

Intel Releases oneAPI Toolkits for XPU Software Development

The toolkits enable developers to build high-performance, cross-architecture applications, accelerating computing in the XPU era

Dec. 8, 2020 —Intel announced the release of [Intel® oneAPI toolkits](#) for developing high-performance, cross-architecture applications across Intel® CPUs, GPUs and FPGAs, collectively described by Intel as XPU. The version 2021.1 toolkits deliver [oneAPI](#), an open, standards-based, unified cross-architecture programming model that provides developers the freedom to choose the best hardware for accelerated computing. The toolkits build on Intel's rich heritage of proven tools developers have relied on for decades to create high-quality performant applications.



The Productive, Smart Path to Freedom for Accelerated Computing

Today's workloads benefit from specialized hardware architectures. However, these architectures have typically required unique languages and tools, which limits code reuse—narrowing hardware choices and hindering adoption of innovative architectures. oneAPI's heterogeneous programming model delivers performance without proprietary lock-in while interoperating with code written in familiar languages such as C, C++, Fortran and Python, and standards such as MPI and OpenMP.

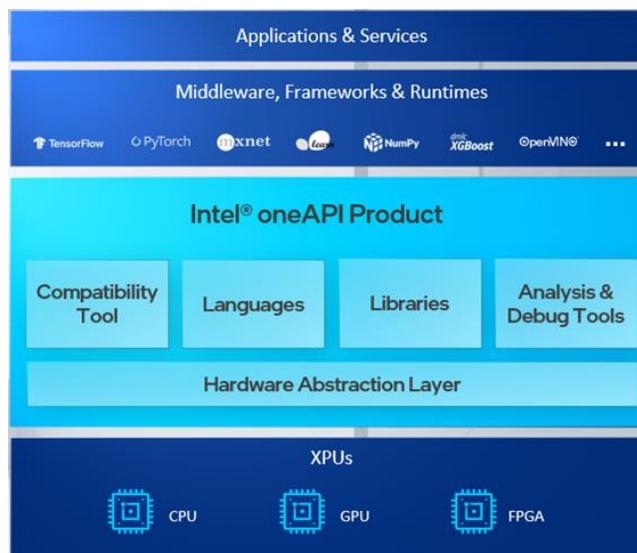
Intel oneAPI toolkits provide developers a quick path to build cross-architecture applications using a single code base across XPU that takes advantage of unique hardware features and lowers software development and maintenance costs. Developers can choose the best architecture for their specific problem without needing to rewrite software for the next architecture and platform.

Intel oneAPI Toolkits Overview

Intel oneAPI toolkits deliver familiar languages such as C, C++, Fortran and Python, and standards such as MPI and OpenMP, providing full continuity with existing code. They enable cutting-edge hardware capabilities and instructions such as Intel® AVX-512 and Intel® DL Boost on CPUs, along with features unique to XPU.

Each Intel oneAPI toolkit is tailored to developers' specific needs.

- [Intel® oneAPI Base Toolkit](#) (Base Kit) is the foundational toolkit to get started. It includes compilers, performance libraries, analysis and debug tools, and a compatibility tool that aids in migrating code written in CUDA to Data Parallel C++ (DPC++).
- Add-on toolkits for [HPC](#), [AI](#), [IoT](#), and [rendering](#) provide tools and components for specialized workloads.



Delivery Options

The toolkits are [free to download](#) and use locally, or access in the [Intel® DevCloud](#). Access options include web download, repositories and containers. Commercial versions providing worldwide support from Intel technical engineers are also offered.

Intel is immediately transitioning Intel® Parallel Studio XE and Intel® System Studio tool suites to its oneAPI products, which are upward compatible and include all current capabilities plus new features and tools.



Intel DevCloud for oneAPI

All Intel oneAPI toolkits are available in the Intel DevCloud, where developers can develop and test code and workloads on a variety of Intel CPUs, GPUs and FPGAs. For GPUs, Intel® Iris® Xe MAX discrete graphics and Intel® Processor Graphics Gen9 (integrated graphics) are available for public access, and select customers are already developing with oneAPI on the Intel® Xe-HP development platform. For FPGAs, Intel® Stratix® 10 and Intel® Arria® devices are available.

