

## PRODUCT BRIEF

High-Performance Computing  
Intel® Parallel Studio XE



# Code that Performs

Intel® Parallel Studio XE helps developers take their HPC, enterprise, 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® and Xeon Phi™ processors, plus support for the latest standards and integrated development environments (IDEs). New features include the combined performance snapshot feature of Intel® VTune™ Amplifier that shows MPI, CPU, and FPU memory use; the roofline analysis feature of Intel® Advisor to find high-impact but under-optimized loops; high-performance Python\* to accelerate HPC—and more.

### Who Needs It?

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

### What it Does

- **Creates faster code<sup>1</sup>.** 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

- **Boost application efficiency and performance** for Intel Xeon and Xeon Phi processors using Intel® AVX-512 instructions.
- **Find high-impact but under-optimized loops** using Intel® Advisor's roofline analysis.
- **Accelerate HPC** with high-performance Python.
- **Easily access** the latest Intel® Performance Libraries and Intel® Distribution for Python via APT GET\*, YUM\*, and Conda\*.
- **Stay up-to-date with the latest standards and IDEs** including full C++14, initial C++17 draft, full Fortran 2008, and initial Fortran 2015 draft language support; Initial OpenMP 5.0 draft; Python 2.7 and 3.6; and Microsoft Visual Studio\* 2017 integration.
- **Quickly spot high-payoff opportunities for faster code** using the combined performance snapshot feature of Intel VTune Amplifier for MPI, CPU, FPU, and memory use.
- **New, broader redistribution rights** for Intel Performance Libraries and Intel Distribution for Python.

## Choose Your Edition

COMPOSER EDITION		PROFESSIONAL EDITION	CLUSTER EDITION
<b>BUILD</b> Compilers and Libraries		<b>ANALYZE</b> Analysis Tools	<b>SCALE</b> Cluster Tools
C / C++ Compiler Optimizing Compiler	Intel® MKL FastMath Kernel Library	Intel® VTune™ Amplifier Performance Profiler	Intel® MPI Library Message Passing Interface Library
Fortran Compiler Optimizing Compiler	Intel® IPP Image, Signal, and Data Processing	Intel® Inspector Memory and Thread Debugger	Intel® Trace Analyzer and Collector MPI Tuning and Analysis
Intel® TBB C++ Threading Library	Intel® DAAL Data Analytics Library	Intel® Advisor Vectorization Optimization and Thread Prototyping	Intel® Cluster Checker Cluster Diagnostic Expert System
Intel® Distribution for Python* High-Performance Scripting			
Intel® Architecture Platforms			
Operating System: Windows*, Linux*, MacOS*			

## 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.<sup>2</sup> 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.


## Build with the Composer Edition

- **Improve performance<sup>1</sup>** 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 HPC** with high-performance Python, powered by native Intel® 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.

