




ONEAPI


SINGLE PROGRAMMING MODEL TO DELIVER CROSS-ARCHITECTURE PERFORMANCE



GETTING STARTED WITH ONEAPI

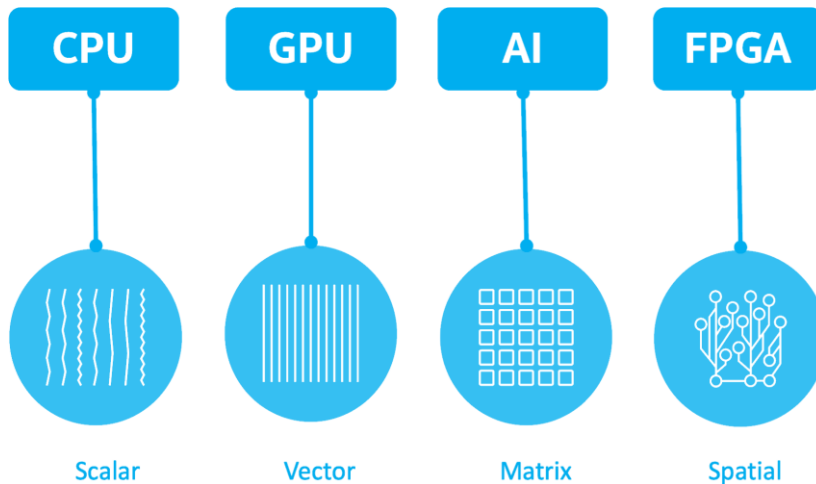
March 2020

HOW ONEAPI ADDRESSES OUR HETEROGENEOUS WORLD?



DIVERSE WORKLOADS DEMAND DIVERSE ARCHITECTURES

The future is a diverse mix of scalar, vector, matrix, and spatial architectures deployed in CPU, GPU, AI, FPGA, and other accelerators.



Optimization Notice

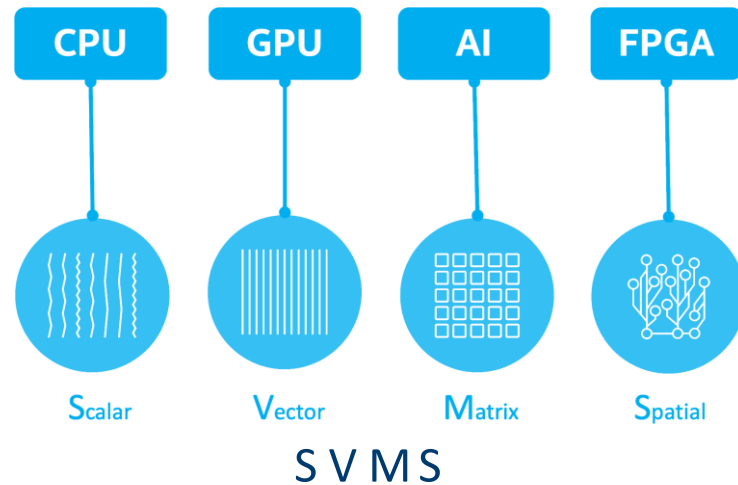
Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Getting started with oneAPI



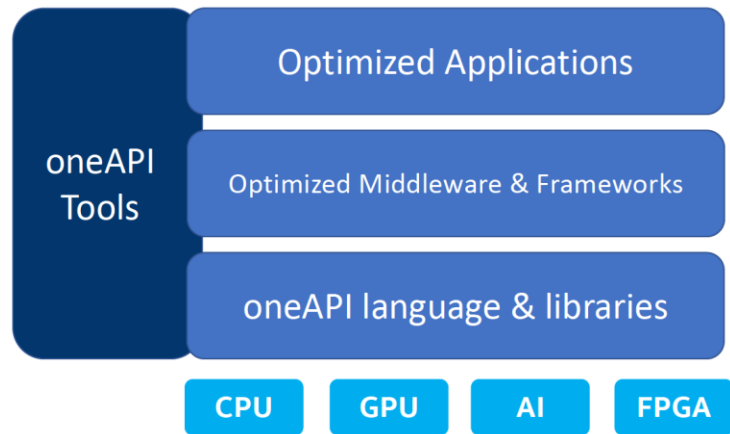
CHALLENGE: PROGRAMMING IN A HETEROGENEOUS WORLD

- ▷ Diverse set of data-centric hardware
- ▷ No common programming language or APIs
- ▷ Inconsistent tool support across platforms
- ▷ Proprietary solutions on individual platforms
- ▷ Each platform requires unique software investment



