2nd March 2022

oneAPI workshop

Joachim Aertebjerg	EMEA Technical Sales Director
Jean-Laurent Philippe	EMEA HPC Director
Shailen Sobhee	EMEA Software Architect for Exascale
Igor Vorobtsov	EMEA Compiler Engineer
Gilles Civario	EMEA Software Architect for Exascale
Dmitry Tarakanov	EMEA Performance Analysis Tools Engineer
Alina Shadrina	EMEA Compiler Engineer

with the participation of Andrey Alekseenko from KTH Royal Institute of Technology

SiPearl + Intel Collaboration

SIPEARL CPU

+

oneAPI

Open Software Environment

+

intel® Xe GPU

Targeting HPC Deployments in Europe

ExaScale & beyond

Collaboration with OEMs, Ecosystem & Supercomputing Centers

oneAPI

Cross-Architecture Programming for Accelerated Compute, Freedom of Choice for Hardware

oneAPI Industry Initiative & Intel® oneAPI Tools

Shailen Sobhee
EMEA HPC Software Architect for Exascale Supercomputers

shailen.sobhee@intel.com

2nd March 2022

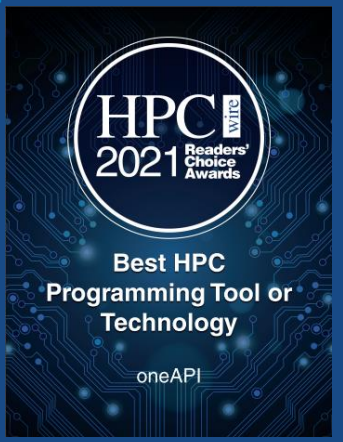

intel

What is oneAPI?

In a nutshell…

oneAPI is a cross-industry, open, standards-based unified programming model that delivers a common developer experience across processor and accelerator architectures.

Programming Challenges
for Multiple Architectures

Growth in specialized workloads

Variety of data-centric hardware required

Separate programming models and toolchains for each architecture are required today

Software development complexity limits freedom of architectural choice

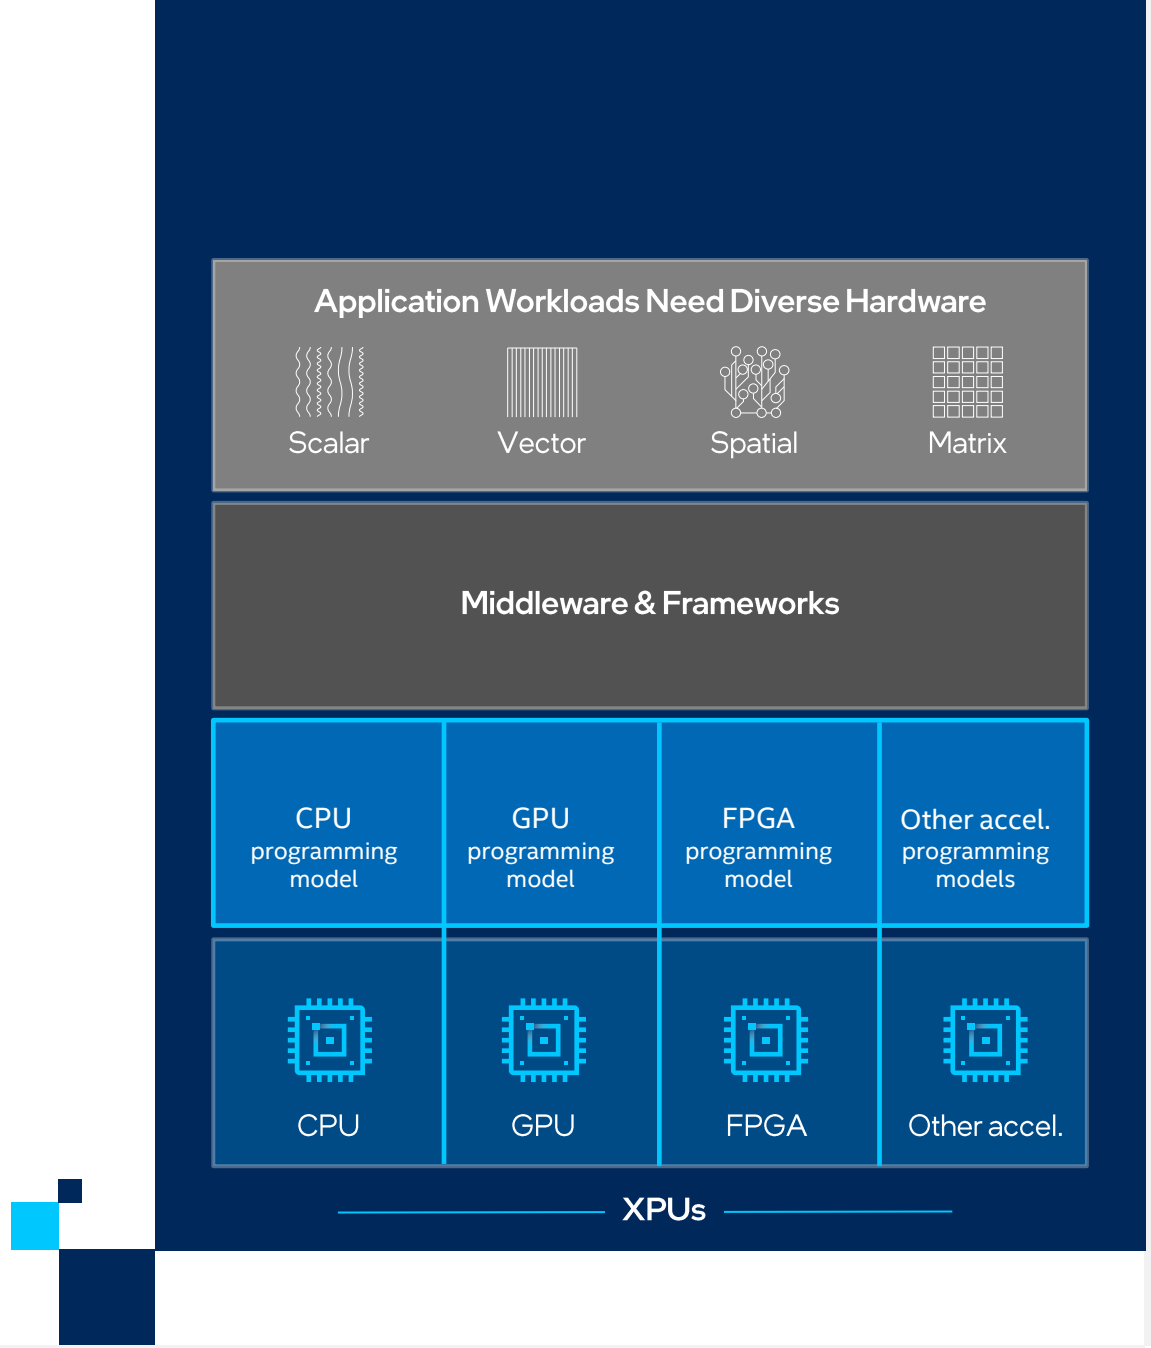

oneAPI Industry Initiative

One Programming Model for Multiple Architectures and Vendors

Freedom to Make Your Best Choice

- Choose the best accelerated technology the software doesn't decide for you

Realize all the Hardware Value

- Performance across CPU, GPUs, FPGAs, and other accelerators

Develop & Deploy Software with Peace of Mind

- Open industry standards provide a safe, clear path to the future
- Compatible with existing languages and programming models including C++, Python, SYCL, OpenMP, Fortran, and MPI

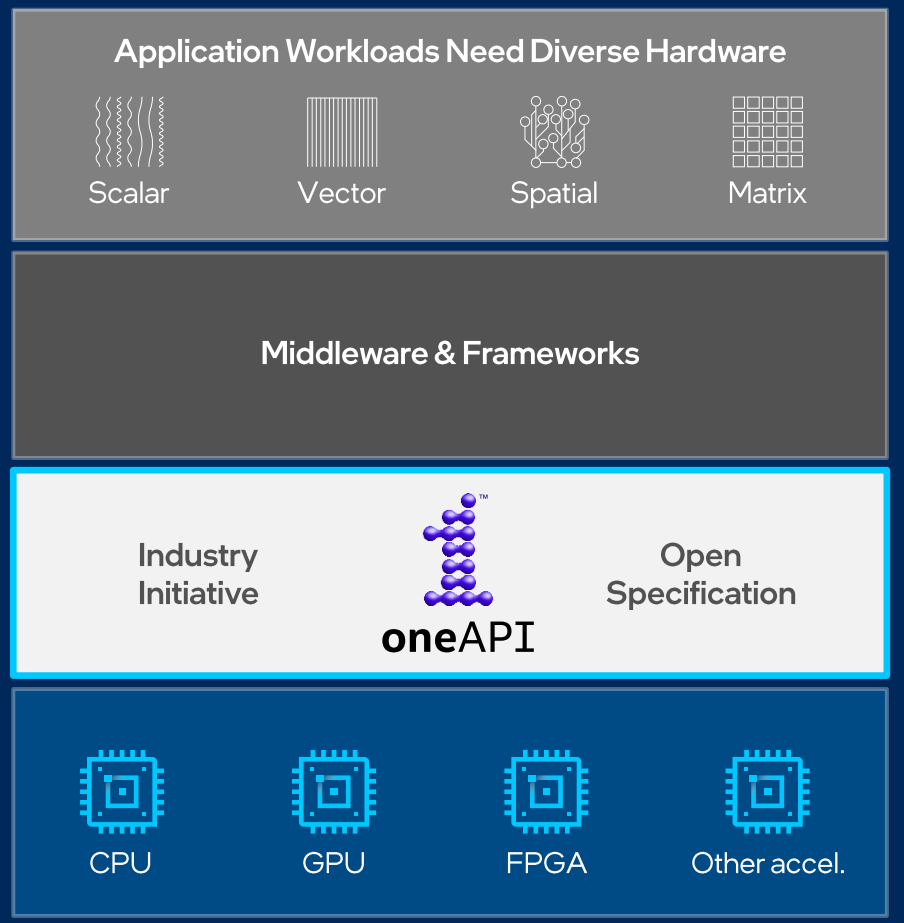

oneAPI Industry Initiative

Break the Chains of Proprietary Lock-in

Open to promote community and industry collaboration

Enables code reuse across architectures and vendors

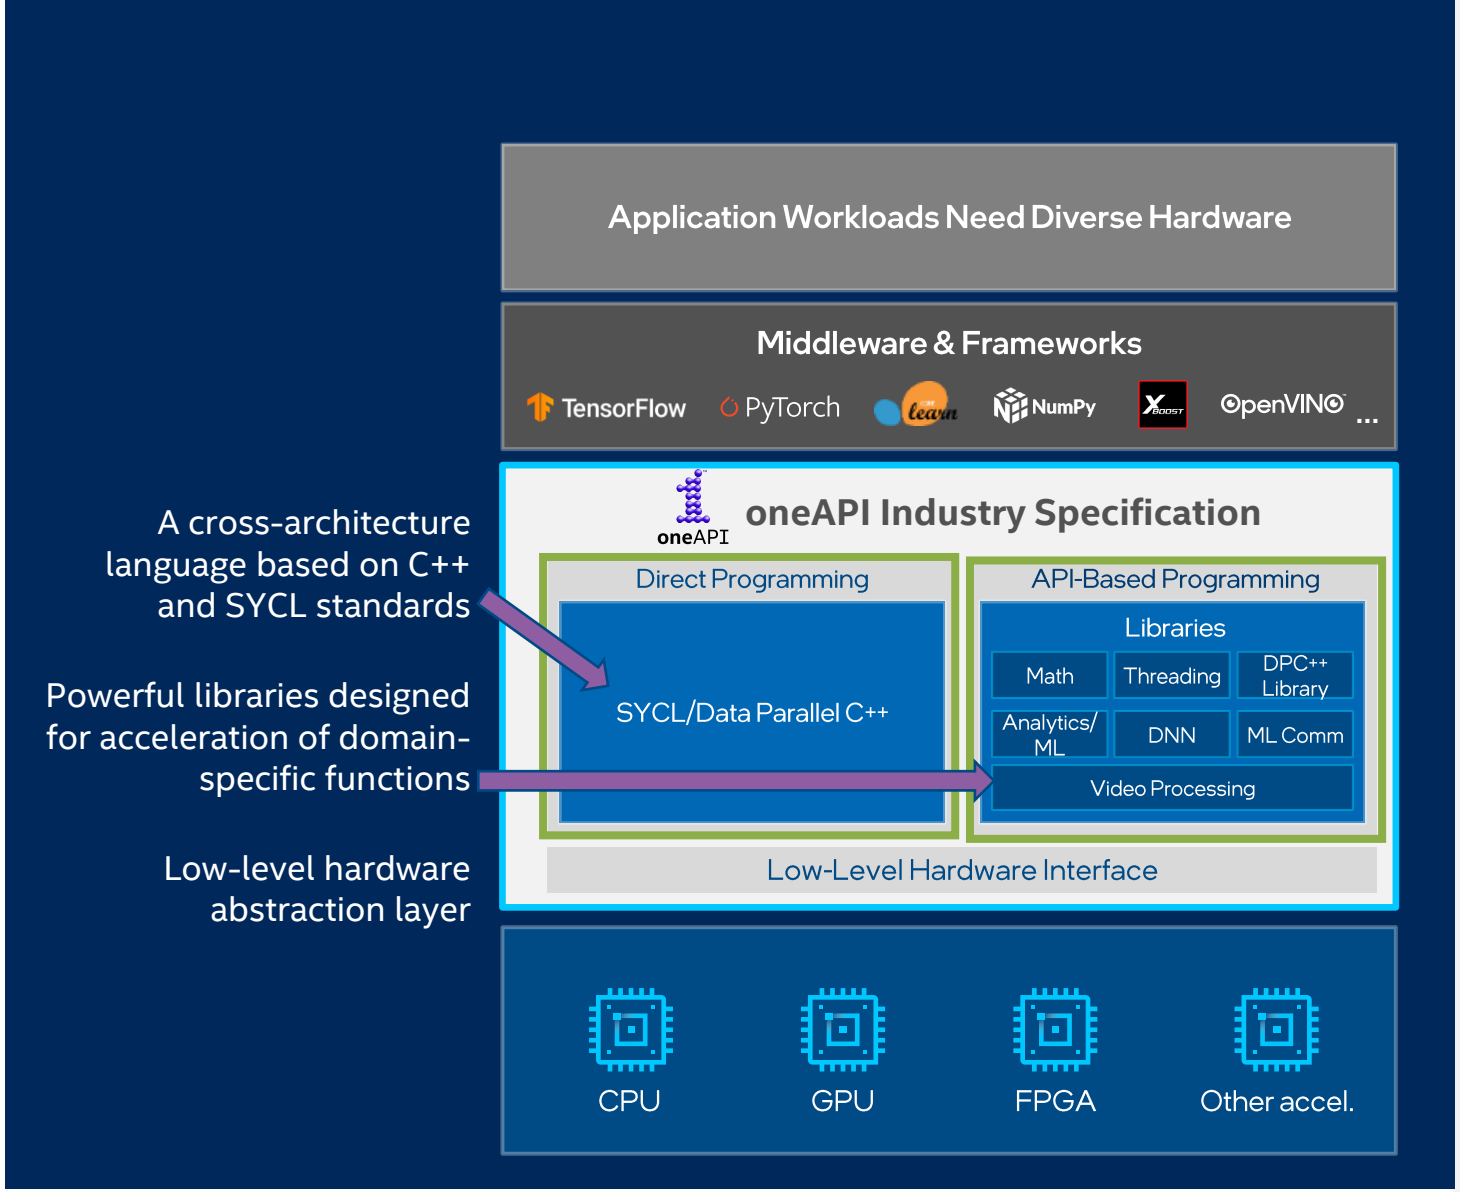

oneAPI

The productive, smart path to freedom for accelerated computing from the economic and technical burdens of proprietary programming models