### Boost Application Performance on Linux\* using Intel® Compilers (higher is better)



Relative geometric performance, SPEC\* benchmark - higher is better

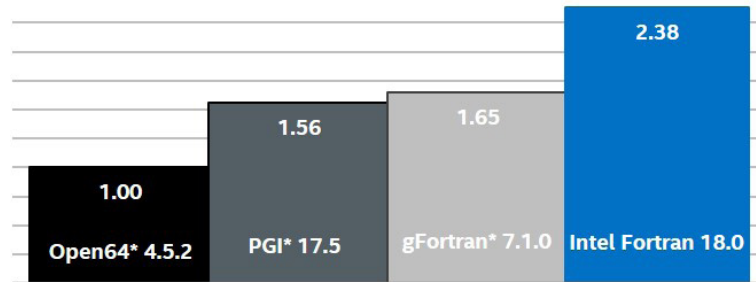
Configuration: Linux\* hardware: 2x Intel® Xeon® Gold 6148 processor @ 2.40GHz, 192 GB RAM, HyperThreading is on. Software: Intel® compilers 18.0, GCC 7.1.0, PGI 15.10, Clang/LLVM 4.0. Linux OS: Red Hat Enterprise Linux Server\* release 7.2 (Maipo), kernel 3.10.0-514.el7.x86\_64. SPEC\* Benchmark (www.spec.org). SmartHeap\* 10 was used for CXX tests when measuring SPECint\* benchmarks. SPECfp\* tests measure C/C++ code performance only. SPECint\_rate\_base\_2006 compiler switches: SmartHeap 10 were used for C++ tests. Intel® C/C++ compiler 18.0: -m32 -xCORE-AVX512 -ipo -O3 -no-prec-div -qopt-prefetch -qopt-mem-layout-trans=3 C++ code adds option -static. GCC 7.1.0: -m32 -Ofast -fno-math-errno -march=core-avx2 -mfpmath=sse -funroll-loops. Clang 4.0: -m32 -Ofast -march=core-avx2 -fno-mfpmath=sse -funroll-loops C++ code adds option -fno-fast-math. SPECfp\_rate\_base\_2006 compiler switches: Intel C/C++ compiler 18.0: -m64 -xCORE-AVX512 -ipo -O3 -no-prec-div -qopt-prefetch -qopt-mem-layout-trans=3 -static. GCC 7.1.0: -m64 -Ofast -fno-math-errno -march=core-avx2 -mfpmath=sse -funroll-loops. Clang 4.0: -m64 -Ofast -march=core-avx2 -fno-mfpmath=sse -funroll-loops. SPECint\_speed\_base\_2006 compiler switches: SmartHeap 10 were used for C++ tests. Intel C/C++ compiler 18.0: -m64 -xCORE-AVX512 -ipo -O3 -no-prec-div -qopt-prefetch -auto-p32. C code adds options -static -parallel. GCC 7.1.0: -m64 -Ofast -fno-math-errno -march=core-avx2 -mfpmath=sse -funroll-loops. C code adds options -fno-fast-math. SPECfp\_speed\_base\_2006 compiler switches: Intel C/C++ compiler 18.0: -m64 -xCORE-AVX512 -ipo -O3 -no-prec-div -qopt-prefetch -static -auto-p32. C code adds option -parallel. Intel Fortran 18.0: -m64 -xCORE-AVX512 -ipo -O3 -no-prec-div -qopt-prefetch -static -parallel. GCC 7.1.0: -m64 -Ofast -fno-math-errno -march=core-avx2 -mfpmath=sse -funroll-loops. C code adds option -fno-parallelize-loops=40. Clang 4.0: -m64 -Ofast -march=core-avx2 -fno-mfpmath=sse -funroll-loops. Benchmark Source: Intel Corporation

## 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)




Relative geomean performance, Polyhedron\* benchmark– higher is better

Configuration: Hardware: 2x Intel® Xeon® Gold 6148 processor @ 2.40GHz, 192 GB RAM, HyperThreading is on. Software: Intel® Fortran compiler 18.0, PGI Fortran\* 17.5, Open64\* 4.5.2, gFortran\* 7.1.0. Linux OS: Red Hat Enterprise Linux Server\* release 7.2 (Maipo), kernel 3.10.0-514.el7.x86\_64 Polyhedron Fortran Benchmark (www.fortran.uk). Linux compiler switches: Gfortran: -Ofast -mpmath=sse -fltto -march=haswell -funroll-loops -ftrue-parallelize-loops=8. Intel Fortran compiler: -fast -parallel -xCORE-AVX2 -nostandard-realloc-lhs. PGI Fortran: -fast -Mipa=fast,inline -Msmatalloc -Mfprelaxed -Mstack\_arrays -Mconcur=bind -tp haswell Open64: -march=bdver1 -mavx -mno-fma4 -Ofast -mso Benchmark Source: Intel Corporation

