PRODUCT BRIEF

High-Performance Computing
Intel® Parallel Studio XE 2020

intel® Software

Code that Performs

Intel® Parallel Studio XE helps developers take their HPC, enterprise, AI, and cloud applications to the max—with fast, scalable, and portable parallel code

Intel® Parallel Studio XE is a comprehensive suite of development tools that make it fast and easy to build modern code that gets every last ounce of performance out of the newest Intel® processors. This tool-packed suite simplifies creating code with the latest techniques in vectorization, multi-threading, multi-node, and memory optimization. Get powerful, consistent programming with Intel® Advanced Vector Extensions 512 (Intel® AVX-512) instructions for Intel® Xeon® Scalable processors, plus support for the latest standards and integrated development environments (IDEs).

Who Needs It?

- **C, C++, Fortran, and Python* software developers and architects** building HPC, enterprise, AI, and cloud solutions
- **Developers looking to maximize their software's performance** on current and future Intel® platforms

What it Does

- **Creates faster code[1].** Boost application performance that scales on current and future Intel® platforms with industry-leading compilers, numerical libraries, performance profilers, and code analyzers.
- **Builds code faster.** Simplify the process of creating fast, scalable, and reliable parallel code.
- **Delivers Priority Support.** Connect directly to Intel's engineers for confidential answers to technical questions, access older versions of the products, and receive free updates for a year. Paid license required.

What's New

- **Speed artificial intelligence inferencing.** Intel® Compilers, Intel® Performance Libraries and analysis tools support Intel® Deep Learning Boost, which includes Vector Neural Network Instructions (VNNI) in 2nd generation Intel® Xeon® Scalable processors (codenamed Cascade Lake/AP platforms)
- **Develop for large memories of up to 512GB DIMMs** with Persistence. Identify, optimize, and tune Intel® platforms for Intel® Optane™ DC Persistent Memory using Intel® VTune™ Profiler
- **Stay up to date with the latest standards support** providing additional Fortran 2018 features, full support of C++17 with initial C++20 support, and expanded OpenMP 4.5/5.0 support.
- **Use extended coarse grain profiling** with platform-level collection and analysis in Intel VTune Profiler to understand and optimize platform configuration for applications.

What's New (cont.)

- **HPC cloud support.** Take advantage of the AWS* Parallel Cluster* and AWS Elastic Fabric Adapter* for low-latency, high-bandwidth communications for MPI applications with Intel® MPI Library.
- **Supports the latest Intel® processors** including Intel Xeon Scalable processors (codenamed Cascade Lake, Cascade Lake AP, Cooper Lake, and Ice Lake).
- **Supports new OS**— Amazon Linux* 2 (supported features of tools and libraries may vary by instances and configurations).

Confidential Support and One Year of Updates Included

Every paid version of Intel® Software Development Products automatically includes priority support at our Online Service Center for at least one year from your date of purchase. You can extend it at a reduced rate.[2] You get:

- **Free access** to all new product updates and continued access to and support for older versions of the product
- **Direct and private interaction** with Intel's engineers. Submit confidential inquiries and code samples
- **Responsive help** with your technical questions and other product needs for both new and older versions
- **Community product forums** covering all of Intel's software development products
- **Access to a vast library** of self-help documents that build off decades of experience creating high-performance code

Choose Your Edition

Build with the Composer Edition

- **Improve performance**[1] 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.

INTEL® C++ COMPILER

- Use industry-leading, standards-based C/C++ tools to speed application performance.
- Experience seamless compatibility with popular compilers, development environments, and operating systems.
- Get superior vectorization and parallelization capabilities (including Intel® AVX 512 instructions) using the latest OpenMP* 5.0 parallel programming model.

Intel® C++ Compiler Boosts Application Performance on Linux*

Performance Advantage Measured by SPEC on Intel® Xeon® Platinum 8180 Processor
Relative Geomean Performance (FP Rate Base and FP Speed Base Higher is Better)

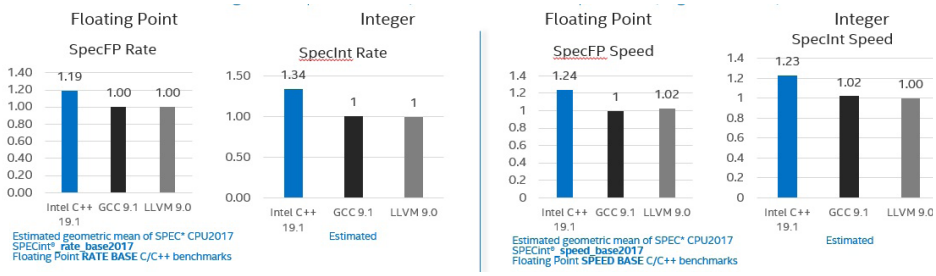

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

Performance results are based on testing as of August 26, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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.