SYCL, OpenCL and SPIR-V, as open industry standards, enable flexible integration and deployment of multiple acceleration technologies

SYCL Source Code

SYCL enables Khronos to influence ISO C++ to (eventually) support heterogeneous compute

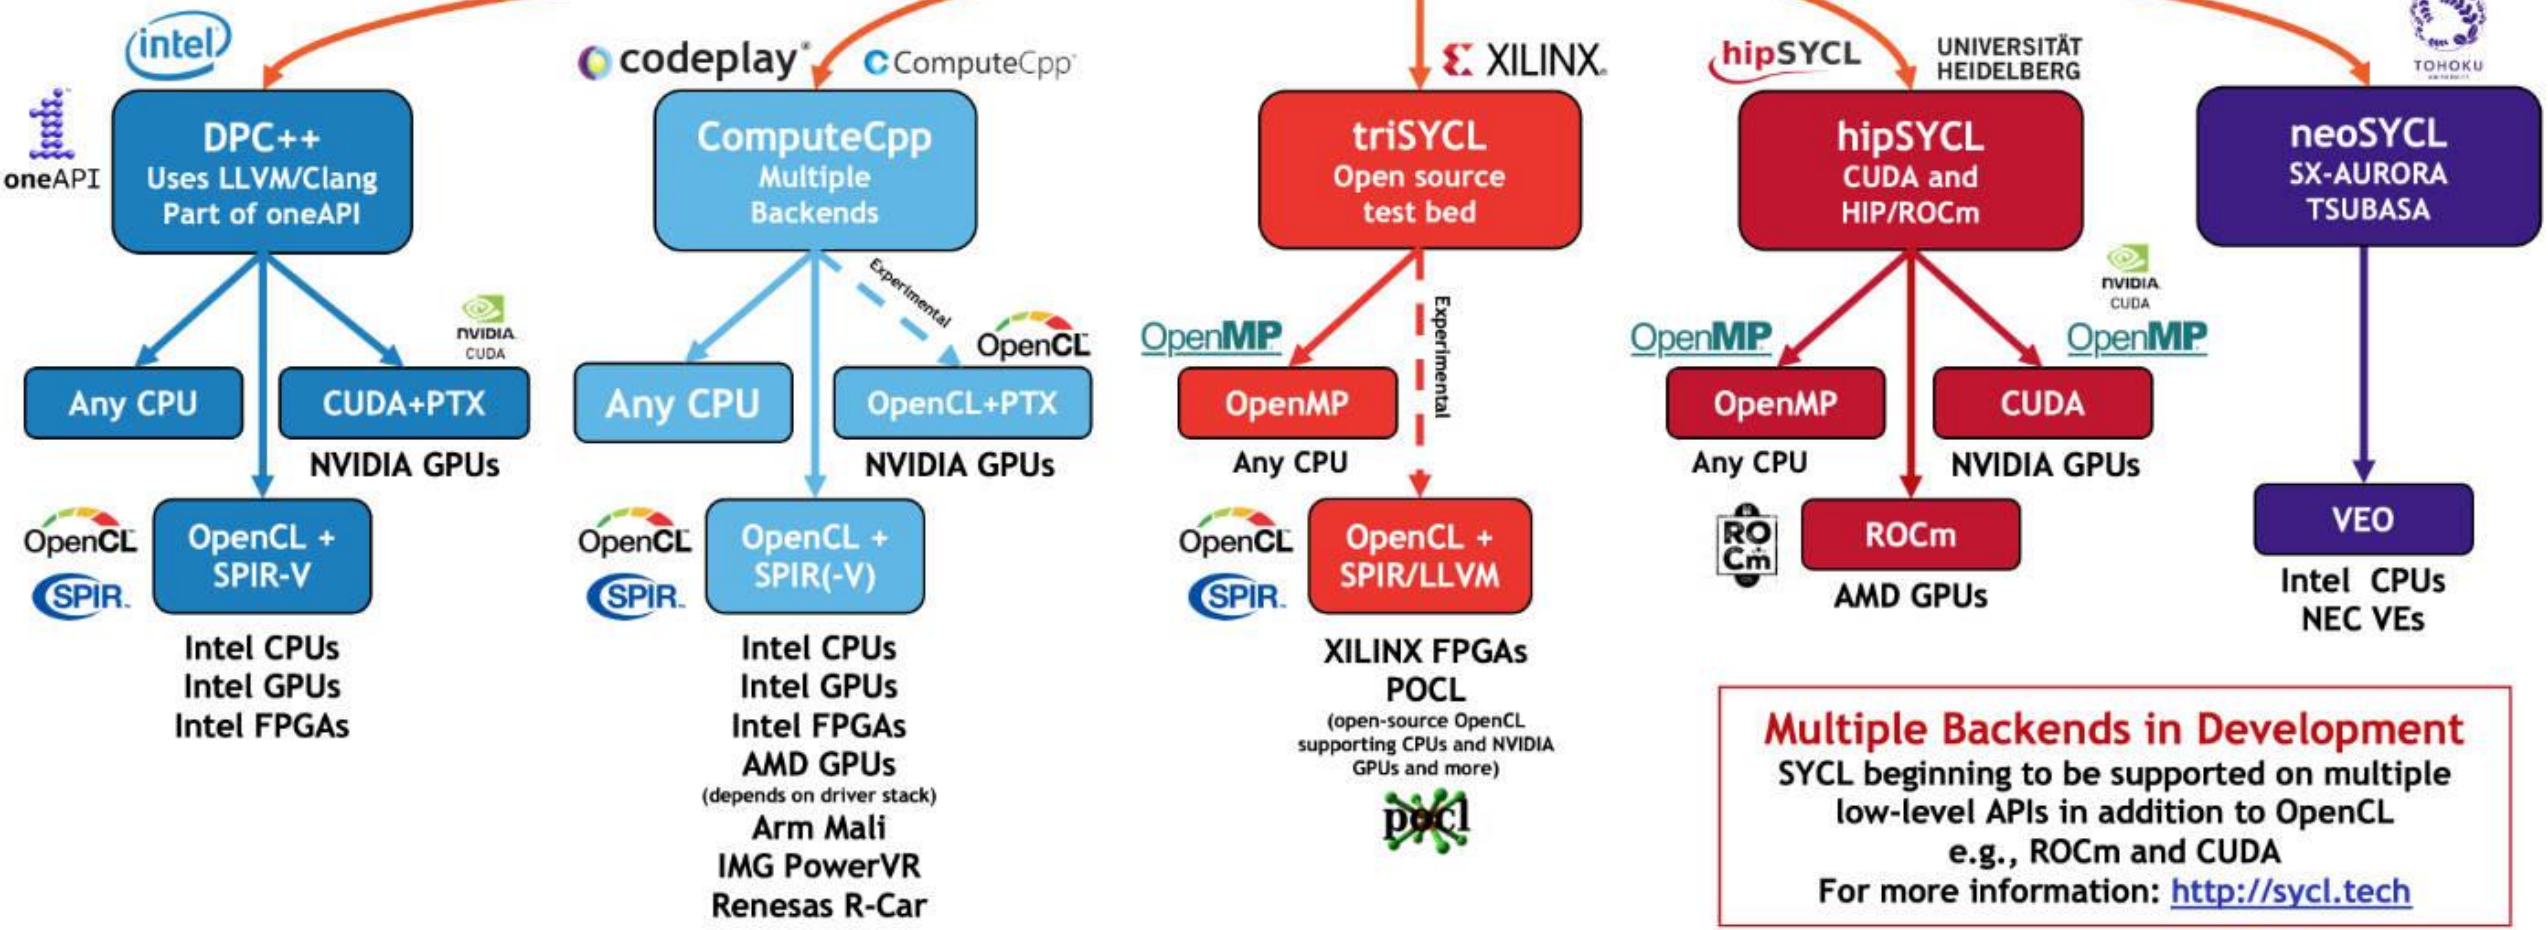

Multiple Backends in Development

SYCL beginning to be supported on multiple low-level APIs in addition to OpenCL
e.g., ROCm and CUDA
For more information: http://sycl.tech

Image source: https://www.khronos.org/sycl/

A64FX Fugaku

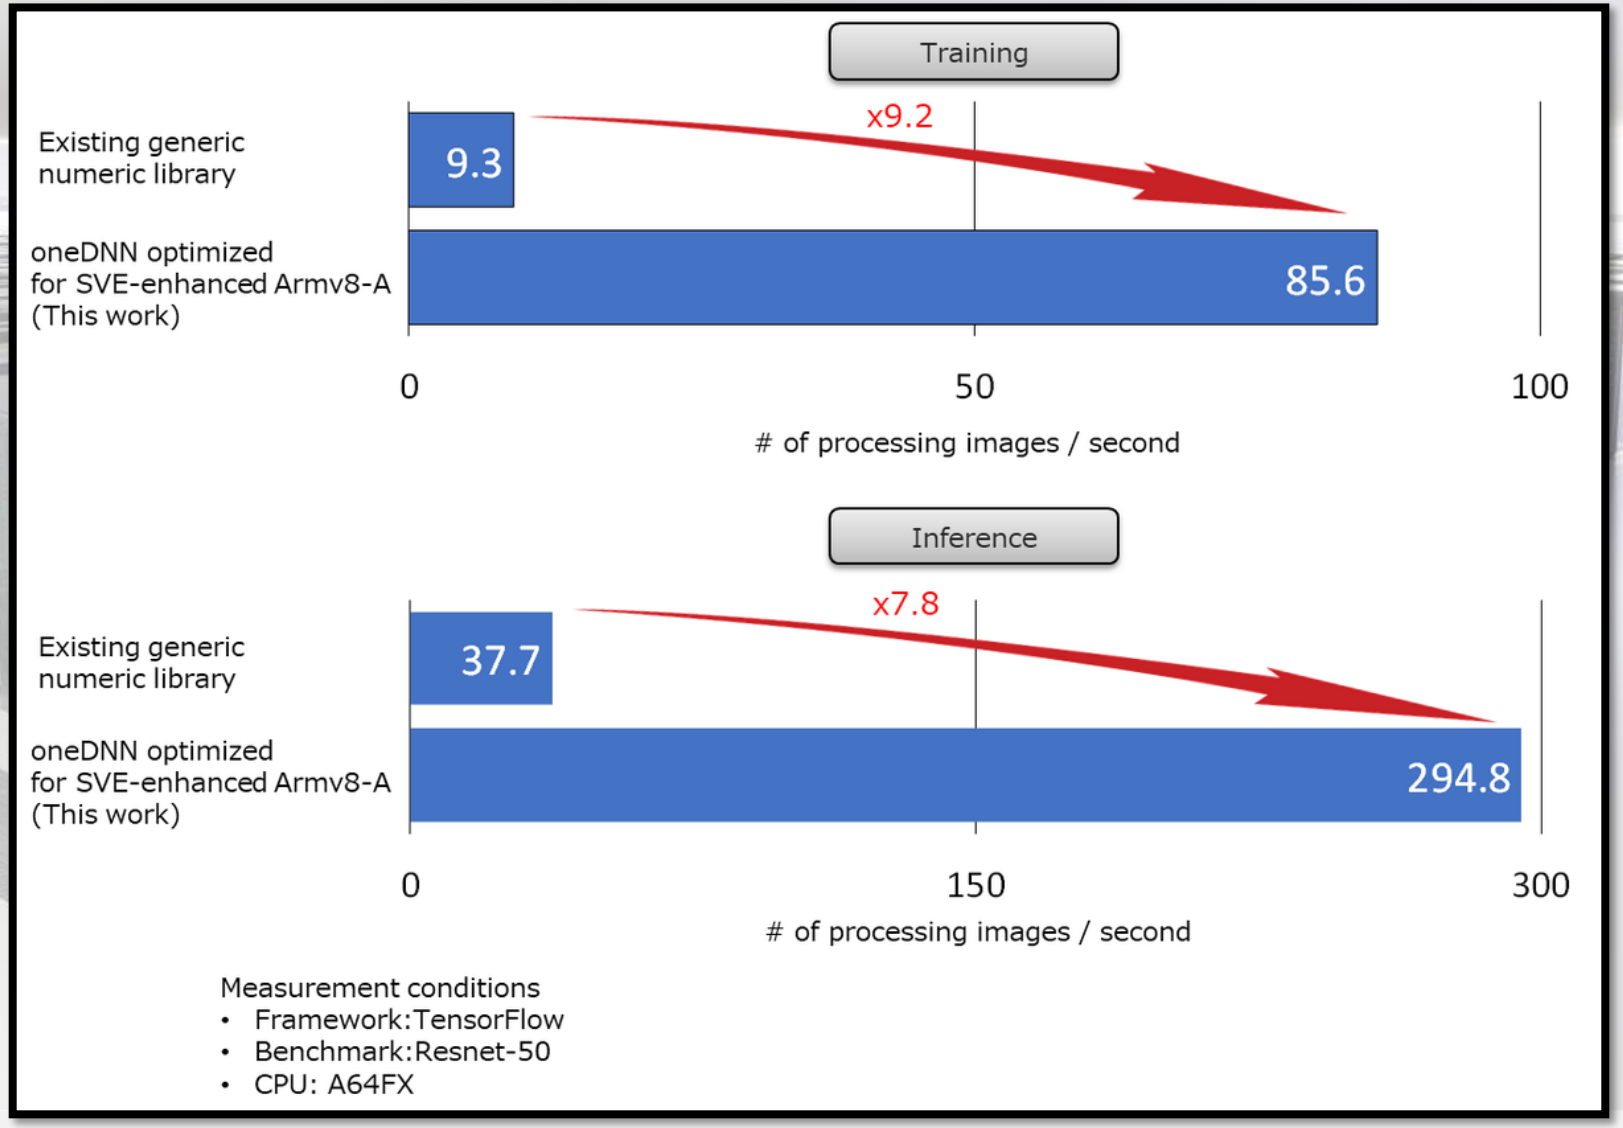

Link: https://www.fujitsu.com/global/about/resources/publications/technicalreview/2020-03/article09.html

oneAPI Ecosystem Support

These organizations support the oneAPI initiative 'concept' 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® oneAPI Tools

Built on Intel's Rich Foundation of CPU Tools Expanded to Accelerators

A complete set of advanced compilers, libraries, and porting, analysis and debugger tools

- Accelerates compute by exploiting cutting-edge hardware features
- Interoperable with existing programming models and code bases (C++, Fortran, Python, OpenMP, etc.), developers can be confident that existing applications work seamlessly with oneAPI
- Eases transitions to new systems and accelerators—using a single code base frees developers to invest more time on innovation

