intel software

Developing Cross-Architecture, High-Performance Applications

Take your HPC, enterprise, AI, and cloud applications to the max— with fast, scalable, and portable parallel code

Intel® oneAPI Base and HPC Toolkit is a comprehensive suite of development tools that make it fast and easy to build modern code that gets every last ounce of performance out of the newest Intel® processors in high-performance computing (HPC) platforms. Combining the core set of tools from the Intel oneAPI Base Toolkit and adding tools focused on HPC simplifies creating code with the latest techniques in vectorization, multi-threading, multi-node, memory optimization, and accelerator offloading. Get powerful, consistent programming with Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instructions for Intel® Xeon® Scalable processors and supported XPUs with standards-driven programming languages and parallel programing models and integrated development environments (IDEs).

Who Needs It

- **C, C++, Data Parallel C++, Fortran, Python, OpenMP, and MPI software developers and architects** building HPC, enterprise, AI, and cloud solutions
- **Developers looking to maximize their software's performance and flexibility** to support cross architectures on current and future Intel® platforms

What it Does

- **Creates fast parallel code.** Boost application performance that scales on current and future Intel platforms with industry-leading compilers, performance libraries, performance profilers, and code and cluster analysis tools.
- **Builds code faster.** Simplify the process of creating fast, scalable, reliable parallel code.
- **Delivers Priority Support.** Connect directly to Intel's engineers for confidential, quick answers to technical questions. Access older versions of the products. And receive updates for a year.

Highlights

Choose Single-Node or Multi-Node

The Intel oneAPI HPC Toolkit provides all the tools you need to target across architectures, whether you're running the applications on shared or distributed memory systems. The target platforms for development and deployment can range from a workstation to a multi-node cluster requiring different support efforts. Choose the paid product with support that best fits the use model targeted:

- **Intel® oneAPI Base and HPC Toolkit Single-Node:** Target platform of shared memory systems including PCs, laptops, or workstations.
- **Intel® oneAPI Base and HPC Toolkit Multi-Node:** Target platform of shared memory systems such as PCs, laptops, workstations, or distributed memory high-performance compute clusters.

Powerful Tools for Cross-Architecture High-Performance Computing

Building, analyzing, and scaling for high-performance computing to get maximum performance from your compute platforms is challenging. The Intel oneAPI HPC Toolkit delivers on performance from standards-driven compilers and performance libraries, finds opportunities for optimization or use of accelerators through analysis tools, and keeps high-performance clusters running optimally.

Tools to Build

- **Improve performance** with a simple recompile using industry-leading, standards-driven C++ and Fortran compilers.

- **Simplify adding parallelism** with built-in, intuitive, parallel models and vectorization support.

- **Drop in advanced libraries** optimized for the latest hardware.

- **Accelerate diverse HPC to AI workloads** with high-performance Python, powered by native performance libraries, in an integrated distribution package.

Tools to Analyze

- **Advanced performance profiler** to tune application performance of the CPU, threading, memory, and storage

- **Vectorization and threading advisor** to optimize vectorization and quickly prototype threading designs

- **Memory and thread debugger** to efficiently find memory errors and intermittent threading errors

Tools to Scale

- **Accelerate** applications' performance on Intel® architecture-based clusters with multiple fabric flexibility.

- **Profile** MPI applications to quickly finding bottlenecks, achieving high performance for parallel cluster applications.

- **Verify** that cluster components continue working together throughout the cluster life cycle.

What You Get

- **Intel® oneAPI DPC++/C++ Compiler:** A standards-based, CPU, GPU and FPGA compiler supporting Data Parallel C++, C++, C, SYCL and OpenMP. that 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® C++ Compiler Classic:** A standards-based C/C++ compiler supporting OpenMP, focused on CPU development. Take advantage of more cores and built-in technologies in platforms based on Intel® CPU architectures. Experience seamless compatibility with popular compilers, development environments, and operating systems.

- **Intel® Fortran Compiler** (Beta) for XPU development: A standards-based, CPU and GPU compiler supporting Fortran and OpenMP. 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® Fortran Compiler Classic:** A standards-based Fortran compiler supporting OpenMP focused on CPU development. Take advantage of more cores and built-in technologies in platforms based on Intel® CPU architectures. Experience seamless compatibility with popular compilers, development environments, and operating systems.

- **Intel® Cluster Checker:** Verify that cluster components work together seamlessly for optimal performance, improved uptime, and lower total cost of ownership.

- **Intel® Inspector:** Locate and debug threading, memory, and persistent memory errors early in the design cycle to avoid costly errors later.

- **Intel® MPI Library:** Deliver flexible, efficient, scalable cluster messaging on Intel® architecture.

- **Intel® Trace Analyzer and Collector:** Understand MPI application behavior across its full runtime.

- **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.

- **Intel® oneAPI Video Processing Library:** Deliver fast, high-quality, real-time video decoding, encoding, transcoding, and processing.

- **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® DPC++ Compatibility Tool:** Migrate legacy CUDA code to a multi-platform program in DPC++ code with this assistant.

- **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 of GDB:** Enables deep, system-wide debug of DPC++, C, C++, and Fortran code.

- **Intel® FPGA Add-On for oneAPI Base Toolkit** (Optional): Use reconfigurable hardware to accelerate data-centric workloads.

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 engineers

Get Started

- **Get the Intel oneAPI Base and HPC Toolkit >**

- **Learn more >**

intel® software