### 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.

### Intel® Distribution for Python\* Performance Speedups for Select Math Functions on Intel® Xeon™ Processors

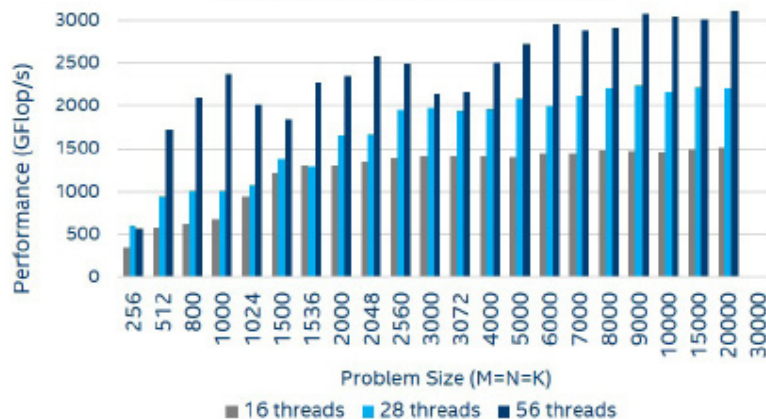


Configuration: Hardware: Intel® Xeon® CPU E5-2699 v4 @ 2.20GHz (2 sockets, 22 cores per socket, 1 thread per core – HT is off), 256GB DDR4 @ 2400MHz. Software: Stock: CentOS Linux release 7.3.1611 (Core), python 3.6.2, pip 9.0.1, numpy 1.13.1, scipy 0.19.1, scikit-learn 0.19.0. Intel® Distribution for Python\* 2018 Gold: mkt 2018.0.0 intel\_4, daal 2018.0.0.20170814, numpy 1.13.1 py36\_intel\_15, openmp 2018.0.0 intel\_7, scipy 0.19.1 np113py36\_intel\_11, scikit-learn 0.18.2 np113py36\_intel\_3. Benchmark Source: Intel Corporation

### 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.

### DGEMM Optimized by Intel® Math Kernel Library 2018 for Intel® Xeon® Platinum Processor (formerly codenamed Skylake Server)




Configuration: Intel® Xeon® Platinum 8180, 2x28 processors, 2.5GHz, 38.5MB L3 cache, 376GB RAM, OS Ubuntu\* 16.04 LTS; Intel® Math Kernel Library (Intel® MKL) 2018. Benchmark source: Intel Corporation.

## Build with the Composer Edition (Continued)

### INTEL® DATA ANALYTICS ACCELERATION LIBRARY (INTEL® DAAL)

- Helps applications deliver better predictions faster; 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.




Configuration: 2x Intel® Xeon® E5-2660 processors @ 2.60GHz, 128 GB, Intel® DAAL 2018; Alternating Least Squares – Users=1M Products=1M Ratings=10M Factors=100 Iterations=1 MLib time=165.9 sec DAAL time=40.5 sec Gain=4.1x; Correlation – N=1M P=2000 size=37 GB MLib time=169.2 sec DAAL=12.9 sec Gain=13.1x; PCA – n=10M p=1000 Partitions=360 Size=75 GB MLib=246.6 sec DAAL (seq)=17.4 sec Gain=14.2x. Benchmark source: Intel Corporation.

### INTEL® INTEGRATED PERFORMANCE PRIMITIVES

- Deliver highly optimized image and signal processing, data compression, and cryptography applications using Intel® Streaming SIMD Extensions and Intel® Advanced Vector Extensions instruction sets.
- Multi-core, multi-OS, and multi-platform ready. Plug in and use APIs to quickly improve application performance.
- Reduces development time and costs.