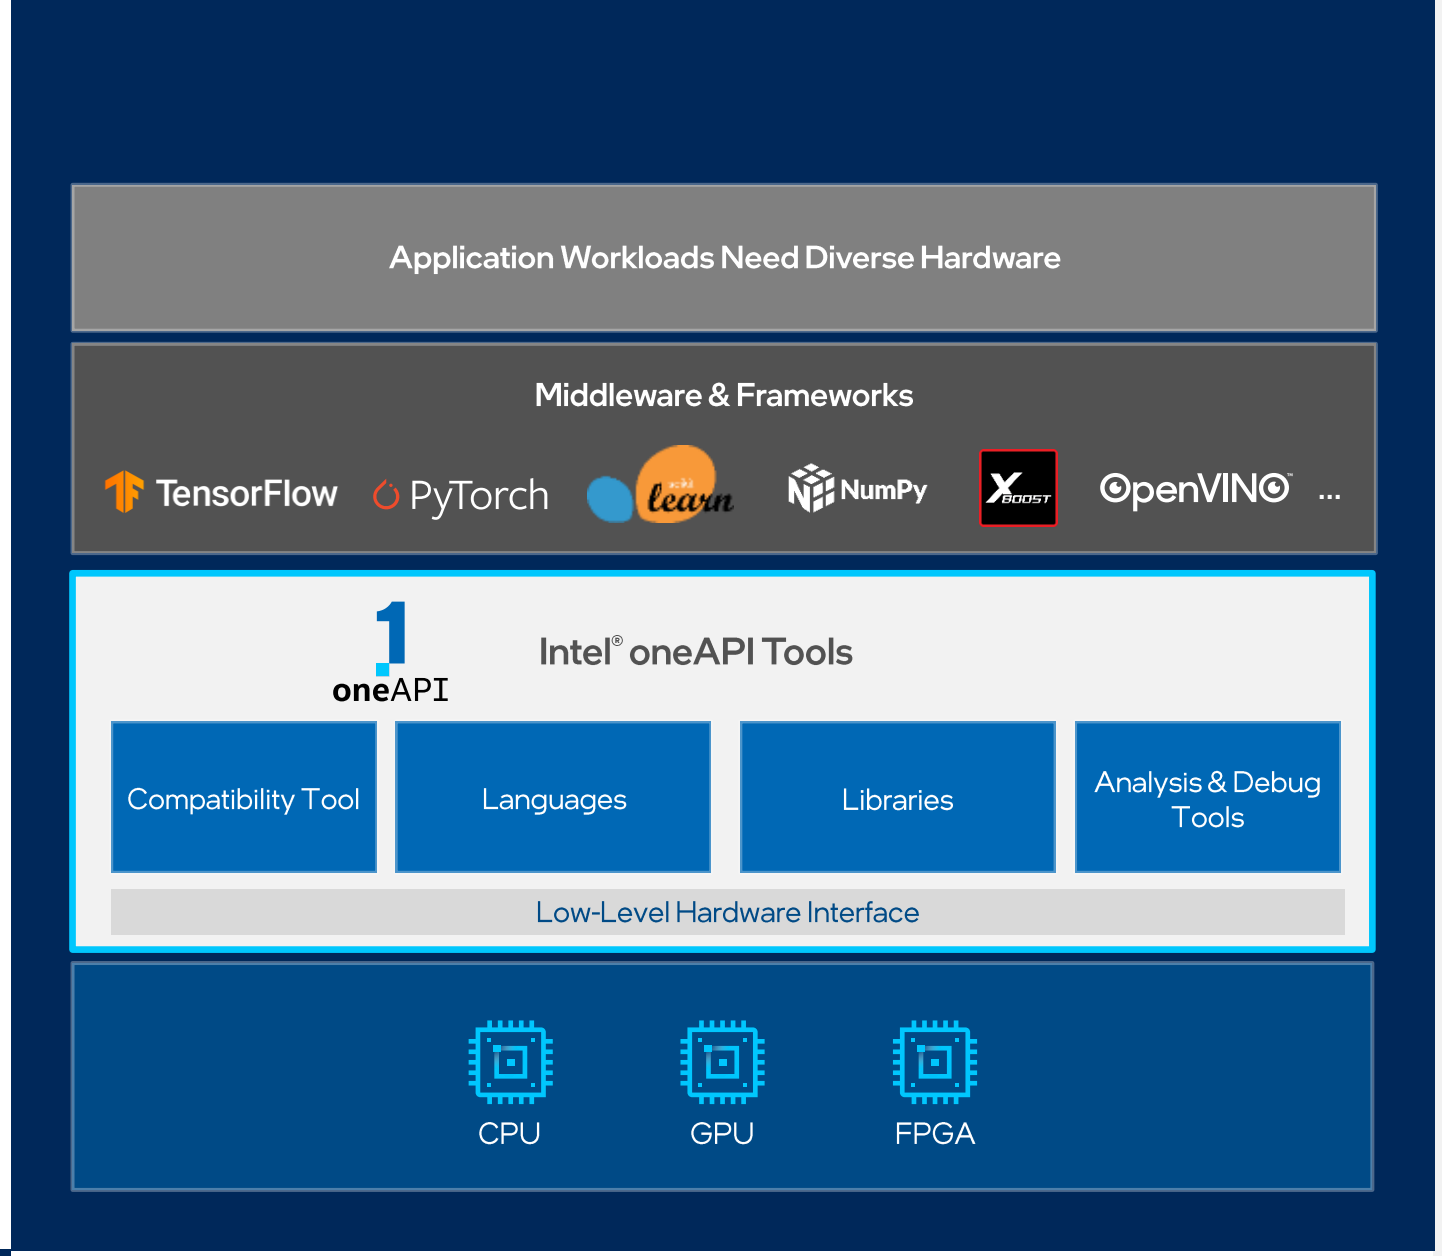

Available Now

Latest version is 2021.1

Visit software.intel.com/oneapi for more details
Some capabilities may differ per architecture and custom-tuning will still be required. Other accelerators to be supported in the future.

Intel® oneAPI Toolkits

A complete set of proven developer tools expanded from CPU to Accelerators

oneAPI

Intel® oneAPI **Base** Toolkit

A core set of high-performance libraries and tools for building C++, SYCL and Python applications

intel oneAPI BASE TOOLKIT

Add-on **Domain-specific** Toolkits

intel oneAPI HPC TOOLKIT

Intel® oneAPI Tools for HPC

Deliver fast Fortran, OpenMP & MPI applications that scale

intel oneAPI IoT TOOLKIT

Intel® oneAPI Tools for IoT

Build efficient, reliable solutions that run at network's edge

intel oneAPI AI ANALYTICS TOOLKIT

Intel® oneAPI AI Toolkit

Accelerate machine learning & data science pipelines end-to-end with optimized DL frameworks & high-performing Python libraries

intel oneAPI RENDERING TOOLKIT

Intel® oneAPI Rendering Toolkit

Create performant, high-fidelity visualization applications

Toolkit **powered by oneAPI**

OpenVINO™

Intel® Distribution of OpenVINO™ Toolkit

Deploy high performance inference & applications from edge to cloud

Latest version available 2022.1

Intel® oneAPI Toolkits Free Availability

Get Started Quickly

Code Samples, Quick-start Guides, Webinars, Training

software.intel.com/oneapi

oneAPI Resources

software.intel.com/oneapi

Get Started

- software.intel.com/oneapi
- Documentation + dev guides
- Code Samples
- Intel® DevCloud

Industry Initiative

- oneAPI.io
- oneAPI open Industry Specification
- Open-source Implementations

Learn

- Training: Webinars & courses
- Intel® DevMesh Innovator Projects
- Summits & Workshops: Live & on-demand virtual workshops, community-led sessions
- Training by certified oneAPI experts worldwide for HPC & AI

Ecosystem

- Community Forums
- Intel® DevMesh Innovator Projects
- Academic Programs: oneAPI Centers of Excellence: research, enabling code, curriculum, teaching

Learn DPC++ for Free

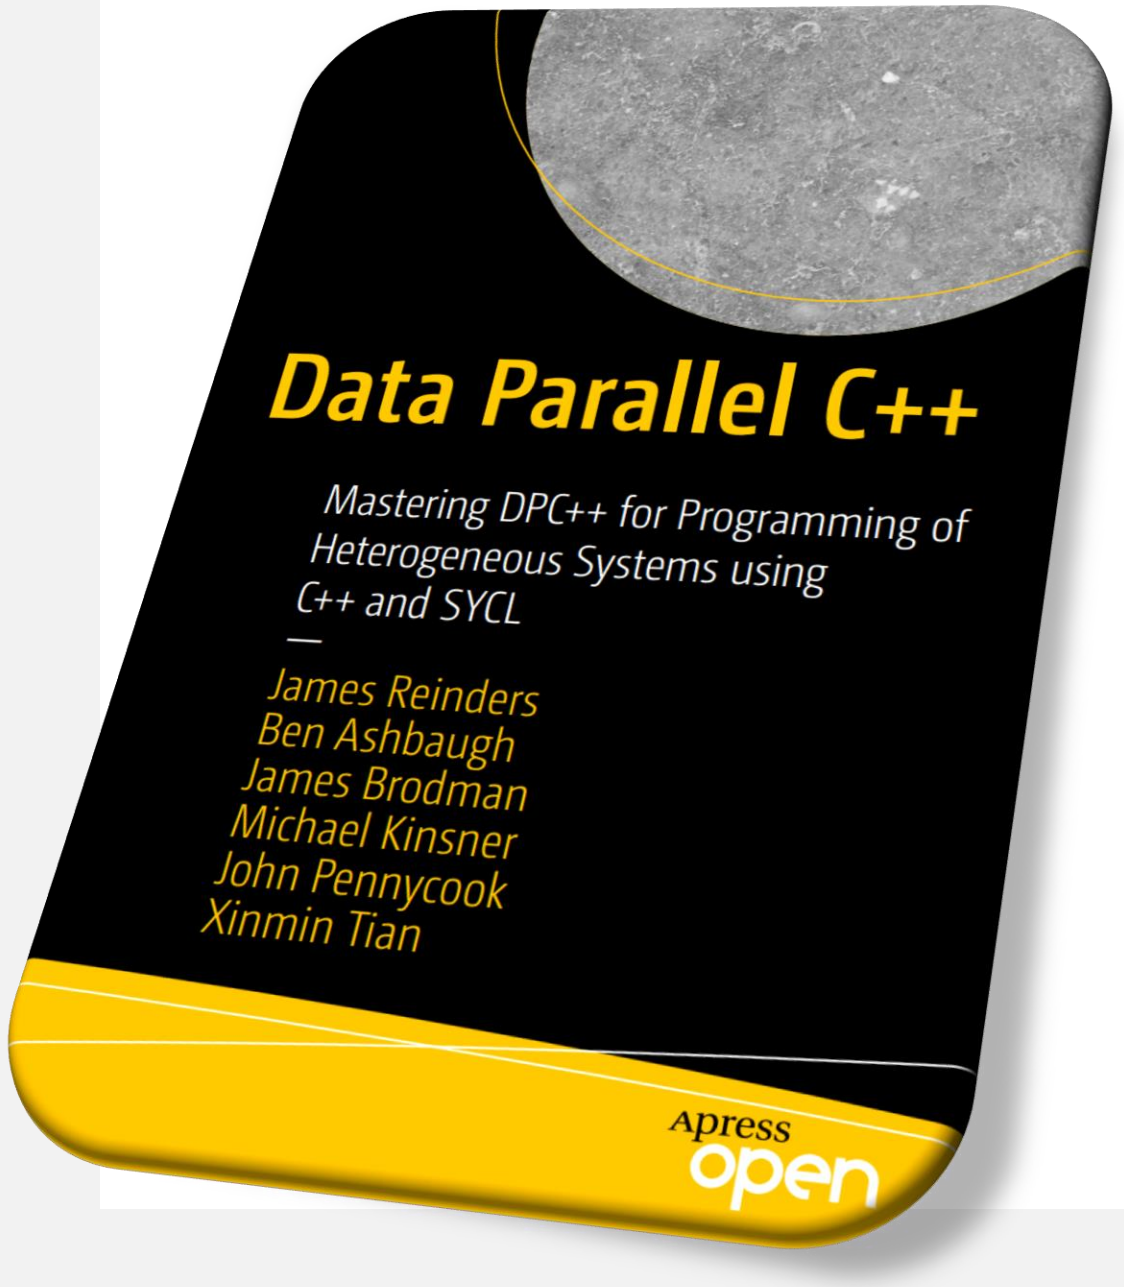

Direct Download: https://www.apress.com/gp/book/9781484255735

Intel® oneAPI Toolkits – Proven Performance

Top Takeaways & Proof Points

- **HPC Cross-architecture** – **Argonne National Labs** is running Exascale-class applications efficiently on current and future generations of Intel CPUs and GPUs
- **HPC Cross-architecture** – **Zuse Institute Berlin (ZIB)** ported the tsunami simulation *easyWave* application from CUDA to SYCL delivering performance across multiple architectures from multiple vendors
- **HPC & AI** – **CERN** uses Intel® DL Boost and oneAPI to speed simulations with inference acceleration by nearly **2x** without accuracy loss*
- **Hyper-real Visualization & AI Using Advanced Ray Tracing** – **Bentley Motors Limited's** AI-based car configurator processes **1.7M+** images with up to **10B** possible configurations per model*
- **IoT** – **Samsung Medison** accelerates ultrasound image processing at the edge on multiple Intel® architectures for improved accuracy and fast diagnosis
- **Major CSPs & Framework** endorse oneAPI – **Microsoft Azure, Google Cloud, TensorFlow**
- **FPGA** – Using oneAPI, **Bittware** had its application running **in days** vs. what typically would take several weeks using Verilog or VHDL*
- And more… 250+ applications developed with oneAPI tools > view catalog

Video [3:45]

*Detailed slides per customer are noted in the oneAPI Customer Use Cases deck. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. See Notices & Disclaimers for more details.

Summary

- Diverse workloads are driving the need for heterogeneous compute architectures, but each architecture has required separate programming models.
- oneAPI cross-architecture programming model provides freedom of choice. Apply your skills to the next innovation, not to rewriting software for the next hardware platform.
- Intel® oneAPI products take full advantage of accelerated compute by maximizing performance across Intel CPUs, GPUs, and FPGAs.
- Develop confidently with a proven set of cross-architecture libraries and advanced tools that interoperate with existing performance programming models.

Call to Action

- Reach out to us for collaboration and to learn more.
- Together, let's make sure the future of software is open, standardized, portable and scalable.


intel

Details about Intel® oneAPI Toolkits

Intel® oneAPI Base Toolkit

Intel® oneAPI Base Toolkit

Accelerate Data-centric Workloads

A core set of core tools and libraries for developing high-performance applications on Intel® CPUs, GPUs, and FPGAs.

Who Uses It?