INTEL'S ONEAPI CORE CONCEPT

- ▶ Project oneAPI delivers a unified programming model to simplify development across diverse architectures
- ▶ Common developer experience across SVMs
- ▶ Uncompromised native high-level language performance
- ▶ Unified language and libraries for expressing parallelism



- ▶ Support for CPU, GPU, AI, and FPGA
- ▶ Based on industry standards and open specifications

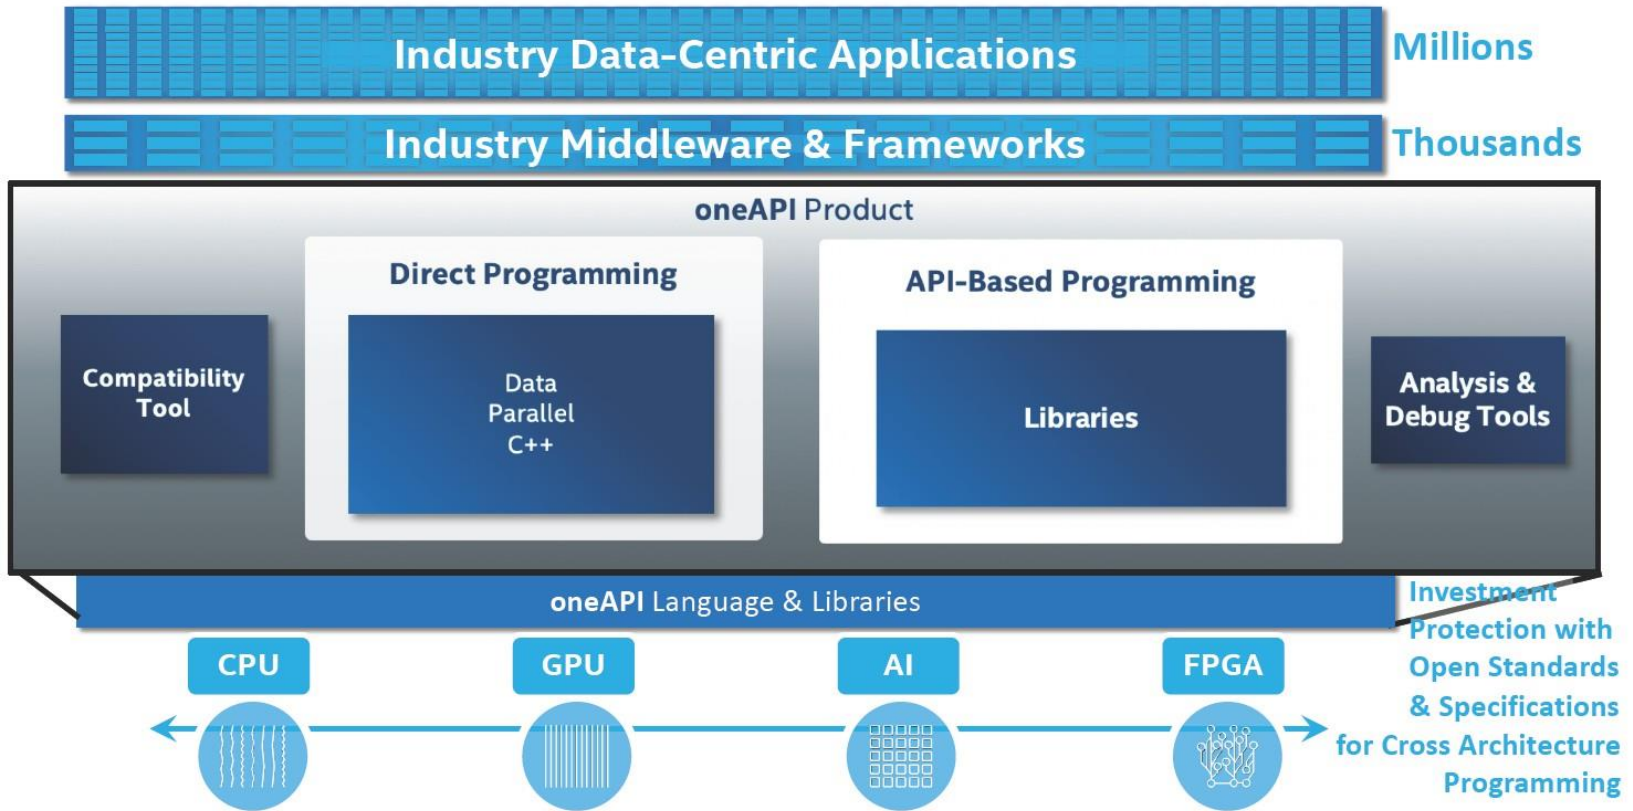
<https://www.oneapi.com/spec/>

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



ONEAPI FOR CROSS-ARCHITECTURE PERFORMANCE




Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Getting started with oneAPI



WHAT IS DATA PARALLEL C++?



WHAT IS DPC++?

The language is:

C++

+

SYCL

<https://www.khronos.org/sycl/>

+

Additional Features


such as...

ndrange subgroups, USM, ordered queue

<https://github.com/intel/llvm/blob/sycl/sycl/doc/extensions/>



INTEL[®] ONEAPI TOOLKITS



INTEL® ONEAPI TOOLKITS(BETA)

One core toolkit:



Intel® oneAPI Base Toolkit

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

- ▶ Additional toolkits targeting specific data-centric workloads
- ▶ Each includes oneAPI components and complementary oneAPI ecosystem components



Intel® oneAPI HPC Toolkit

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



Intel® oneAPI IoT Toolkit