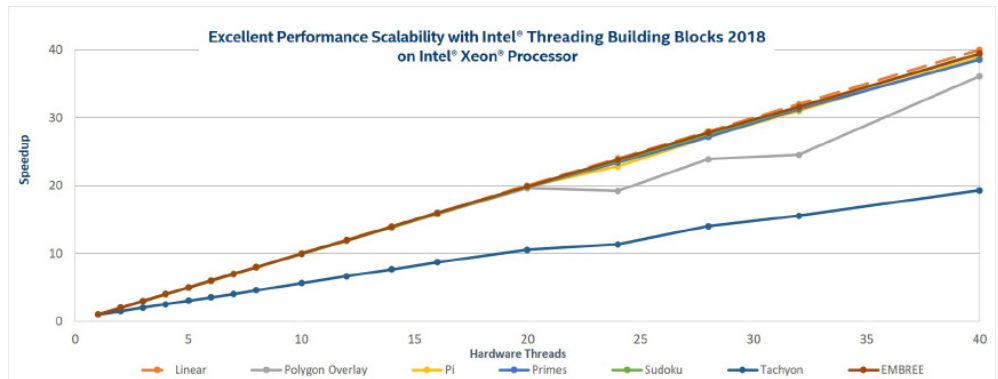
#### Take Advantage of Intel® AVX-512 with Intel® Integrated Performance Primitives (Intel® IPP) 2018



Configuration: Intel® Xeon® Platinum 81xx processor, Intel Xeon Platinum 8168 processor @ 2.70GHz, L3=33 MB, 2x24 cores + HT, Ubuntu\*-64, 109 GB, Intel® Compiler 18, Intel® IPP 2018. Benchmark source: Intel Corporation.

### 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.



Configuration: Software versions: Intel® C++ Intel® 64 Compiler, Version 17.4, Intel® Threading Building Blocks 2018 (Intel® TBB); Hardware: 2x Intel® Xeon® E5-2699 v4 processors @ 2.20GHz 44/T, 128GB Main Memory; Operating System: Red Hat Enterprise Linux Server\* 7.2 (Maipo), kernel 3.10.0-327.4.5.el7.x86\_64; Note: sudoku, primes and tachyon are included with Intel® TBB. Benchmarks Source: Intel Corporation.

## Build with the Composer Edition (Continued)

### ROGUE WAVE\* IMSL FORTRAN NUMERICAL LIBRARY

- Numerical analysis functions for Fortran applications.
- Includes 1,000+ mathematics and statistics algorithms.
- Available as an add-on for any Intel Parallel Studio XE suite.

```

vliberf  x  vliberf  vliberf
G (Global Scope)  -  vliberf0
4  c  Computer:  PC64MS/SINGLE
5  c
6  c  Revised:  June 1, 1990
7  c
8  c  Purpose:  Environment and installation assurance test for
9  c  subroutine OMR.
10 c
11 c  Copyright: 1990 by IMSL, Inc. All Rights Reserved.
12 c
13 c  Warranty:  IMSL warrants only that IMSL testing has been applied
14 c  to this code. No other warranty, expressed or implied,
15 c  is applicable.
16 c
17 c
18 c
19 c  subroutine vliber
20 c  SPECIFICATIONS FOR LOCAL VARIABLES
21 integer  ickk(1), lev, nout
22 real    alpha, ermax
23 complex ans, conans, sd(5,5)
24 character aname(1)*6, uplo*1
25 c
26 integer  ij, incx, lda, llick, n, ntest
27 complex sa(5,5), sx(5)
28 c  SPECIFICATIONS FOR SUBROUTINES
29 external ccopy, cher, eiclk, elpos, rltite, tckke, umach, wcrmn
30 c  SPECIFICATIONS FOR FUNCTIONS
31 external amach, nlrty
32 integer  nlrty
33 real    amach
34 c
35 data  n/5/, lda/5/, incx/1/
36 c
37 sa(1,1) = (3.0,0.0)
38 sa(2,1) = (3.0,0.0)
39 sa(3,1) = (3.0,0.0)
40 sa(4,1) = (3.0,0.0)
41 sa(5,1) = (3.0,0.0)
    
```