Developers can get started easily in the DevCloud—no downloads, installations or configurations required. It supports Jupyter notebooks and Visual Studio Code and provides samples and tutorials. [Access Intel DevCloud today to get started.](#)

Details about Intel oneAPI Toolkits: Unleash Performance across Architectures

In addition to the Base Kit, additional oneAPI toolkits for specialized workloads include:

- **Intel® oneAPI HPC Toolkit** helps developers deliver fast Fortran, OpenMP and MPI applications that scale.
- **Intel® oneAPI IoT Toolkit** helps developers bring the power of oneAPI to global IoT edge innovations for applications used in healthcare, smart homes, industrial, retail, aerospace and more.
- **Intel® oneAPI Rendering Toolkit** enables creation of high-performance, high-fidelity, extensible, and cost-effective visualization applications and solutions used in entertainment, scientific visualization and other segments.
- **Powered by oneAPI toolkits** include frameworks and middleware that have been optimized using DPC++ and/or any of the oneAPI libraries.
 - **Intel® AI Analytics Toolkit** provides drop-in acceleration for end-to-end data science and machine learning pipelines; it is used by data scientists, AI developers and researchers.
 - **Intel® Distribution of OpenVINO™ Toolkit** helps developers deliver high-performance deep learning inference and computer vision.



Developer Resources: To help developers build high-performance, cross-architecture applications using oneAPI and the toolkits, many resources are available.

- **Key Links:** [Documentation](#), [oneAPI programming guide](#), [Code samples](#)
- **Training:** Access [free training](#) through webinars, deep dive workshops, full learning paths or even short videos. Explore the [training catalog](#).
 - A new [Intel[®] oneAPI Technology Partner program](#) includes a group of specially trained consulting companies to help customers accelerate HPC and AI solutions using oneAPI cross-architecture programming.
- **Technical Support:** Developers can access the free [community forums](#). [Priority Support](#), with direct, private interactions with Intel technical consulting engineers is included in [toolkit commercial packages](#).
- Developers can join [Intel[®] DevMesh Innovator Projects](#) to connect with developers using oneAPI.

General Technical Specifications

Processors	<p>CPUs – Intel[®] Xeon[®], Core™, Atom, and compatible processors</p> <p>Integrated GPUs – Intel[®] Processor Graphics Gen9, Gen11, and Intel[®] Iris[®] Xe</p> <p>Discrete GPUs – Intel[®] Iris[®] Xe MAX graphics</p> <p>FPGAs – Intel[®] Arria[®] and Intel[®] Stratix[®] 10 FPGAs</p>
Operating Systems	<p>Linux</p> <p>Windows</p> <p>macOS (partial support)</p>

Specific toolkit, tools or component support by processor are available in the latest release notes.

oneAPI Industry Momentum

oneAPI ecosystem support has been steadily growing since 2019. More than 60 leading research organizations, companies and universities support the oneAPI initiative and some call out success using Intel oneAPI toolkits. See [oneAPI ecosystem support](#) and [reviews](#) for details.

- **Argonne National Laboratory:** Researchers at the U.S. Department of Energy's [Argonne National Laboratory](#) are using Intel oneAPI toolkits to test code performance and functionality using programming models that will be supported on Aurora. Aurora is set to be one of the nation's first exascale systems and will be used to dramatically advance scientific research and discovery.
- **Codeplay builds oneAPI support:** Codeplay Software [announced](#) the first release of its Data Parallel C++ (DPC++) compiler for Nvidia GPUs.
- **oneAPI Centers of Excellence (CoE):** In 2020, four oneAPI CoEs were established to help expand oneAPI adoption by accelerating key technology focus areas and widely used life sciences and other applications to take advantage of heterogeneous architectures. These include [Stockholm University](#), [Heidelberg University](#) and the [University of Illinois](#). **New: Lobachevsky University of Nizhny Novgorod (UNN) [announced today a oneAPI CoE](#)** to facilitate studies in contemporary physics using the power of CPUs, GPUs and other accelerators with oneAPI cross-architecture programming.
- **oneAPI 1.0 specification release:** It's now easier to embrace accelerator programming and address data-intensive workloads. The oneAPI 1.0 specification spans a language and many domains that benefit from accelerators, including math libraries, deep learning and machine learning interfaces, video analytics APIs and a low-level hardware abstraction interface or runtime API. An open source

stack is publicly available to use and port applications. The release marks the culmination of a year of collaboration among leading technologists from industry, academia and government, paving an open road for cross-architecture development.

