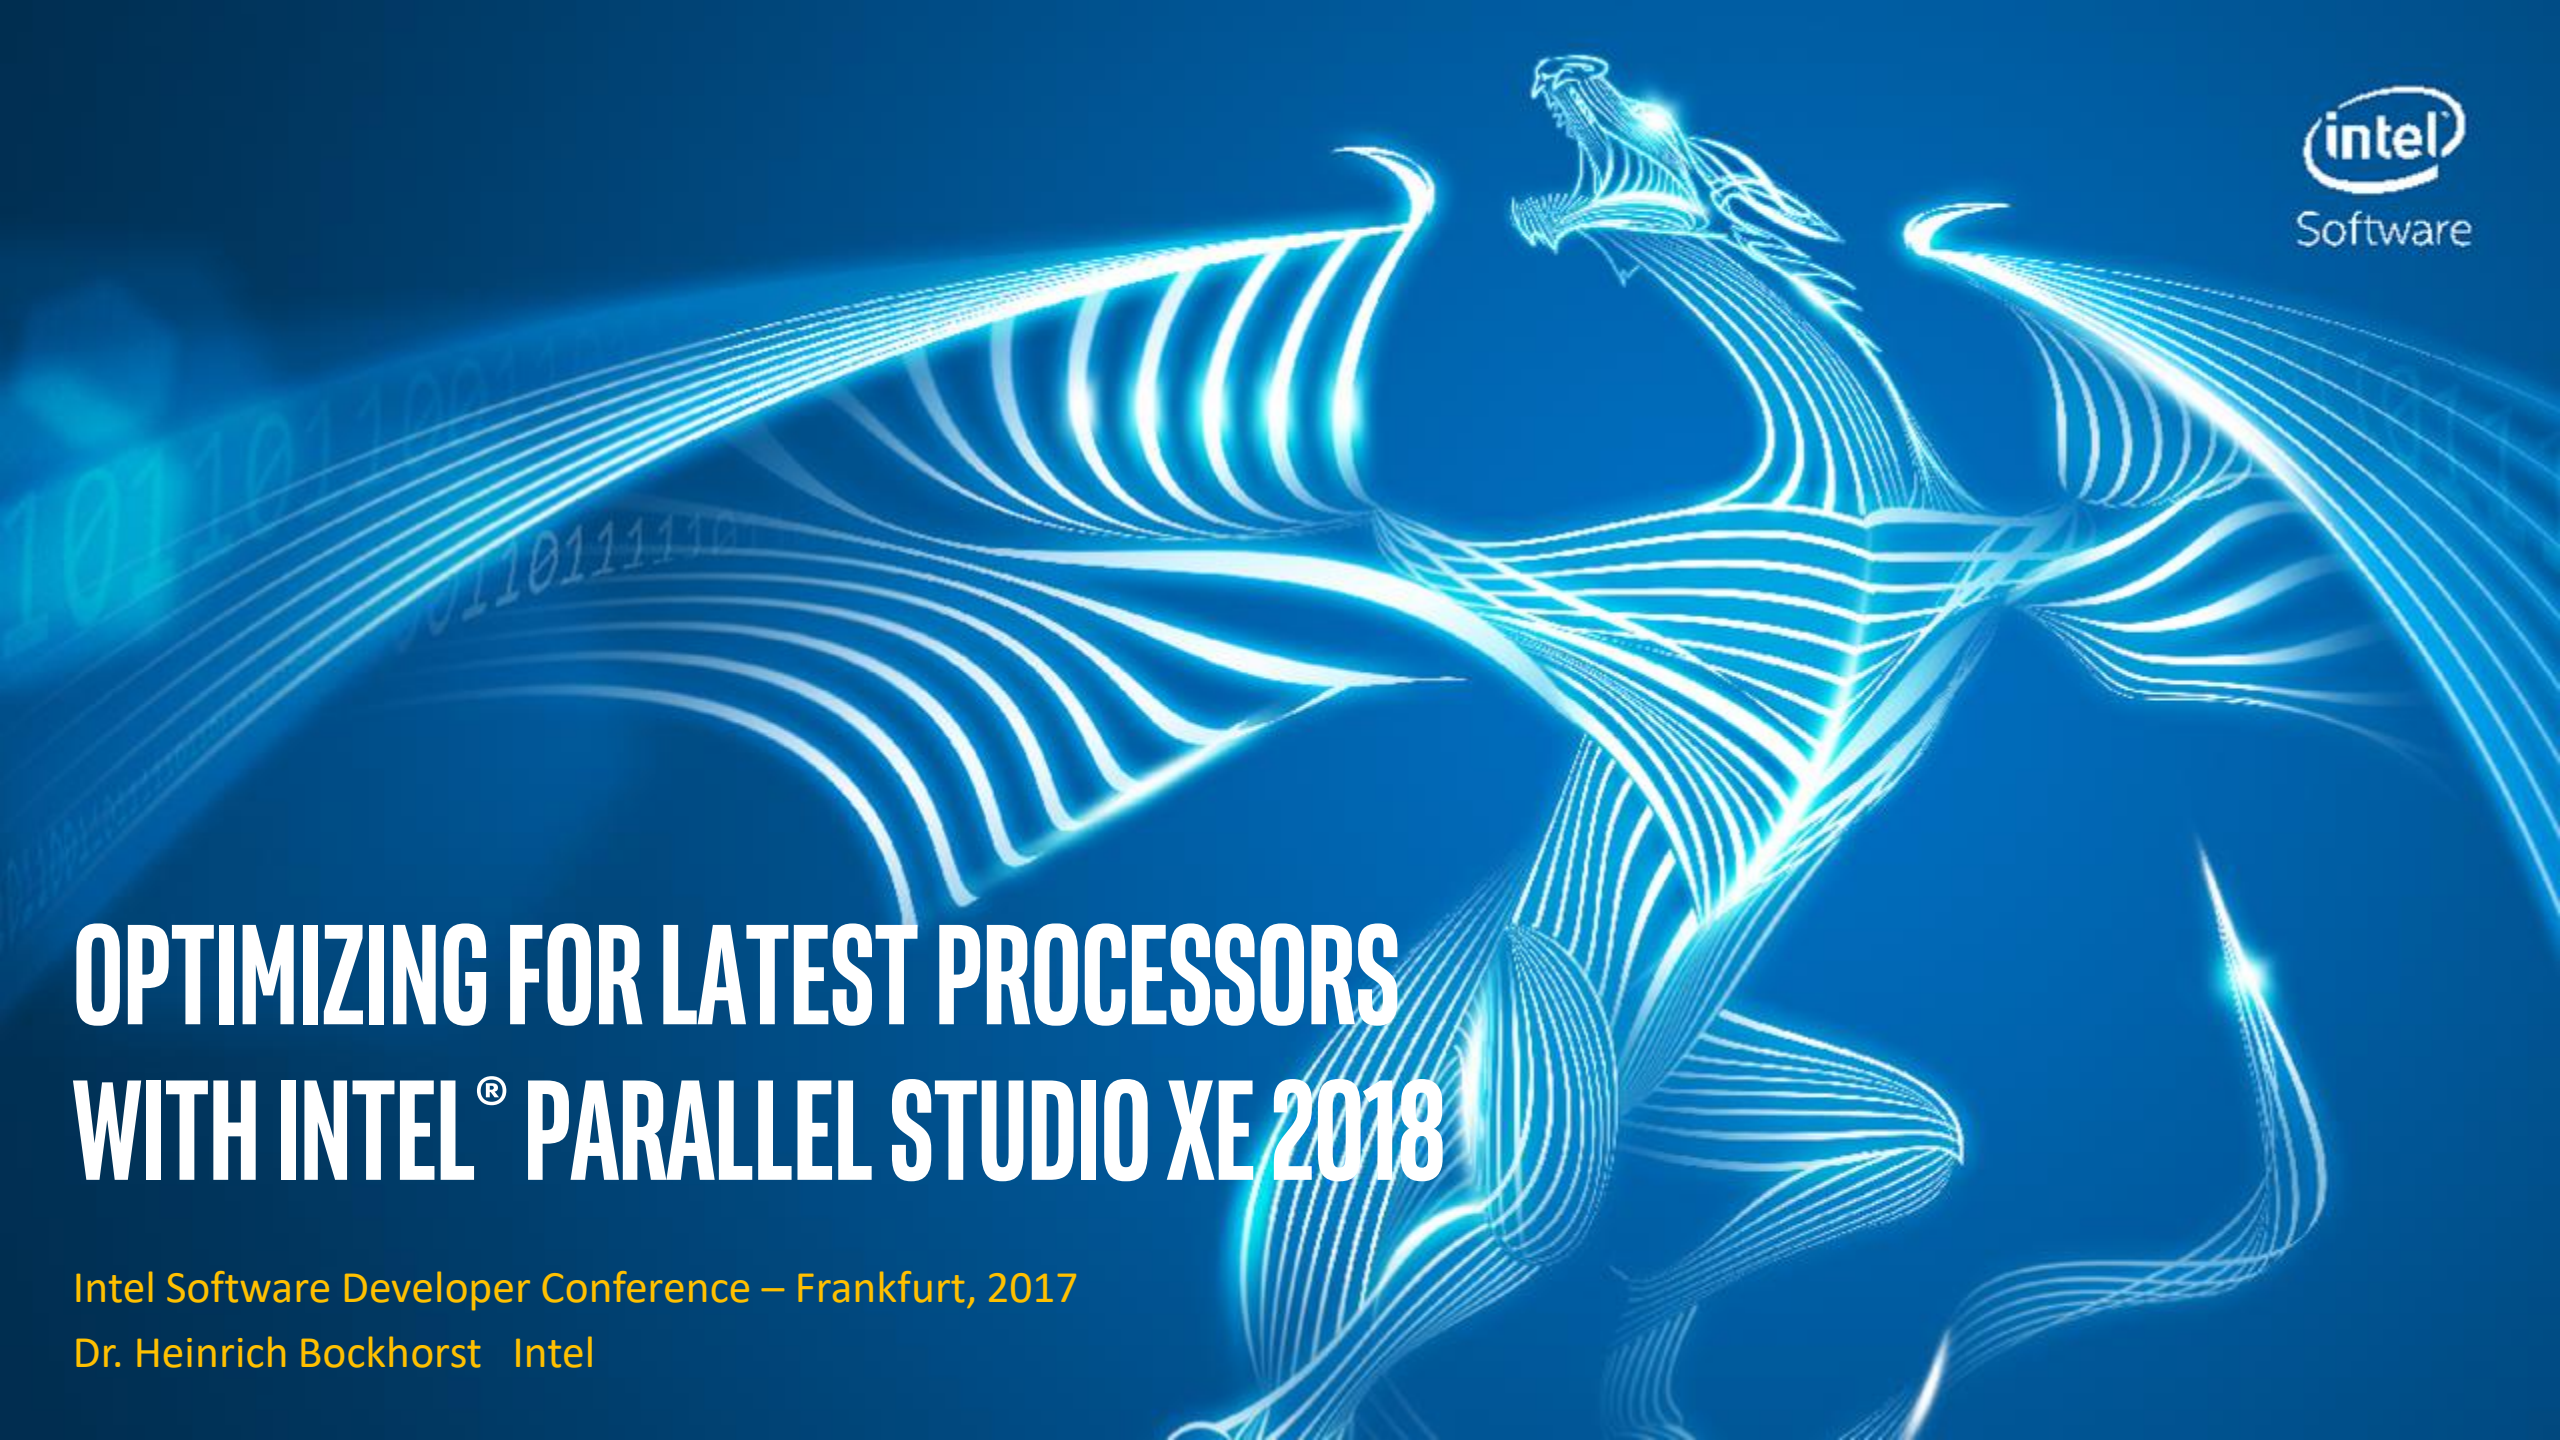

OPTIMIZING FOR LATEST PROCESSORS WITH INTEL® PARALLEL STUDIO XE 2018

Intel Software Developer Conference – Frankfurt, 2017

Dr. Heinrich Bockhorst Intel

AGENDA

- What's Inside Intel® Parallel Studio XE
- Which tool should I use – tuning workflow
- Intel® Parallel studio XE component tools (Build, Analyze, Scale)

