

oneAPI

Single Programming Model
to Deliver Cross-Architecture Performance

Industry initiative, Intel® oneAPI Beta Products

francisco.perez@intel.com

January 2020

INTEL DATA-CENTRIC HARDWARE: HIGH PERFORMANCE, FLEXIBLE OPTIONS

General Purpose
CPU



SCALAR

Programmable Data
Parallel Accelerator
GPU




VECTOR

FPGA



SPATIAL

Domain Optimized
Accelerator
NNP



MATRIX

GENERAL PURPOSE

HARDWARE

WORKLOAD OPTIMIZED

Provide optimal performance
over the widest variety of
workloads

Deliver highest performance
per \$/Watt/U/Rack for critical
applications

Growth in specialized workloads
Diverse set of data-centric hardware required

PROGRAMMING CHALLENGES FOR MULTIPLE ARCHITECTURES

Today, each kind of data-centric HW:

No common programming language or APIs

Inconsistent tool support across platforms
(ie. profiling, debugging)

Each platform requires unique software
investment (time and \$)

No reuse to target a different architecture

Application Workloads Need Diverse Hardware



SCALAR



VECTOR



MATRIX



SPATIAL

Middleware / Frameworks

Language & Libraries



CPU



GPU



FPGA



OTHER ACCEL.

INTRODUCING ONEAPI

A project to deliver a unified software development environment across CPU and accelerator architectures.

Unified and simplified language and libraries for expressing parallelism

Delivering native high-level language performance

Based on industry standards and open specifications

Application Workloads Need Diverse Hardware



SCALAR



VECTOR



MATRIX



SPATIAL

Middleware / Frameworks

Industry Initiative



Intel Product



CPU



GPU



FPGA



OTHER ACCEL.

ONEAPI INDUSTRY INITIATIVE

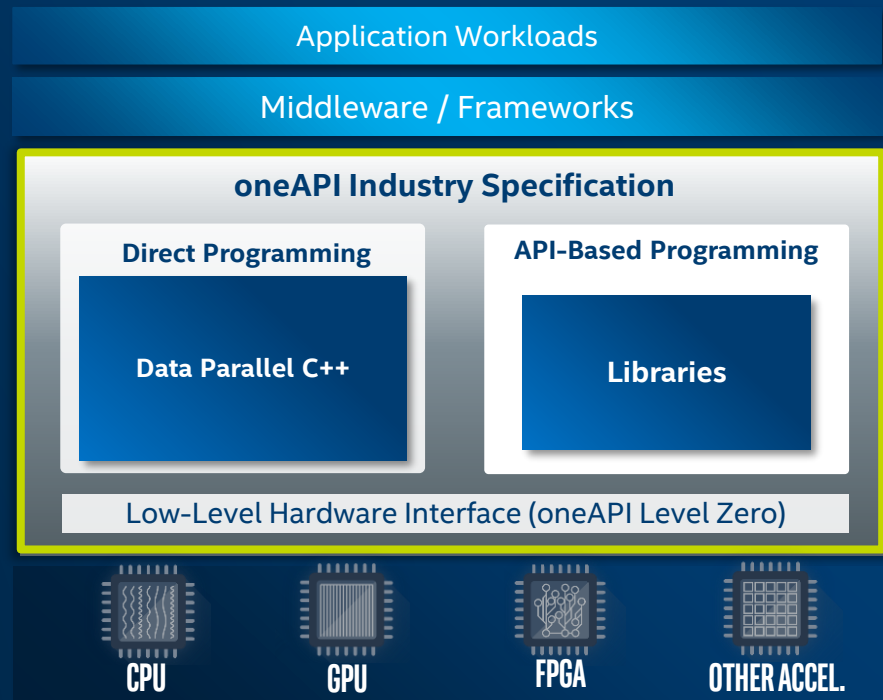
ALTERNATIVE TO SINGLE-VENDOR SOLUTION

Open standard specification to promote community and industry vendors support and includes:

Direct programming flow with an open, unified language: DPC++ based on C++ with SYCL extensions

API-based programming flow with a set of powerful libraries designed for each hardware to accelerate key domain-specific functions, most of them open sourced

Specification of Low-level interface to provide a hardware abstraction layer to vendors