Configuration: Testing by Intel as of August 26, 2019. Intel® Atom™ Processor C3958 @ 2.00GHz, 64 GB RAM, HyperThreading is not supported. Software: Intel® C++ Compiler 19.1., GCC* 9.1.0, Clang/LLVM 9.0. Red Hat Enterprise Linux Server* release 7.5 (Maipo), 3.10.0-862.el7.x86_64. CoreMark-Pro* Benchmark (www.eembc.org). Compiler flags: Intel C++ 19.1: '-Ofast -ipo -xATOM_SSE4.2 -mtune=goldmont. GCC 9.1' -Ofast -mfpmath=sse -flto -march=native -funroll-loops -ffat-lto-objects. Clang/LLVM 9.0: ' -Ofast -mfpmath=sse -flto -march=native -funroll-loops. GCC and clang/LLVM 32-bit modes have additional flag –m32.

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

Build with the Composer Edition (Continued)

INTEL® FORTRAN COMPILER

- Deliver superior Fortran application performance.
- Get extensive support for the latest Fortran standards (including full Fortran 2008 and initial Fortran 2015), with backwards compatibility to FORTRAN 77.
- Boost SIMD vectorization and threading capabilities (including Intel® AVX 512 instructions) using the latest OpenMP parallel programming model.

Boost Fortran Application Performance
on Linux* using Intel® Fortran Compiler (Higher is Better)

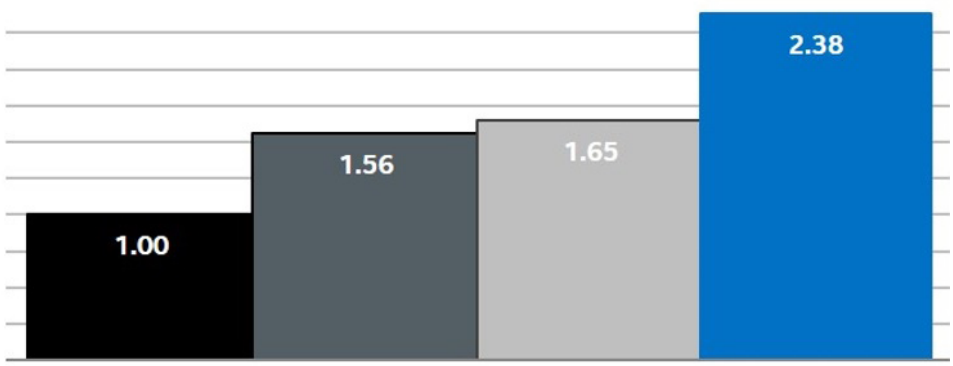

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

Performance results are based on testing as of December 12, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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.

Configuration: Testing by Intel as of December 12, 2019. : Hardware: Intel(R) Core(TM) i7-6700 CPU @ 3.40GHz, HyperThreading is on. RAM 8G. Software: Intel® Fortran Compiler 19.1, PGI Fortran* 19.1, Open64* 4.6.2, gFortran* 9.2.0. Red Hat Enterprise Linux Server release 7.4 (Maipo), Kernel 3.10.0-693.11.6.el7.x86_64 Polyhedron Fortran Benchmark (www.fortran.uk). Linux compiler switches: Auto-parallel: Gfortran:gfortran -Ofast -mfpmath=sse -flto -march=skylake -funroll-loops -ftree-parallelize-loops=8. Intel Fortran compiler: -fast -parallel -xCORE-AVX2 -nostandard-realloc-lhs. PGI Fortran: pgf95 -fast -Mipa=fast,inline -Msmartalloc -Mfprelaxed -Mstack_arrays -Mconcur –mp=bind. Open64: openf95 -march=auto -Ofast -mso -apo. Non-auto parallel (NP): ifort -fast -xCORE-AVX2 -nostandard-realloc-lhs. open64:openf95 -march=auto -Ofast -mso. gcc: gfortran -Ofast -mfpmath=sse -flto -march=native -funroll-loops. pgi: pgf95 -fast -Mipa=fast,inline -Msmartalloc -Mfprelaxed -Mstack_arrays. aocc: flang -flto -Wl,-mllvm -Wl,-function-specialize -Wl,-mllvm -Wl,-region-vectorize -Wl,-mllvm -Wl,-reduce-array-computations=3 -ffast-math -Wl,-mllvm -Wl,-inline-recursion=4 -Wl,-mllvm -Wl,-lsr-in-nested-loop -Wl,-mllvm -Wl,-enable-iv-split -O3 -flto -march=znver2 -funroll-loops -Mrecursive -mllvm -vector-library=LIBMVEC -z muldefs -lamdlibm -lflang -lamdlibm -lm

INTEL® DISTRIBUTION FOR PYTHON*

- Delivers faster Python application performance in an easy, integrated distribution for Windows*, macOS*, and Linux*.
- Accelerates NumPy*/SciPy*/ scikit-learn* packages with native Intel® Performance Libraries such as Intel® Math Kernel Library for multi-threaded performance benefits.