What's Inside Intel® Parallel Studio XE

Comprehensive Software Development Tool Suite

COMPOSER EDITION

BUILD
Compilers & Libraries

C / C++ Compiler
Optimizing Compiler

Fortran Compiler
Optimizing Compiler

Intel® Threading Building Blocks
C++ Threading Library

Intel® Math Kernel Library

Intel® Integrated Performance Primitives
Image, Signal & Data Processing

Intel® Data Analytics Acceleration Library

Intel® Distribution for Python*
High Performance Scripting

PROFESSIONAL EDITION

ANALYZE
Analysis Tools

Intel® VTune™ Amplifier
Performance Profiler

Intel® Inspector
Memory & Thread Debugger

Intel® Advisor
Vectorization Optimization & Thread Prototyping

CLUSTER EDITION

SCALE
Cluster Tools

Intel® MPI Library
Message Passing Interface Library

Intel® Trace Analyzer & Collector
MPI Tuning & Analysis

Intel® Cluster Checker
Cluster Diagnostic Expert System

Intel® Architecture Platforms

Operating System: Windows*, Linux*, MacOS[1]*

More Power for Your Code - software.intel.com/intel-parallel-studio-xe

intel Software

Take Advantage of Intel Priority Support

- Paid licenses of Intel® Software Development Tools include Priority Support for one year from your date of purchase, with options to extend support at a highly discounted rate.
- Benefits

- **Direct & private** interaction with Intel engineers. Submit confidential inquiries & code samples via the Online Service Center.
- **Responsive help** with your technical questions & other product needs.
- **Free access** to all new product updates & access to older versions.

Additional Resources

- **Learn from other experts via community product forums**
- **Access to a vast library** of self-help documents that build off decades of experience with creating high performance code.

WHICH TOOL SHOULD I USE?