- A broad range of developers across industries
- Add-on toolkit users since this is the base for all toolkits

Top Features/Benefits

- Data Parallel C++ compiler, library and analysis tools
- DPC++ Compatibility tool helps migrate existing code written in CUDA
- Python distribution includes accelerated scikit-learn, NumPy, SciPy libraries
- Optimized performance libraries for threading, math, data analytics, deep learning, and video/image/signal processing

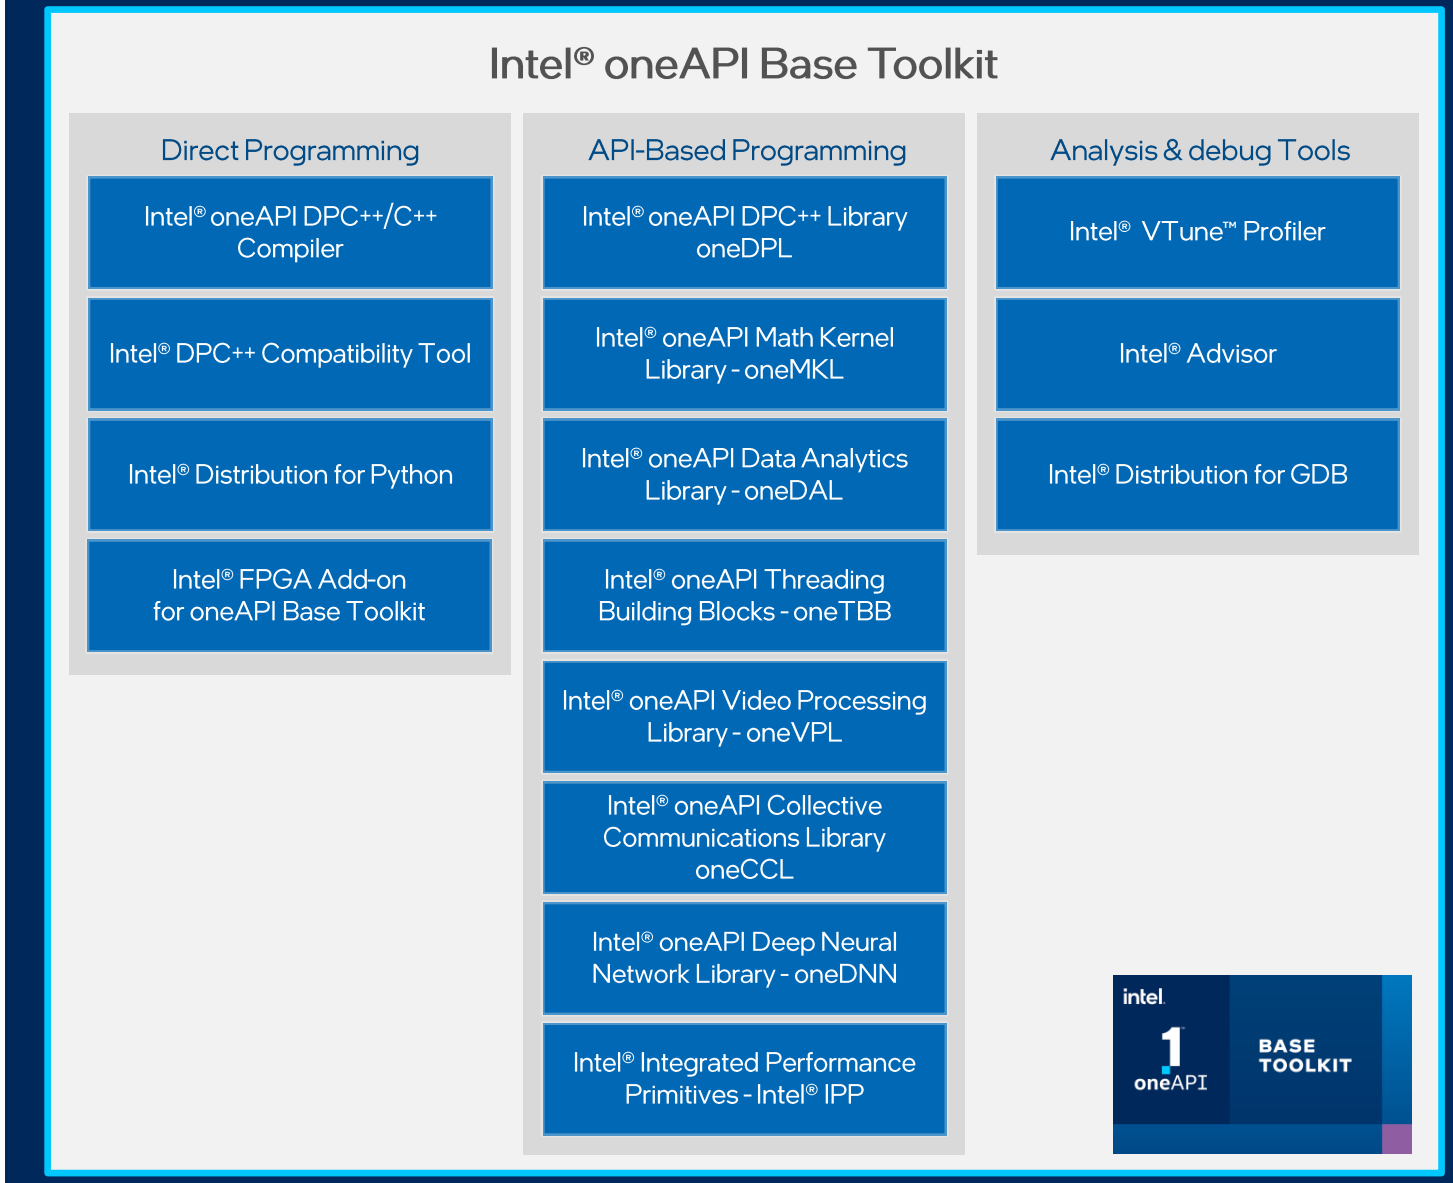

Learn More: intel.com/oneAPI-BaseKit

Intel® oneAPI DPC++/C++ Compiler

Parallel Programming Productivity & Performance

Compiler to deliver uncompromised parallel programming productivity and performance across CPUs and accelerators

- Allows code reuse across hardware targets, while permitting custom tuning for a specific accelerator
- Open, cross-industry alternative to single architecture proprietary language

DPC++ is the oneAPI implementation of Khronos SYCL

- Delivers C++ productivity benefits, using common and familiar C and C++ constructs
- SYCL is a Khronos Group standard to support data parallelism and heterogeneous programming

Builds upon Intel's decades of experience in architecture and high-performance compilers

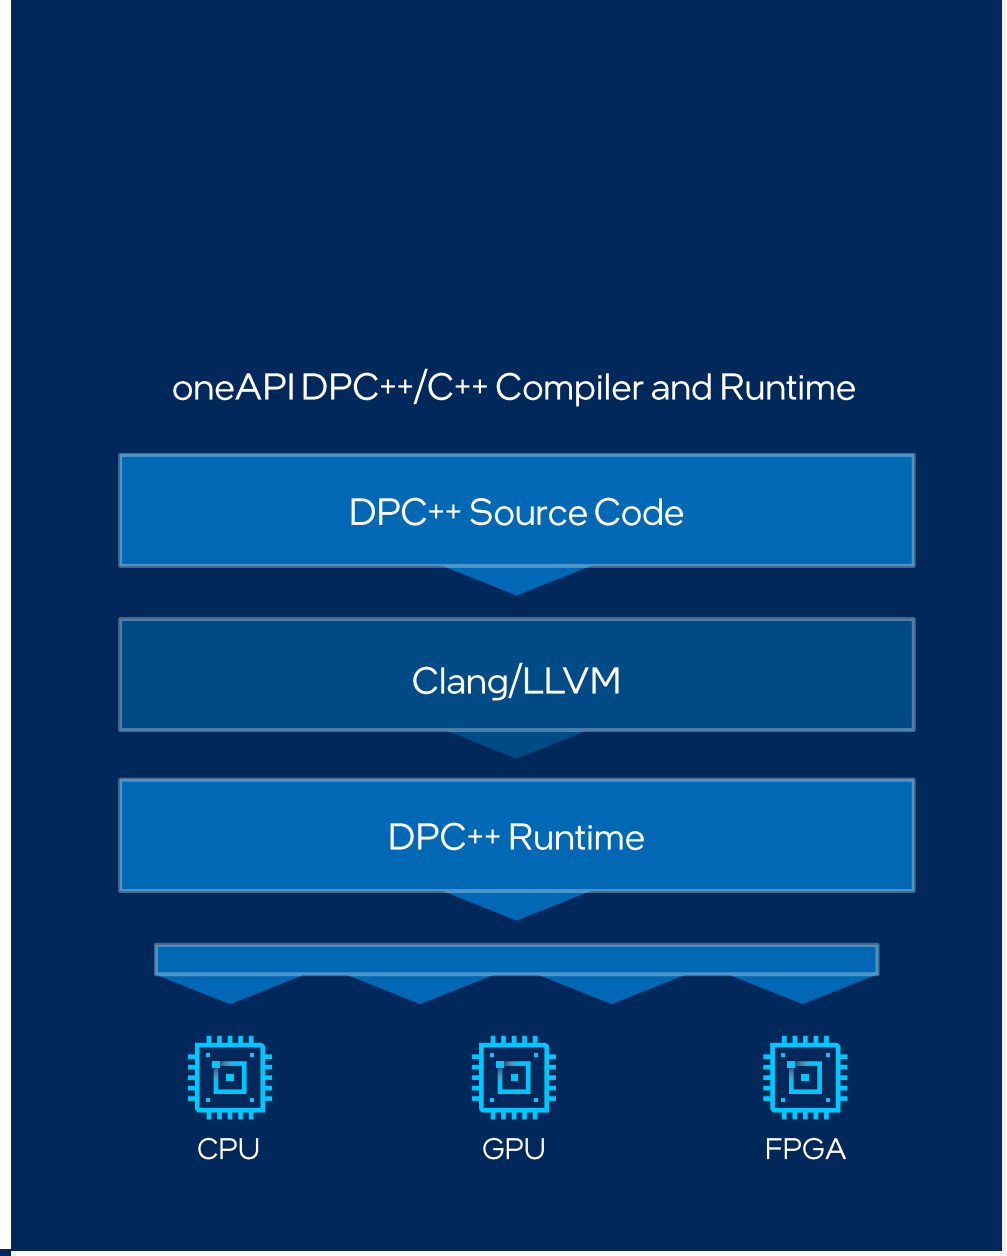

There will still be a need to tune for each architecture.

Intel® DPC++ Compatibility Tool

Minimizes Code Migration Time

Assists developers migrating code written in CUDA to SYCL once, generating **human readable** code wherever possible

~90-95% of code typically migrates automatically[1]

Inline comments are provided to help developers finish porting the application

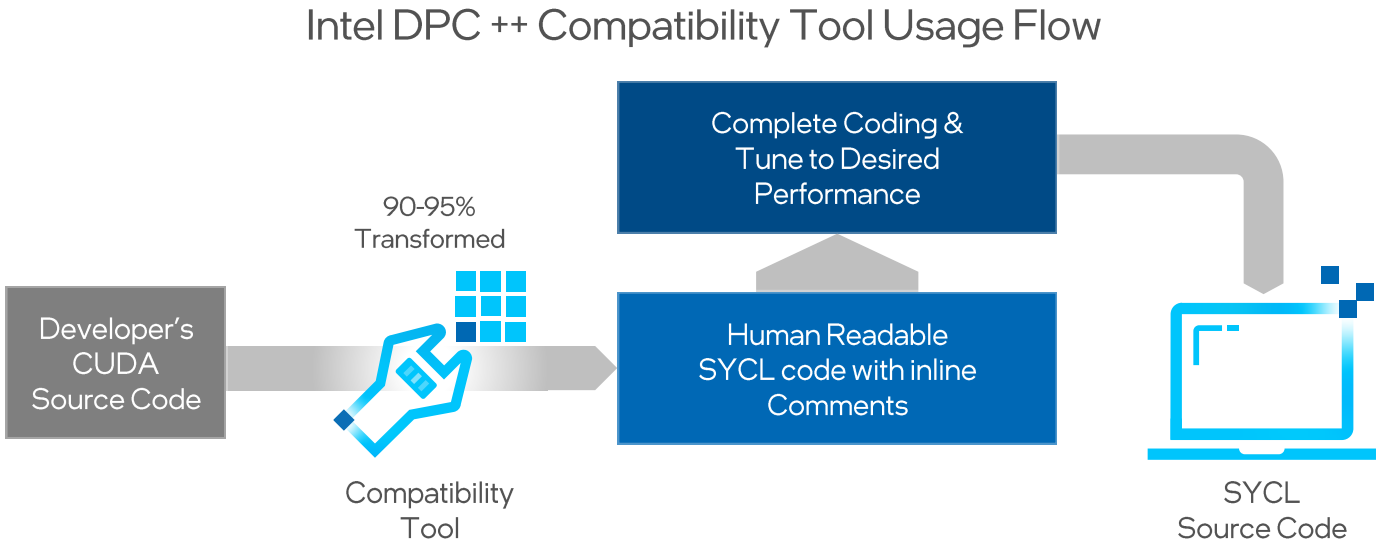

[1]Intel estimates as of September 2021. Based on measurements on a set of 70 HPC benchmarks and samples, with examples like Rodinia, SHOC, PENNANT. Results may vary.

Intel® oneAPI DPC++ Library

Accelerate DPC++ Kernels on Intel CPUs, GPUs & FPGAs

Optimized C++ Standard Algorithms

Contains 75 parallelized C++17 algorithms and utilities for efficient application development and deployment on a variety of hardware.

Based on parallel libraries that C++ developers are already familiar with

Incorporates popular libraries Parallel STL and Boost. Compute for easier developer adoption.

Integrated with Intel® DPC++ Compatibility Tool

Complements all oneAPI DPC++ components to simplify migration of developers' CUDA* code to DPC++ code.

Intel® oneAPI Deep Neural Network Library

Deliver High Performance Deep Learning

Helps developers create high performance deep learning frameworks

Abstracts out instruction set & other complexities of performance optimizations

Same API for both Intel CPUs and GPUs, use the best technology for the job

Supports Linux, Windows

Open sourced for community contributions

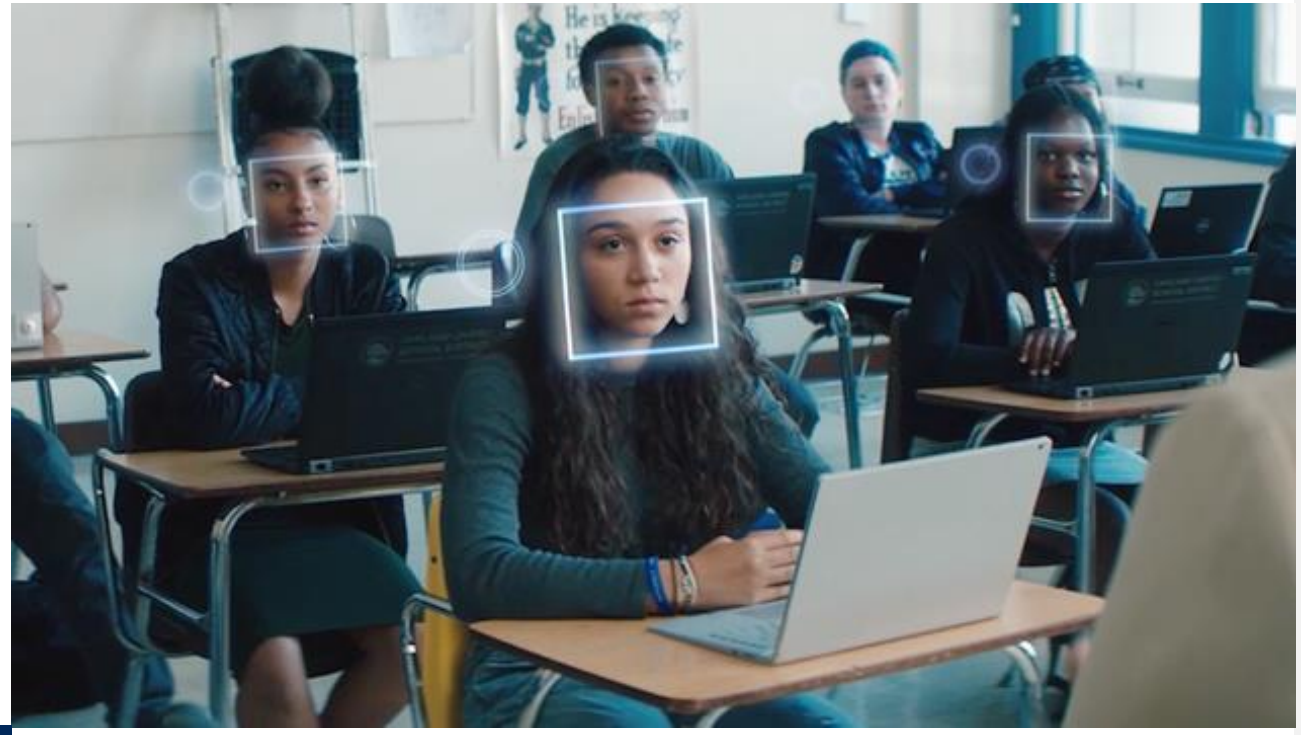

Intel® oneAPI Video Processing Library

Boost Media Performance

Boost media and video application performance with hardware-accelerated codecs and programmable graphics on Intel CPUs and GPUs

Simple API that works the same on CPU and GPU

Using the API, developers have full control over codec visual quality and performance

Intel® oneAPI Collective Communications Library

Optimize Communication Patterns

Provides optimized communication patterns for high performance on Intel CPUs & GPUs to distribute model training across multiple nodes

Transparently supports many interconnects, such as Intel® Omni-Path Architecture, InfiniBand, & Ethernet

Built on top of lower-level communication middleware-MPI & libfabrics

Enables efficient implementations of collectives used for deep learning training-all-gather, all-reduce, & reduce-scatter