Faster Python* with Intel® Distribution for Python*

Intel's Optimizations Improve Python scikit-learn Efficiency Closer to Native Code Speeds on Intel® Xeon® Processors

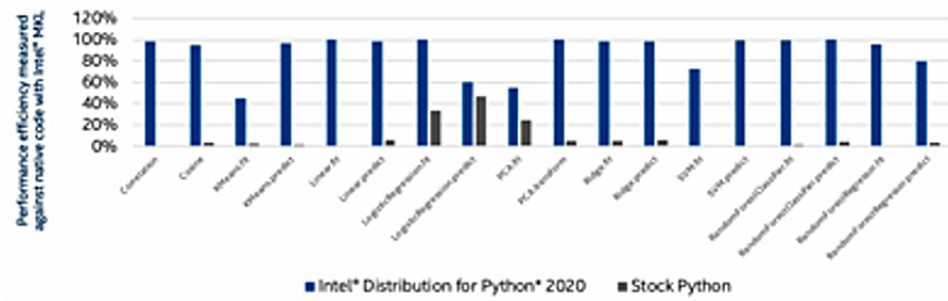

Configuration: Testing by Intel as of November 27, 2019. Stock Python: python 3.7.5 h0371630_0 installed from conda, numpy 1.17.4, numba 0.46.0, llvmlite 0.30.0, scipy 1.3.2, scikit-learn 0.21.3 installed from pip; Intel Python: Intel® Distribution for Python* 2020 Gold: python 3.7.4 hf484d3e_3, numpy 1.17.3 py37ha68da19_4, mkl 2020 intel_133, mkl_fft 1.0.15 py37ha68da19_3, mkl_random 1.1.0 py37ha68da19_0, numba 0.45.1 np117py37_1, llvmlite 0.29.0 py37hf484d3e_9, scipy 1.3.1 py37ha68da19_2, scikit-learn 0.21.3 py37ha68da19_14, daal 2020 intel_133, daal4py 2020 py37ha68da19_4; CentOS Linux 7.3.1611, kernel 3.10.0-514.el7.x86_64; Hardware: Intel(R) Xeon(R) Platinum 8280 CPU @ 2.70GHz (2 sockets, 28 cores/socket, HT: off), 256 GB of DDR4 RAM, 16 DIMMs of 16 GB@2666MHz

INTEL® MATH KERNEL LIBRARY

- Fastest and most-used math library for Intel® and compatible processors.
- Highly tuned for best performance on today's and future Intel® platforms.
- De facto standard APIs for simple code integration.

Intel® Math Kernel Library 2020 on 2nd Generation Intel® Xeon® Scalable Processor

Intel® Math Kernel Library 2020 Gold versus Competitors
DGEMM on 56 Threads

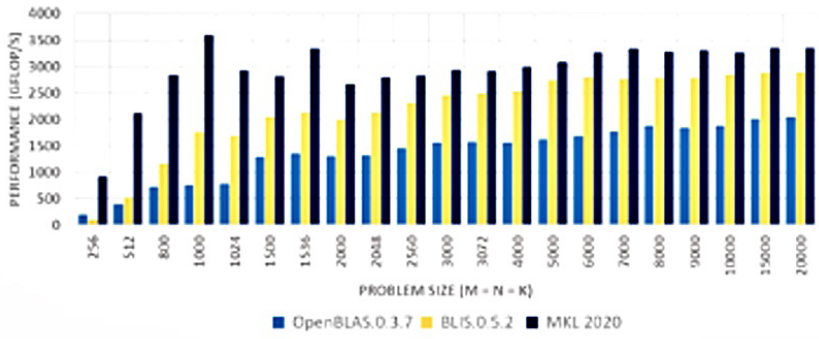

Configuration: Testing by Intel as of 11/7/2019. 2x Intel® Xeon® Platinum 8280L H0 205W 2x28@2.7GHz 192GB DDR4-2666 using Intel® Math Kernel Library 2020. Benchmark Source: Intel® Corporation.

Build with the Composer Edition (Continued)

INTEL® DATA ANALYTICS ACCELERATION LIBRARY (INTEL® DAAL)

- Helps applications deliver better predictions faster and analyzes larger data sets with the same compute resources.
- Optimizes data ingestion and algorithmic compute together for highest performance.
- Supports offline, streaming, and distributed usage models to meet a range of application needs.

