



Unified, Cross-Architecture Programming Model

**Simplify development and save time across multiple architectures—
with uncompromised performance for diverse workloads**

Modern workloads are incredibly diverse—and so are architectures. No single architecture is best for every workload. Maximizing performance takes a mix of scalar, vector, matrix, and spatial (SVMS) architectures deployed in CPU, GPU, FPGA, and other future accelerators.

Intel® oneAPI products will deliver the tools you need to deploy your applications and solutions across SVMS architectures. Its set of complementary toolkits—a base kit and specialty add-ons—simplify programming and help you improve efficiency and innovation.

Use it for:

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

Highlights

Data Parallel C++ Language for Direct Programming

Data Parallel C++ (DPC++) is an evolution of C++ that incorporates SYCL*. It allows code reuse across hardware targets and enables high productivity and performance across CPU, GPU, and FPGA architectures, while permitting accelerator-specific tuning.

Libraries for API-Based Programming

Powerful libraries—including deep learning, math, and video processing—include pre-optimized, domain-specific functions to accelerate compute-intensive workloads on Intel® CPUs and GPUs.

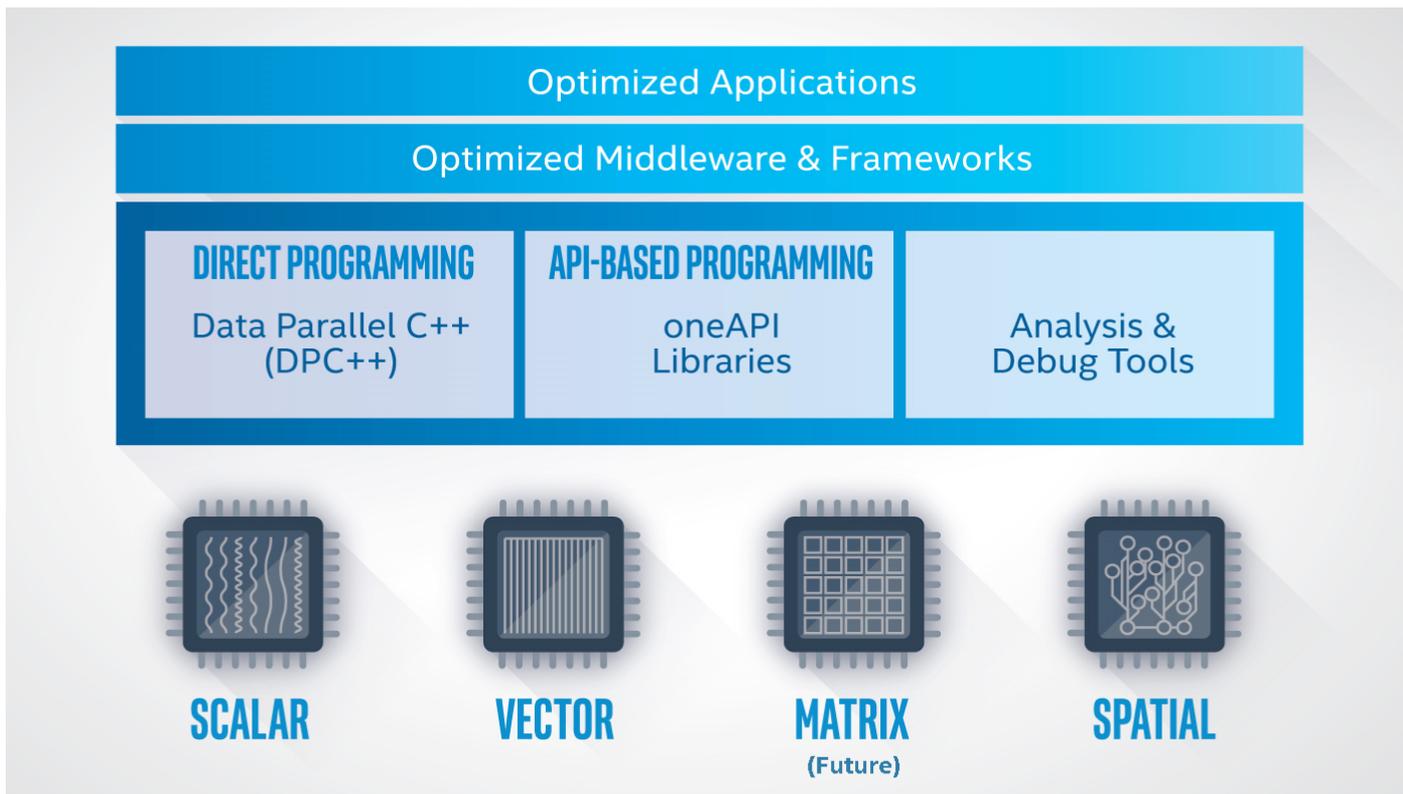
Advanced Analysis and Debug Tools

For profiling, design advice, and debug, Intel oneAPI products include leading analysis tools:

- **Intel® VTune™ Profiler^(Beta)** to find performance bottlenecks fast in CPU, GPU, and FPGA systems
- **Intel® Advisor^(Beta)** for vectorization, threading, and accelerator offload design advice
- **GDB*** for efficient code troubleshooting



**Get the Intel® oneAPI
Base Toolkit Now >**



Toolkits Tailored to Your Needs

Start with the Intel® oneAPI Base Toolkit

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.

It features the Data Parallel C++ (DPC++) language, an evolution of C++ that:

- **Allows code reuse across hardware targets**—CPUs, GPUs, and FPGAs
- **Permits custom tuning** for individual accelerators
- **Domain-specific libraries and the Intel® Distribution for Python*** to provide drop-in acceleration across relevant architectures
- **Enhanced profiling, design assistance, and debug tools** to complete the kit'

Here's what you get:

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

Add Domain-Specific Toolkits for Your Specialized Workloads

Besides the Intel oneAPI Base Toolkit that serves a broad set of developers' needs, there are four add-on toolkits that combine it to give you the specialized tools you need:

