

Gold Intel® oneAPI Toolkits Ship December 2020

Intel oneAPI Toolkits enable developers to code across architectures, accelerate computing in the XPU era

Nov. 11, 2020 — Intel today announced key milestones in the multiyear journey to accelerate deployment of diverse workloads across a mix of architectures with a unified software abstraction. This includes the gold shipment of Intel® oneAPI Toolkits in December, broad public availability of Intel® Iris® X^e MAX graphics and availability of [Intel® X^e-HP](#) to select developers in Intel® DevCloud, and new oneAPI ecosystem advancements and endorsements. These milestones are important steps in delivering hardware and software solutions to address the growth in specialized workloads from artificial intelligence (AI) to high performance computing (HPC) and graphics.


Background: At SuperComputing 2019, Intel announced the [oneAPI](#) industry's initiative to create an open, standards-based, unified programming model that provides developers the freedom to choose the best hardware for accelerated computing—CPUs, GPUs, FPGAs other accelerators, collectively described by Intel as XPU, or a combination thereof. Intel released beta implementations of its [oneAPI toolkits](#) that help developers efficiently build, analyze and optimize cross-architecture applications for Intel® CPUs, GPUs and FPGAs. Throughout 2020, Intel delivered regular updates to its beta toolkits, bringing additional capabilities to maximize XPU hardware performance. See [Intel oneAPI Updates](#) for details.



Intel oneAPI Toolkits: What's New

Intel's gold oneAPI product is free for developers to use locally or in the Intel® DevCloud and ships in December 2020. Commercial versions with expert worldwide support through Intel technical consulting engineers will also be available. Intel is immediately transitioning its prior tool suites, Intel® Parallel Studio XE and Intel® System Studio, to its oneAPI products.

Intel is redefining the future of programming built on its rich heritage of proven CPU developer tools now expanded for XPU to help developers deliver high performance, cross-architecture applications. Intel's gold oneAPI product release provides a



complete set of advanced libraries, compilers and porting and analysis tools. The tools enable:

- Accelerated compute by maximizing performance using the full value of cutting-edge features of Intel CPUs, GPUs and FPGAs.
- Interoperability with existing programming models and code bases like C++, Fortran, Python and OpenMP, etc.; developers can be confident that existing applications work seamlessly with oneAPI.
- Easy transitions to new generations of systems and accelerators — using a single code base across all architectures frees developers to invest on innovating and optimizing features that advance exascale and AI (vs. code maintenance).
- Acceleration for specialized workloads via domain-centric toolkits for HPC, AI, IoT and rendering.

Intel DevCloud Expands to Include New Intel Iris X^e Graphics

Intel DevCloud provides free access to state-of-the-art Intel software and hardware, allowing developers to start writing code using oneAPI and to test code and workloads across a combination of Intel[®] CPUs and accelerators.

What's New: Intel is announcing the broad public availability of [Intel Iris X^e MAX graphics](#) in the Intel DevCloud, and availability of [Intel[®] X^e-HP](#) to select developers. Developers can now test code and workloads on these architectures, paving the road for Intel's pipeline of future graphics products.

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.](#)

oneAPI Ecosystem Advancements

- **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.
- **What's New: University of Illinois (UI):** The Beckman Institute for Advanced Science and Technology at UI [today announced a new oneAPI center of excellence \(CoE\)](#). They are bringing 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 COVID 19.
- **Heidelberg University Computing Center (URZ):** URZ [announced](#) it is establishing a oneAPI CoE focused on bringing oneAPI support to AMD GPUs.
- **Swedish e-Science Research Center (SeRC):** Hosted at Stockholm University and the KTH Royal Institute of Technology, the [SeRC's oneAPI academic CoE](#) 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.

New oneAPI Endorsements



"The industry needs a programming model where developers can take advantage of an array of innovative hardware architectures. The goal of oneAPI is to provide increased choice of hardware vendors, processor architectures and faster support of next-generation accelerators. Microsoft has been using oneAPI elements across Intel hardware offerings as part of its initiatives and supports the open standards-based specification. We are excited to support our customers with choice and accelerate the growth of AI and machine learning."



"With the growth of AI, machine learning and data-centric applications, the industry needs a programming model that allows developers to take advantage of rapid innovation in processor architectures. TensorFlow supports the oneAPI industry initiative and its standards-based open specification. oneAPI complements TensorFlow's modular design and provides increased choice of hardware vendor and processor architecture, and faster support of next-generation accelerators. TensorFlow uses oneAPI today on Xeon processors and we look forward to using oneAPI to run on future Intel architectures."

More than 50 leading companies, research organizations and universities now endorse oneAPI, and some call out success using Intel oneAPI Toolkits. See also [oneAPI Ecosystem Support](#)

Details About Intel oneAPI Toolkits: The Smart Path to Accelerated Computing

Intel oneAPI Toolkits include:

- The **Intel oneAPI Base Toolkit** (Base Kit) is the starting point for developers with a core set of tools, including best-in-class compilers, powerful performance libraries, a compatibility tool, and analysis and debug tools.
- **Domain-Specific Toolkits** aid developers in optimizing applications and solutions for specialized workloads in HPC, AI, internet of things and rendering. These complement the Base Toolkit.
 - **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 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.



- **Toolkits powered by oneAPI** include oneAPI's cross-architecture components or 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.

Intel oneAPI Toolkits General Technical Specifications

Architecture Supported	CPU – Intel [®] Xeon [®] , Core™, Atom, and compatible processors GPU – Intel [®] Processor Graphics Gen9, Intel [®] Iris [®] Xe MAX graphics, Intel [®] Xe-HP FPGAs – Intel [®] Arria [®] , Intel [®] Stratix [®] 10 FPGAs
Operating Systems	Linux and Windows

Please view specific toolkit, tools or component support by processor detailed in the latest release notes.

Customer Usage of Intel[®] oneAPI Tools

	<p>“By integrating Intel’s Intel[®] oneAPI Data Analytics Library (oneDAL) and Intel[®] AI Analytics Toolkit tools into Allegro Trains, Allegro AI offers better performance, and optimized use of cloud instances.”</p> <p><i>Moses Guttman, CEO</i></p>
	<p>“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.”</p> <p><i>Khaled Elamrawi, president</i></p> <p>InsideHPC article white paper</p>
	<p>“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 VTune Profiler and Advisor to help better understand performance optimizations and how to best migrate your workloads to the cloud.”</p> <p><i>Ilias Katsardis, HPC Solution Lead, Google</i></p>
	<p>“We compared the TBB+OpenCL vs. the oneTBB+oneAPI implementations of our heterogeneous schedulers, observing that oneAPI versions result in up to 5x less programming effort and only incur in 3–8% of overhead.”</p> <p><i>Dr. Rafael Asenjo</i></p> <p>Intel[®] oneAPI Toolkits (Beta) Enable Easier Programming of Machine Learning Algorithms on Robots</p>
	<p>“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.”</p> <p><i>Dr. Thomas Steinke, head of the Supercomputing Department</i></p>

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.

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

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.
