

Joint INTEL / EuroCC-Castiel oneAPI Workshop

February 16th and 17th, 2022



Goals and Objectives

Get deeper insights into heterogeneous programming with the Intel oneAPI Development tools

Share knowledge from Intel specialists and from external speaker

Interactive Sessions with Q+A / Use Chat Room

This is a ,public‘ event –
No Confidential Material to be shared

Agenda – Day 1

TOPIC				Presenter
09:00	09:10	00:10	Welcome and Introduction to Day 1	EuroCC, Edmund Preiss (Intel)
09:10	09:40	00:30	oneAPI – Introduction to a new Development Environment - Concept and oneAPI Standardization initiative - Intel's Tools Implementation – Intel oneAPI Toolkits and libs - Transition from Intel Parallel Studio XE to Intel oneAPI toolkits	Edmund Preiss (Intel)
09:40	10:00	00:20	Introduction to the DevCloud - Purpose: Demoing, testing and porting applications - Hardware and Software offerings - How to onboard & how to get an DevCloud account	Klaus-Dieter Oertel (Intel)
10:00	10:05	00:05	Bio Break	
10:05	12:05	02:00	Direct programming with oneAPI Compilers – with Demos - Intro to DPC++ heterogenous programming model and SYCL 2020 - Fortran and offloading with OpenMP5.0 - DPC++ features and examples <ul style="list-style-type: none">o “Hello World” Exampleo Device Selectiono Compilation and Execution Flowo Execution Modelo Memory Model; Buffers, Unified Shared Memory (USM)o Kernel Model	Igor Vorobtsov (Intel)
12:05	13:15	01:10	Lunch Break	
13:15	14:15	01:00	Compatibility tool for porting purposes - with demo - Migration Cuda based GPU Applications to oneAPI (DPC++)	Igor Vorobtsov (Intel)
14:15	15:00	00:45	Intel oneAPI libraries (oneMKL, oneDAL, ...) - with demos - Performance optimized libraries for AI, numerical simulations and other purposes	Gennady Fedorov (Intel)
15:00	15:05	00:05	Bio Break	
15:05	15:35	00:30	Debugging on GPU /Debugging on LLVM (TBD)	Alina Shadrina (Intel)
15:35	16:35	01:00	Intel's oneAPI AI Analytics Toolkit solutions for Artificial Intelligence (AI) - Intel performance optimized Python Distribution - Low level libs with basic oneDAL, py4Dal - Optimized ML libraries such as XGBoost and SciKit Learn, Modin - Optimized Tools Frameworks for Deep Learning	Shailen Sohbee (Intel)
16:35	16:40	00:05	- Questions and Answers - Wrap up	Intel

Agenda – Day 2

TOPIC				Presenter
09:00	09:10	00:10	Welcome and Introduction to Day 1	EuroCC, Edmund Preiss (Intel)
09:10	09:40	00:30	oneAPI – Introduction to a new Development Environment - Concept and oneAPI Standardization initiative - Intel's Tools Implementation – Intel oneAPI Toolkits and libs - Transition from Intel Parallel Studio XE to Intel oneAPI toolkits	Edmund Preiss (Intel)
09:40	10:00	00:20	Introduction to the DevCloud - Purpose: Demoing, testing and porting applications - Hardware and Software offerings - How to onboard & how to get an DevCloud account	Klaus-Dieter Oertel (Intel)
10:00	10:05	00:05	Bio Break	
10:05	12:05	02:00	Direct programming with oneAPI Compilers – with Demos - Intro to DPC++ heterogenous programming model and SYCL 2020 - Fortran and offloading with OpenMP5.0 - DPC++ features and examples <ul style="list-style-type: none">o “Hello World” Exampleo Device Selectiono Compilation and Execution Flowo Execution Modelo Memory Model; Buffers, Unified Shared Memory (USM)o Kernel Model	Igor Vorobtsov (Intel)
12:05	13:15	01:10	Lunch Break	
13:15	14:15	01:00	Compatibility tool for porting purposes - with demo - Migration Cuda based GPU Applications to oneAPI (DPC++)	Igor Vorobtsov (Intel)
14:15	15:00	00:45	Intel oneAPI libraries (oneMKL, oneDAL, ...) - with demos - Performance optimized libraries for AI, numerical simulations and other purposes	Gennady Fedorov (Intel)
15:00	15:05	00:05	Bio Break	
15:05	15:35	00:30	Debugging on GPU /Debugging on LLVM (TBD)	Alina Shadrina (Intel)
15:35	16:35	01:00	Intel's oneAPI AI Analytics Toolkit solutions for Artificial Intelligence (AI) - Intel performance optimized Python Distribution - Low level libs with basic oneDAL, py4Dal - Optimized ML libraries such as XGBoost and SciKit Learn, Modin - Optimized Tools Frameworks for Deep Learning	Shailen Sohbee (Intel)
16:35	16:40	00:05	- Questions and Answers - Wrap up	Intel

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

oneAPI Industry Initiative & Intel® oneAPI Tools

Software & Advanced Technologies Group (SATG)
Software Products & Ecosystem
February 2021