- **Intel® oneAPI HPC Toolkit^(Beta)**: 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 IoT Toolkit^(Beta)**: Accelerate development of smart, connected devices for healthcare, smart homes, aerospace, security, and more.
- **Intel® oneAPI Rendering Toolkit^(Beta)**: 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.
- **Intel® oneAPI DL Framework Developer Toolkit^(Beta)**: Develop new—or customize existing—deep learning frameworks using common APIs. Optimize for high-performance on Intel CPUs and GPUs for either single-node or multi-node distributed processing.

There are three more toolkits closely related to oneAPI:

- **Intel® Distribution of OpenVINO™ Toolkit**: Accelerate deep learning inference and seamlessly deploy intelligent solutions across Intel® platforms and

accelerators through this toolkit powered by oneAPI components.

- **Intel® AI Analytics Toolkit^(Beta)**: Achieve end-to-end performance for AI workloads with this toolkit powered by oneAPI. Accelerate each step in the pipeline—training deep neural networks, integrating trained models into applications for inference, and executing functions for data science and analytics.
- **Intel® System Bring-Up Toolkit^(Beta)**: Strengthen system reliability and optimize system power and performance with this collection of debug, trace, and power and performance analysis tools to let you quickly debug and analyze the entire platform.

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 >](#)



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

**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. For more complete information visit <http://www.intel.com/performance>.

Intel's 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.

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 non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade. This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps. The products and services described may contain defects or errors known as errata which may cause deviations from published specifications. Current characterized errata are available on request. Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.htm. For more information regarding performance and optimization choices in Intel® Software Development Products, see our Optimization Notice: <https://software.intel.com/articles/optimization-notice#opt>
Copyright © 2019, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Inside, Intel Atom, Intel Core, Intel V Tune, and Intel Xeon are trademarks of Intel Corporation in the U.S. and/or other countries.