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



Intel® oneAPI DL Framework Developer Toolkit

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



Intel® oneAPI Rendering Toolkit

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

Toolkits Powered by oneAPI

Intel® System Bring-Up Toolkit

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

Intel® Distribution of OpenVINO™ Toolkit

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

Intel® AI Analytics Toolkit

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

software.intel.com/oneapi

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Getting started with oneAPI



INTEL® ONEAPI BASE TOOLKIT (BETA)

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

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*

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

Optimization Notice

Copyright © Intel Corporation. All rights reserved.

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

Getting started with oneAPI



POWERFUL ONEAPI TOOLS

Intel® VTune™ Profiler


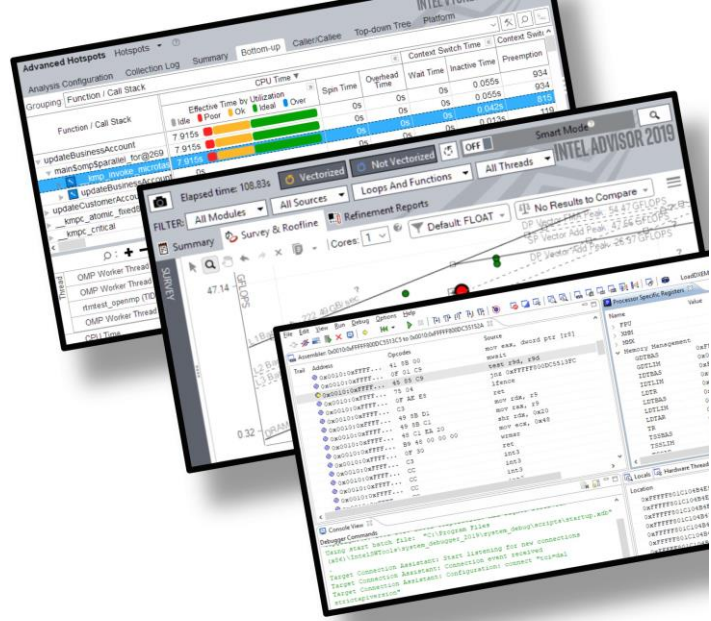
- ▶ Profiler to analyze CPU and accelerator performance of compute, threading, memory, storage, and more

Intel® Advisor

- ▶ Design assistant to provide advice on threading, and vectorization
- ▶ Offload Advisor and GPU Roofline*

