

Product Brief

High-Performance Computing
Intel® oneAPI Base and Rendering Toolkit



Powerful Libraries for High-Fidelity Visualization Applications



Intel® oneAPI Base and Rendering Toolkit is an award-winning, comprehensive suite of development tools that enable developers and content creators to build high-performance, high-fidelity, extensible, and cost-effective visualization applications and solutions. Using open source rendering libraries, your solution is optimized on Intel® Xeon® Scalable processors, other Intel® CPUs, Intel® Xe architecture GPUs. Combine the core set of tools from the Intel oneAPI Base Toolkit and add tools focused on advanced ray tracing and rendering, and developers can achieve high performance and high image quality in the most data-intensive use cases.

Who needs it

Developers working on rendering applications. (Figure 1) including:

- Digital content creation
- Professional rendering
- Animation
- Scientific visualization
- Computer-aided design, architectural engineering
- Gaming VR and AR

What it does

Meets the demands of the highest-quality ray tracing use cases without the limitations and compromises of rasterization. Users can interactively visualize huge data sets (terabytes), embracing full system memory beyond today's memory limits of GPU add-in cards. Plus, your existing investments in graphics and rendering solutions based on Intel oneAPI Rendering Toolkit libraries will seamlessly scale to gain the exponential performance benefits of future, flexible CPU-plus-accelerator platforms.

Highlights

Intel oneAPI Rendering Toolkit is a set of advanced, open source libraries that deliver high performance and high image quality for data-intensive use cases on CPU platforms of all sizes, including workstation, data center/cloud, and high performance computing (HPC) clusters. It provides both scalable and interactive ray tracing and OpenGL visualization.

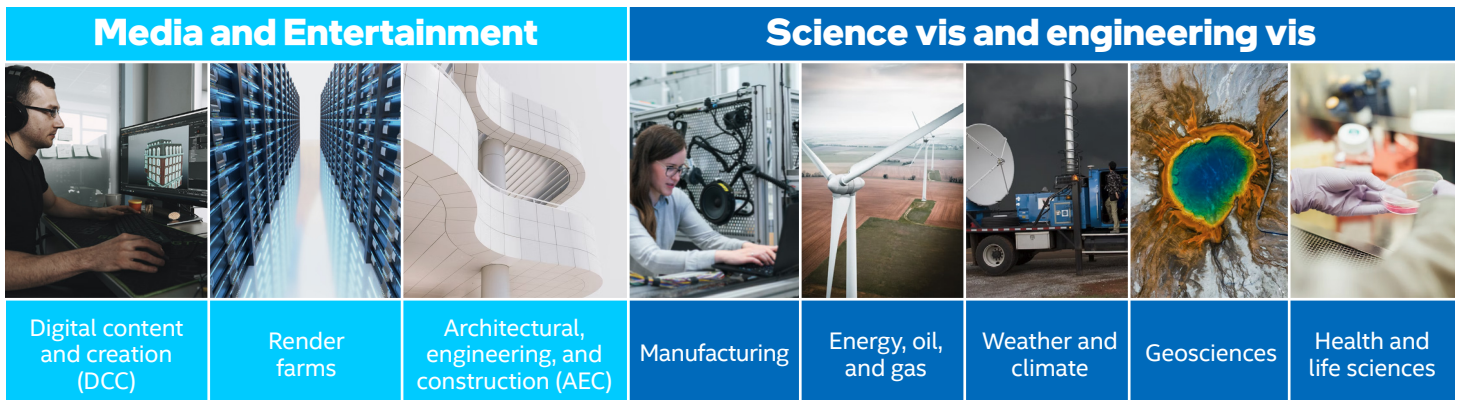


Figure 1. Applications for Intel® oneAPI Rendering Toolkit

Intel® oneAPI Base + Rendering Toolkit

Direct Programming

- Intel® oneAPI DPC++/ C++ Compiler
- Intel® DPC++ Compatibility Tool
- Intel® Distribution for Python*
- Intel® Embree
- Intel® FPGA Add-On for oneAPI Base Toolkit
- Intel® Open Image Denoise
- Intel® OpenSWR
- Intel® Open VKL
- Intel® OSPRay
- Intel® OSPRay Studio
- Intel® OSPRay for Hydra*
- Rendering Toolkit Utilities


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 and Debugging Tools

- Intel® Advisor
- Intel® VTune™ Profiler
- Intel® Distribution for GDB*


*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


What you get

The Intel oneAPI Rendering Toolkit is a set of advanced, open source libraries that deliver high performance and both scalable and interactive ray tracing and OpenGL image quality for data-intensive use cases on CPU platforms visualization.

Intel leadership in ray tracing for high-performance graphics and compute




Intel® Embree: High-performance, feature-rich ray tracing and photorealistic rendering.




Intel® OSPRay: Scalable, portable, distributed rendering API.



Intel® Open Image Denoise: AI-accelerated denoiser for superior visual quality.



Intel® Open Volume Kernel Library: Render and simulate 3D spatial data processing.



Intel® OSPRay Studio: Real-time rendering via graphical user interface-based, scene graph application add-on Intel OSPRay.

Ray Tracing and Rendering Software Toolkit Capabilities	Intel® oneAPI Rendering Toolkit	Competitive Tools
Runs at scale from laptop or workstation across enterprise HPC, cloud, and wherever you have compute infrastructure.	✓	No
Visualize huge data sets (terabytes) interactively, embracing full system memory, create freely beyond limits to 32 GB card VRAM.	✓	No
Persistent memory support for data loss protection, reduced restart times, and I/O performance for ever-growing geometric and volumetric data.	✓	No
Combines volume and geometry rendering.	✓	No
Render in high-fidelity, path-traced, scientific visualization: Shadows, ambient occlusion, global illumination, denoising.	✓	✓
Supports model complexity beyond triangles: Triangle meshes; quad meshes; flat or round oriented curves; discs, spheres, and subdivision surfaces.	✓	Limited
Motion blur support for linear, time-segmented objects and quaternion.	✓	✓
Open source libraries. Permissive Apache 2 license with 100+ ISVs delivering content creation tools using common scalable compute infrastructure.	✓	No
Future-proofs code with built-in support for hybrid rendering on Intel® CPUs and future Intel® Xe architecture/GPUs.	✓	No

The **Intel oneAPI Base Toolkit** is a core set of advanced compilers, libraries, and analysis and debug tools, such as Intel® VTune™ Profiler and Intel® Advisor, for building and deploying high-performance, data-centric applications across a variety of Intel® architectures. Data Parallel C++ (DPC++) language for direct programming, which is also included, is an evolution of C++ that incorporates SYCL, allowing code reuse across hardware targets. Enables high productivity and performance across CPU, GPU, and FPGA architectures while permitting accelerator-specific tuning.

- **Intel® oneAPI DPC++/C++ Compiler:** A standards-based, cross-architecture 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® 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.
- **Intel® oneAPI Threading Building Blocks:** Simplify parallelism with this advanced threading and memory management template library.
- **Intel® Advisor:** Design code for efficient vectorization, threading, and offloading to accelerators.
- **Intel® VTune Profiler:** Find and optimize performance bottlenecks across CPU, GPU, and FPGA systems.

And more

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.

Get started

Get the **Intel® oneAPI Base & Rendering Toolkit** ›

Learn more about **Priority Support for Intel® oneAPI Toolkits** ›

Check out the **Intel® Developer Cloud** ›



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.