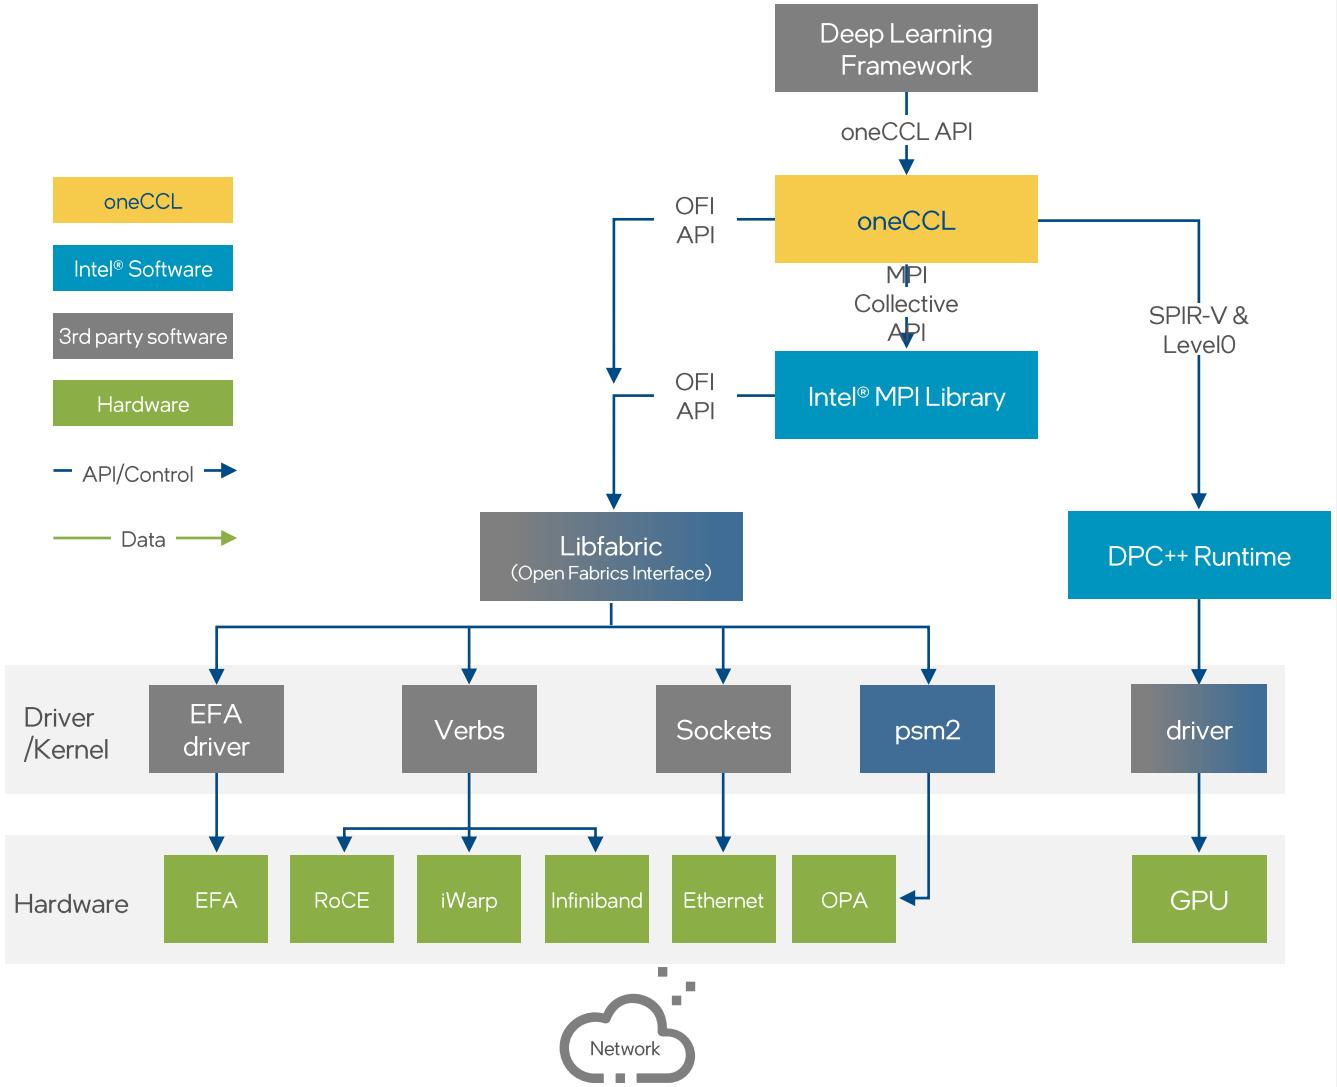

Intel® VTune™ Profiler

DPC++ Profiling—Tune for CPU, GPU & FPGA

Analyze SYCL code

See the lines of SYCLthat consume the most time

Tune for Intel CPUs, GPUs & FPGAs

Optimize for any supported hardware accelerator

Optimize Offload

Tune OpenMP offload performance

Wide Range of Performance Profiles

CPU, GPU, FPGA, threading, memory, cache, storage...
Flame graph display improves visualization of hot spots

Supports Popular Languages

SYCL, C, C++, Fortran, Python, Go, Java, or a mix

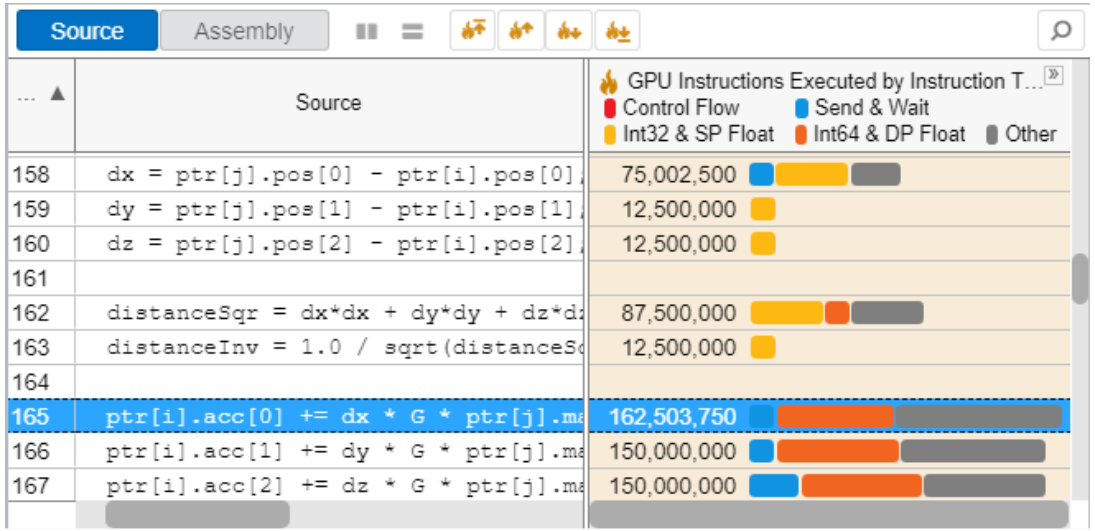

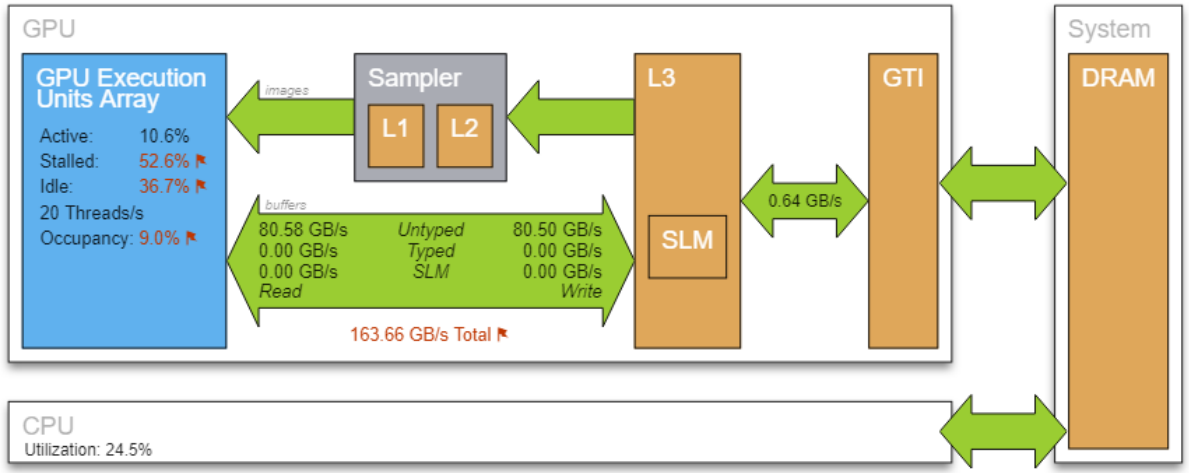

Images above show analysis of SYCL code and GPU Offload profiling.

There will still be a need to tune for each architecture.

Intel® Advisor

Design Assistant—Design for Modern Hardware

Offload Advisor

Estimate performance of offloading to an accelerator

Roofline Analysis

Optimize CPU/GPU code for memory and compute

Vectorization Advisor

Add and optimize vectorization

Threading Advisor

Add effective threading to unthreaded applications

Flow Graph Analyzer

Create and analyze efficient flow graphs

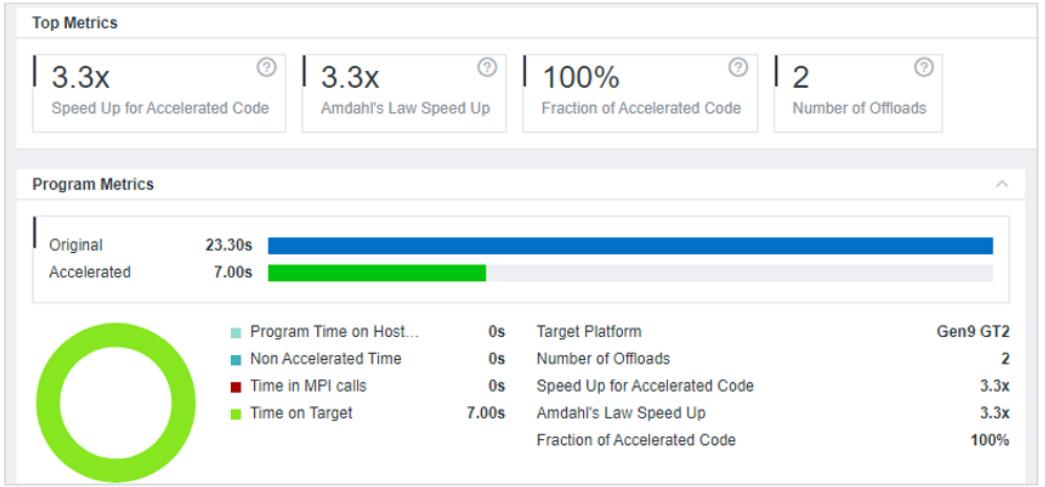

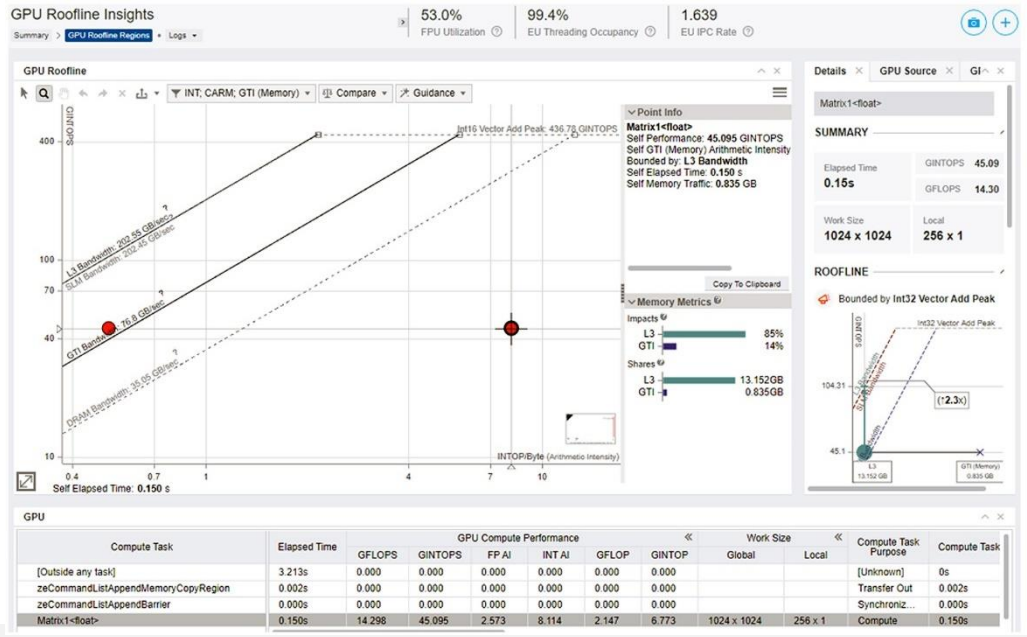

Intel® Distribution of GDB

Data Parallel C++ Debug—Heterogeneous Application Debug

High-level language debug support

Multiple accelerator support:
Intel CPU, GPU, FPGA emulation

Auto-detect accelerator architecture during application runtime

Non-proprietary open-source solution based on GDB

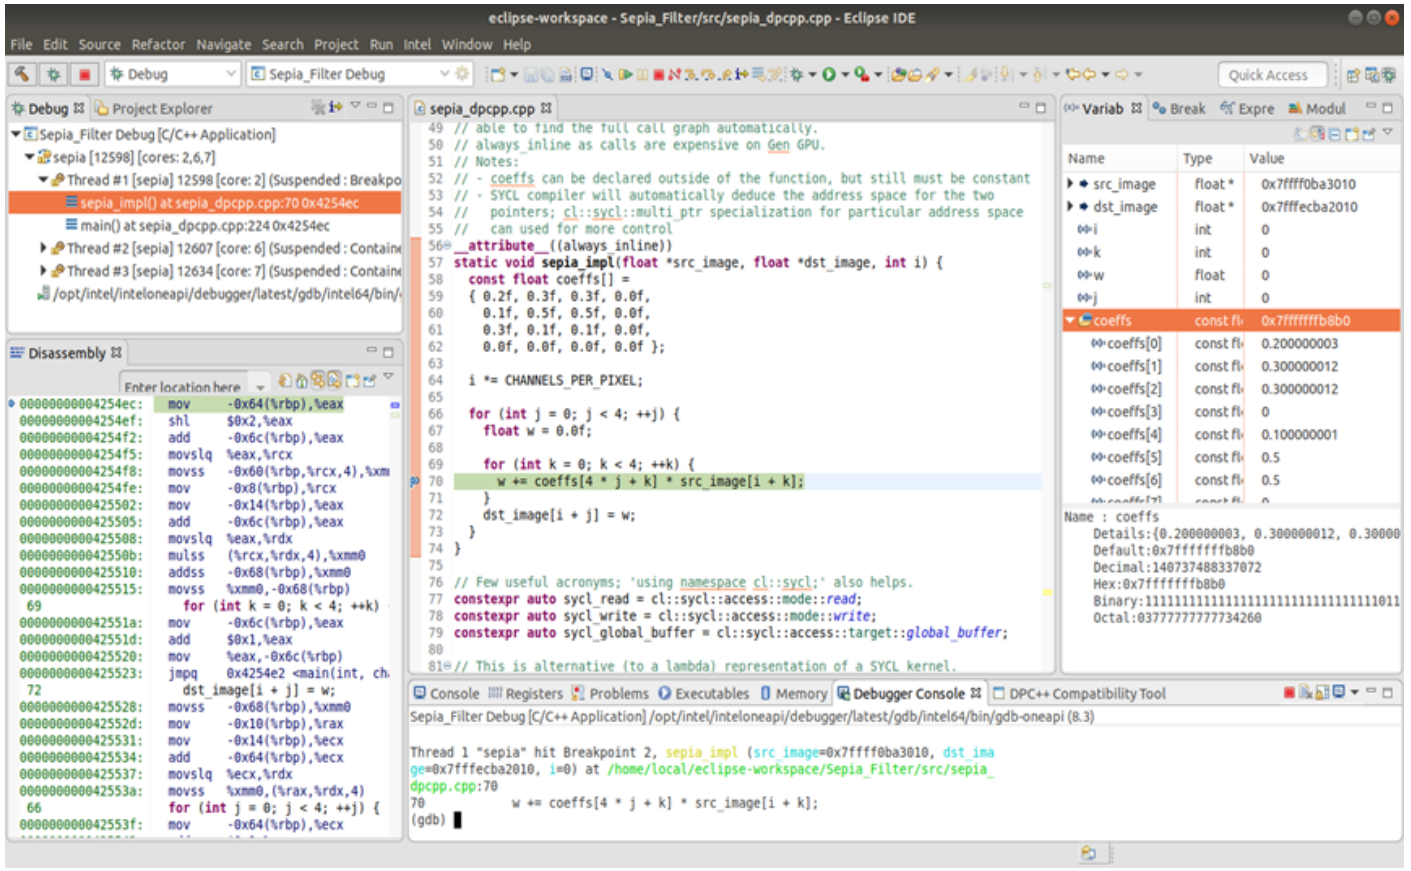

oneAPI for FPGA

SYCLCoding for Spatial Architecture

For Experienced FPGA Developers

Ease of Use

Experienced FPGA users can take advantage of a streamlined programming model using Data Parallel C++

