>(n, q);

q.parallel_for( sycl::range<1>{n},
  [=] (sycl::id<1> i) {
      B[i] += A[i] * A[i];
  }
).wait();
```

```
sycl::queue q(gpu_selector{});

auto A = sycl::malloc_shared<float>(n, q);
auto B = sycl::malloc_shared<float>(n, q);

q.parallel_for( sycl::range<1>{n},
  [=] (sycl::id<1> i) {
      B[i] += A[i] * A[i];
  }
).wait();
```



*Example uses static **targets in bold**, but can be programmed to be dynamic as well

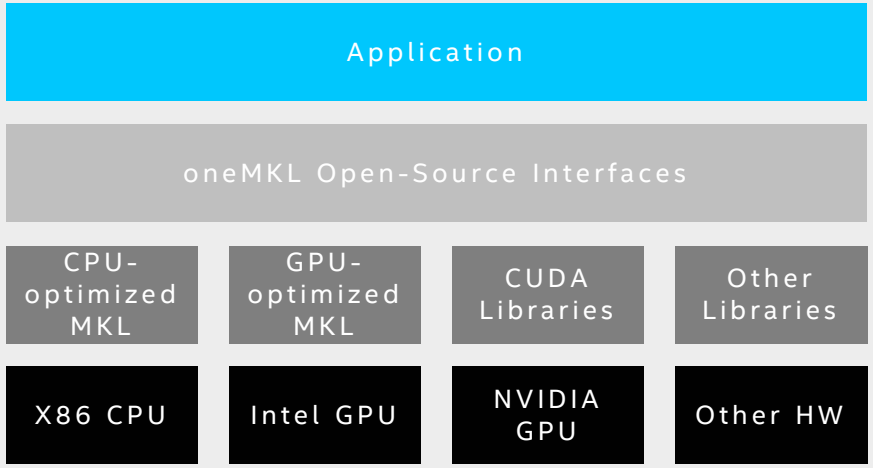
- Standard C++17 aids time to developer productivity
- Syntax for accelerators (device selection, offload, memory transfer)
- Unified shared memory
- Single source (host and device code)
- Multi-Architecture (CPU, GPU, FPGA, and other targets)
- Stack based on standards and open specifications (CLANG, LLVM, SPIR-V, Level Zero)

© INTEL CORPORATION

©Intel Corporation

Standardized Library Interfaces

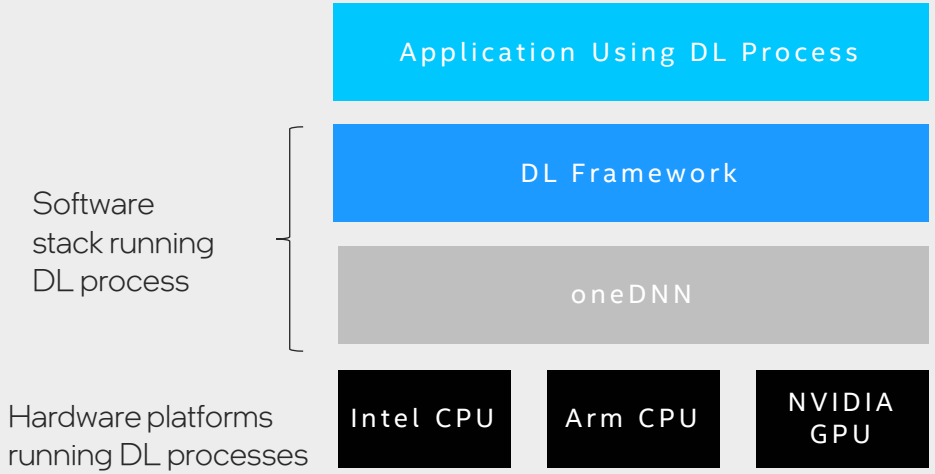
Our Vision



High-Level Overview of oneMKL Open-Source Interfaces

Source: https://www.intel.com/content/www/us/en/developer/articles/technical/a-vendor-neutral-path-to-math-acceleration.html#gs.v4bfu6

In Practice



Software Stack for Deep Learning Processes

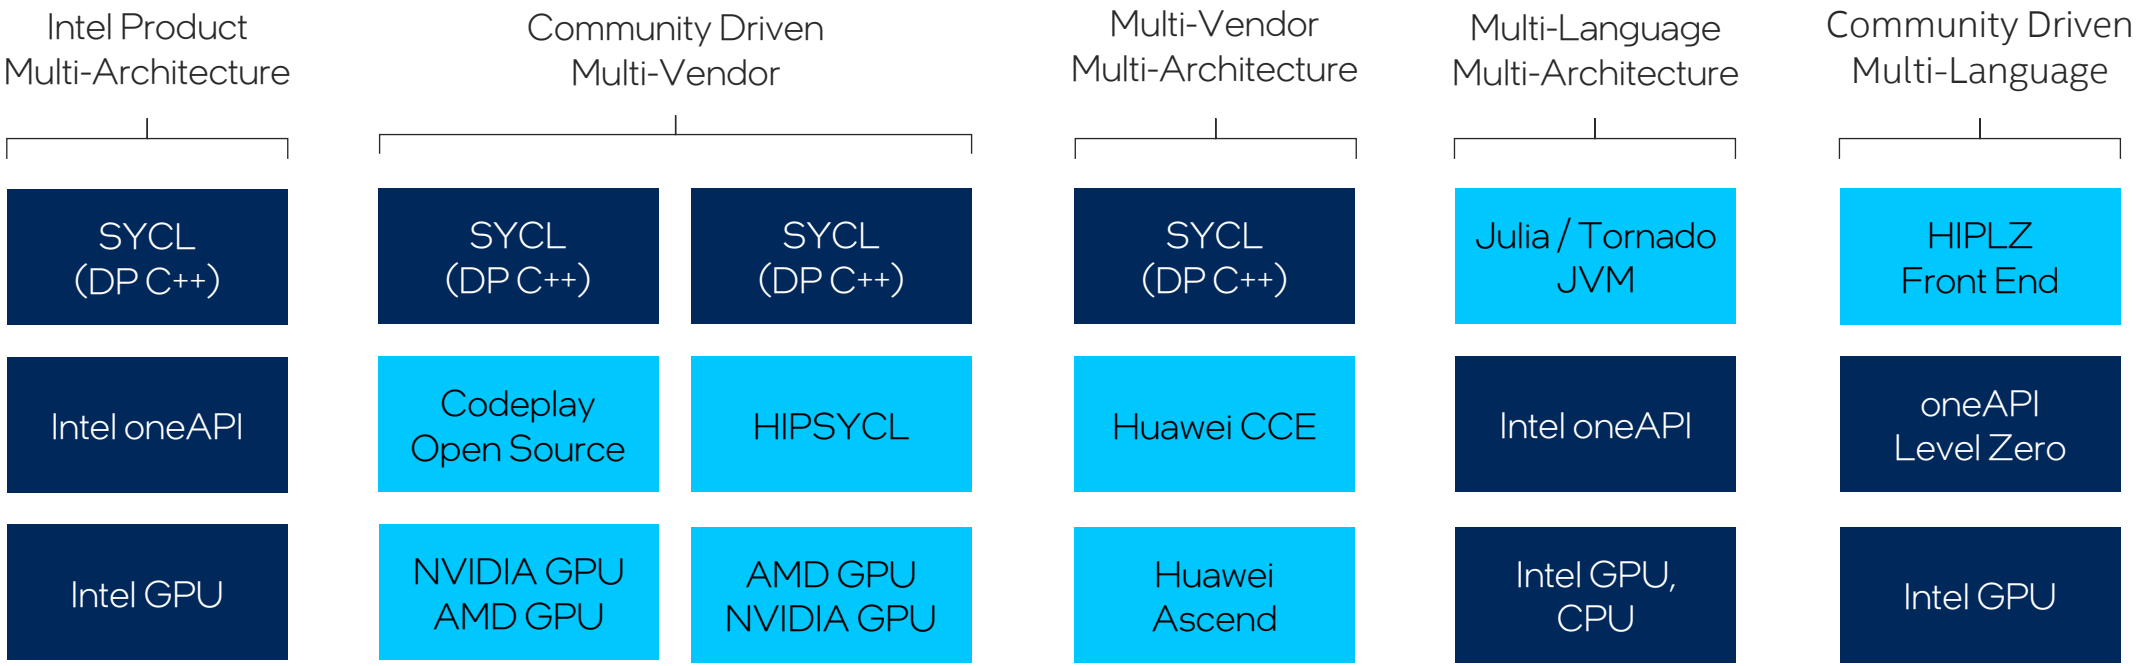
Source: https://blog.fltech.dev/entry/2020/11/19/fugaku-onednn-deep-dive-en

- Faster time to applications
- Ability for HW companies to provide unique value-add underneath standard interfaces

© INTEL CORPORATION

©Intel Corporation

Ecosystem Driven Innovation



Standards-based architecture allows others to freely develop new language front ends and support new hardware targets.

Intel Provided
Ecosystem Content

©Intel Corporation

Status and How to Join In the Fun

vision: true multivendor and multiarchitecture

- github – open source
  - oneAPI is a collection of activities generally grounded in open-source projects (with additions planned)