- New [oneAPI applications catalog](#) details more than 230 applications powered by oneAPI.

Customer Usage of oneAPI & Intel[®] oneAPI Tools



“By integrating Intel’s Intel[®] **oneAPI Data Analytics Library** (oneDAL) and Intel[®] **AI Analytics Toolkit** into Allegro Trains, Allegro AI offers **better performance, and optimized use** of cloud instances.”
Moses Guttman, chief executive officer



“**Customizing, deploying and evaluating oneAPI was a breeze.** I was not expecting the process to be so smooth and frictionless.”
Dr. Jernej Zidar, senior application engineer



“Analytics Zoo and the Intel[®] **AI Analytics Toolkit** with the Intel[®] **oneAPI Data Analytics Library** helped reduce end-to-end data processing time and improved our prediction model’s accuracy significantly for AsialInfo 5G network intelligence, including customer satisfaction analysis, power saving for 5G base station and user location analysis.”
Duozhi Zhu, general manager, 5G Network Product R&D



“Using Intel technology, we’re now able to render more than **10 billion images in 16 weeks**...by working with Intel on the **oneAPI Rendering Toolkit**, we’ve been able to develop proofs of concept on **cost-efficient methods** for creation-rendered images and user experiences for the future.”
Paul Chapman, manager, Virtual Media



“Some application codes which have historically been implemented on CPUs or GPUs actually run more efficiently on FPGAs. Until now, porting these applications to FPGAs has been a significant investment, requiring expert hardware programmers. oneAPI is a welcome, bold initiative that introduces a unified software programming model capable of supporting Intel[®] Stratix[®] 10 and Agilex™ FPGA devices. Fundamentally, oneAPI opens up the compelling benefits of FPGAs to customers who are software-oriented.”
Craig Petrie, vice president, Marketing



“With a smooth learning curve from C++, minor code divergence, and a common codebase, the leading-edge Intel[®] oneAPI technology **reduces the barriers of programming on different architectures**. This allows **maximum flexibility** to harness all the computational capacity of HPC clusters.”
Khaled Elamrawi, president



“oneAPI enabled us to rapidly develop scalable, single-source, data-parallel algorithms for DNA data storage that can **target CPUs and GPUs** (integrated and discrete), using a unified programming model.”
Raja Appuswamy, assistant professor



GeoEast

"We have been using Intel® Parallel Studio XE for years. It is exciting to see that Intel is building a comprehensive and powerful ecosystem with **Intel® oneAPI Base Toolkit and Intel® oneAPI HPC Toolkit**, and we are going to **migrate our whole development environment to oneAPI**. It unifies the programming language on different platforms and will definitely **simplify the development progress and downgrade the developing difficulty.**"

Mr. Zhang, manager



"Collaborating with Intel and using the **Intel® oneAPI HPC Toolkit** has been instrumental in helping our customer engineers understand in depth our customers' HPC workloads and performance on GCP instances. We recommend using **Intel® MPI for best performance**, and tools such as Intel® **VTune™ Profiler** and **Advisor** to help better **understand performance optimizations** and how to best migrate your workloads to the cloud."

Ilias Katsardis, lead, HPC Solution



"We are integrating Samsung's image processing technology, semiconductors ergonomic mechanics and... AI technology into our ultrasound systems for efficient and confident diagnosis. One tool...to accelerate for pursuing this effort more efficiently and with flexibility is Intel's oneAPI solution. Samsung Medison is working on a medical imaging proof of concept using oneAPI to write one source code implementation for performance acceleration on different kinds of hardwares... The **Intel® DPC++ compatibility tool** made it easy to port our existing code to DPC++ and Intel's training... and technical resources helped us use **Intel® VTune™ Profiler** to analyze code performance and further enhance it to run optimally on our products. We look forward to our continued collaboration with Intel, as using the oneAPI solution will allow us to respond quickly to new requests from our healthcare professionals."