Edmund Preiss
Intel Business Development Manager



Programming Challenges

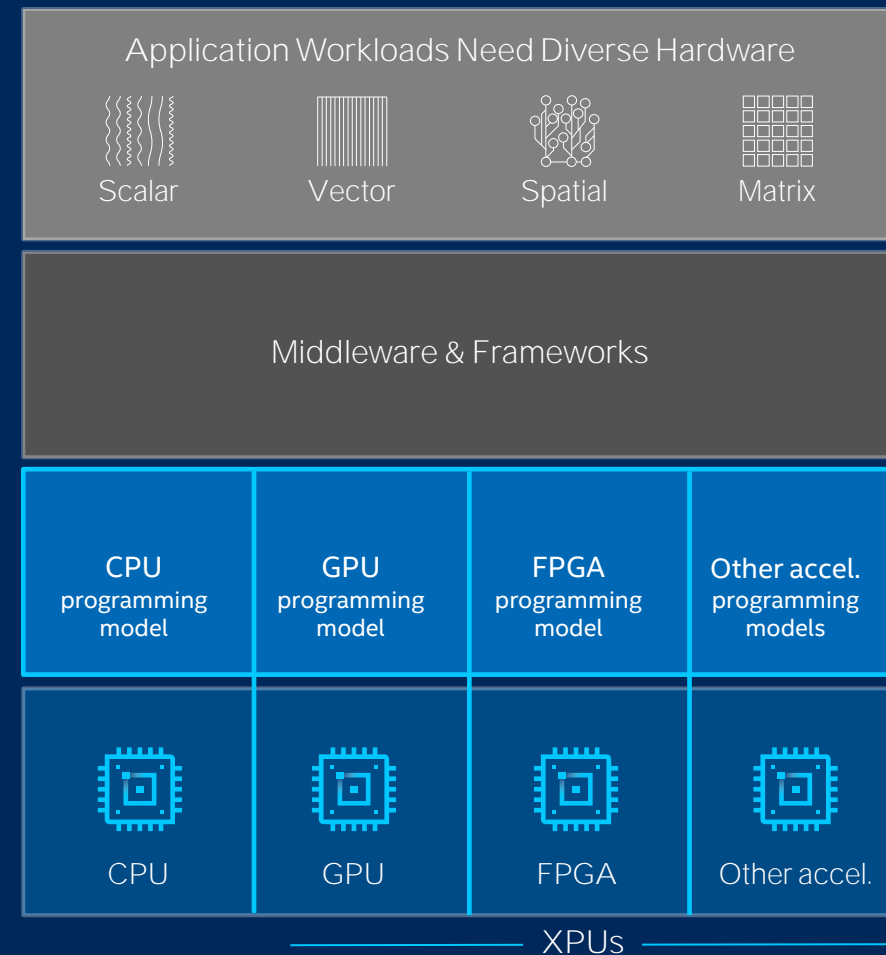
for Multiple Architectures

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

One Programming Model for Multiple Architectures and Vendors



Freedom to Make Your Best Choice

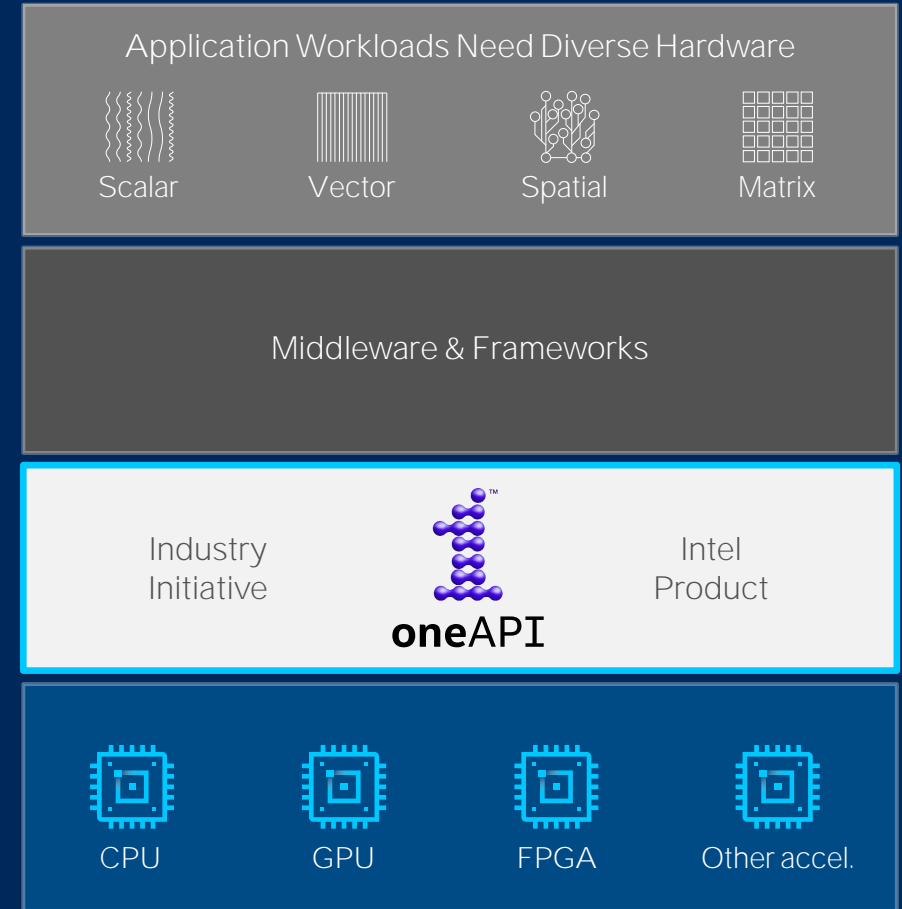
- Choose the best accelerated technology the software doesn't decide for you

