

Intel Empowers Developers with the Release of oneAPI 2022 Toolkits

Open your code to more hardware choices, move beyond proprietary barriers, and unleash your code's best performance.

Dec. 22, 2021 —Intel released [oneAPI 2022 Toolkits](#). The newly enhanced toolkits expand cross-architecture features to provide developers greater utility and architectural choice to accelerate computing. The [oneAPI cross-architecture programming](#) model provides developers with tools that aim to improve the productivity and velocity of code development when building cross-architecture applications. (See page 2 for more 2022 version details.)



Why oneAPI Matters: According to an Evans Data survey, 40% of developers target heterogeneous systems that use more than one type of processor, processor core or coprocessor.¹ The pace of innovation is only accelerating with cross-architecture computing driven forward with oneAPI across heterogeneous fabrics of CPUs, GPUs, and FPGAs.

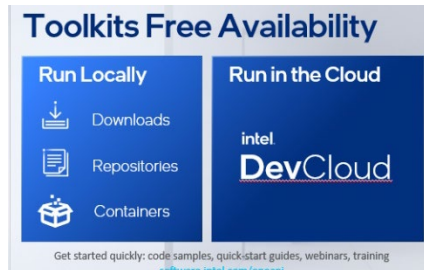
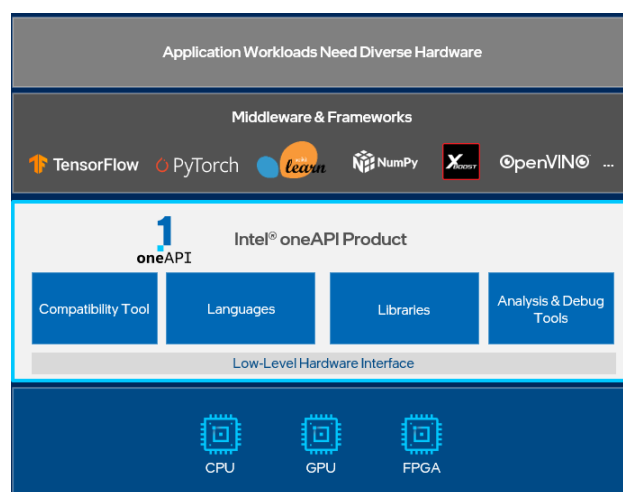
oneAPI offers freedom for developers to choose the best hardware for a given solution without needing to rewrite software for the next architecture and platform, overcoming the barriers of proprietary programming models.

Intel oneAPI Toolkits Overview

The [Intel® oneAPI Toolkits](#) deliver performance and productivity through a complete set of advanced tools including compilers, libraries, pre-optimized frameworks, analyzers and debuggers. The toolkits support familiar languages such as C, C++, SYCL, Fortran, and Python, and standards like MPI and OpenMP, providing full continuity with existing code. They enable Intel cutting-edge hardware capabilities and instructions such as Intel® AVX-512, Intel® AMX, and Intel® DL Boost on CPUs, along with features unique to accelerators.

The oneAPI toolkits are tailored for specific developer needs. The foundational [Base Kit](#) includes oneAPI's advanced performance libraries, compilers and analysis and debug tools, along with a porting tool that aids in migrating code written in CUDA to SYCL. Add-on toolkits for [HPC](#), [AI](#), [IoT](#), and [rendering/ray tracing](#) provide tools and components for those workloads.

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



What's New in the 2022 Version: There are more than 900 new and enhanced features added over the past year that strengthen every tool in the foundational and domain-specific toolkits. Top highlights 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
- Performance acceleration for AI and ray tracing visualization workloads

Intel[®] oneAPI
Toolkits 2022.1
Now Available

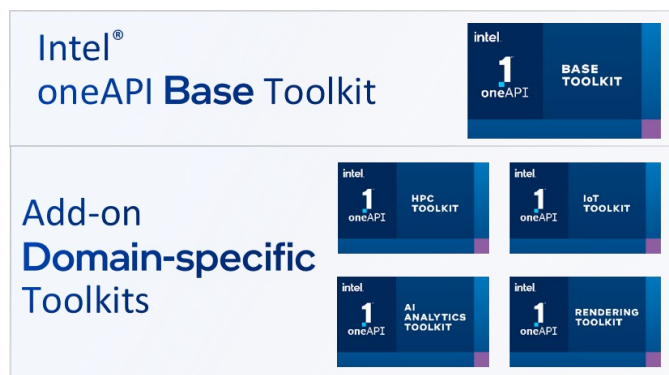
[GET THEM NOW >](#)

Get more details: Read [announcement blog](#) | Access [product release notes](#)

Toolkit Details: The Smart Path to Accelerated Computing

Intel[®] oneAPI Toolkits include:

- **Intel[®] oneAPI Base Toolkit** (Base Kit) is the starting point for developers with a core set of advanced compilers, performance libraries, analysis and debug tools, and a CUDA code porting tool to build high-performance cross-architecture applications.
- **Add-on Toolkits** aid developers in optimizing applications and solutions for HPC, AI, rendering and IoT workloads. These complement the Base Kit:
 - **Intel[®] oneAPI HPC Toolkit** helps developers deliver fast Fortran, OpenMP and MPI applications that scale.
 - **Intel[®] oneAPI 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[®] oneAPI Rendering Toolkit** enables creation of high-performance, high-fidelity, extensible and cost-effective visualization applications and solutions used in entertainment, scientific visualization, 3D/architectural design, and other segments.
 - **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[®] Distribution of OpenVINO™ toolkit**, powered by oneAPI libraries, helps developers deliver high-performance deep learning inference and computer vision.