Intel® DAAL versus Apache Spark* MILib Performance (Higher is Better)

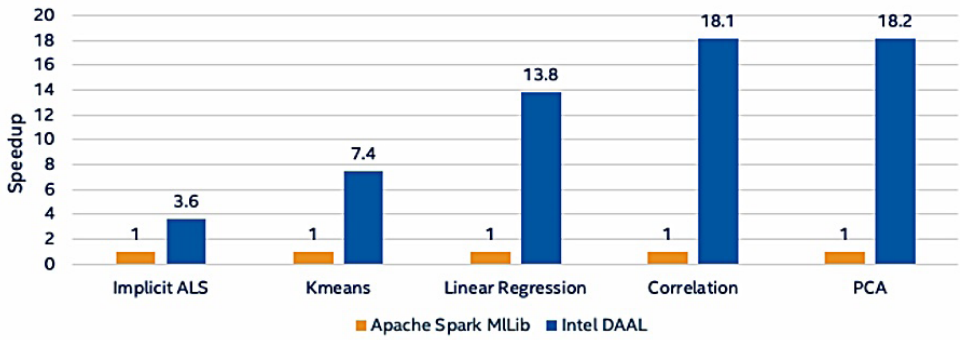

Configuration: Testing by Intel as of 11/11/2019. 7 x m5.2xlarge AWS instances, Intel® Data Analytics Acceleration Library 2020 (Intel® DAAL); Correlation (# samples = 10M, # features = 1000, (Intel® DAAL=35.2s, MLLib=638.2s)), PCA (# samples = 10M, # features = 1000 (Intel® DAAL=35.2s, MLLib=639.8s)), implicit ALS (# users = 1M, # items = 1M, # factors = 100, # Iterations = 1 (Intel® DAAL=37.6s, MLLib=134.9s)), Linear Regression (# samples = 100M, # features = 50 (Intel® DAAL=16.3s, MLLib=224.5s)), k-means (#

INTEL® INTEGRATED PERFORMANCE PRIMITIVES

- Helps applications deliver better predictions faster and analyzes larger data sets with the same compute resources.
- Optimizes data ingestion and algorithmic compute together for highest performance.
- Supports offline, streaming, and distributed usage models to meet a range of application needs.

Intel® IPP Image Resize Functions Performance Boost
Intel IPP Optimization Code versus Compiled C Code

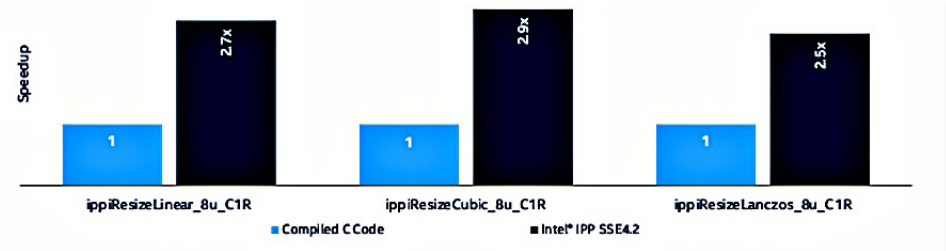

Configuration: Testing by Intel as of 11/08/2019. Intel® Xeon® Platinum 8260L CPU @ 2.30GHz, 64 GB RAM, Intel HyperThreading Technology is supported. Software: Intel® C++ Compiler 18.0.0, GCC 4.8.2. Linux OS: Red Hat Enterprise Linux* 7.4, Kernel 3.10.0-693.11.6.el7.x86_64. Intel® Integrated Performance Primitives v.2020.0.0. Compiler flags: Intel C++ 18.0.0: '-O3 –fstack-protector -ffreestanding –xATOM_SSE4.22'. GCC 4.8.2 '-O2'.

INTEL® THREADING BUILDING BLOCKS

- Specify tasks instead of manipulating threads. Intel® Threading Building Blocks (Intel® TBB) maps your logical tasks onto threads with full support for nested parallelism.
- Intel TBB uses proven, efficient parallel patterns and work-stealing to load balance and cut task execution time.
- Licensed versions are available for Linux, Windows, and macOS. Compatible with multiple compilers and Intel processors.

Excellent Performance Scalability with Intel® Threading Building Blocks on Intel® Xeon® Processors

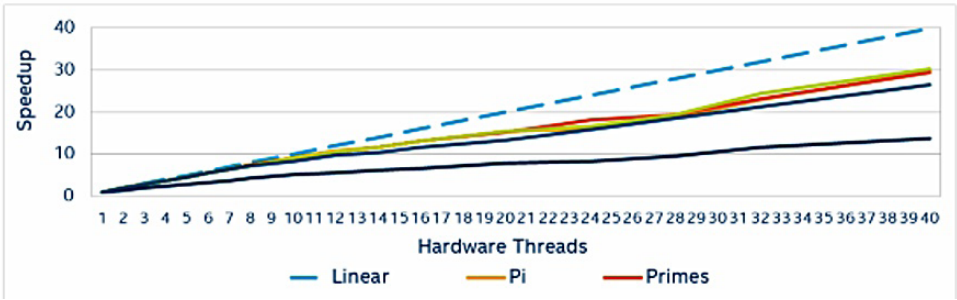

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer. Performance results are based on testing as of December 4, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product can be absolutely secure.