Intel®-enhanced gdb

- ▶ Application debugger for fast code debug on CPUs and accelerators



Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

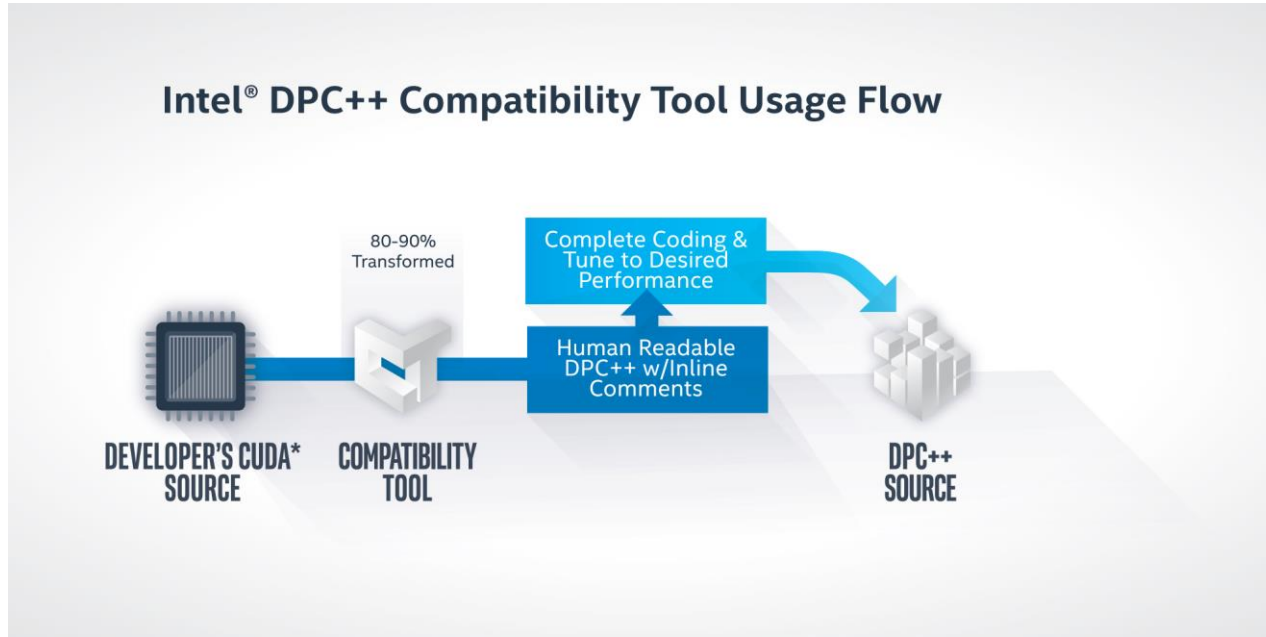
Getting started with oneAPI



INTEL® DPC++ COMPATIBILITY TOOL^(BETA)

MINIMIZES CODE MIGRATION TIME

- Assists developers migrating code written in CUDA* to DPC++ **once**, generating **human readable** code wherever possible
- Inline comments are provided to help developers complete their code



<https://software.intel.com/en-us/sycl-compatibility-tool-get-started>


Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.

Getting started with oneAPI



WHERE AND HOW TO GET AND USE DPC++?



EASIEST - USE THE PREBUILT DPC++ WITH COMPLETE ONEAPI TOOLKITS

▶ Download Toolkits

You'll want oneAPI toolkits, even if you build your own DPC++ compiler.

▶ DevCloud

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



INTEL® DEVCLLOUD FOR ONEAPI PROJECTS

A development sandbox to develop, test, and run your workloads across a range of Intel®-based CPUs, GPUs, and FPGAs using oneAPI^(Beta) software

[Sign Up for Beta](#)

What You Can Do



Learn Data Parallel C++



Learn about Intel® oneAPI Toolkits



Evaluate Workloads



Prototype Your Project




Build Heterogeneous Applications

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



RESOURCES



RESOURCES

- ▶ Get oneAPI Toolkit(s) now:

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

- ▶ oneAPI Spec:


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

- ▶ Webinars:

<https://techdecoded.intel.io/topics/oneapi/>

<https://techdecoded.intel.io/essentials/dpc-part-2-programming-best-practices/>

- ▶ Book (Chapters 1-4 Preview)