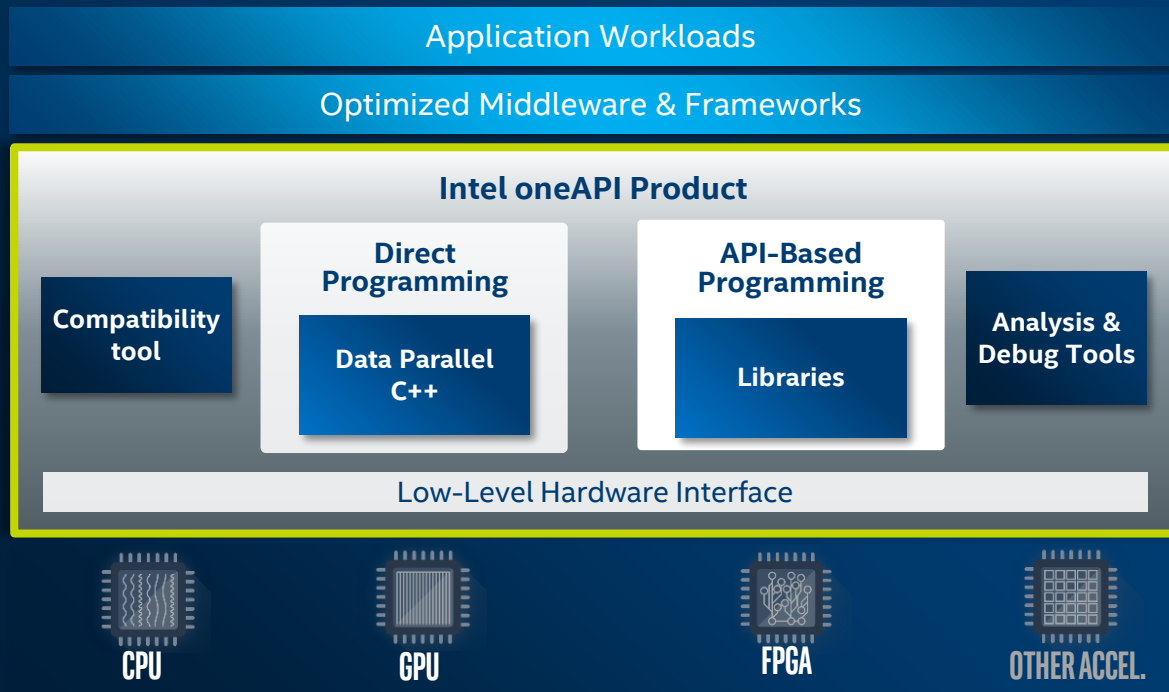
Together, these components allow Intel and other companies to build their own implementations of oneAPI

INTEL[®] ONEAPI PRODUCTS^(BETA)

Intel's reference implementation of oneAPI

Distributed through a core toolkit and a complementary set of add-on domain-specific toolkits

Includes DPC++ compatibility tool for code migration along with advanced performance analysis and debug tools



Visit software.intel.com/oneapi for more details

Some capabilities may differ per architecture and custom-tuning will still be required. Other accelerators to be supported in the future.

[Optimization Notice](#)

Copyright © 2019, Intel Corporation. All rights reserved.

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

INTEL® ONEAPI TOOLKITS^(BETA)

TOOLKITS TAILORED TO YOUR NEEDS

Domain-specific sets of tools to get your job done quickly.



Intel® oneAPI Base Toolkit

A core set of high-performance tools for building Data Parallel C++ applications and oneAPI library based applications

[Learn More](#)



Intel® oneAPI HPC Toolkit

Everything HPC developers need to deliver fast C++, Fortran, & OpenMP* applications that scale

[Learn More](#)



Intel® oneAPI IoT Toolkit

Tools for building high-performing, efficient, reliable solutions that run at the network's edge

[Learn More](#)



Intel® oneAPI DL Framework Developer Toolkit

Tools for developers & researchers who build deep learning frameworks or customize existing ones so applications run faster

[Learn More](#)



Intel® oneAPI Rendering Toolkit

Powerful rendering libraries to create high-performance, high-fidelity visualization applications

[Learn More](#)

Toolkits Powered by oneAPI

Intel® System Bring-Up Toolkit

Tools to debug & tune power & performance in pre- & post-silicon development

[Learn More](#)