Real Time Processing

Process data faster with deterministic low latency and high throughput

Runtime Analysis Support

Collect profiling data at runtime to analyze CPU and FPGA interaction with Intel® VTune™ Profiler

Device Specific Optimizations

One-day class provides experienced FPGA developers training to begin optimizing oneAPI code for FPGA

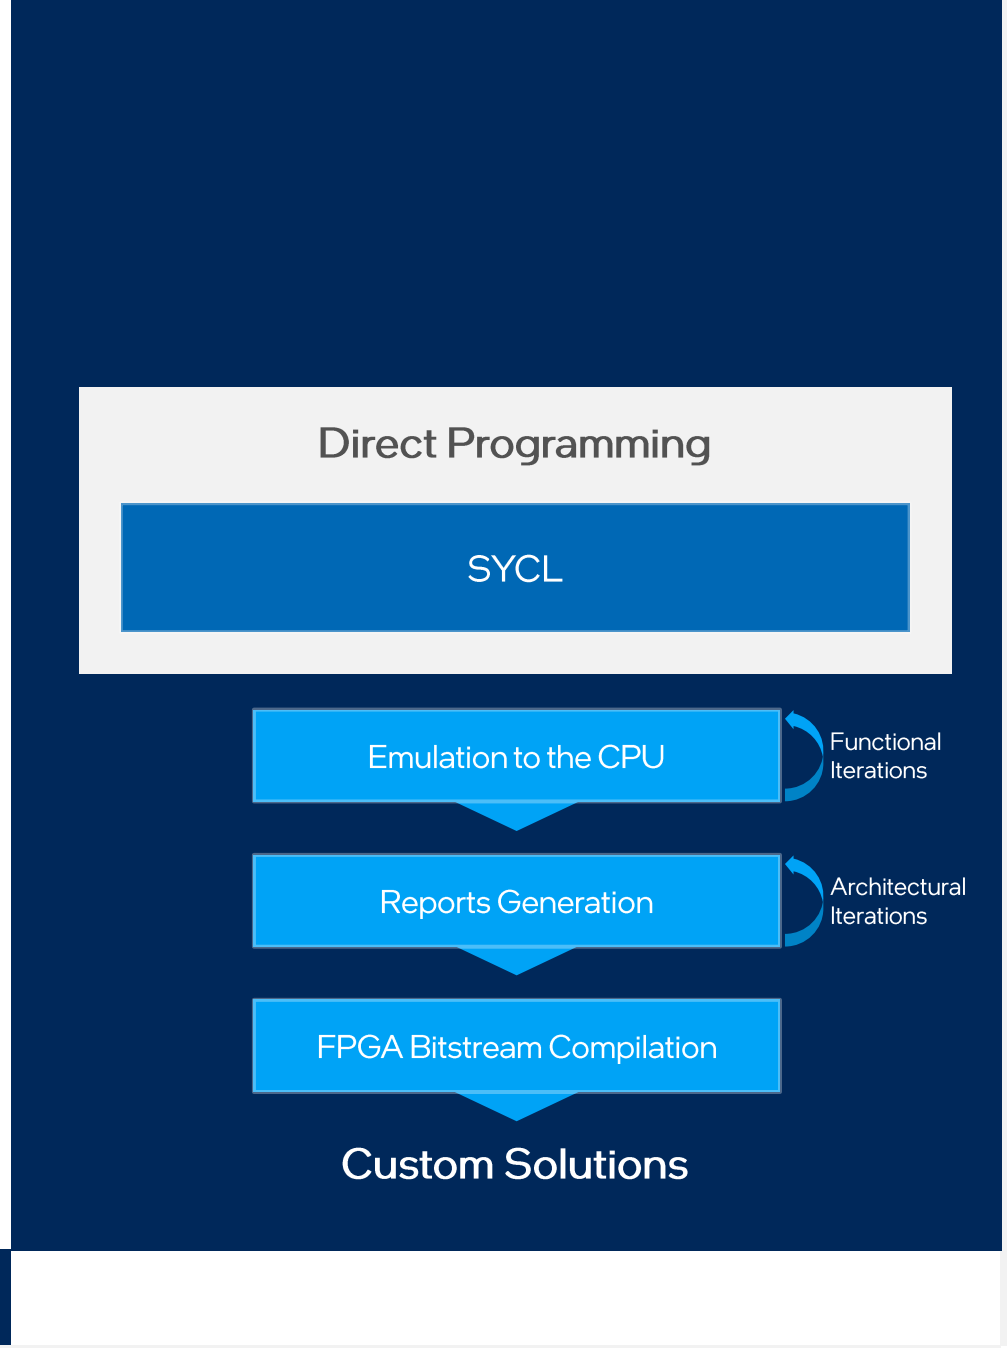

Domain-specific Toolkits for Specialized Workloads

- Intel® oneAPI HPC Toolkit
- Intel® oneAPI AI Analytics Toolkit
- Intel® oneAPI Rendering Toolkit
- Intel® oneAPI IoT Toolkit
- Intel® Distribution of OpenVINO™ toolkit, powered by oneAPI

Intel® oneAPI Tools for HPC

Intel® oneAPI HPC Toolkit

Deliver Fast Applications that Scale

What is it?

A toolkit that adds to the Intel® oneAPI Base Toolkit for building high-performance, scalable parallel code on C++, Fortran, SYCL, OpenMP & MPI from enterprise to cloud, and HPC to AI applications.

Who needs this product?

- OEMs/ISVs
- C++, Fortran, OpenMP, MPI Developers

Why is this important?

- Accelerate performance on Intel® Xeon® & Core™ processors and Intel accelerators
- Deliver fast, scalable, reliable parallel code with less effort built on industry standards

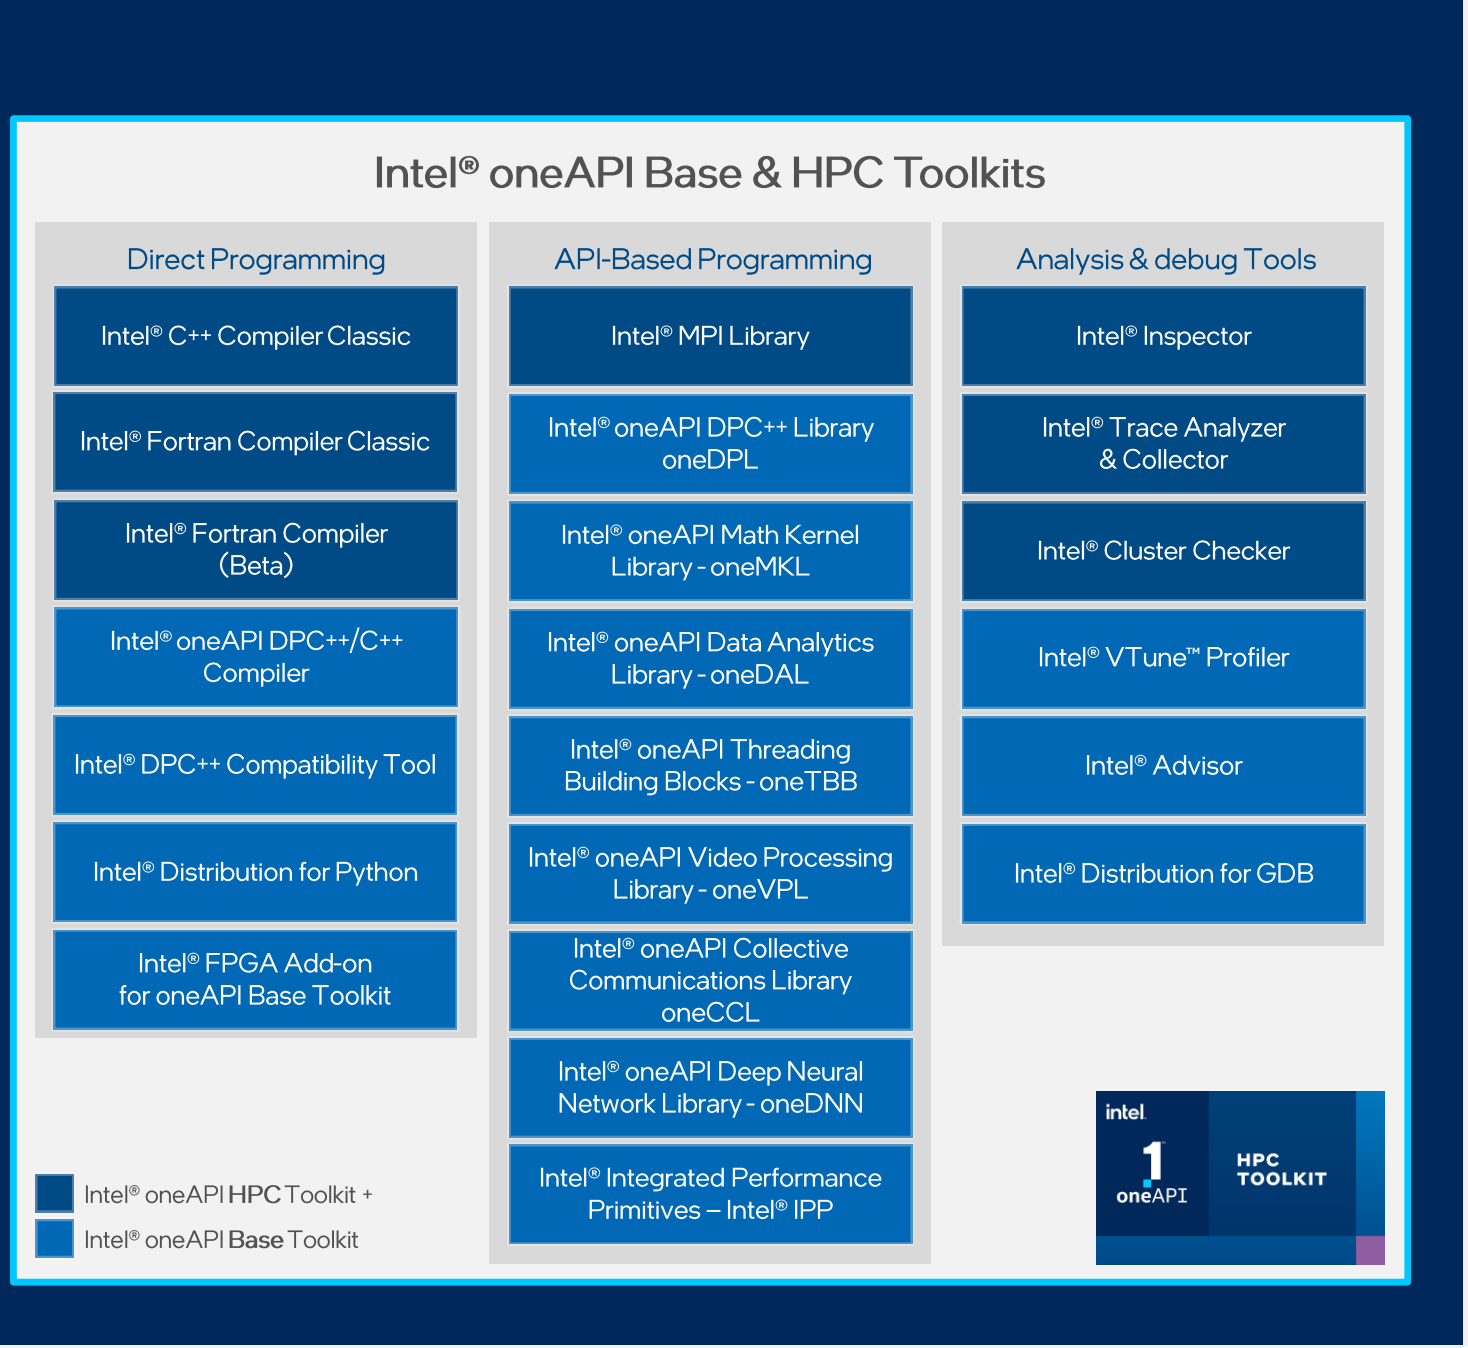

Learn More: intel.com/oneAPI-HPCKit

[Back to Domain-specific Toolkits for Specialized Workloads](#)

Deliver Fast HPC Applications that Scale

Customer Use Cases – Intel® oneAPI Base & HPC Toolkits

SAMPLE USE CASES & PROOF POINTS

Intel oneAPI tools help prepare code for Aurora.
Aurora, Argonne Leadership Computing Facility's Intel-HPE/Cray supercomputer, will be one of the U.S.'s 1st exascale systems

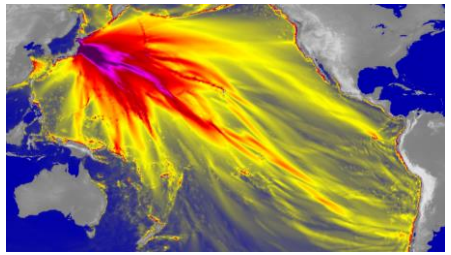

Zuse Institute Berlin (ZIB) ported the *easyWave* tsunami simulation application from CUDA to Data Parallel C++ (DPC++) **delivering performance on Intel CPUs, GPUs, FPGAs**, & Nvidia P100

Accelerating Google Cloud for HPC
C2 provides great performance for HPC workloads: **40% higher performance/core.** Runs on Intel® Xeon® processors + AMD, optimized by Intel® oneAPI **Base** & **HPC** Toolkits. Video
Video | Podkast

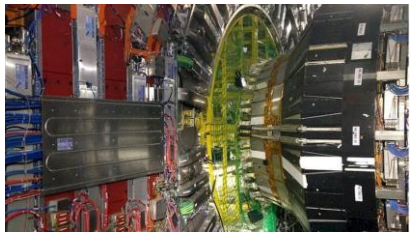

Acceleration for HPC & AI Inferencing
CERN, SURFsara, and Intel are investigating approaches driving **breakthrough performance on simulations** used in scientific, engineering, and financial applications*.

Texas Advanced Computing Center (TACC)
Frontera SuperComputer Visualization & Filesystem Use Cases Show Value of Large Memory Fat Nodes on Intel® Xeon® processors & Intel® Optane Persistent Memory*

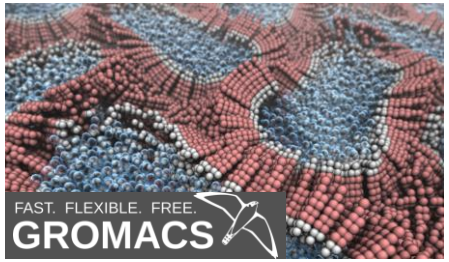

University of Stockholm/KTH
GROMACS, a simulation application used to design new drugs, was optimized by oneAPI. CUDA code was migrated to oneAPI to create new cross-architecture code targeting Intel CPUs and multiple accelerators.