<https://tinyurl.com/book-dpcpp>

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



INSTALLATION

oneAPI main entry <https://software.intel.com/en-us/oneapi>

- Links to installation of oneAPI kits
- Installation guide <https://software.intel.com/en-us/articles/installation-guide-for-intel-oneapi-toolkits>
- Installation using online installer, full packages or repositories!
- Hardware and software requirements <https://software.intel.com/en-us/articles/intel-oneapi-dpcpp-system-requirements-beta>



PUBLIC DPC++ COMPILER

DPC++ is SYCL with Intel extensions

Public compiler from Intel <https://github.com/intel/llvm>

- Sometimes fixes are included faster

SYCL extensions <https://github.com/intel/llvm/tree/sycl/sycl/doc/extensions>



DEVCLLOUD

DevCloud - a free development sandbox with access to the latest SVMS hardware from Intel and Intel oneAPI software

DevCloud link on <https://software.intel.com/en-us/oneapi>

Get started (after sign-in) <https://devcloud.intel.com/oneapi/get-started/>



SPECIFICATION, GUIDE AND BOOK

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

oneAPI Programming Guide <https://software.intel.com/en-us/oneapi-programming-guide> (also as PDF)

James Reinder's oneAPI book (chapters 1-4) <https://jamesreinders.com/dpcpp/>



TRAINING AND WEBINARS

Get started <https://software.intel.com/en-us/get-started-with-intel-oneapi-linux>

- oneAPI for Intel FPGAs
<https://www.intel.com/content/www/us/en/programmable/products/design-software/high-level-design/one-api-for-fpga-support.html>

Colfax training <https://www.colfax-intl.com/training/intel-oneapi-training>

Webinars (not oneAPI exclusive) <https://techdecoded.intel.io/webinar-registration/upcoming-webinars>

WEDNESDAY,
MARCH 25, 2020
9:00 AM PDT

#oneAPI

Intel® oneAPI Threading Building Blocks: Optimizing for NUMA Architectures

For over a decade, TBB has been a popular C++ library that delivers composable, scalable performance on multicore systems. And in today's new cross-architecture paradigm, the oneAPI version of TBB (oneTBB) is being tapped for multi-socket NUMA systems ... which presents new optimization challenges. Find out how to solve them.

[Read More +](#)



DPC++ COMPATIBILITY TOOL

Migration of CUDA code to DPC++

DPC++ Compatibility Tool is included in the oneAPI Base Toolkit

Get Started <https://software.intel.com/en-us/get-started-with-intel-dpcpp-compatibility-tool>

DPC++ Compatibility Tool User Guide <https://software.intel.com/en-us/intel-dpcpp-compatibility-tool-user-guide>

- Includes an example in section “Usage Workflow Overview”



SUPPORT OF NVIDIA CUDA DEVICES

First SYCL compiler available from Codeplay <https://github.com/codeplaysoftware/sycl-for-cuda>

Now also included in the public Intel compiler
<https://github.com/intel/llvm/blob/sycl/sycl/doc/GetStartedWithSYCLCompiler.md>

Build SYCL toolchain with support for NVIDIA CUDA

There is experimental support for SYCL for CUDA devices.

To enable support for CUDA devices, the following arguments need to be added to the CMake command when building the SYCL compiler.

```
-DCUDA_TOOLKIT_ROOT_DIR=/usr/local/cuda/ \  
-DLLVM_ENABLE_PROJECTS="clang;llvm-spirv;sycl;libclc"\  
-DSYCL_BUILD_PI_CUDA=ON\  
-DLLVM_TARGETS_TO_BUILD="X86;NVPTX"\  
-DLIBCLC_TARGETS_TO_BUILD="nvptx64-;nvptx64--nvidiacl"
```

Enabling this flag requires an installation of [CUDA 10.1](#) on the system, refer to [NVIDIA CUDA Installation Guide for Linux](#).

Currently, the only combination tested is Ubuntu 18.04 with CUDA 10.2 using a Titan RTX GPU (SM 71), but it should work on any GPU compatible with SM 50 or above.

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.



NOTICES & DISCLAIMERS

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, Pentium, Xeon, Core, VTune, OpenVINO, Cilk, 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

Optimization Notice

Copyright © 2020, Intel Corporation. All rights reserved.
*Other names and brands may be claimed as the property of others.





Software