Optimizing Performance on Parallel Hardware
Intel® Parallel Studio XE

It's an Iterative Process…

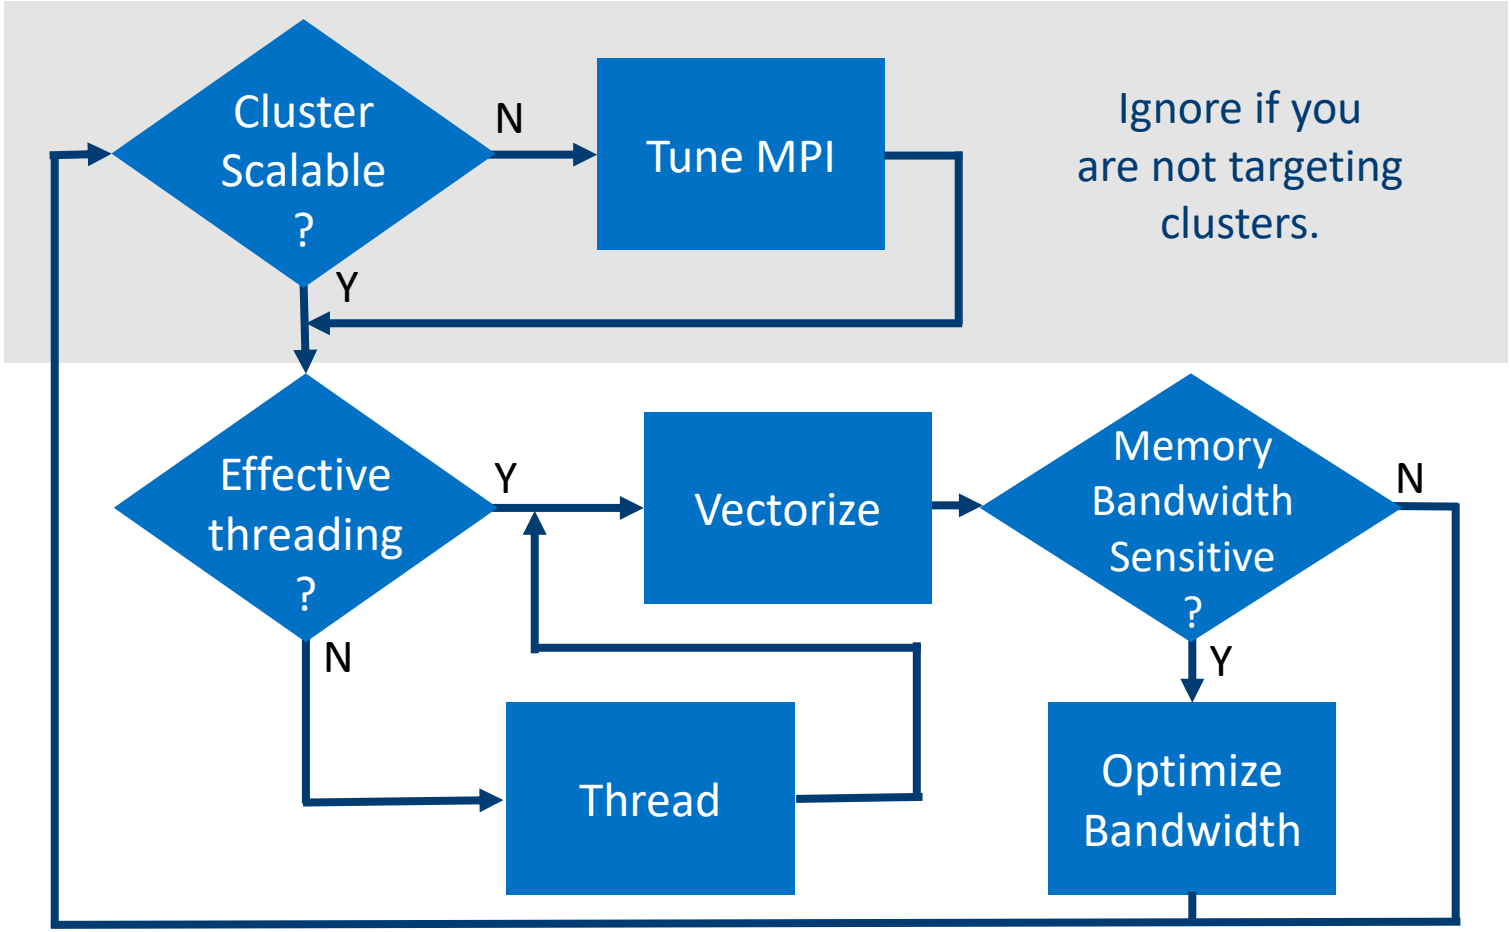

Performance Analysis Tools for Diagnosis
Intel® Parallel Studio XE

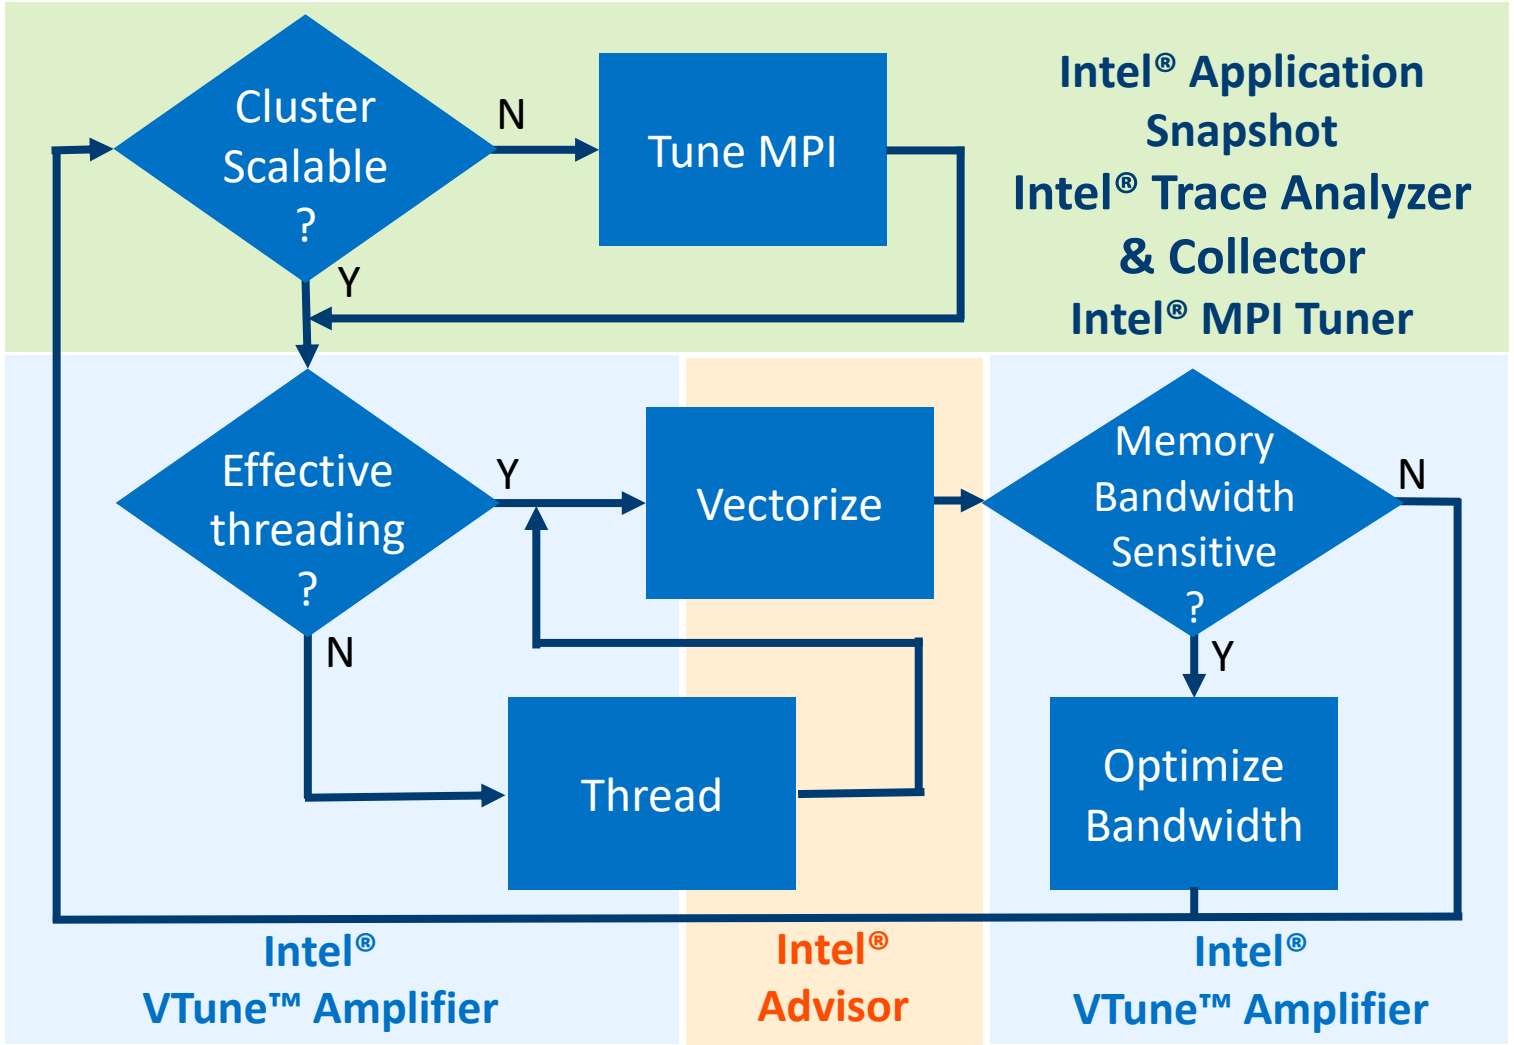

INTEL® PARALLEL STUDIO XE COMPONENT TOOLS

BUILD

Intel® C++ Compiler
Intel® Fortran Compiler
Intel® Distribution for Python*
Intel® Math Kernel Library
Intel® Integrated Performance Primitives
Intel® Threading Building Blocks
Intel® Data Analytics Acceleration Library
Included in Composer Edition

ANALYZE

Intel® VTune™ Amplifier XE
Intel® Advisor
Intel® Inspector

Part of the Professional Edition

SCALE

Intel® MPI Library
Intel® Trace Analyzer & Collector
Intel® Cluster Checker

Part of the Cluster Edition

What's New in Intel® Compilers 2018

Updates to All Versions

- **Advance Support for Intel® Architecture** – Use Intel compiler to generate optimized code for Intel Atom® through Intel® Xeon® Scalable and Xeon Phi™ processor families
- **Achieve Superior Parallel Performance** – Vectorize & thread your code (using OpenMP*) to take full advantage of the latest SIMD-enabled hardware, including AVX-512 instructions
- **Develop Smart Code with Confidence** – Access extensive compiler diagnostics to study code generation characteristics, use with Intel® VTune™ Amplifier & Intel® Advisor for further analysis
- **Faster Compile Time** – Memory management improvements reduce application compile time without sacrificing runtime performance
- **Lightweight Hardware-based Profile-guided Optimization** alternative – Experience many benefits of profile information without the overhead of instrumentation[1]