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.

Configuration: Testing by Intel as of December 4, 2019. Hardware: Intel® Xeon® Gold 6152 CPU @ 2.10GHz, 192GB Main Memory. Operating System: CentOS Linux release 7.6, kernel 3.10.0-957.el7.x86_64. Benchmark Source: Intel Corp. Note: sudoku, primes and tachyon are included within Intel TBB.

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

Analyze with the Professional Edition

Includes everything in the Composer Edition, plus:

- **Advanced performance profiler** to tune application performance of the CPU, GPU, threading, memory, cache, and storage
- **Design advisor** to optimize vectorization, quickly prototype threading and create and analyze flow graphs
- **Memory and thread debugger** to efficiently find memory errors and intermittent threading errors

Analyze with the Professional Edition (Continued)

INTEL® VTUNE™ PROFILER

- Accurately profile C, C++, Fortran, Python, Go*, Java*, or a mix of coding languages.
- Provides diverse data to optimize for CPU/GPU, threading, memory, cache, MPI, and storage.
- Delivers fast answers. Rich analysis turns data into insight that saves time optimizing code.

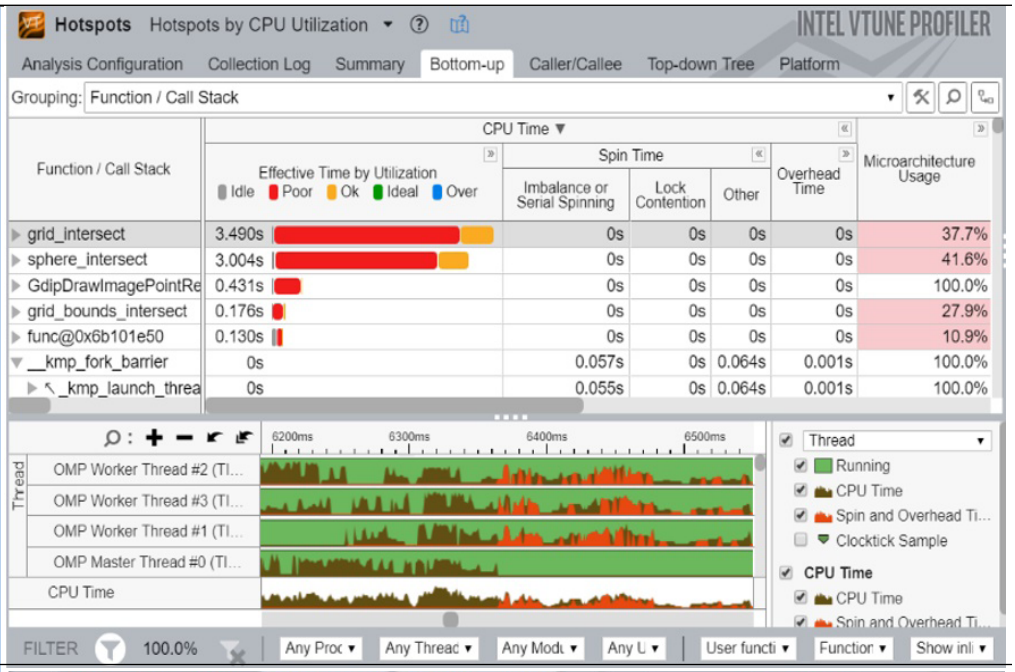

INTEL® ADVISOR

- Add and optimize vectorization
- Find high-impact but under-optimized loops using powerful Roofline analysis
- Analyze memory access patterns
- Add effective threading to unthreaded applications
- Create and analyze efficient flow graphs