Intel® Distribution of OpenVINO™ Toolkit

Tools to build high performance deep learning inference & computer vision applications (production-level tool)

[Learn More](#)

Intel® AI Analytics Toolkit

Tools to build applications that leverage machine learning & deep learning models

[Learn More](#)

DETAILS ABOUT

INTEL[®] ONEAPI TOOLKITS^(BETA)

INTEL[®] ONEAPI BASE TOOLKIT

INTEL® ONEAPI BASE TOOLKIT (BETA)

Core set of frequently used tools and libraries for developing high-performance applications across diverse architectures—CPU, GPU, FPGA.

Top Features/Benefits

Data Parallel C++ compiler, library, and analysis tools

DPC++ Compatibility tool helps migrate existing code written in CUDA*

Python distribution includes accelerated scikit-learn, NumPy, SciPy libraries

Optimized performance libraries for threading, math, data analytics, deep learning, and video/image/signal processing

Intel® oneAPI Base Toolkit

DIRECT PROGRAMMING

Intel® oneAPI DPC++ 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 Comms. Library

Intel® oneAPI Deep Neural Network Library

Intel® Integrated Performance Primitives

ANALYSIS & DEBUG TOOLS

Intel® VTune™ Profiler

Intel® Advisor

GDB*

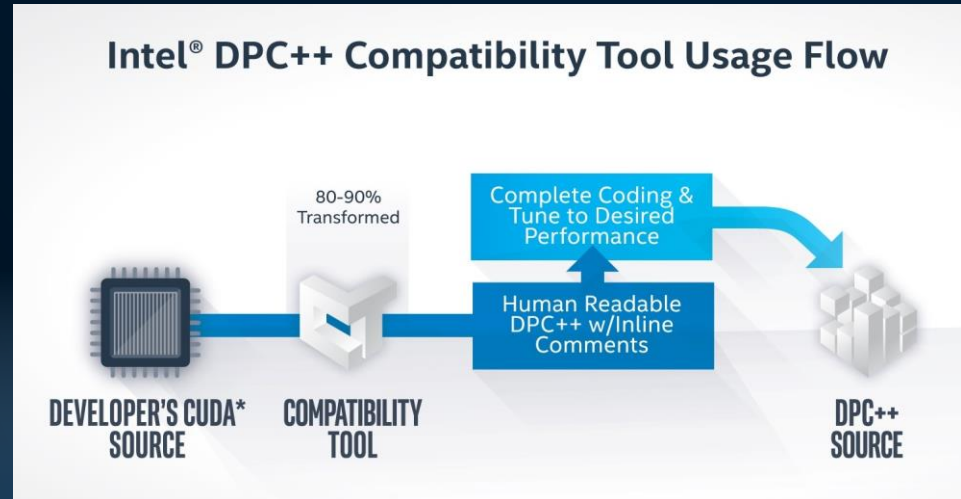
INTEL® DPC++ COMPATIBILITY TOOL^(BETA)

MINIMIZES CODE MIGRATION TIME

Assists developers migrating code written in CUDA* to DPC++ bridging the gap between the languages

Migrated result includes:

- DPC++ code and inline comments
- Unchanged CUDA code and hints to assist developers complete the rest manually



INTEL® VTUNE™ PROFILER (BETA)

DPC++ PROFILING — TUNE FOR CPU, GPU & FPGA

Analyze Data Parallel C++ (DPC++)

See the lines of DPC++ that consume the most time

Tune for CPU, GPU & FPGA


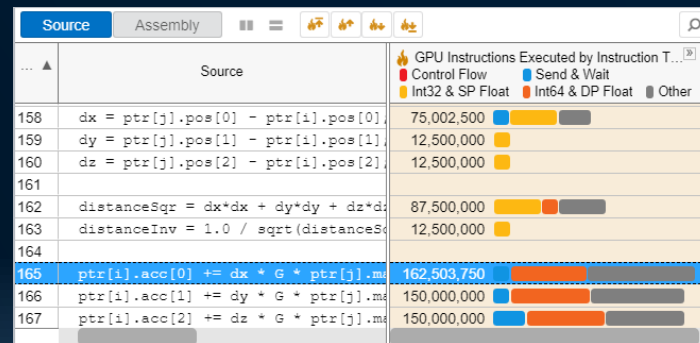
Optimize for any supported hardware accelerator