Dr. Won-Chul Bang, vice president and head of Product Strategy



"**Intel® oneAPI Base Toolkit and oneAPI HPC Toolkit** are helping to build readable source code, **improving development, maintenance** and our accelerated Reverse Time Migration (RTM) application."

Dr. Ph.D. Clícia Pinto, lead researcher, HPC, Performance Engineering



"The **Intel® DPC++ Compatibility Tool's** debug interoperability with Microsoft Visual Studio IDE helped to **seamlessly migrate CUDA code to DPC++**. The Intel® oneAPI DPC++ Compiler, Intel® Integrated Performance Primitives, the Intel® oneAPI Math Kernel Library and the Intel® VTune™ Profiler from the **Intel® oneAPI Base Toolkit and Intel® oneAPI IoT Toolkit** are all critical to our product line."

Tian Wang, manager, Healthcare & Magnetic Resonance Software R&D



"We compared the TBB + OpenCL vs. the **oneTBB + Intel® oneAPI** implementations of our heterogeneous schedulers, observing that oneAPI versions result in **five times less programming effort** and only incur in 3–8% of overhead."

Dr. Rafael Asenjo, professor



"The **Intel® DPC++ Compatibility Tool** greatly supported our porting efforts to create a DPC++ version of the tsunami simulation code 'easyWave,' which we **executed on CPUs, GPUs and FPGAs.**"

Dr. Thomas Steinke, Supercomputing head

“We are pleased to see the SYCL standard used as the foundation of Intel[®] oneAPI toolkits. This drives the collaboration on open-source implementations including upstreaming to Clang/LLVM and motivates further community input to the standards body at Khronos SYCL.”

Ronan Keryell, principal software engineer, Xilinx Research Labs

oneAPI COE

- [Heidelberg University Computing Center \(URZ\)](#) is focused on bringing oneAPI support to AMD GPUs.
- **New: Lobachevsky University of Nizhny Novgorod (UNN)** [announced a oneAPI CoE](#) to facilitate studies in contemporary physics using the power of CPUs, GPUs and other accelerators with oneAPI cross-architecture programming. UNN The first software to be ported to oneAPI by the UNN researchers is the open-source framework Hi-Chi (High-Intensity Collisions and Interactions), which is intended for the simulation of high-intensity laser-matter interactions.
- **Stockholm University/ KTH Royal Institute of Technology:** The [Swedish e-Science Research Center](#) is using oneAPI's unified and heterogeneous programming model to accelerate research conducted with GROMACS, a widely-used free and open source application designed for molecular dynamics simulation.
- **University of Illinois:** The [Beckman Institute for Advanced Science and Technology](#) will bring the oneAPI programming model to life sciences and application NAMD to additional heterogeneous computing environments. NAMD, which simulates large biomolecular systems, is helping to tackle real-world challenges such as COVID19.

Intel[®] Graphics and Visualization Institutes of XeLLENCE

- Computer Graphics Charles University & Chaos Research
- Northern Illinois University
- Texas Advanced Computing Center & Kitware, Inc. (TACC)
- University of Cambridge/The Stephen Hawking Center for Theoretical Cosmology
- University of Tennessee Knoxville
- University of Utah/Scientific Computing & Imaging Institute
- Visualisierungsinstitut der Universität Stuttgart

Learn More

[Intel oneAPI Toolkits](#) | [Intel oneAPI Updates](#) | [oneAPI.com](#)

Notices & Disclaimers

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

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.

Copyright ©, Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, Core, VTune, and OpenVINO are trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries.

About Intel

Intel (Nasdaq: INTC) is an industry leader, creating world-changing technology that enables global progress and enriches lives. Inspired by Moore's Law, we continuously work to advance the design and manufacturing of semiconductors to help address our customers' greatest challenges. By embedding intelligence in the cloud, network, edge and every kind of computing device, we unleash the potential of data to transform business and society for the better. To learn more about Intel's innovations, go to newsroom.intel.com and intel.com.