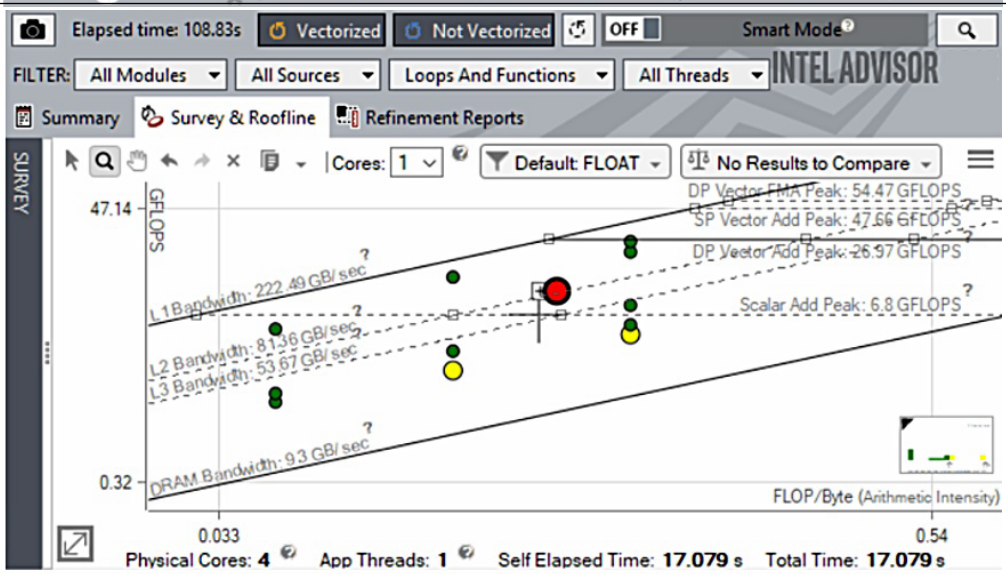

INTEL® INSPECTOR

- Debug threading, memory and persistent memory errors.
- Save money. Locate the root cause of memory and threading errors before you release.
- Save time. Quickly debug intermittent races and deadlocks.
- Save data. Find errors such as missing or redundant cache flushes for persistent memory implementations.
- Save effort. No special compilers or builds are required..

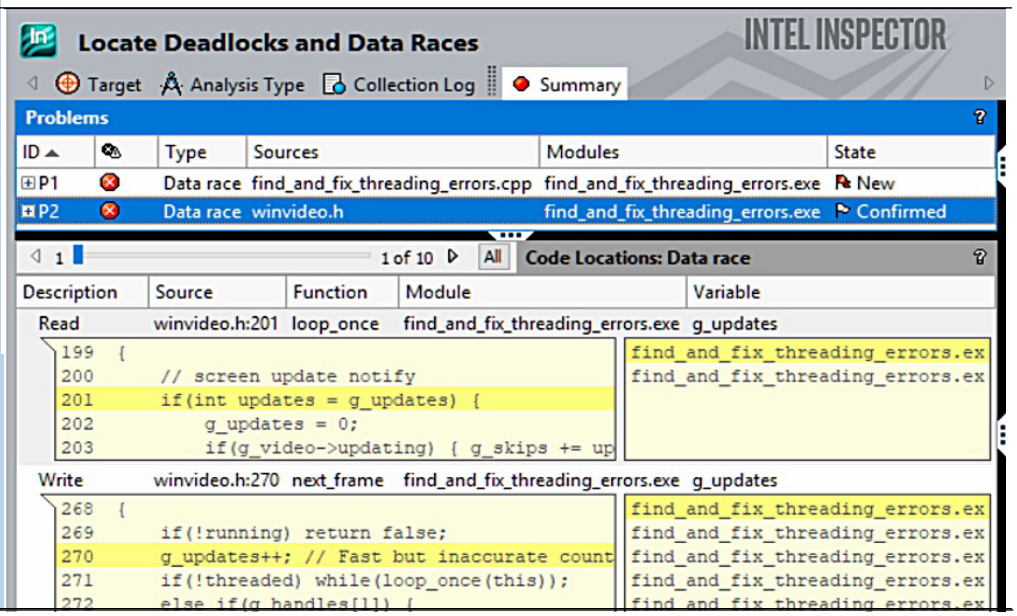

Scale with the Cluster Edition

Includes everything in the Professional Edition, plus tools to:

- **Accelerate** applications' performance on Intel® architecture-based clusters with multiple fabric flexibility.
- **Profile** MPI applications to quickly finding bottlenecks, achieving high performance for parallel cluster applications.
- **Verify** that cluster components continue working together throughout the cluster life cycle

INTEL® MPI LIBRARY

- **Boost** distributed application performance.
- **Enable** your MPI applications to perform better on Intel® architecture-based clusters with multiple-fabric flexibility.
- **Delivers** sustained scalability—low latencies, higher bandwidth, and increased processes.
- **Supports** Intel® multicore and many-core systems.