Optimize Offload

Tune OpenMP* offload performance

Supports Popular Languages

DPC++, C, C++, Fortran, Python*, Go*, Java*, or a mix



INTEL® ADVISOR (BETA)

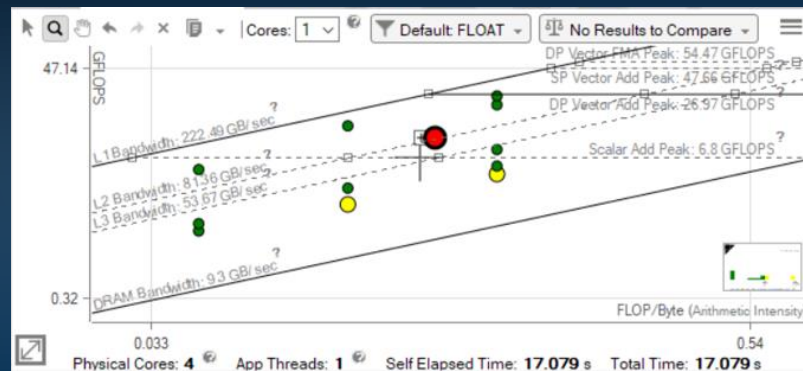

DESIGN ASSISTANT — DESIGN FOR MODERN HARDWARE

Determine if your code would benefit from offloading to an accelerator – even before you have the hardware

Projects performance on accelerators

Estimates overhead from data transfers and kernel launch costs

Pinpoint accelerator performance bottlenecks (memory, cache, compute and data transfer)



GDB* (BETA)