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

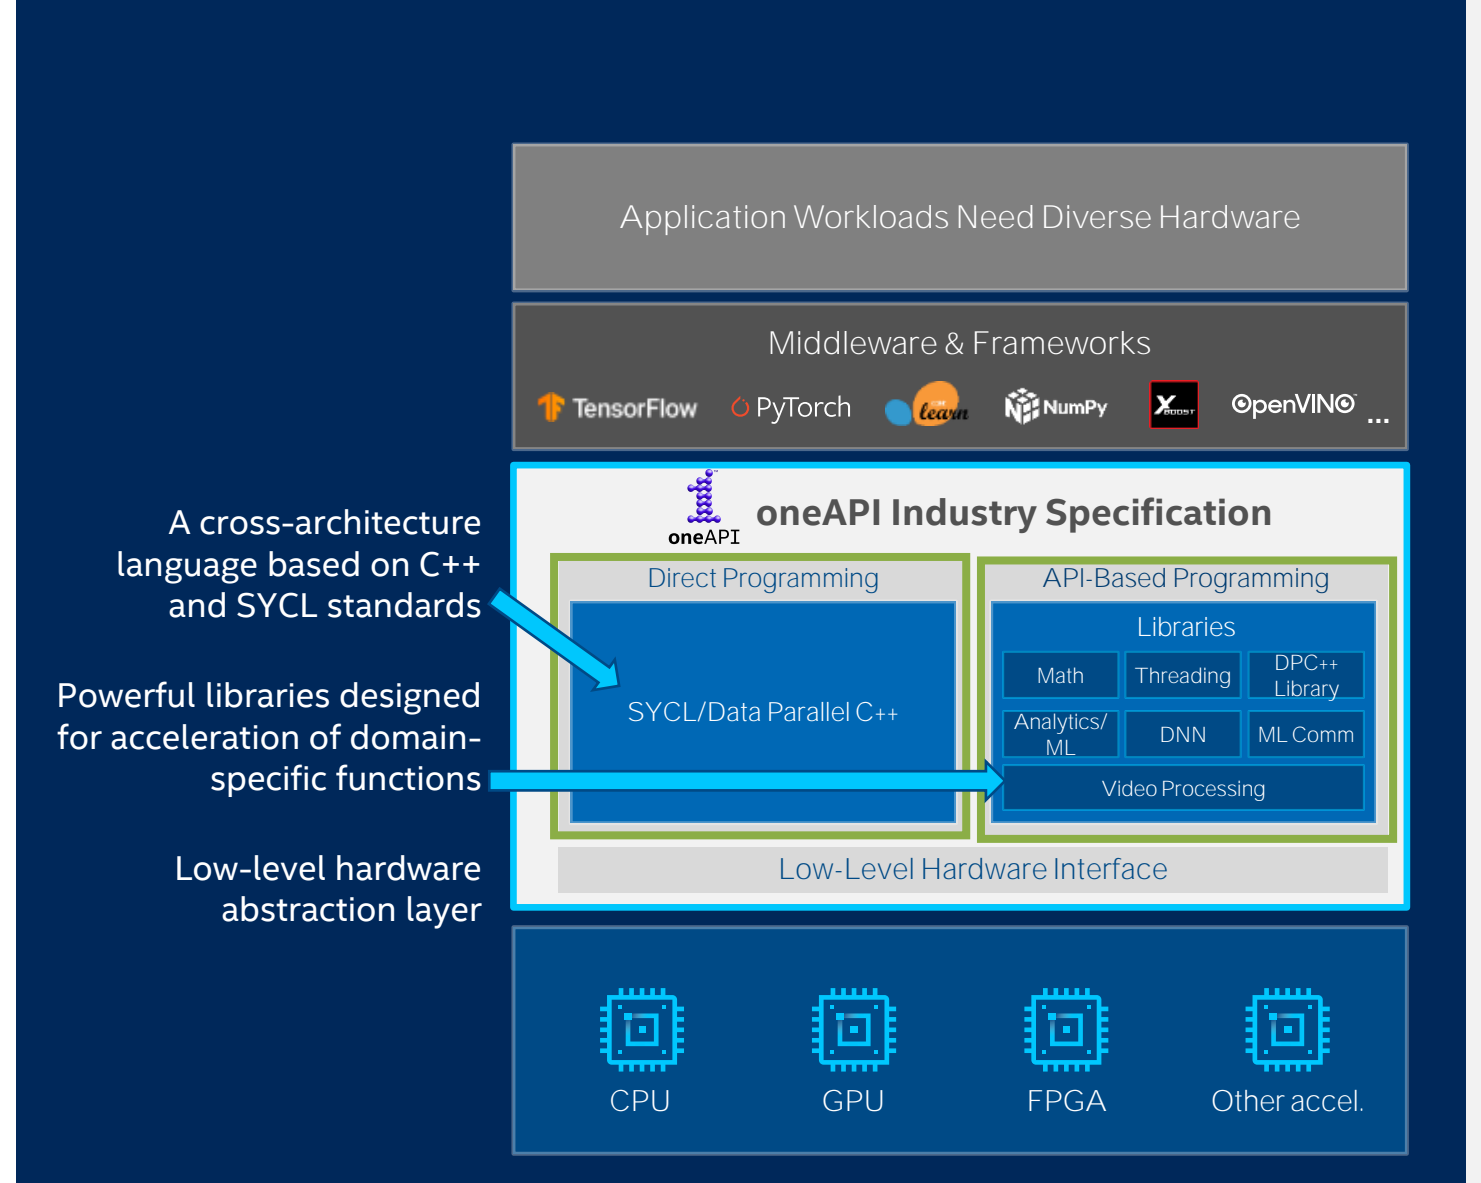
Open to promote community and industry collaboration


Enables code reuse across architectures and vendors



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

Visit oneapi.com for more details





oneAPI

Industry Specification

spec.oneapi.com/oneAPI/

- [Notices and Disclaimers](#)
- [Contribution Guidelines](#)
- [Introduction](#)
- [Software Architecture](#)
- [Library Interoperability](#)
- [oneAPI Elements](#)
- [Data Parallel C++ \(DPC++\)](#)
- [oneAPI Data Parallel C++ Library \(oneDPL\)](#)
- [oneAPI Deep Neural Network Library \(oneDNN\)](#)
- [oneAPI Collective Communications Library \(oneCCL\)](#)
- [oneAPI Level Zero \(Level Zero\)](#)
- [oneAPI Data Analytics Library \(oneDAL\)](#)
- [oneAPI Threading Building Blocks \(oneTBB\)](#)
- [oneAPI Video Processing Library \(oneVPL\)](#)
- [oneAPI Math Kernel Library \(oneMKL\)](#)
- [Contributors](#)

Data Parallel C++

Standards-based, Most Comprehensive,
Cross-architecture Implementation of SYCL

DPC++ = ISO C++ and Khronos SYCL and community extensions

Freedom of Choice: Future-Ready Programming Model

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

DPC++ = ISO C++ and Khronos SYCL and community extensions

- Designed for data parallel programming productivity
- Provides full native high-level language performance on par with standard C++ and broad compatibility
- Adds SYCL from the Khronos Group for data parallelism and heterogeneous programming

Community Project Drives Language Enhancements

- Provides extensions to simplify data parallel programming
- Continues evolution through open and cooperative development