Learn more: *oneAPI Discussions with HPC Thought Leaders* Video [2.20]

*Uses Intel® oneAPI Rendering Toolkit

Intel® oneAPI AI Analytics Toolkit

Accelerate end-to-end AI and data analytics pipelines with libraries optimized for Intel® architectures

Who Uses It?

Data scientists, AI researchers, ML and DL developers, AI application developers

Top Features/Benefits

- Deep learning performance for training and inference with Intel optimized DL frameworks and tools
- Drop-in acceleration for data analytics and machine learning workflows with compute-intensive Python packages

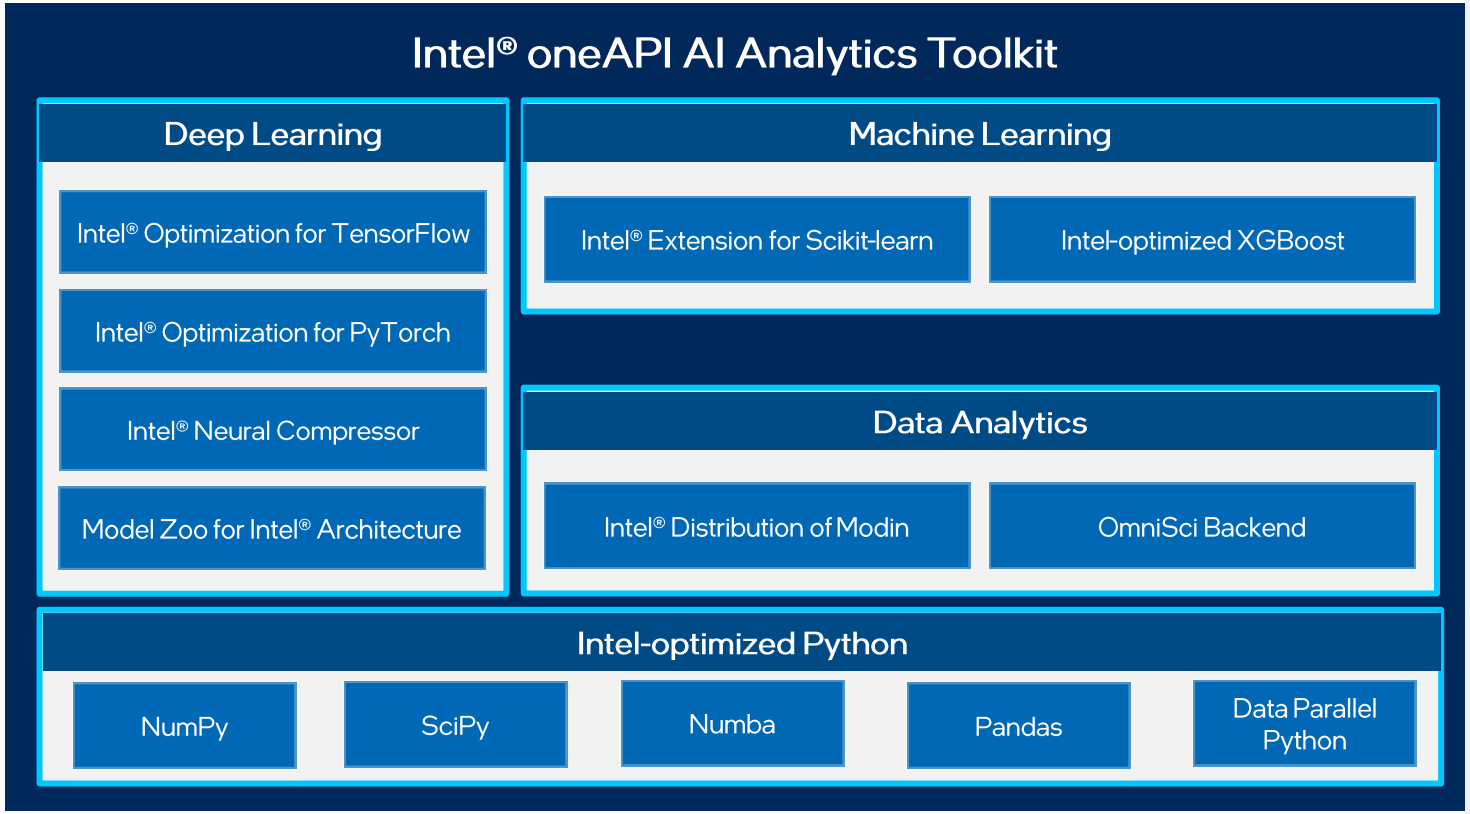

CPU GPU

Hardware support varies by individual tool. Architecture support will be expanded over time.

Get the Toolkit HERE or via these locations

Intel Installer | Docker | Apt, Yum | Conda | Intel® DevCloud

Learn More: software.intel.com/oneapi/ai-kit

Back to Domain-specific Toolkits for Specialized Workloads

Achieve End-to-End Performance for AI Workloads–Customer Use Cases

Accelerate Training + Inference - Most are optimized by Intel® oneAPI AI Analytics Toolkit

May also use Intel® Distribution of OpenVINO™ toolkit & Intel® oneAPI Base Toolkit, or Intel® oneAPI Rendering Toolkit

SAMPLE USE CASES & PROOF POINTS > View More Case Studies

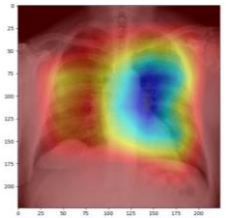

Accrad AI-based Solution Helps Accelerate Lung Disease Diagnosis – Acceleration for training + inference

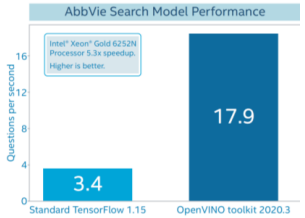

AbbVie Machine Translation Solution accelerates natural language processing inference models using processor optimized capabilities

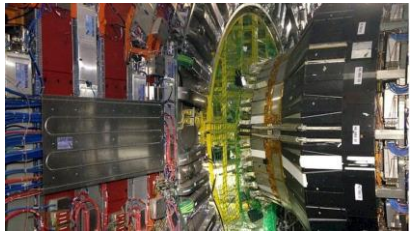

Acceleration for HPC & AI Inferencing - CERN, SURFsara, & Intel are driving breakthrough performance on scientific, engineering, and financial simulations. Includes strong inference benchmark.

AI Machine 3D Imaging – For many machine learning & AI tasks, **Daspatial** achieved significant performance improvements. The company is a design win – moving from Nvidia HW to Intel's, HW, & successfully migrated CUDA code to DPC++.

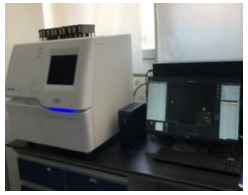

KFBIO AI-based Solution Helps Accelerate Tuberculosis Diagnosis demonstrating Intel inference performance leadership

Optimized by Intel® oneAPI Analytics Toolkit & Intel® Distribution of OpenVINO™ toolkit

AsiaInfo AI-based Solution optimized the AI end-to-end workflow for performance, and helps Accelerate 5G Network Intelligence

Bentley Motors Limited's Car Configurator uses real-time, accurate, visualization & AI for **1.7M+** rendered images in providing options to customers.
PR | Video | Demo | Session

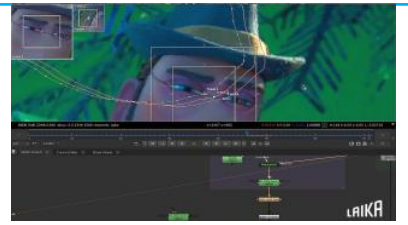

LAIKA Studios & Intel Join Forces to Expand What's Possible in Stop-Motion Filmmaking – decreases production time by transitioning manual work to AI-based processing
Article | Video

Intel & Facebook Accelerate PyTorch Performance on training workloads using Intel® Deep Learning Boost on 3rd gen Intel® Xeon® Scalable processors

Red Hat Optimizing Data Science Workflows
Red Hat Optimizing Data Science Workflows

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Render Your Vision in Highest Fidelity

Intel® oneAPI Rendering Toolkit

Powerful Libraries for High-Fidelity Visualization Applications

- Deliver high-performance, high-fidelity visualization applications on Intel® architecture
- Create amazing visual, hyper-realistic renderings via ray tracing with global illumination
- Access all system memory space to create renderings using the largest data sets
- Flexible, cost efficient development using open source libraries

Intel® Embree 2021 Awards

EuroGraphics Symposium on Rendering - Software Award

Academy Award® Technical Achievement Award

Award recipients: Sven Woop, Carsten Benthin, Attila Afra, Manfred Ernst, & Ingo Wald

Intel oneAPI Rendering & Ray Tracing Libraries

Intel® Embree
High-Performance, Feature-Rich Ray Tracing & Photorealistic Rendering

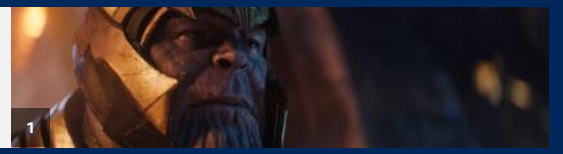

Intel® Open Image Denoise
AI-Accelerated Denoiser for Superior Visual Quality

Intel® OpenSWR
High-Performance, Scalable, OpenGL*-Compatible Rasterizer

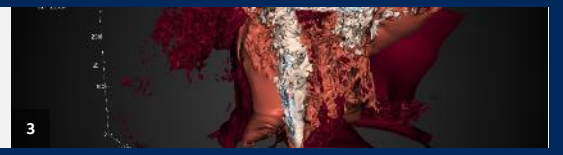

Intel® Open Volume Kernel Library
Render & Simulate 3D Spatial Data Processing

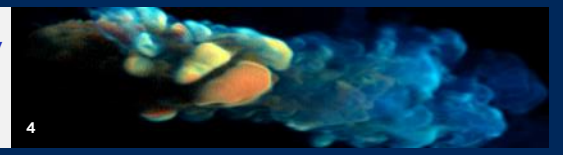

Intel® OSPRay
Scalable, Portable, Distributed Rendering API

Intel® OSPRay Studio
Real-time rendering through a graphical user interface with this new scene graph application

Intel® OSPRay for Hydra
Connect the Rendering Toolkit libraries to Universal Scene Description Hydra Rendering subsystem via plugin

Learn More: intel.com/oneAPI-RenderKit

intel oneAPI RENDERING TOOLKIT

[1] Avengers: Infinity War - Digital Domain, Marvel Studios, Chaos Group V-Ray
[2] Scene courtesy of Frank Meinl
[3] Model from Leigh Orf at University of Wisconsin. For more tornado visualization, visit Leigh Orf's site
[4] Smoke volume, data courtesy OpenVDB example repository
[5] Moana Island Scene, Walt Disney Animation Studios , publicly available dataset: 15fps+,~160 billion prims

Intel® oneAPI Rendering Toolkit

Render Your Vision in Highest Fidelity: Your Open Path to Advanced Ray Tracing

SAMPLE USE CASES & PROOF POINTS > View more case studies

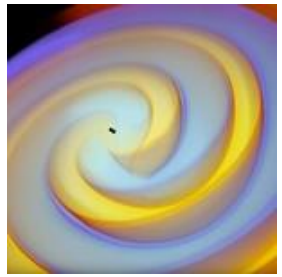

Stephen Hawking Centre for Cosmology Visualizes Cosmos Physics

Chaos Group
Case Study | Video [2:22] | Session [17:42]

Maxon's Cinema 4D Enhances CG Workflows with AI-Trained Denoising
Case Study | Session CG Workflow [7:03] | Tutorial Scene Creation/3D Scans [16:33]

Bentley Motors Limited's Car Configurator uses visualization & AI for **1.7M+** rendered images
PR | Video [2:31] | Demo [6:58] | Session [10:37]

For The Addams Family 2, Cinesite gained 10% to up to 25% efficiency in rendering.

Tangent Studios gained 5X-6X reduction in renders using Intel® Embree & AI-based Intel® Open Image Denoise
Video [2:35]

Texas Advanced Computing Center (TACC) Frontera SuperComputer Visualization & Filesystem Use Cases Show Value of Large Memory Fat Nodes

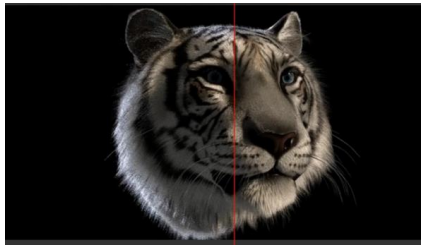

Mercenaries Engineering is delivering AI-based Animation & VFX through its production-ready tools and **improved rending times** by up to **138%** in renders.*

View a **Customer Showcase Reel** [Video 2:34]

[1]Courtesy Baozou Production in association with Tangent Animation using Blender with Intel® Embree. Media courtesy of Netflix, Inc. Now streaming on Netflix. Netflix subscription required.
*See slide notes for configuration details.
Refer to software.intel.com/articles/optimization-notice for more information regarding performance & optimization choices in Intel software products.
For workloads and configurations visit www.Intel.com/PerformanceIndex. Results may vary. Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy. .

Intel® oneAPI Tools for IoT

Intel® IoT Toolkit

Accelerate development of IoT applications for smart, connected devices that run at the networks edge

Who Uses It?

A broad range of application developers creating highly reliable IoT devices running on Intel CPUs, GPUs, & FPGAs

Top Features/Benefits

- Leverage more cores & built-in technologies in Intel® architecture-based platforms through optimized compilers & libraries
- Easily connect sensors to devices, and devices to cloud with IoT Connection Tools
- Speed development & maintenance of Yocto Project platform projects
- Develop with confidence with powerful analysis tools to identify threading, memory & offloading optimization opportunities
- DPC++ compatibility tool helps migrate existing code written in CUDA