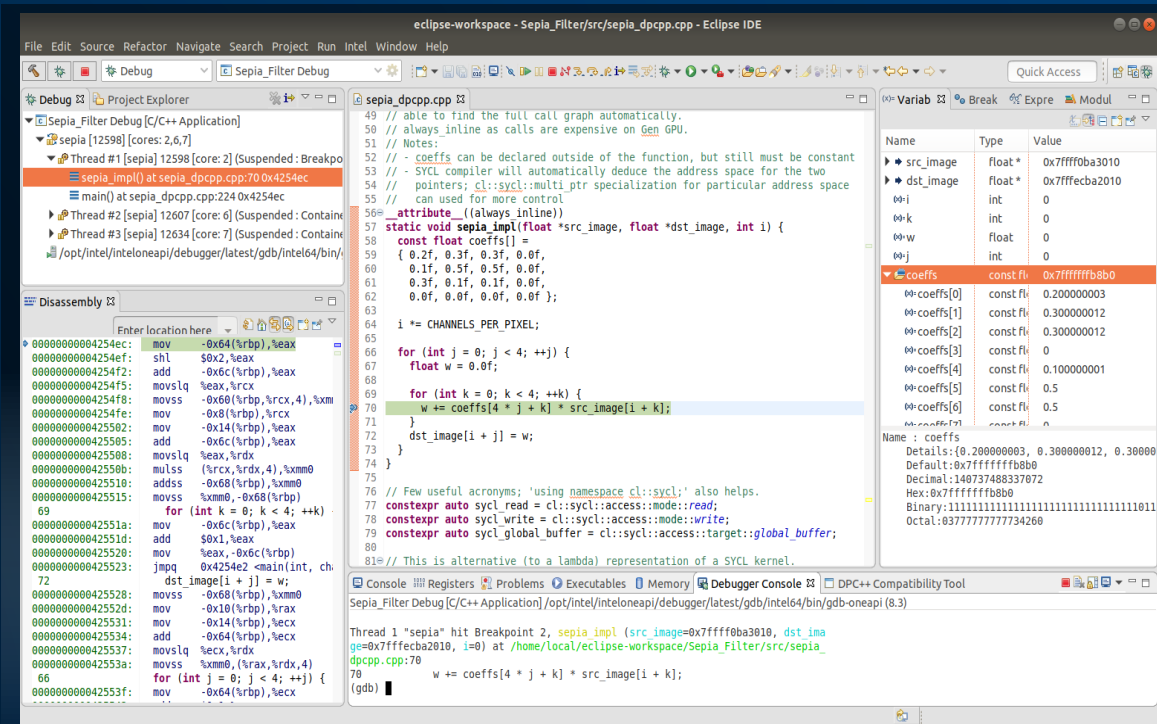
DPC++ DEBUG — HETEROGENEOUS APPLICATION DEBUG

High-level language debug support

Multiple accelerator support: CPU, GPU and FPGA in emulation

Auto-detect accelerator architecture during application runtime

Non-proprietary open-source solution based on GDB*



ONEAPI FOR FPGA

DPC++ CODING FOR SPATIAL ARCHITECTURE

For Experienced FPGA Developers

Ease of Use

Experienced FPGA users can take advantage of a streamlined programming model using DPC++

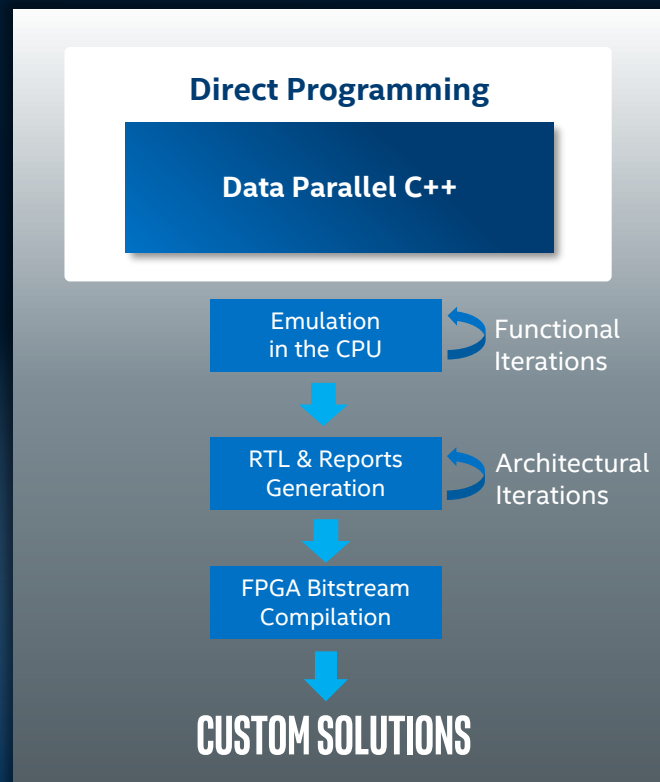
Runtime Analysis Support

Collect profiling data at runtime to analyze CPU & FPGA interaction with Intel® VTune™ Profiler

Device Specific Optimizations

Developers are expected to make use of FPGA specific pragmas & attributes to optimize their designs

Process data faster with deterministic low latency & high throughput in FPGAs



DETAILS ABOUT INTEL® ONEAPI TOOLKITS^(BETA)

DOMAIN-SPECIFIC TOOLKITS FOR SPECIALIZED WORKLOADS

INTEL® ONEAPI HPC TOOLKIT (BETA)

A toolkit that makes easier to build, analyze, optimize & scale HPC applications for Intel® Xeon® Scalable, Intel® Core™ processors & Intel® Accelerators.

Who Uses It?

C/C++, Fortran, OpenMP & MPI application developers

Top Features/Benefits

Optimized compilers & performance libraries for Intel® architectures

Powerful analysis tools to identify optimization opportunities for threading, memory & offloading

Intel oneAPI Tools for HPC

DIRECT PROGRAMMING

Intel® C++ Compiler with OpenMP*

Intel® Fortran Compiler with OpenMP*

Intel® oneAPI DPC++ Compiler

Intel® DPC++ Compatibility Tool

Intel® Distribution for Python*

Intel® FPGA Add-on for oneAPI Base Toolkit

API-BASED PROGRAMMING

Intel® MPI Library

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® Inspector

Intel® Trace Analyzer & Collector

Intel® Cluster Checker

Intel® VTune™ Profiler

Intel® Advisor

GDB*

■ Intel® oneAPI HPC Toolkit +

■ Intel® oneAPI Base Toolkit

LEARN DPC++ AND ONEAPI

<https://software.intel.com/en-us/oneapi/documentation>

 Intel Software Developer Zone

 INTEL® ONEAPI TOOLKITS(BETA)

oneAPI Home

Featured Documentation

Get Started with the Intel® oneAPI Toolkits

[Linux*](#) | [Windows*](#) | [Code Samples](#) | [Download Documentation](#)