DPC++ aims to be the best implementation of SYCL

Direct Programming:
SYCL/Data Parallel C++

Community Extensions

Khronos SYCL

ISO C++

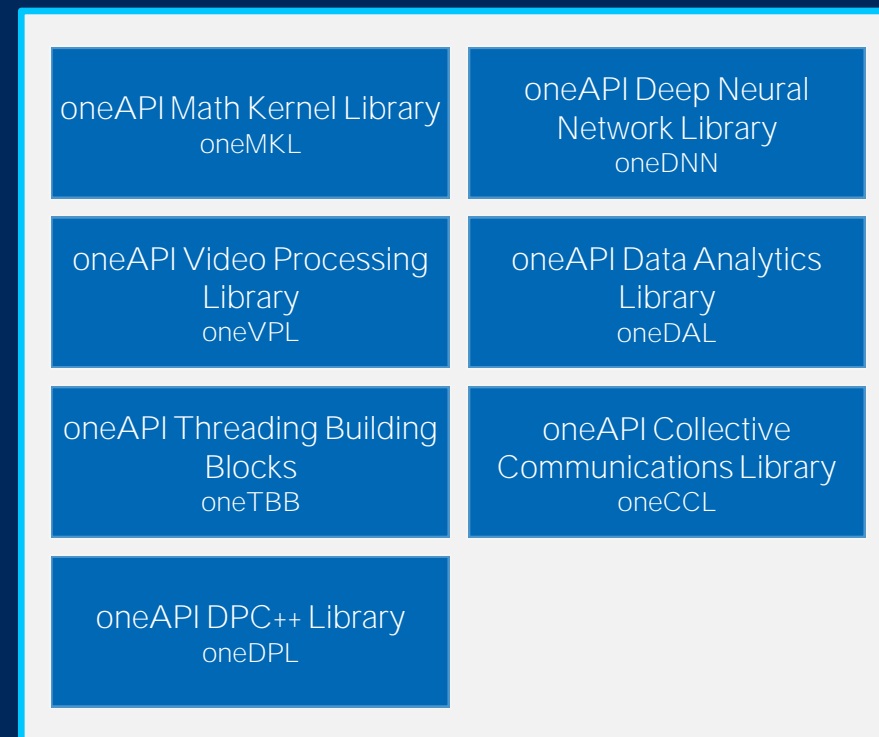
Powerful oneAPI Libraries

Realize all the Hardware Value

Designed for acceleration of key domain-specific functions

Freedom of Choice

Pre-optimized for each target platform for maximum performance



oneAPI: Open Accelerator Ecosystem

Freedom of Choice in Hardware Drives Productivity

Codeplay contribution to DPC++ brings SYCL support for **NVIDIA GPUs**

oneAPI oneDNN on Arm for A64FX Fugaku

Extending DPC++ with Support for **Huawei AI Chipset**

NERSC, ALCF, CODEPLAY PARTNER ON SYCL FOR NEXT-GENERATION SUPERCOMPUTERS

on Nvidia

ARGONNE, ORNL AWARD CODEPLAY CONTRACT TO STRENGTHEN SYCL SUPPORT FOR **AMD GPUS**


European exascale combines **SiPearl's CPU RHEA** with Intel's X^e GPU PVC

"DPC++ and oneAPI helped us to develop much faster the accelerators for machine learning algorithms."

– Chris Kachris, co-founder, InAccel

"If you like modern, standard C++ and you want to target GPUs or other accelerators, you will love SYCL!"

– Marcel Breyer



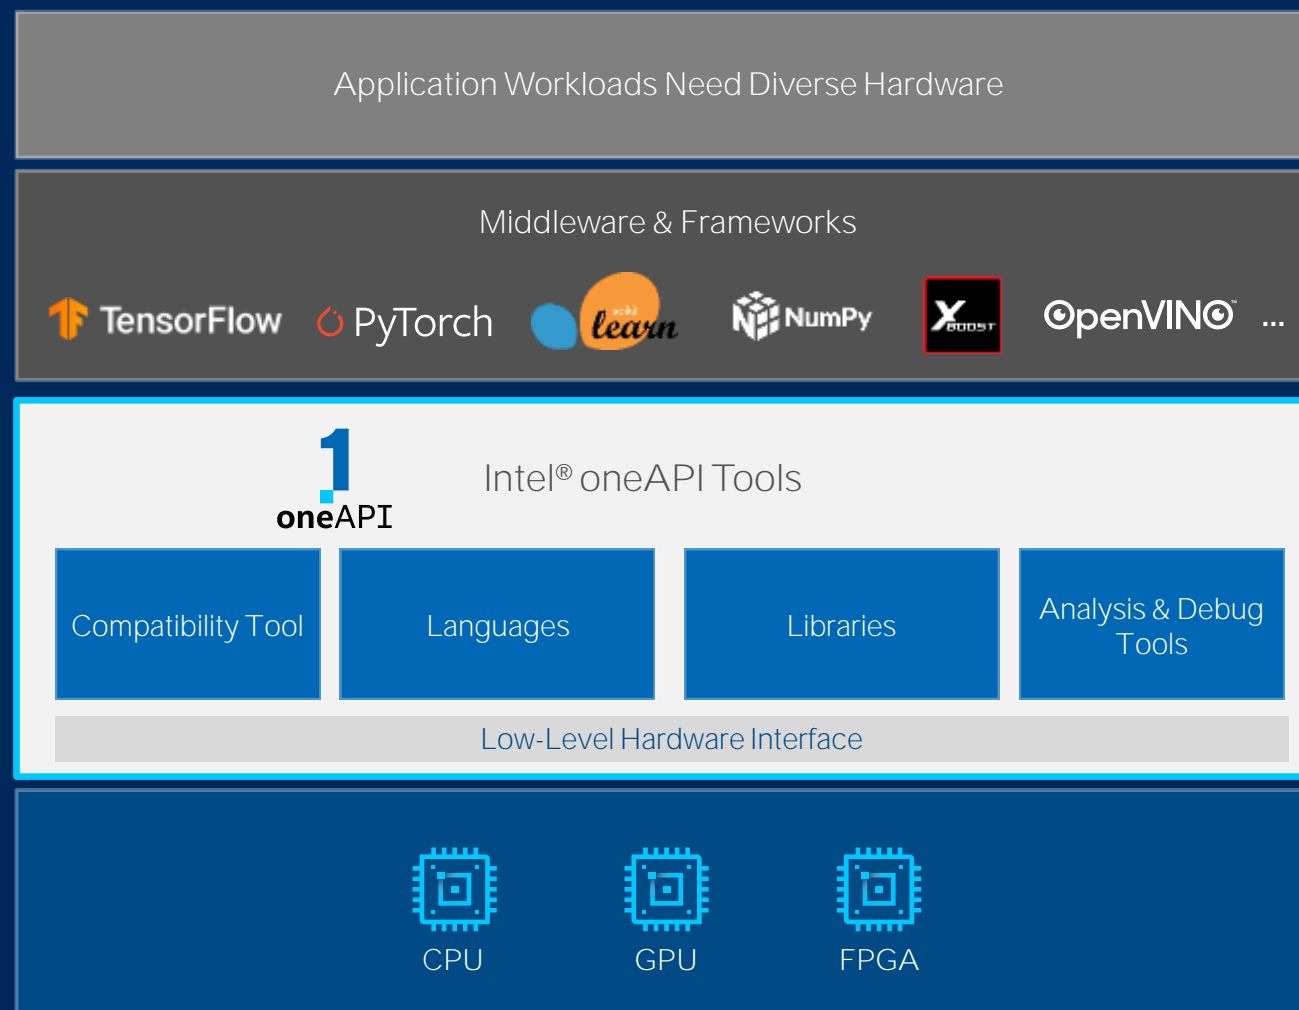
Visualization of easyWave tsunami simulation application -
Courtesy Zuse Institute Berlin (ZIB)