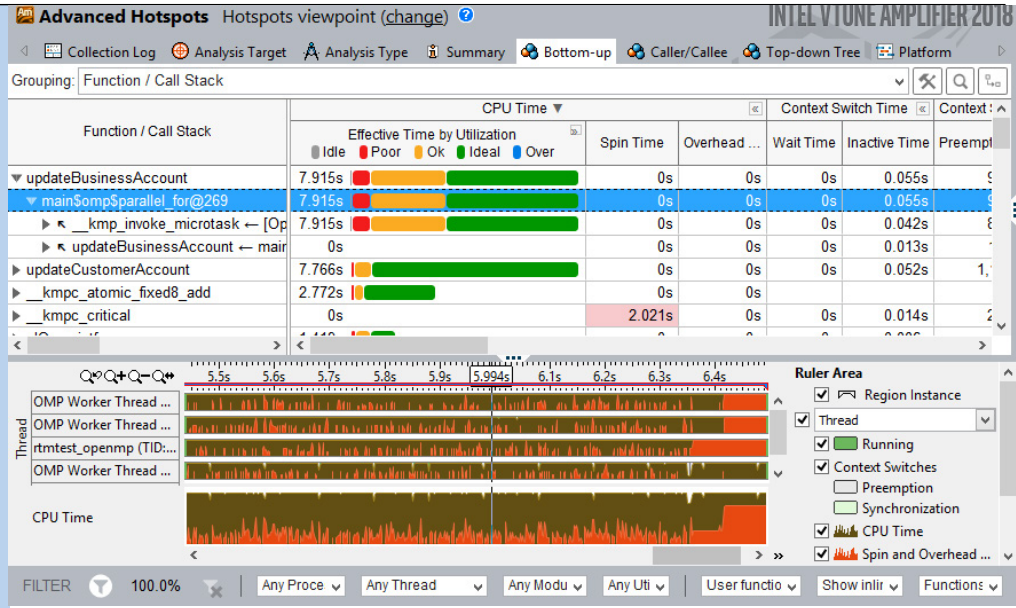
## Analyze with the Professional Edition

Includes everything in the Composer Edition, plus:

- **Advanced performance profiler** to tune application performance of the CPU, threading, memory, and storage
- **Vectorization and threading advisor** to optimize vectorization and quickly prototype threading designs
- **Memory and thread debugger** to efficiently find memory errors and intermittent threading errors

### INTEL® VTUNE AMPLIFIER

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



## Analyze with the Professional Edition (Continued)

### INTEL® ADVISOR

- Vectorize and thread your code—or performance dies on modern processors.
- Get trip counts, data dependencies, memory access patterns, and more.
- Follow an easy optimization workflow with tips for faster code.
- Find high-impact but under-optimized loops using powerful roofline analysis.

Function Call Sites and Loops	Vector Issues	Self Time	Total Time	Type	FLOPS	Why No Vectorization?	Vectorized Loops	Trip Counts
[loop in S252 at loops90.f:1172]	1 Possible...	3.129s	3.129s	Vectorized	0.1911 0.115	1 vectorizat...	AVX2 17%	99: 6: 1: 1
[loop in S2101 at loops90.f:1749]	2 Possible...	2.765s	2.765s	Scalar	0.1421 0.067	vectorizatio...		12
[loop in s442_SompParallel_for...]	1 Ineffecti...	1.492s	1.492s	Vectorized	0.5861 0.165		AVX2 14%	1.09x 8 30: 1: 3
[loop in S126 at loops90.f:447]	2 Assume...	1.079s	1.079s	Scalar	0.3671 0.179	vector depe...		49
[loop in S343 at loops90.f:2300]	1 Assume...	1.029s	1.029s	Scalar		vector depe...		50
[loop in s471_SompParallel_for...]	2 Assume...	1.017s	1.047s	Threaded (O...	0.4471 0.119	vector depe...		250
[loop in S353 at loops90.f:2381]	1 Possible...	0.989s	0.989s	Vectorized (...)	2.0231 0.134		AVX2 27%	2.16x 8 6: 4: 1

### INTEL® INSPECTOR

- Debug memory and threading errors.
- Save money. Find the root cause of memory and threading errors early—before you release.
- Save time. Quickly debug intermittent races and deadlocks.
- Save effort. No special builds—just use your normal build.