Get Started: Documentation, Training & Support: To help developers build high-performance cross-architecture applications, many resources are available.

- [Documentation](#), [oneAPI programming guide](#), and [code samples](#) help developers get started.
- [Free training](#) is available through webinars, deep dive workshops, full learning paths, or even 3-minute quick hit videos. Explore the [training catalog](#). Training by certified oneAPI experts is also available worldwide for HPC and AI through the Intel[®] oneAPI Technology Partner program.
- For technical support using Intel oneAPI tools, developers can access the free [community forums](#). [Priority Support](#) with direct, private interactions with Intel engineers to accelerate development is included in [toolkit commercial packages](#).
- Developers can join [Intel[®] DevMesh Innovator Projects](#) to connect with other developers already using oneAPI.

General Technical Specifications

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

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

Note that former tool suites Intel[®] Parallel Studio XE and Intel[®] System Studio were transitioned into Intel oneAPI products, which are upward compatible and include all current capabilities + new capabilities and tools.

Intel[®] DevCloud

All oneAPI Toolkits are available in the Intel DevCloud, an environment 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 platforms 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.](#)

oneAPI Momentum

oneAPI Ecosystem Adoption: oneAPI has made great progress since its launch in 2019. It's an open project, getting community input and contributions into the oneAPI specification. More than 100 national labs, research organizations, companies and universities are embracing the value of oneAPI's cross-architecture and cross-vendor programming over proprietary programming models that restrict adoption of other vendors' architectures. Many third-party vendors and open-source communities have also integrated one or more libraries or components from the Intel oneAPI Rendering Toolkit into their rendering tools.

Learn more: oneAPI [Open Source Projects](#) | [News](#) | [Community](#) | [Ecosystem Support](#)

Intel oneAPI Implementation/Tools Adoption: Usage of Intel oneAPI toolkits and the DevCloud continue to increase. A few customer examples follow.


Learn more: [Ecosystem Support](#) (incl. tools usage testimonials) | [Intel oneAPI Projects](#) | [oneAPI Application Catalog](#)



Innovation Leaders Using Intel® oneAPI Cross-Architecture Tools


ANSYS, Brightskies, Dell, Google Cloud, Samsung Medison, Stephen Hawking Center for Theoretical Cosmology (CTC), and Zuse Institute Berlin (ZIB) share how Intel oneAPI cross-architecture tools deliver performance and productivity for heterogeneous development.

[Video](#)



Accelerate Workloads in Google Cloud Using Intel® oneAPI Toolkits

Find out how Google Cloud Platform service sped up its HPC workloads using Intel oneAPI Tools and Intel® Xeon® processors. [Video](#)




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

[The Next Platform](#) Article



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

[Article](#)



The Addams Family 2: [Intel Denoiser Delivers Up to 25% Rendering Efficiency for The Addams Family 2](#)

Find out how [Intel® Open Image Denoise](#) helped Cinesite Studios deliver incredible visual storytelling.

[Video \[3:18\]](#) | [Case Study](#)



Samsung Medison Uses oneAPI to Power IoT 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 and improved workflows.

[Intel PR Sept. 10, 2020](#) | [Video \[1:45\]](#)

Product Innovations Using Intel oneAPI Tools

- **AI Development** - [SberBank](#) built its own oneAPI cloud for AI development for its Russia customers and is planning expansion to other countries in Europe. It includes the oneAPI Analytics Toolkit.
- **AI Platforms** - [Dell, HP, and Lenovo are delivering Linux-based mobile and desktop AI/data science workstations](#) built on Intel[®] Core™ and Intel[®] Xeon[®] architecture and integrated the oneAPI Analytics Toolkit to enable “out of the box” AI development to a broad set of users. See also [Intel press release](#).
- **AI Tools** - Alibaba and Intel partnered to build an end-to-end recommendation engine toolkit called DeepRec to facilitate deep learning training and deployment of recommendation systems powered by oneAPI with AVX-512, VNNI and BF16 acceleration. See also [Intel press release](#).
- **Rendering Tools** – Many third-party vendors and open-source communities have integrated one or more libraries or components from the Intel oneAPI Rendering Toolkit into their rendering tools. This includes many of the world’s most popular creator, gaming, scientific and product design tools—such as Blender, Chaos Group’s Corona and V-Ray, Cinema 4D, Kitware ParaView, Unity and Unreal engine — with the most recent being [AutoDesk’s Arnold renderer](#), which is widely used in film, and [Open3D Engine](#).

Intel oneAPI Centers of Excellence: To help bring oneAPI heterogeneous programming and tools to developers, Intel invests in research, technologies and education through Intel oneAPI Centers of Excellence (CoEs) to accelerate oneAPI ecosystem adoption for HPC, AI and Visualization. In 2021, 13 Intel oneAPI Centers of Excellence were added in the U.S., Europe and China, with more to come. These CoEs will deliver strategic code ports, additional hardware support, new technologies and services and curriculum.



For more details visit: [Intel oneAPI Centers of Excellence Fact Sheet](#) | [Intel[®] Academic Program for oneAPI](#)

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.