Intel® oneAPI Tools

Built on Intel's Rich Heritage of CPU Tools Expanded to XPU

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++, SYCL, 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® DPC++ Compatibility Tool

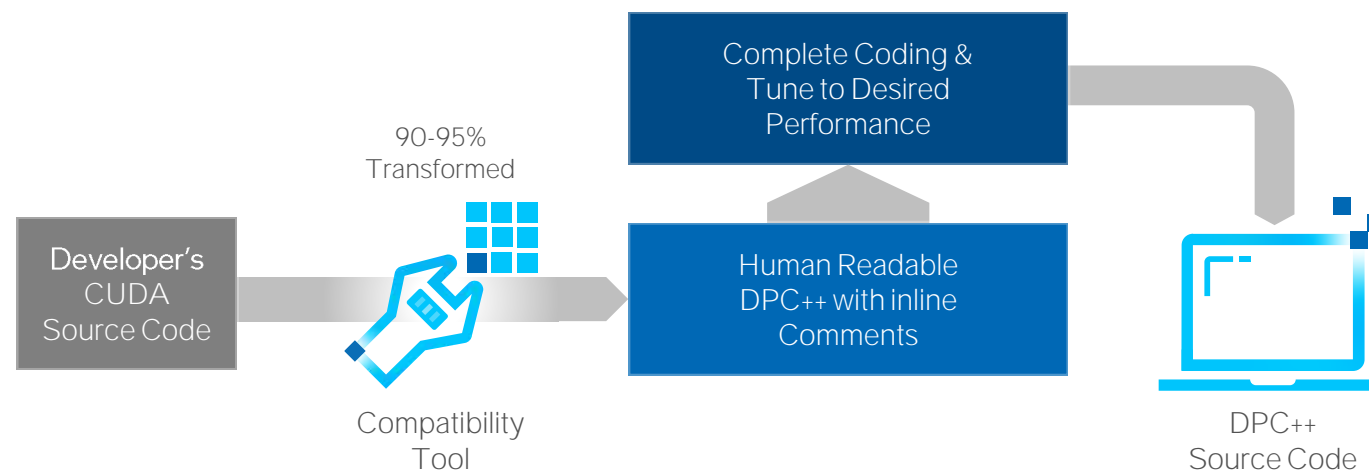
Minimizes Code Migration Time

Assists developers migrating code written in CUDA to Data Parallel C++ once, generating **human readable** code wherever possible

~90-95% of code typically migrates automatically¹

Inline comments are provided to help developers finish porting the application

Intel DPC ++ Compatibility Tool Usage Flow



¹Intel 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.

Analysis & Debug Tools

Get More from Diverse Hardware



Design

Intel® Advisor

- Efficiently offload code to GPUs
- Optimize your CPU/GPU code for memory and compute
- Enable **more** vector parallelism and improve efficiency
- Add effective threading to unthreaded applications



Debug

Intel® Distribution for GDB

- Multiple accelerator support with CPU, GPU and FPGA
- Enables deep, system-wide debug of Data Parallel C++ (DPC++), C, C++, and Fortran code



Tune

Intel® VTune™ Profiler

- Tune for GPU, CPU, and FPGA
- Optimize offload performance
- Supports DPC++, C, C++, Fortran, Python, Go, Java or a mix of languages




Intel® oneAPI Toolkits

A complete set of proven developer tools expanded from CPU to XPU (accelerators)