ID	Type	Sources	Object Size	State	Modules
P5	Mismatched allocation/deallocati...	gdvideo.cpp		New	find_an ...
P6	Memory leak	find_and_fix_memor...	28672	Confirmed	find_an ...
P7	Memory leak	gdplusplusgraphics.h	507904	New	find_an ...
P9	Invalid memory access	dynamic_link.cpp; fi ...		Fixed	find_an ...
P.	Memory not deallocated	api.cpp; util.cpp; vid ...	10376	New	find_an ...

## 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.

## Superior MPI Performance with Intel® MPI Library 2018 on Linux\* 64

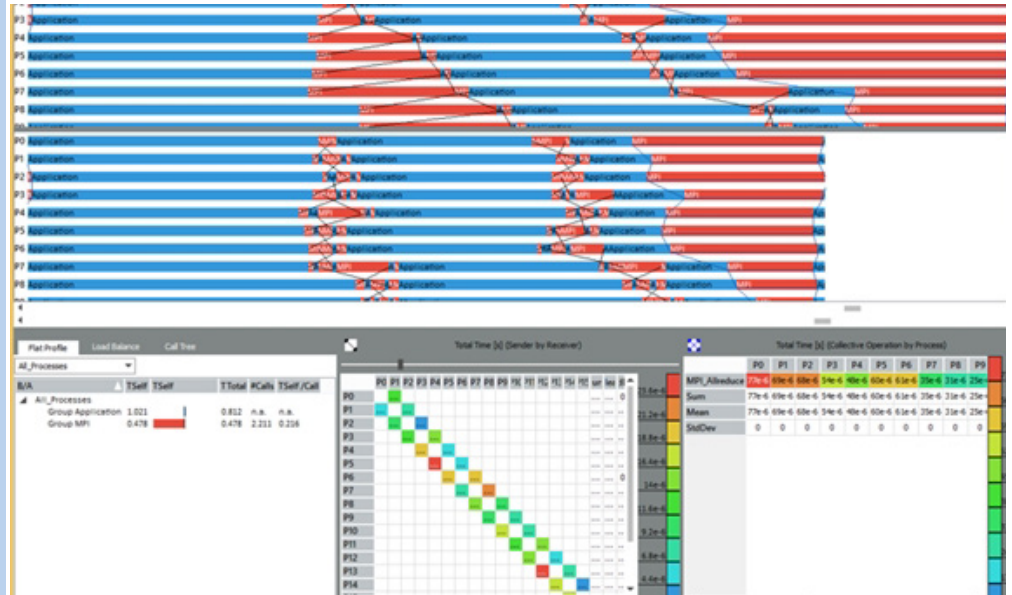
1,280 Processes, 32 Xeon nodes (Omni-Path) Linux\* 64 Relative (Geomean) MPI Latency Benchmarks (Higher is Better)



Configuration: Hardware: Intel® Xeon® Gold 6148 processor @ 2.40GHz; 192 GB RAM. Interconnect: Intel® Corporation Omni-Path HFI Silicon 100 Series [discrete]. Software: RHEL\* 7.3; IFS 10.2.0.0.158; Libfabric 1.3.0; Intel® MPI Library 2018 (I\_MPI\_FABRICS=shm:ofi); Intel® MPI Benchmarks 2018 (built with Intel® C++ Compiler XE 18.0.0 for Linux\*). Benchmark Source: Intel Corporation.

## 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.

```

root@jarry:~# clck@jarry:~
Intel(R) Cluster Checker 2018 (build 20170609)
Database: clck_default, SQLite file /home/clck/2018.0/clck.db.

38 checks requested: all_to_all, cpu, datconf, dgemm, environment, ethernet, heartbeat, hpl, icr_cluster, icr_version, imb_pingpong, infiniband, intel_cluster_runtimes, iozone, java, kernel, kernel_param, libraries, lshw, lustre, memory, miccheck, micinfo, mount, mpi_internode, mpi_local, ntp, offload_phi, opa, perl, process, python, rpm, shells, storage, stream, tcl, x11_tools

Reading the database for the following checks:
all_to_all... done (0.0264 seconds)
cpu... done (0.0176 seconds)
datconf... done (0.0342 seconds)
dgemm... done (0.0144 seconds)
environment... done (0.12 seconds)
ethernet... done (0.0265 seconds)
heartbeat... done (0.0317 seconds)
hpl... done (0.32 seconds)
icr_cluster... done (0.00015 seconds)
icr_version... done (0.00616 seconds)
imb_pingpong... done (0.171 seconds)
infiniband... done (0.0612 seconds)
intel_cluster_runtimes... done (0.00753 seconds)
iozone... done (0.0162 seconds)
java... done (0.00336 seconds)
    
```

