intel software

Develop Cross-Architecture Applications for CPUs, GPUs, and FPGAs

Take the smart path to accelerated computing—without the economic and technical burdens of proprietary programming models

Many data-centric workloads run best when they're deployed across a mix of heterogeneous architectures—CPU, GPU, FPGA, and other accelerators. But different architectures typically require unique languages, tools, and libraries—adding complexity for developers and limiting code reuse. This makes it hard to take advantage of cross-architecture solutions and inefficient to optimize application performance.

oneAPI is an industry initiative creating an open, standards-based, cross-architecture programming model to simplify development for a wide range of data-centric workloads across a variety of architectures. It includes a cross-architecture language, Data Parallel C++ (DPC++), based on ISO C++ and Khronos Group's SYCL, plus advanced libraries and performance analysis tools.

Use it for:

- High-performance computing (HPC)
- Machine learning, deep learning, and analytics
- IoT applications
- Video processing
- Rendering
- And more

Who Needs It

- **Developers looking to maximize performance and productivity and freedom of architectural choice** by building cross-architecture applications and solutions that take advantage of a variety of Intel® CPUs, GPUs, and FPGAs.

- **Developers building data-centric applications to run on all types of platforms** (edge to cloud) and architectures (Intel CPUs, GPUs, and FPGAs) across a range of domains including HPC, AI, IoT, and visual computing.

- **Developers with existing CUDA code** who want to take advantage of other architectures by porting their code to DPC++ code.

What It Does

oneAPI and Intel® oneAPI toolkits together bring developers productive and performant heterogenous programming.

Lets You Get All of Your Hardware's Value

- **Expose and exploit all the cutting-edge features of the latest hardware**. For instance, get the most from 3rd-gen Intel® Xeon® Scalable processors, which enable Intel® Advanced Vector Extensions 512 (Intel® AVX-512) including Intel® Deep Learning Boost for AI acceleration.

Optimized Applications

Intel oneAPI Base Toolkit

DIRECT PROGRAMMING

- Intel® oneAPI DPC++/C++ Compiler
- Intel® DPC++ Compatibility Tool
- Intel® Distribution for Python
- Intel® FPGA Add-On for oneAPI Base Toolkit

API-BASED PROGRAMMING

- Intel® oneAPI DPC++ Library
- Intel® oneAPI Math Kernel Library
- Intel® oneAPI Data Analytics Library
- Intel® oneAPI Threading Building Blocks
- Intel® oneAPI Video Processing Library
- Intel® oneAPI Collective Communications Library
- Intel® oneAPI Deep Neural Network Library
- Intel® Integrated Performance Primitives

ANALYSIS TOOLS

- Intel® vTune™ Profiler
- Intel® Advisor
- Intel® Distribution for GDB
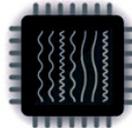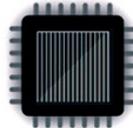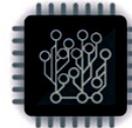
CPU   **GPU**   **FPGA**

- **Get top performance for accelerated architectures**. Take full advantage of accelerated compute by maximizing performance across Intel CPUs, GPUs, and FPGAs.

Lets You Quickly and Confidently Develop Performant Code

- **Make development fast and efficient** with a complete set of cross-architecture libraries and advanced tools.
- **Easily integrate with legacy code** including migrating CUDA code to DPC++.

Future-Ready Programming Model Provides Freedom of Choice

- **Get an open alternative to single-vendor/proprietary lock-in** for easy architecture retargeting.
- **Apply your skills to the next innovation**, not to rewriting software for the next hardware platform.

Highlights

Intel® oneAPI DPC++/C++ Compiler for Direct Programming

DPC++ is an evolution of C++ that incorporates SYCL and community enhancements. It allows code reuse across hardware targets and enables high productivity and performance across CPU, GPU, and FPGA architectures while permitting accelerator-specific tuning. The unified Intel® oneAPI DPC++/C++ Compiler also includes full modern C++.

Libraries for API-Based Programming

Powerful libraries—including deep learning, math, and video and media processing—are preoptimized for domain-specific functions and custom-coded to accelerate compute-intense workloads.