Intel® oneAPI Base Toolkit

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



Add-on Domain-specific Toolkits



Intel® oneAPI Tools for HPC

Deliver fast Fortran, OpenMP & MPI applications that scale



Intel® oneAPI Tools for IoT

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



Intel® oneAPI AI Analytics Toolkit

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



Intel® oneAPI Rendering Toolkit

Create performant, high-fidelity visualization applications

Toolkit
powered by oneAPI



Intel® Distribution of OpenVINO™ Toolkit

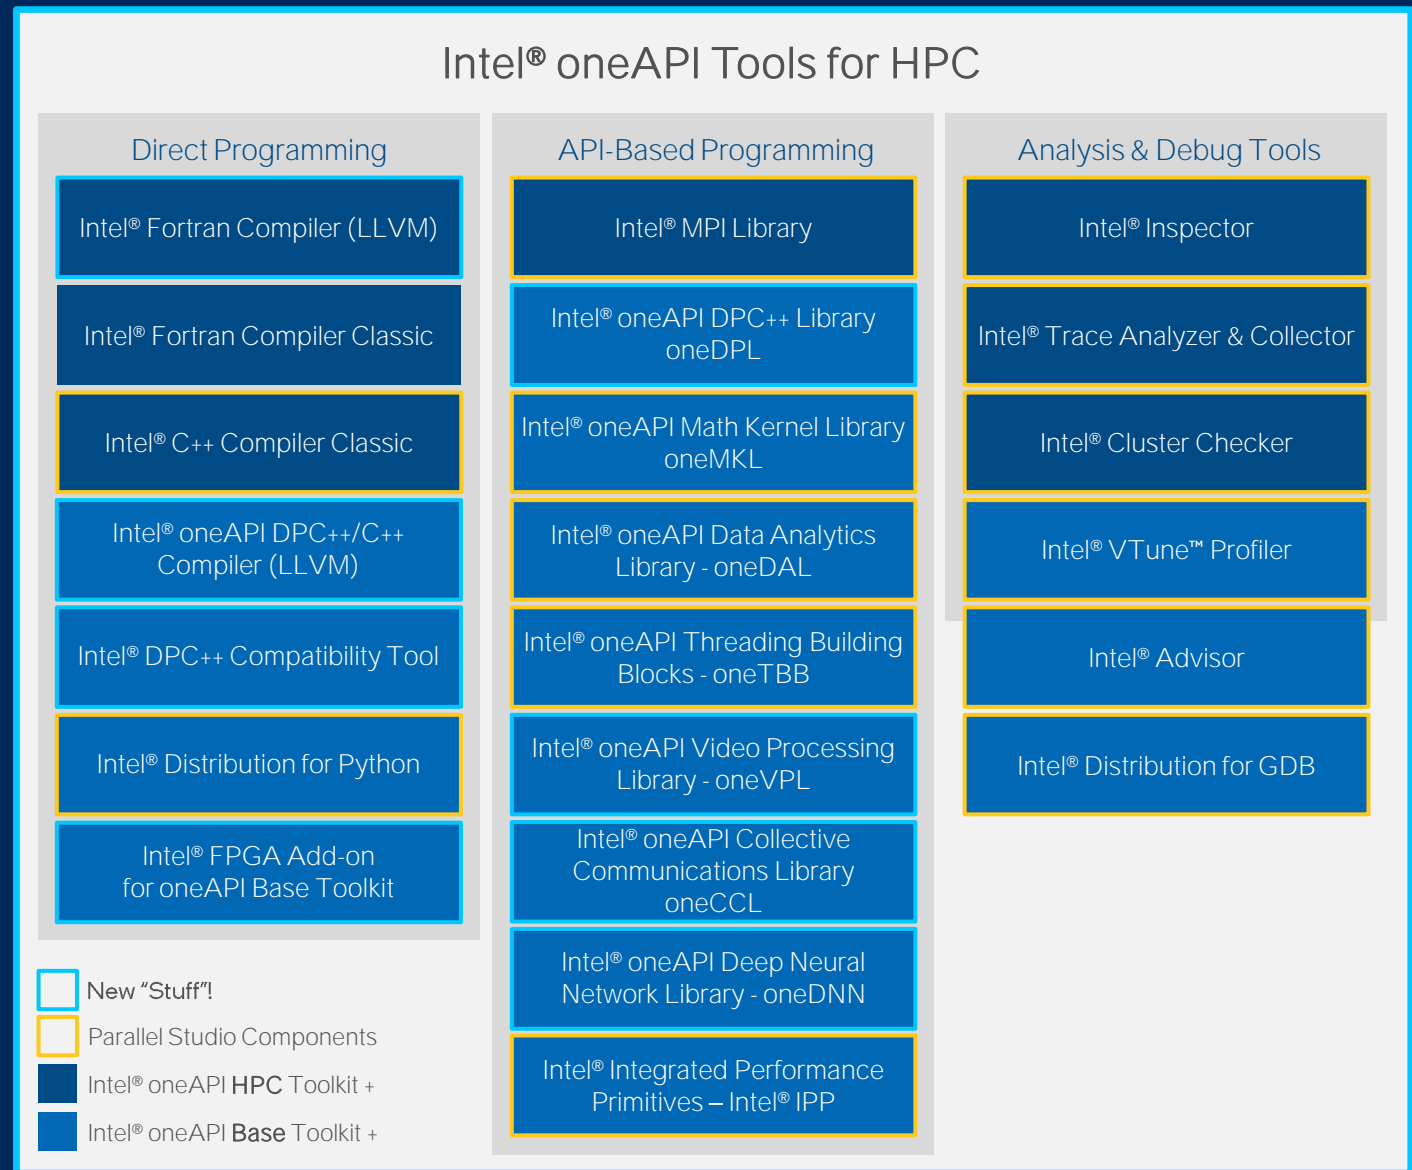
Deploy high performance inference & applications from edge to cloud

Latest version available 2022.1

Intel® Parallel Studio XE Product Transition

All components delivered today will be available in oneAPI Toolkits along with **new “Stuff”**

- Intel Priority Support continues without interruption
- Transition goal from Parallel Studio to the oneAPI Toolkits is to make it as smooth as possible for customers
- Existing customers will be getting Intel® oneAPI Base & HPC Toolkits as if they are the next release of what they have
- Protects and extends customers investment in Intel Software development tools



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*

¹ Avengers: Infinity War - Digital Domain, Marvel Studios, Chaos Group V-Ray

² Scene courtesy of Frank Meinel

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 & 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 AI Analytics Toolkit