Guides


[Installation Instructions for oneAPI Toolkits Programming Guide](#)
[Intel® DPC++ Compatibility Tool User Guide](#)

<https://jamesreinders.com/dpcpp/>


learn DPC++ and oneAPI


 Preview: Chapters 1-4 from the new DPC++ Book

 Download: DPC++ Code from DPC++ Book

 View Currently Known (Book) Errata

 Online Training (Great Complement to Book!)

 Download: DPC++ Code and Labs for Online Training

 Download: Slides used in Online Training

 Intel oneAPI (and DPC++) website for Software Developers

 DPC++ Open Source (github)

 oneAPI project home

ONEAPI AVAILABLE NOW ON INTEL[®] DEVCLOUD

A development sandbox to develop, test and run your workloads across a range of Intel CPUs, GPUs, and FPGAs using Intel's oneAPI beta software

software.intel.com/devcloud/oneapi



NO DOWNLOADS | NO HARDWARE ACQUISITION | NO INSTALLATION | NO SET-UP & CONFIGURATION


GET UP & RUNNING IN SECONDS!

ONEAPI DEVCLLOUD

What you can do


- Learn DPC++
- Learn about Intel oneAPI toolkits
- Evaluate workloads
- Prototype your projects
- Build heterogeneous applications

Hardware




CPU:

Intel® Xeon® Scalable 6128 processors
Intel® Xeon® Scalable 8256 processors



GPU:

Intel® Xeon® E-2176 P630 processors
(with Intel® Processor Graphics Gen9)



FPGA:

Intel® Arria® 10 FPGAs

Included Toolkits

- Intel® oneAPI Base Toolkit
- Intel® oneAPI HPC Toolkit
- Intel® oneAPI Deep Learning Framework Developer Toolkit
- Intel® AI Analytics Toolkit
- Intel® Distribution of OpenVINO™ Toolkit
- + more

ONEAPI DEV CLOUD

What you get

- Access to Intel® oneAPI software and hardware
- 200 GB of file storage
- 192 GB RAM
- 120 days of access (extensions available)
- Terminal Interface (Linux*) and web (jupyter notebook)

How it works



SUMMARY



Diverse workloads are driving the need for heterogeneous compute architectures

oneAPI is an open industry initiative & an Intel reference product

oneAPI unifies & simplifies programming of heterogeneous architectures
delivering developer productivity & maximum performance for each hardware

Get Started – test code & workloads using the Intel® DevCloud

<https://www.oneapi.com/>

NOTICES & DISCLAIMERS

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. No product or component can be absolutely secure. Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

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 www.intel.com/benchmarks.

INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS". NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO THIS INFORMATION INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Copyright ©, Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, Core, VTune, and OpenVINO are trademarks of Intel Corporation or its subsidiaries in the U.S. and other countries.


Optimization Notice

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.

Notice revision #20110804

ONEAPI INDUSTRY SPECIFICATION

[SPEC.ONEAPI.COM/ONEAPI/](https://spec.oneapi.com/oneapi/)



[Notices and Disclaimers](#)

[Contribution Guidelines](#)

[Introduction](#)

[Software Architecture](#)

[Library Interoperability](#)

[oneAPI Elements](#)

[Data Parallel C++ \(DPC++\)](#)

[oneAPI Data Parallel C++ Library \(oneDPL\)](#)

[oneAPI Deep Neural Network Library \(oneDNN\)](#)

[oneAPI Collective Communications Library \(oneCCL\)](#)

[oneAPI Level Zero \(Level Zero\)](#)

[oneAPI Data Analytics Library \(oneDAL\)](#)

[oneAPI Threading Building Blocks \(oneTBB\)](#)

[oneAPI Video Processing Library \(oneVPL\)](#)


[oneAPI Math Kernel Library \(oneMKL\)](#)

[Contributors](#)

Refer to software.intel.com/articles/optimization-notice for more information regarding performance & optimization choices in Intel software products.

Copyright © Intel Corporation. All rights reserved.

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



Software