Advanced Analysis and Debug Tools

Get what you need for profiling, design advice, and debug:

- **Intel® VTune™ Profiler** to find performance bottlenecks fast in CPU, GPU, and FPGA systems
- **Intel® Advisor** for efficient vectorization, threading, and offloading to accelerators
- **Intel® Distribution for GDB** for efficient code troubleshooting

What You Get

The Intel® oneAPI Base Toolkit is a core set of tools and libraries for building and deploying high-performance, data-centric applications across diverse architectures:

- **Intel oneAPI DPC++/C++ Compiler**: A standards-based, cross architecture compiler supporting Data Parallel C++, C++, C, SYCL, and OpenMP, it leverages well-proven LLVM compiler technology and Intel's history of compiler leadership for performance. Experience seamless compatibility with popular compilers, development environments, and operating systems.

- **Intel® DPC++ Compatibility Tool**: Migrate CUDA source code to DPC++ code with this assistant.
- **Intel® oneAPI DPC++ Library**: Speed up data parallel workloads with these key productivity algorithms and functions.
- **Intel® oneAPI Threading Building Blocks:** Simplify parallelism with this advanced threading and memory-management template library.
- **Intel® oneAPI Math Kernel Library**: Accelerate math processing routines including matrix algebra, fast Fourier transforms (FFT), and vector math.
- **Intel® oneAPI Data Analytics Library:** Boost machine learning and data analytics performance.
- **oneAPI Video Processing Library:** Deliver fast, high-quality, real-time video decoding, encoding, transcoding, and processing for broadcasting, live streaming and VOD, cloud gaming, and more.
- **Intel® Advisor:** Design code for efficient vectorization, threading, and offloading to accelerators.
- **Intel® Distribution for Python:** Achieve fast math-intensive workload performance without code changes for data science and machine learning problems.
- **Intel® Integrated Performance Primitives:** Speed performance of imaging, signal processing, data compression, cryptography, and more.
- **Intel® VTune™ Profiler:** Find and optimize performance bottlenecks across CPU, GPU, and FPGA systems.
- **Intel® Distribution for GDB:** Enables deep, system-wide debug of DPC++, C, C++, and Fortran code.
- **Intel® FPGA Add-On for oneAPI Base Toolkit (Optional):** Program these reconfigurable hardware accelerators to speed specialized, data-centric workloads. Requires installation of the Intel oneAPI Base Toolkit.
- **Intel® oneAPI Deep Neural Network Library:** Develop fast neural networks on Intel® CPUs and GPUs with performance-optimized building blocks.
- **Intel® oneAPI Collective Communications Library:** Implement optimized communication patterns to distribute deep learning and machine learning model training across multiple nodes.

Priority Support

Every paid version of Intel® Software Development Products automatically includes priority support at our Online Service Center for a duration associated with your purchase, typically one year. You get:

- **Direct and private interaction with Intel's support engineers** and ability to submit confidential support requests

- **Accelerated response time** for technical questions and other product needs
- **Priority support** for escalated defects and feature requests
- **Free download access** to all new product updates and continued access to older versions of the product
- **Access to a vast library** of self-help documentation that builds off decades of experience with creating high-performance code
- **Access to Intel public community forums** supported by community technical experts and monitored by Intel engineers
- **Optional services** at additional cost including on-site/online training and consultation by Intel technical consulting engineer

Related Products: Add Domain-Specific Toolkits for Your Specialized Workloads

Besides the Intel oneAPI Base Toolkit, you can choose from three add-on toolkits that combine to give you the specialized tools you need:

- **Intel® oneAPI Base and HPC Toolkit:** Deliver fast applications that scale with tools to build, analyze, optimize, and scale HPC applications with the latest techniques in vectorization, multithreading, multi-node parallelization, and memory optimization.
- **Intel® oneAPI Base and IoT Toolkit:** Accelerate development of smart, connected devices for healthcare, smart homes,
- **Intel® oneAPI Base and Rendering Toolkit:** Get powerful rendering and ray-tracing libraries for high-fidelity visualization applications—for medical research, geophysical exploration, movie-making, and more—that require massive amounts of raw data to be quickly rendered into rich, realistic visuals.

Try Your Code in the Intel® DevCloud

Develop, run, and optimize your Intel oneAPI code in the Intel® DevCloud—a free development sandbox with access to the latest Intel CPU, GPU, and FPGA hardware and Intel oneAPI software.

Get Started

- **Learn More about Intel oneAPI Products >**
- **Get the Intel oneAPI Base Toolkit >**
- **Check out the Intel DevCloud >**

intel software