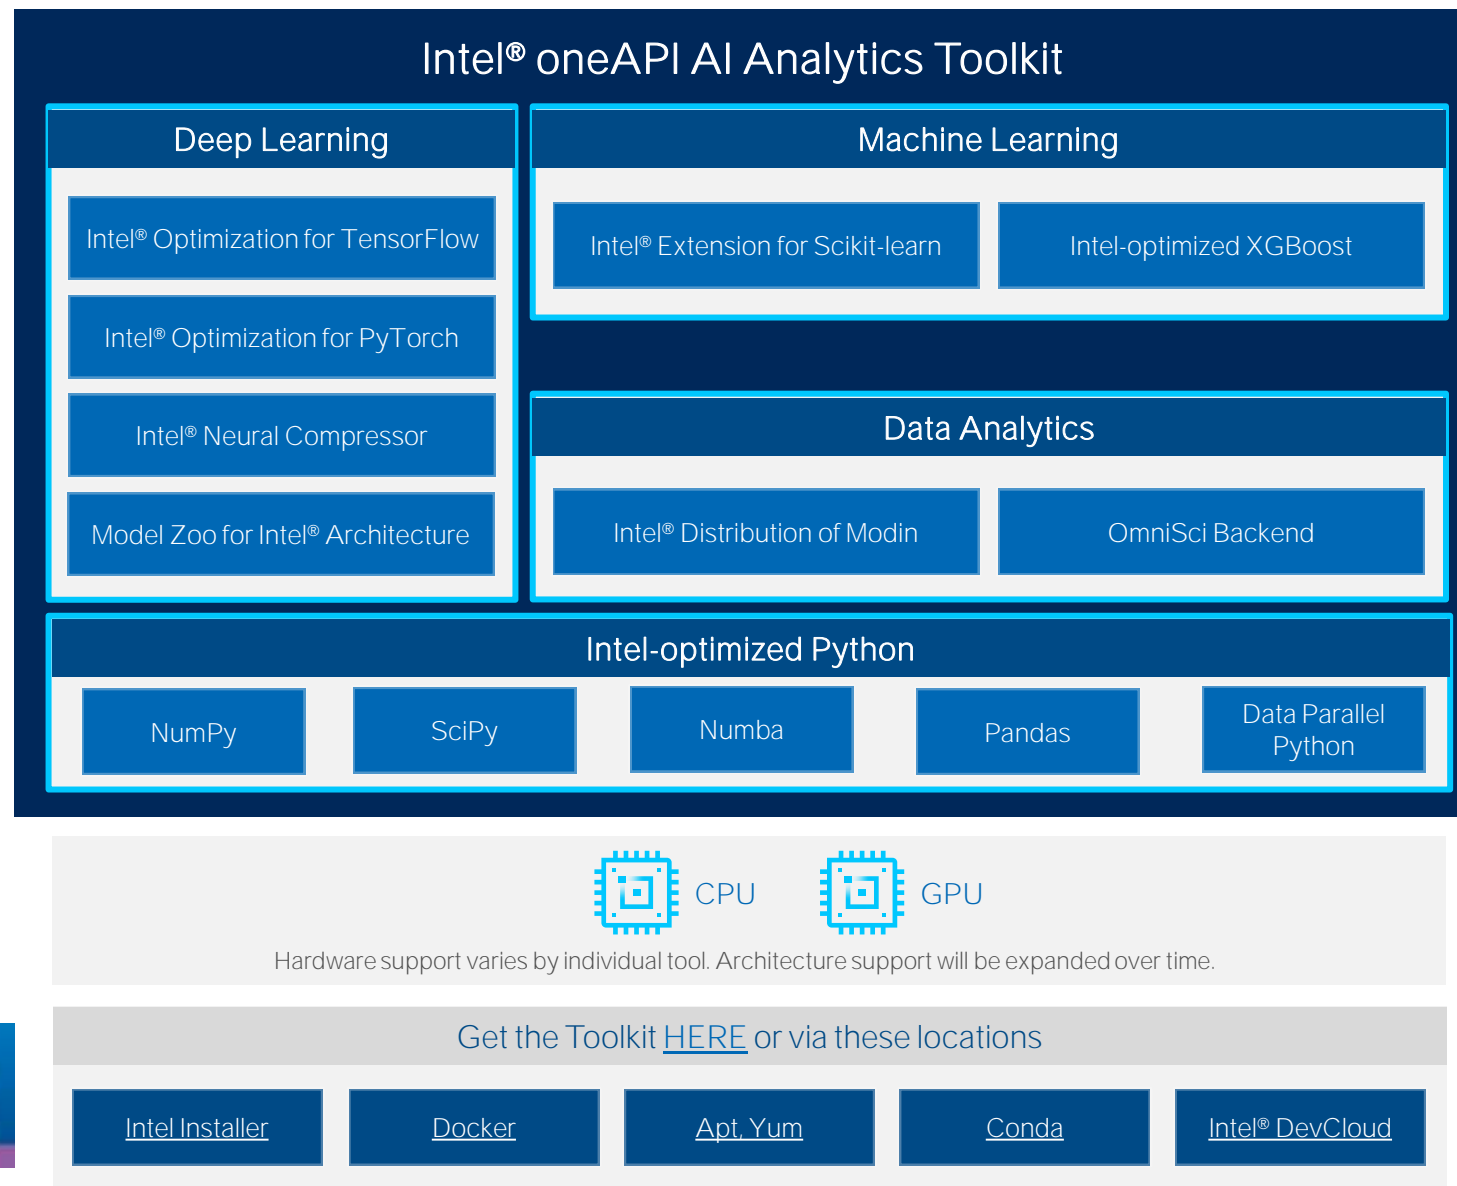
Accelerate end-to-end AI and data analytics pipelines with libraries optimized for Intel® architectures