SPEC MPI Medium Benchmarks with Intel® MPI Library
384 Processes, 8 Nodes (Mellanox Technologies Family), LInux* 64 (Higher is Better)

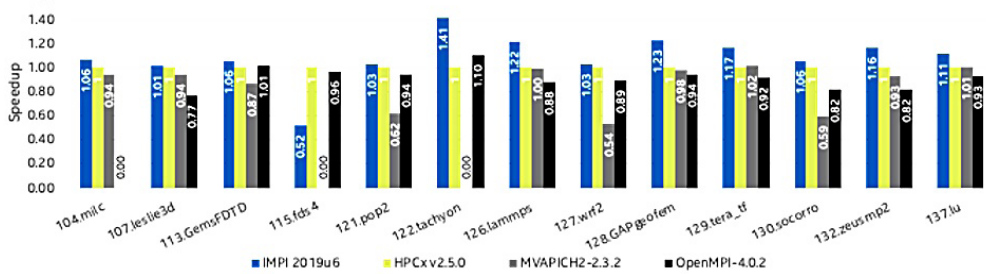

Configuration: Testing by Intel as of November 8, 2019. Intel® Xeon™ Gold 6252 CPU @ 2.10GHz, 196 GB RAM, Intel HyperThreading Technology is supported but not enabled. Mellanox Technologies* MT28908 Family [ConnectX-6]. MLNX_OFED_LINUX-4.6-1.0.1.1(OFED-4.6-1.0.1). Software: Intel® C/C++ Compiler 19.0.0, Intel® Math Kernel Library 2016 Update 1. Linux OS: Red Hat Enterprise Linux 8.0 (Ootpa). Kernel 4.18.0-80.el8.x86_64. Compiler flags: '-O3 -no-prec-div -xCORE-AVX512'. Application: SPEC MPI® 2007

INTEL® TRACE ANALYZER AND COLLECTOR

- **Profile and analyze** MPI applications for performance.
- **Scalable**, with low overhead and effective visualization.
- **Flexible** to fit your workflow: Compile, link, or run.
- **Support** for OpenSHMEM*.

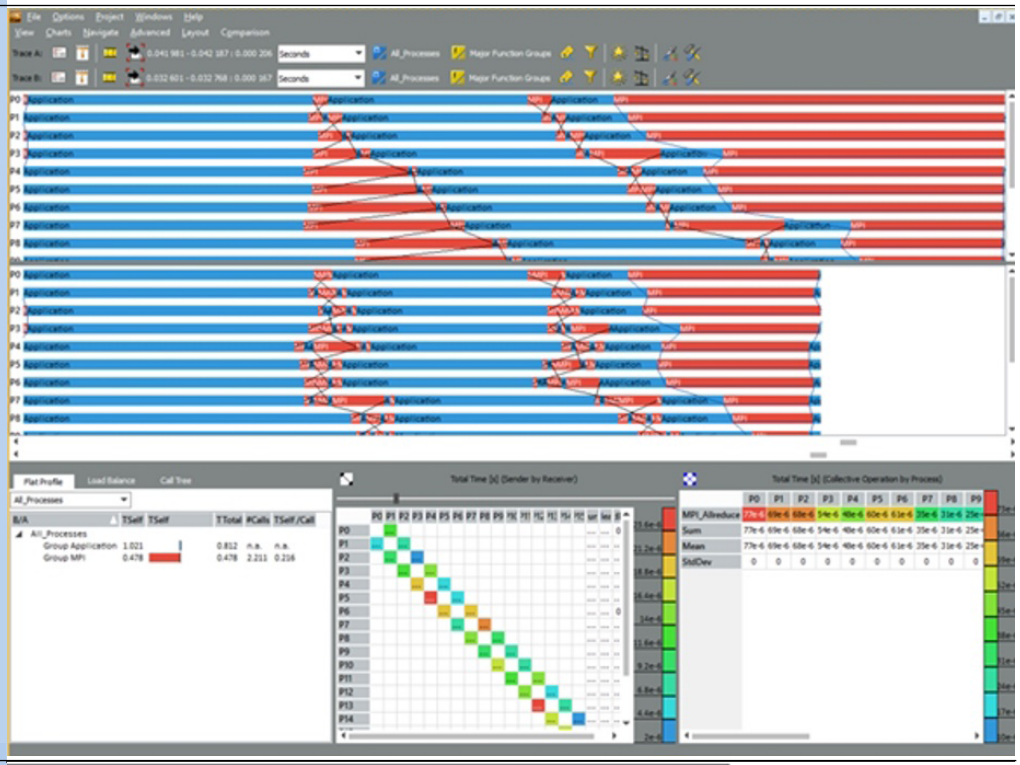

INTEL® CLUSTER CHECKER

- Ensure high-performance, reliable HPC platforms with an advanced cluster diagnostic expert system tool.
- Simpler diagnosis of issues to improve cluster functionality and performance.
- Integrates into other software using an API.
- Comprehensive cluster environment checking, extensible with custom tests.

