

Product Brief

High-Performance Computing
Intel® oneAPI Base and HPC Toolkit

Developing Multi-Architecture, High-Performance Computing

intel
1
oneAPI
HPC
TOOLKITintel
1
oneAPI
BASE
TOOLKIT

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

Intel® oneAPI Base Toolkit together with the Intel® oneAPI HPC Toolkit create a comprehensive suite of development tools that make it fast and easy to build modern code. You can obtain all possible 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 HPC-focused tools simplifies creating code with the latest techniques in vectorization, multithreading, multinode, memory optimization, and accelerator offloading. Get powerful, consistent programming for Intel® Xeon® Scalable processors and supported GPUs and FPGAs with standard programming languages, parallel programming models, and integrated development environments (IDEs).

Who needs it

- C, C++, Data Parallel C++, SYCL*, 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 multiple architectures across current and future Intel® platforms

What it does

- **Creates fast parallel code:** Boosts application performance that scales on current and future Intel platforms with industry-leading compilers, performance libraries, performance profilers, code, and cluster analysis tools.
- **Builds code faster:** Simplifies the process of creating fast, scalable, and reliable parallel code.
- **Benefits from Priority Support:** Connect directly to Intel's engineers for confidential, quick answers to technical questions; get access to—and support for—older versions of the products; and receive updates for a year.

Highlights

Choose single node or multinode

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 multinode cluster, requiring different support efforts. Choose the product that best fits your use model:

- **Intel oneAPI Base and Intel oneAPI HPC Toolkit single node:** Target platform of shared memory systems including PCs, laptops, or workstations.
- **Intel oneAPI Base and Intel oneAPI HPC Toolkit multinode:** Target platform of shared memory systems such as PCs, laptops, workstations, or high-performance compute clusters.

Intel® oneAPI Base + HPC Toolkit

Direct Programming

Intel® oneAPI DPC++/C++ Compiler
Intel® C++ Compiler Classic
Intel® Fortran Compiler
Intel® Fortran Compiler Classic

API-Based Programming

Intel® MPI Library


Analysis Tools

Intel® Cluster Checker
Intel® Inspector
Intel® Trace Analyzer and Collector


*Hardware support varies by individual oneAPI tool.

Additional architecture support will be expanded over time.


*Other names and brands may be claimed as the property of others.



CPU



GPU



FPGA

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 adviser to optimize vectorization and quickly prototype threading designs.
- Memory and thread debugger to find memory errors and intermittent threading errors efficiently.

Tools to scale

- Accelerate applications' performance on Intel® architecture-based clusters with multiple fabric flexibility.
- Profile MPI applications to quickly find 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 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:** 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, 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 multiplatform 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 FPGA hardware to accelerate data-centric workloads.
- **Free download access** to all new product updates and continued access to older versions of the product.
- **Access to Intel public community forums** supported by community technical experts and monitored by Intel engineers.

Priority Support

Every paid version of Intel® oneAPI Software Development Toolkits automatically includes Priority Support at our Online Service Center for a duration starting with your purchase, typically one year.

You get:

- **Direct and private interaction with Intel's support engineers**, including the ability to submit confidential support requests.
- **Accelerated response time** for toolkit-related technical questions and other product needs.
- **Priority oneAPI Toolkit Support** for older versions, defect escalation, 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.
- **Additional services** at a reduced cost, including on-site or online training and consultation by Intel technical consulting engineers.

Try your code in the Intel® Developer Cloud

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.



Intel® technologies may require enabled hardware, software, or service activation. Learn more at intel.com or from the OEM or retailer. Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Optimization notice: Intel® compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel® microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel® microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product user and reference guides for more information regarding the specific instruction sets covered by this notice. Notice Revision #20110804. <https://software.intel.com/en-us/articles/optimization-notice>

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations, and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. See backup for configuration details. For more complete information about performance and benchmark results, visit intel.com/benchmarks.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See configuration disclosure for details. No product or component can be absolutely secure.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and noninfringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

© 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.