  - LLVM, libraries, tools, and the specification itself
  - input and contributions are always welcome
- Adoption
  - oneAPI brings advantages that are leading many to evaluate and commit to oneAPI adoption
- Intel oneAPI support – pioneering multivendor "plugin" in product offerings
  - In late 2022, Intel introduced a "plug in" model for its C++ (with SYCL) compiler that allows seamless (no overhead) addition of non-Intel support into the same compiler, with merged single-binary outputs supporting multiple vendors
  - Codeplay announced product support for NVIDIA and AMD GPUs
- oneAPI specification – governance – https://oneapi.io
  - oneAPI is transitioning to a fully independent process
    - Intel shepherded oneAPI in early years, with independent technical advisory boards (all meeting notes are openly online)
    - In late 2022, Intel announced formation of independent steering committees
    - Organization interested in being "founding members" are participating in decisions on the exact details of this independent foundation
      - Goal: formalize governance and membership, agreement, and announce final decisions before end of 2023

©Intel Corporation

oneAPI Ecosystem Support



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.

©Intel Corporation

Summary

- oneAPI: multi-vendor, multi-architecture programming for accelerators
- Standards and open specification based
  - SYCL language
  - Standard library interface
  - LLVM/CLANG, SPIR-V, Level Zero
- Available today in both open-source GitHub and as finished Intel product

Visit oneapi.io for specifications, source repository

Please join the community, provide input

Intel platforms – download for free at software.intel.com

© INTEL CORPORATION

©Intel Corporation


intel