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



Intel® oneAPI Toolkits Availability

Get Started Quickly

Code Samples, Quick-start Guides, Webinars, Training

software.intel.com/oneapi

Run the tools locally

 Downloads

 Repositories

 Containers

Or run the tools in intel. DevCloud

1 Minute to Code

No Hardware Acquisition

No Download, Install or Configuration

Samples & Tutorials

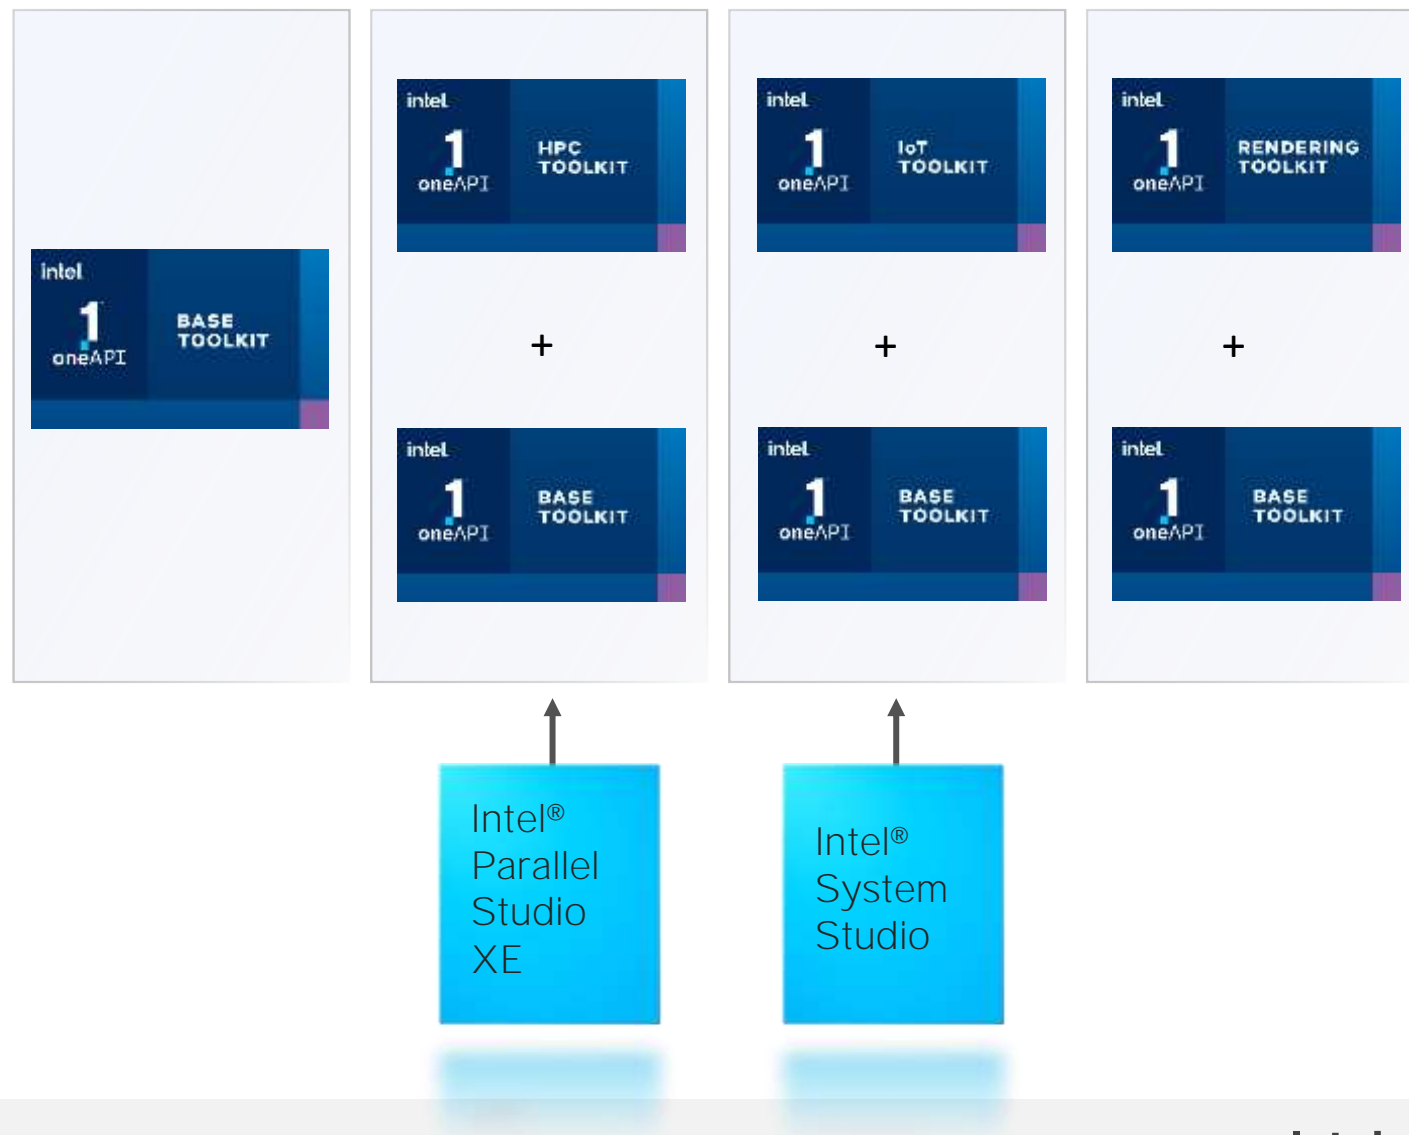
Supports Jupyter Notebooks, Visual Studio Code

Get Up & Running In Seconds!

Commercial Toolkits Deliver Priority Support

Next Generation of Commercial Intel® Software Development Products

- Worldwide support from Intel technical consulting engineers
- Prior commercial tool suites, Intel® Parallel Studio XE and Intel® System Studio, transition to oneAPI products



What's New in Intel® oneAPI Toolkits 2022.1

New & enhanced features released Dec. 22, 2021, learn more: [Read the blog & product release notes](#)



Intel [oneAPI 2022 Toolkits](#) expand cross-architecture features to provide developers greater utility and architectural choice to accelerate computing. New capabilities include the world's first unified compiler implementing C++, SYCL and Fortran, data parallel Python for CPUs and GPUs, advanced accelerator performance modeling and tuning, and performance acceleration for AI and ray tracing visualization workloads. There are **more than 900** new and enhanced features added over the last year that strengthen every tool in the foundational and domain-specific toolkits. Top highlights:

Cross-architecture programming

- Intel created the world's first unified compiler implementing C++, SYCL, and Fortran for CPUs and GPUs utilizing a common LLVM backend.
- Accelerated compute on CPUs and GPUs for Python, today's most popular programming language.
- The Intel® DPC++ Compatibility Tool was improved to automatically migrate 90-95% of CUDA code to SYCL/DPC++. ²

Performance on the latest hardware

- **Hardware support** – Intel oneAPI Toolkits are optimized to enable advanced features of the latest and upcoming new hardware including: 12th Gen Intel® Core™ processors (code named Alder Lake) with AVX-VNNI, Sapphire Rapids with Intel® Advanced Matrix Extension (Intel® AMX), and upcoming Xe client and datacenter GPUs.
- **AI performance optimizations**
 - Deep learning framework performance is accelerated up to 10x over earlier versions with the latest Intel® Optimization for TensorFlow and Intel® Optimization for PyTorch. ³
 - New [Intel® Extension for Scikit-learn](#) speeds up machine learning algorithms over 100x on Intel CPUs over the stock open source version. ⁴
 - Introduced [Intel® Neural Compressor](#) to achieve increased inference performance through post-training optimization techniques across multiple deep learning frameworks.

Advanced tools for development productivity

- **Expert performance analysis for CPUs and accelerators** – Intel® VTune™ Profiler's flame graph display improves the ability to visualize performance hot spots. Intel® Advisor's accelerator performance modeling allows developers to estimate performance benefits of offloading to a GPU before making code changes.
- **Advanced ray tracing** – New features including cone telemetry, auxiliary feature denoising and FP16 support provide more robust shapes and shorten rendering times improving overall rendering performance. Support for upcoming Intel Xe GPUs as well as real-time denoising further improve industry-leading, final frame, production-quality rendering.
- **Expanded development environment support** including deeper Microsoft Visual Studio Code integration, support for Microsoft Visual Studio 2022, and Microsoft WSL2 for Linux development on Windows.

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 Data Parallel C++ 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.

GROMACS – Using oneAPI



Intel oneAPI Tools: Empowering
GROMACS Cross-Architecture
Development

[@IntelDevTools](#)

Click on the image to run the video

oneAPI Resources

software.intel.com/oneapi


Get Started

- software.intel.com/oneapi
- [Documentation](#) + dev guides
- [Code Samples](#)
- Intel® DevCloud



Industry Initiative

- [oneAPI.io](https://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



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.

Notices & Disclaimers

Performance varies by use, configuration and other factors. Learn more at [www.Intel.com/PerformanceIndex](https://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.

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 Corporation. Intel, the Intel logo, Xeon, Core, VTune, OpenVINO, Agilex, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.