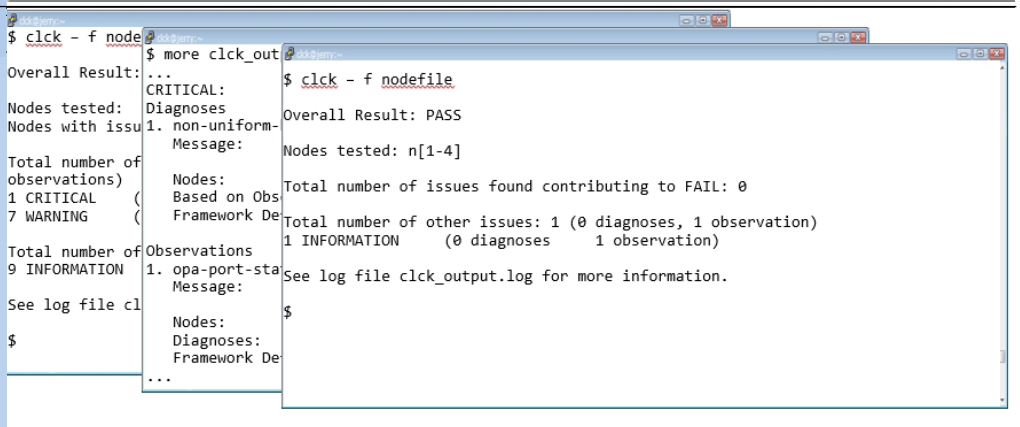

License Options

Each software purchase has a perpetual license with no timeout. Two licensing models are available:

- **Named user licenses** price products per named user.
- **Floating licenses** can be shared by multiple users simultaneously on several systems, managed from a licensing server. Two- or five-seat licenses are available. When a license is released from one user, another user can request it.

Discounted pricing for academia and free versions for students, educators, and open source contributors are available.

Support Services Renewal Options

- **Renewal before subscription expiration.** You can extend your serial number for 12 months after the expiration date. You'll enjoy a price benefit if you renew within a one-month grace period after the expiration date. Support ends at 12 months.
- **Renewal after subscription expiration.** You can extend the same serial number any time within 12 months after the expiration date. Your new subscription will begin from your date of purchase. After the 12-month extension period, you can purchase a new license and get a new serial number.

Specifications at a Glance

Processors	Supports multiple generations of Intel and compatible processors including, but not limited to, Intel® Core™ processors, Intel Xeon Scalable processors
Languages	• Supports processors including, but not limited to, Intel® Core™ and Xeon® Scalable, and Xeon Phi™ processor families. • C, C++, Fortran, Python*[1], C#[2], Go3, and OpenSHMEM*[4]
Operating Systems	• Windows, Linux, and macOS
Development Environment	• Compatible with compilers from Microsoft, GCC, Intel, and others that follow established language standards. • Integrates with Microsoft Visual Studio* (Windows*), Eclipse (Linux*) and XCode* (macOS).
Details	See **software.intel.com/articles/intel-parallel-studio-xe-release-notes**

[1]Intel® Data Analytics Acceleration Library and Intel® VTune™ Amplifier only.
[2]Intel® Advisor and Intel VTune Profiler only.
[3]Intel VTune Profiler only
[4] Intel Trace Analyzer and Collector only.

What's Included

	Feature	Composer Edition[1]	Professional Edition[1]	Cluster Edition
Build	Intel® C++ Compiler	•	•	•
	Intel® Fortran Compiler	•	•	•
	Intel® Distribution for Python[2]	•	•	•
	Intel® Math Kernel Library	•	•	•
	Intel® Data Analytics Acceleration LIbrary	•	•	•
	Intel® Threading Building Blocks	•	•	•
	Intel® Integrated Performance Primitives	•	•	•
Analyze	Intel® VTune™ Profiler[3]		•	•
	Intel® Advisor		•	•
	Intel® Inspector		•	•
Scale	Intel® MPI Library			•
	Intel® Trace Analyzer and Collector			•
	Intel® Cluster Checker			•
	Operating System (Development Environment)	Windows (Visual Studio), Linux (GNU), macOS[5] (XCode*)[6]	Windows (Visual Studio), Linux (GNU)	Windows (Visual Studio), Linux (GNU)

[1]Available with a single language (C++ or Fortran) or both languages.
[2]Available on Windows*, Linux*, and macOS.
[3]Available bundled in a suite or as a standalone.
[4]Available as an add-on to any Windows* Fortran suite or bundled with a version of the Composer Edition.
[5]Available only in the Composer Edition.
[5]Available as a single language suite.

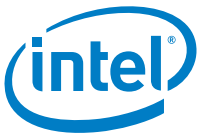

Learn more and get started with Intel Parallel Studio XE >

[1] 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.

[2] Priority support is available only for paid licenses.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation.

Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or learn more at www.intel.com.

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 information go to www.intel.com/performance.

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NONINFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications.

* Other names and brands may be claimed as the property of others. Printed in USA 1219/SS Please Recycle 20110804