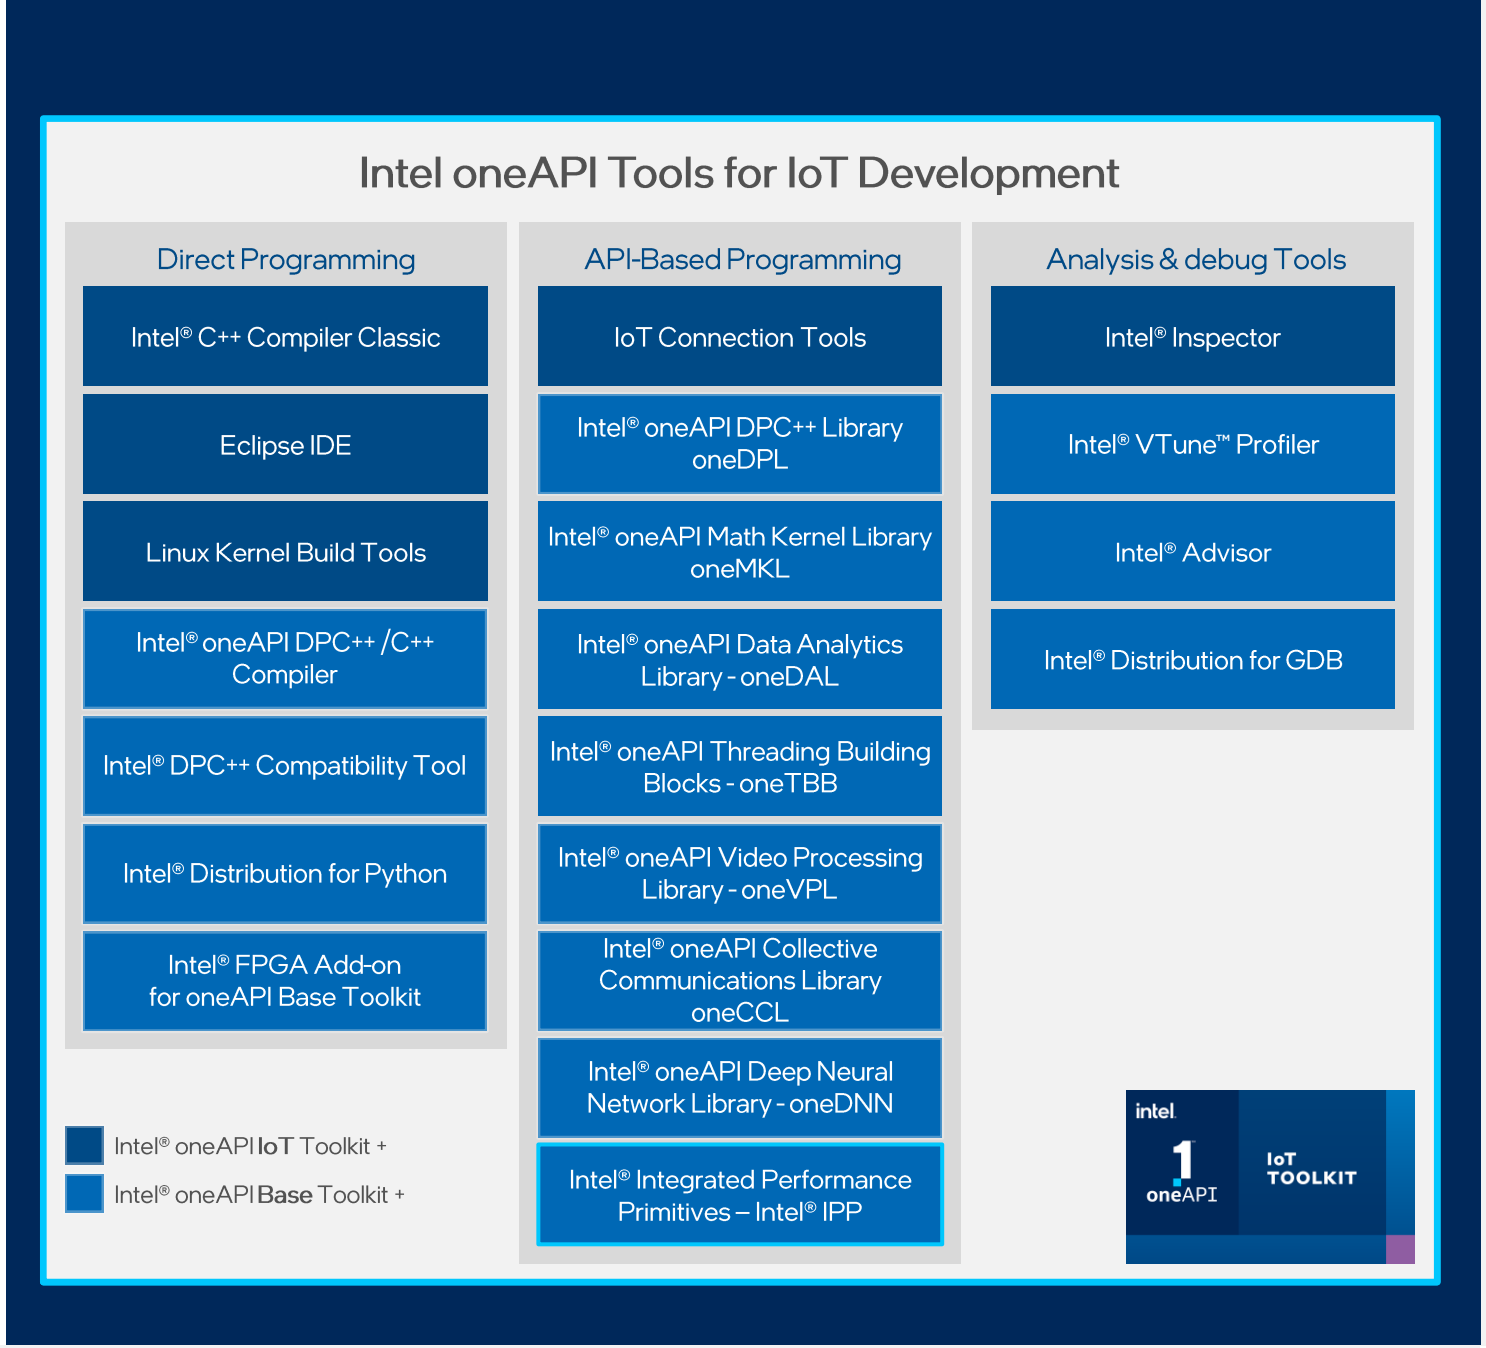

Learn More: intel.com/oneAPI-IoTKit

Back to Domain-specific Toolkits for Specialized Workloads

Accelerate Development of Smart, Connected Devices

Customer Use Cases

May be optimized one or a combination of the Intel® oneAPI Base, IoT, & AI Analytics Toolkits, & Intel® Distribution of OpenVINO™ toolkit

oneAPI

SAMPLE USE CASES & PROOF POINTS

Samsung Medison Uses oneAPI to Power Obstetric Ultrasound Systems
Intel® oneAPI Base Toolkit & Intel® Distribution of OpenVINO™ toolkit, powered by oneAPI, help accelerate image processing at the edge for consistent measurement accuracy & improved workflows.[1]

Intel PR News Byte Sept. 10, 2020 | Video [1.45]

United Imaging Successfully Ported Code Written in CUDA to oneAPI
United Imaging develops advanced medical products covering the entire process of imaging diagnosis and treatment; and offers innovative medical IT solutions. It used Intel oneAPI Base Toolkit for code migration and optimizations.

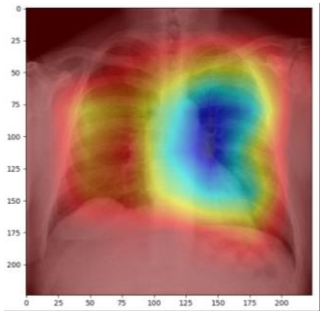

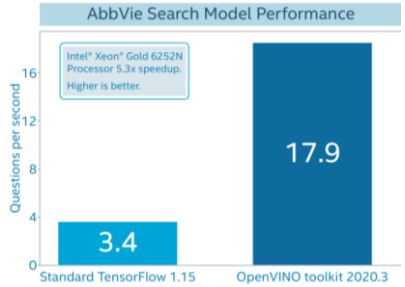

Optimized by Intel® oneAPI Analytics Toolkit & Intel® Distribution of OpenVINO™ toolkit

Accrad AI-based Solution Helps Accelerate Lung Disease Diagnosis – Acceleration for training + inference.

Learn more in the solution brief

AbbVie Machine Translation Solution accelerates natural language processing inference models using processor optimized capabilities.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel® Distribution of OpenVINO™ toolkit

Powered by oneAPI

Deliver High-Performance Deep Learning Inference

A toolkit to accelerate development of high-performance deep learning inference & computer vision in vision/AI applications used from edge to cloud. It enables deep learning on hardware accelerators & easy deployment across Intel® CPUs, GPUs, FPGAs, VPUs.

Who needs this product?

- Computer vision, deep learning software developers
- Data scientists
- OEMs, ISVs, System Integrators

Usages

Security surveillance, robotics, retail, healthcare, AI, office automation, transportation, non-vision use cases (speech, NLP, Audio, text) & more

Edge AI & Vision Alliance 2019 VISION PRODUCT OF THE YEAR

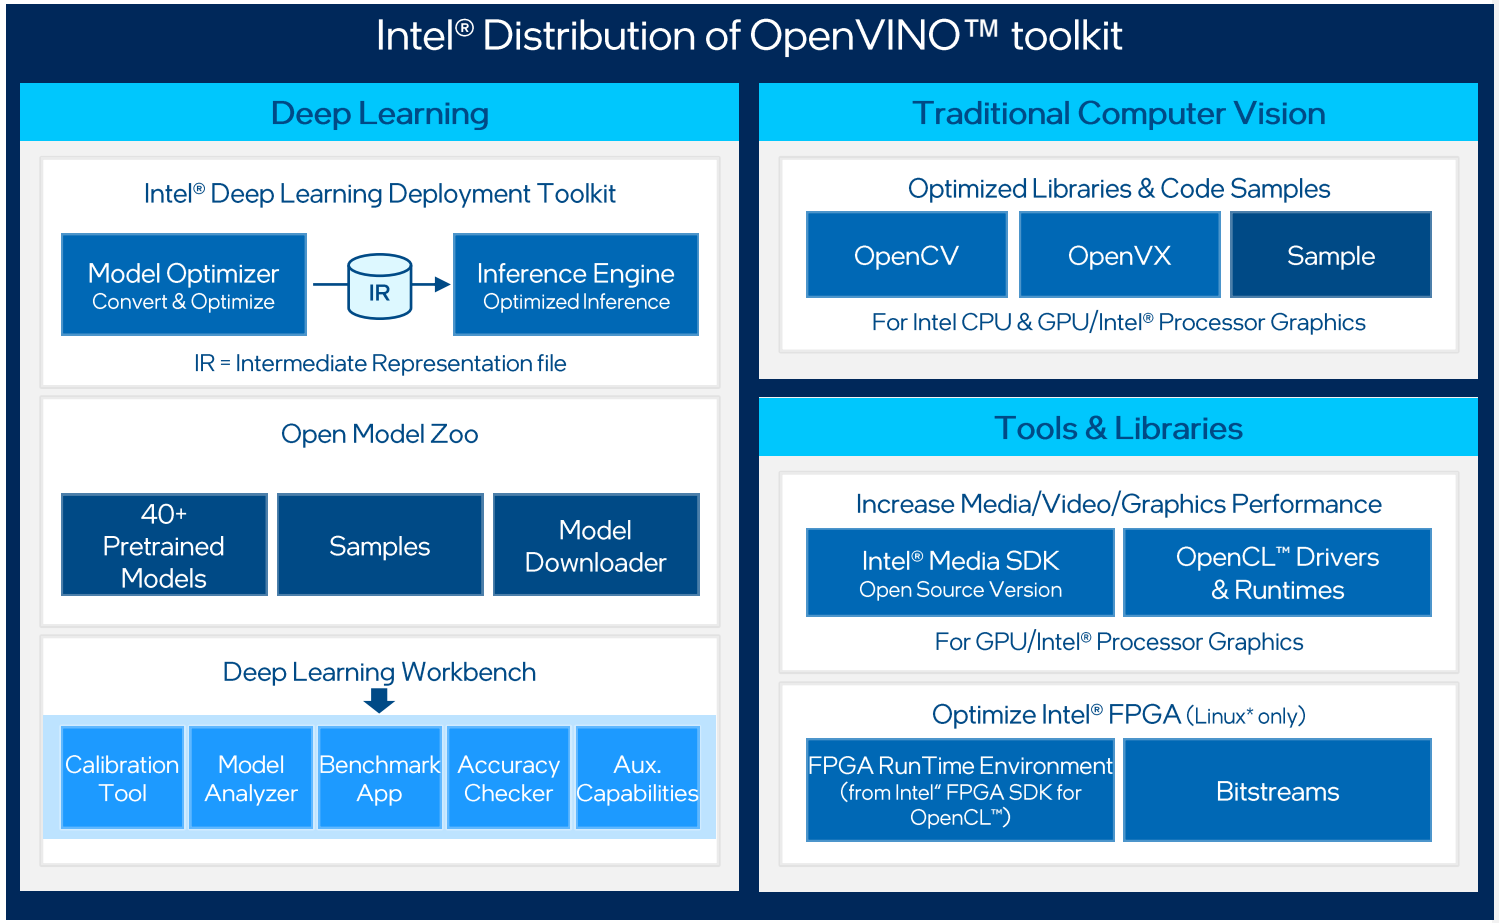

Back to Domain-specific Toolkits for Specialized Workloads

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex. Results may vary.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details.
No product or component can be absolutely secure.

Slide 50 - **Texas Advanced Computing Center (TACC) Frontera references**
Article: *HPCWire: Visualization & Filesystem Use Cases Show Value of Large Memory Fat Notes on Frontera*.
www.intel.com/content/dam/support/us/en/documents/memory-and-storage/data-center-persistent-mem/Intel-Optane-DC-Persistent-Memory-Quick-Start-Guide.pdf
software.intel.com/content/www/us/en/develop/articles/introduction-to-programming-with-persistent-memory-from-intel.html
wreda.github.io/papers/assise-osdi20.pdf
KFBIO
KFBIO m. tuberculosis screening detectron2 model throughput performance on 2nd Intel® Xeon® Gold 6252 processor: NEW: Test 1 (single instance with PyTorch 1.6: Tested by Intel as of 5/22/2020. 2-socket 2nd Gen Intel® Xeon® Gold 6252 Processor, 24 cores, HT On, Turbo ON, Total Memory 192 GB (12 slots/16 GB/2666 MHz), BIOS: SSE5C620.86B.02.01.0008.031920191559 (ucode: 0x500002c), Ubuntu 18.04.4 LTS, kernel 5.3.0-51-generic, mitigated Test 2 (24 instances with PyTorch 1.6: Tested by Intel as of 5/22/2020. 2-socket 2nd Gen Intel Xeon Gold 6252 Processor, 24 cores, HT On, Turbo ON, Total Memory 192 GB (12 slots/16 GB/2666 MHz), BIOS: SSE5C620.86B.02.01.0008.031920191559 (ucode: 0x500002c), Ubuntu 18.04.4 LTS, kernel 5.3.0-51-generic, mitigated BASELINE: (single instance with PyTorch 1.4): Tested by Intel as of 5/22/2020. 2-socket 2nd Gen Intel Xeon Gold 6252 Processor, 24 cores, HT On, Turbo ON, Total Memory 192 GB (12 slots/16 GB/2666 MHz), BIOS: SSE5C620.86B.02.01.0008.031920191559 (ucode: 0x500002c), Ubuntu 18.04.4 LTS, kernel 5.3.0-51-generic, mitigated.

Tangent Studios
Configurations for Render Times with Intel® Embree, testing conducted by Tangent Animation Labs. Render farm: 8x Intel® Core™ processors +hyperthread*2 + 128gig. In-office workstations: Intel® Xeon® processors HP blade c7000 chassis, with HP460 gen8 blades - 2x Intel Xeon E5-2650 V2, Eight Core 2.6GHz-128GB. Software: Blender 2.78 with custom build using Intel® Embree. For more information on Tangent's work with Embree, watch this video:
www.youtube.com/watch?time_continue=251&v=_2Ia4h8q3xs&feature=emb_logo
Recreation of the performance numbers can be recreated using Agent327, Blender and Embree.

Chaos Group - Up to 90% Memory Reduction for Displacement
Testing conducted by Chaos Group with Intel® Embree 2020. Software Corona Renderer 5 with Intel Embree. Up to 90% memory reduction calculated using Corona Renderer 5 with regular displacement grids per triangle of 154 bytes versus Corona Renderer 5 with Intel Embree, which has a displacement capability grid of 12 bytes per grid triangle. (12/154 = 7.8% usage or >90% memory reduction.) Recreation of the performance numbers can be accomplished using Corona Renderer 5 and Embree. For more information, visit the Corona Renderer Blog: blog.corona-renderer.com/corona-renderer-5-for-3ds-max-released/

The Addams Family 2 - Gained a 10% to 20%—and sometimes 25%—efficiency in rendering, saving thousands of hours in rendering production time.
Testing Date: Results are based on data conducted by Cinesite 2020-21. 10% to up to 25% rendering efficiency/thousands of hours saved in rendering production time/15 hrs per frame per shot to 12-13 hrs.
Cinesite Configuration: 18-core Intel® Xeon® Scalable processors (W-2295) used in render farm, 2nd gen Intel Xeon processor-based workstations (W-2135 and -2195) used. Rendering tools: Gaffer, Arnold, along with optimizations by Intel® Open Image Denoise.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.


intel