## 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

<b>Processors</b>	Supports multiple generations of Intel and compatible processors including, but not limited to, Intel® Core™ processors, Intel Xeon processors, and Intel Xeon Phi processors and coprocessors.
<b>Languages</b>	<ul style="list-style-type: none"> <li>• Supports processors including, but not limited to, Intel® Core™, Xeon®, and Xeon Phi™ processor/co-processor families.</li> <li>• C, C++, Fortran, Python*<sup>1</sup>, C#<sup>2</sup>, Go<sup>3</sup>, and OpenSHMEM<sup>4</sup></li> </ul>
<b>Operating Systems</b>	Windows, Linux, and macOS
<b>Development Environment</b>	<ul style="list-style-type: none"> <li>• Compatible with compilers from Microsoft, GCC, Intel, and others that follow established language standards.</li> <li>• Integrates with Microsoft Visual Studio* (Windows*), Eclipse (Linux*) and XCode* (macOS).</li> </ul>
<b>Details</b>	See <a href="https://software.intel.com/articles/intel-parallel-studio-xe-release-notes">software.intel.com/articles/intel-parallel-studio-xe-release-notes</a>

<sup>1</sup>Intel® Data Analytics Acceleration Library and Intel® VTune™ Amplifier only.

<sup>2</sup>Intel® Advisor and Intel VTune Amplifier only.

<sup>3</sup>Intel VTune Amplifier only

<sup>4</sup>Intel Trace Analyzer and Collector only.

## What's Included

	Feature	Composer Edition <sup>1</sup>	Professional Edition <sup>1</sup>	Cluster Edition
<b>Build</b>	Intel® C++ Compiler	•	•	•
	Intel® Fortran Compiler	•	•	•
	Intel® Distribution for Python <sup>2</sup>	•	•	•
	Intel® Math Kernel Library	•	•	•
	Intel® Data Analytics Acceleration Library	•	•	•
	Intel® Threading Building Blocks	•	•	•
	Intel® Integrated Performance Primitives	•	•	•
<b>Analyze</b>	Intel® VTune™ Amplifier <sup>3</sup>		•	•
	Intel® Advisor		•	•
	Intel® Inspector		•	•
<b>Scale</b>	Intel® MPI Library			•
	Intel® Trace Analyzer and Collector			•
	Intel® Cluster Checker			•
	Rogue Wave* ISML Library <sup>4</sup>	Bundled and Add-On	Add-on	Add-on
	Operating System (Development Environment)	Windows* (Visual Studio*), Linux* (GNU), macOS* (XCode*) <sup>5</sup>	Windows (Visual Studio), Linux (GNU)	Windows (Visual Studio), Linux (GNU)

<sup>1</sup>Available with a single language (C++ or Fortran) or both languages.

<sup>2</sup>Available on Windows\*, Linux\*, and macOS.

<sup>3</sup>Available bundled in a suite or as a standalone.

<sup>4</sup>Available as an add-on to any Windows\* Fortran suite or bundled with a version of the Composer Edition.

<sup>5</sup>Available as a single language suite.





Learn more and get started with Intel Parallel Studio XE >

<sup>1</sup> 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

<sup>2</sup> 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](http://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](http://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 NON-INFRINGEMENT 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.

Copyright © 2017 Intel Corporation. All rights reserved. Intel, Xeon, Xeon Phi, VTune, and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

\* Other names and brands may be claimed as the property of others. Printed in USA 0917/SS Please Recycle 20110804