What's New in C++

Initial C++17, OpenMP* 5; full C++ 14 support

- Standards-driven parallelization for C++ developers

What's New in Fortran

Full Fortran 2008 support

- Submodules, BLOCK, superior coarray performance

Initial Fortran 2015 support (draft standard)

- Further C interoperability (ISO/IEC TS 29113:2012)

Full OpenMP* 4.5 support; initial OpenMP 5

- Thread & vectorize your code using standard APIs

[1]Requires Intel® VTune™ Amplifier

Python* Landscape

Intel® Distribution for Python*

- **Challenge#1**
- Domain experts are not professional software programmers

Adoption of Python continues to grow among domain experts & developers for its productivit benefits

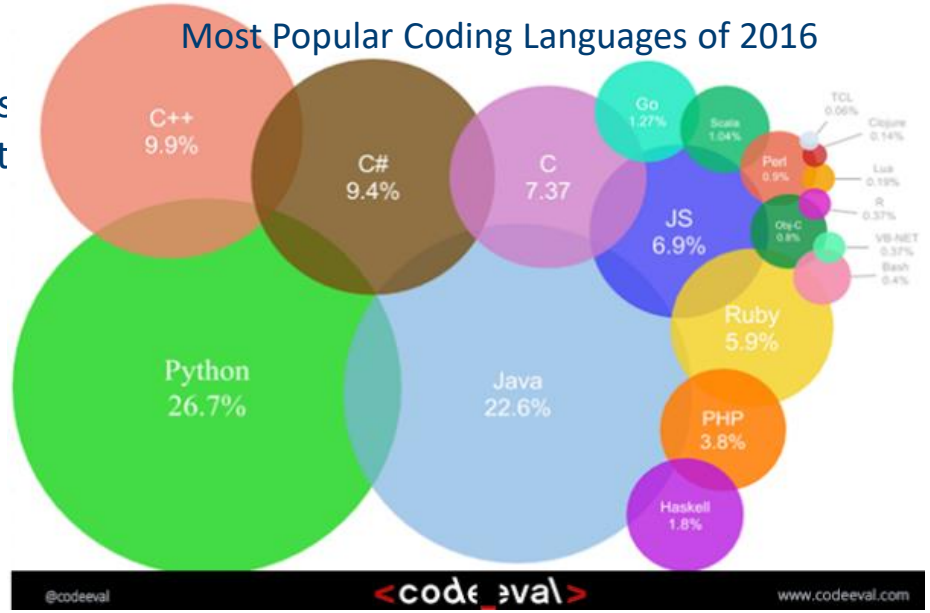

- **Challenge#2**
- Python performance limits migration to production systems

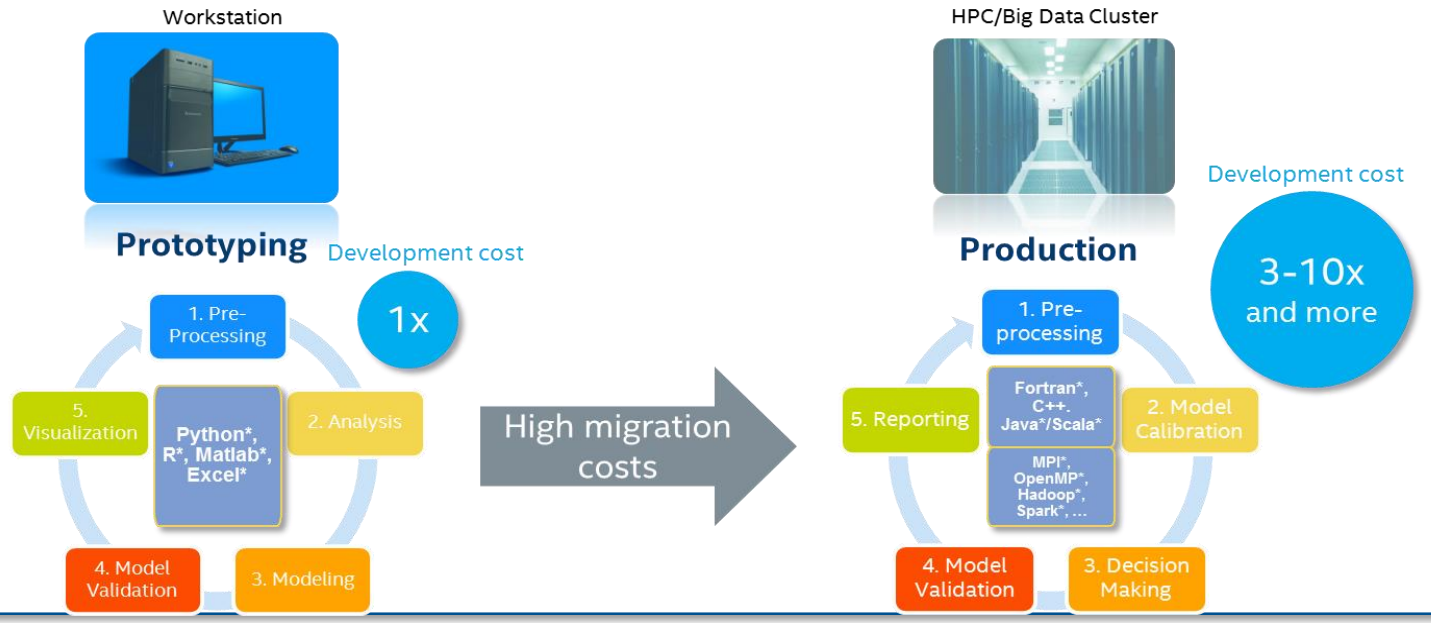

intel Software

Python* Landscape

Intel® Distribution for Python*

Adoption of Python continues to grow among domain experts & developers for its productivit benefits

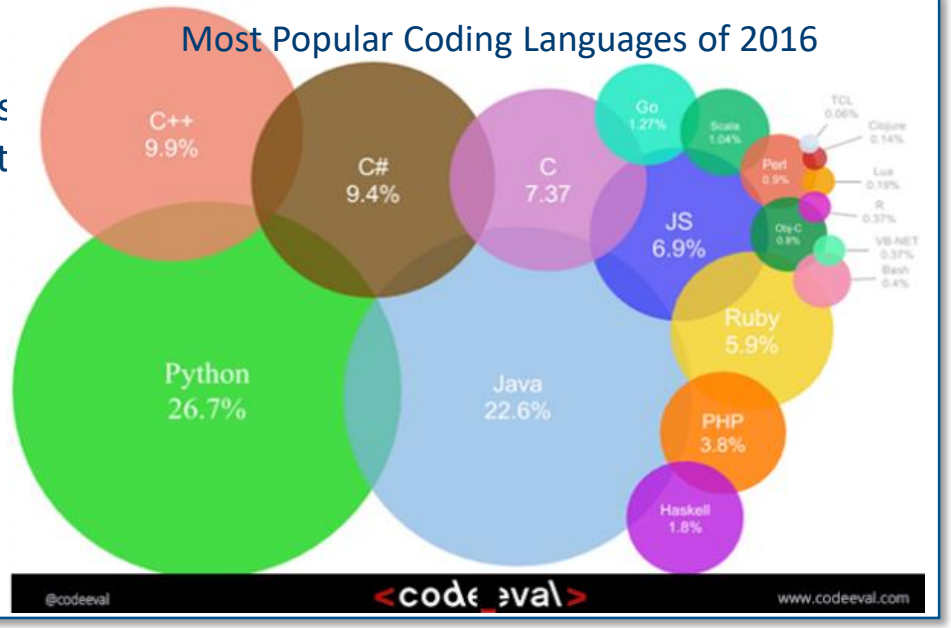

- **Challenge#1**
- Domain experts are not professional software programmers

Intel's Python Tools

- Accelerate Python performance
- Enable easy access
- Empower the community

- **Challenge#2**
- Python performance limits migration to production systems

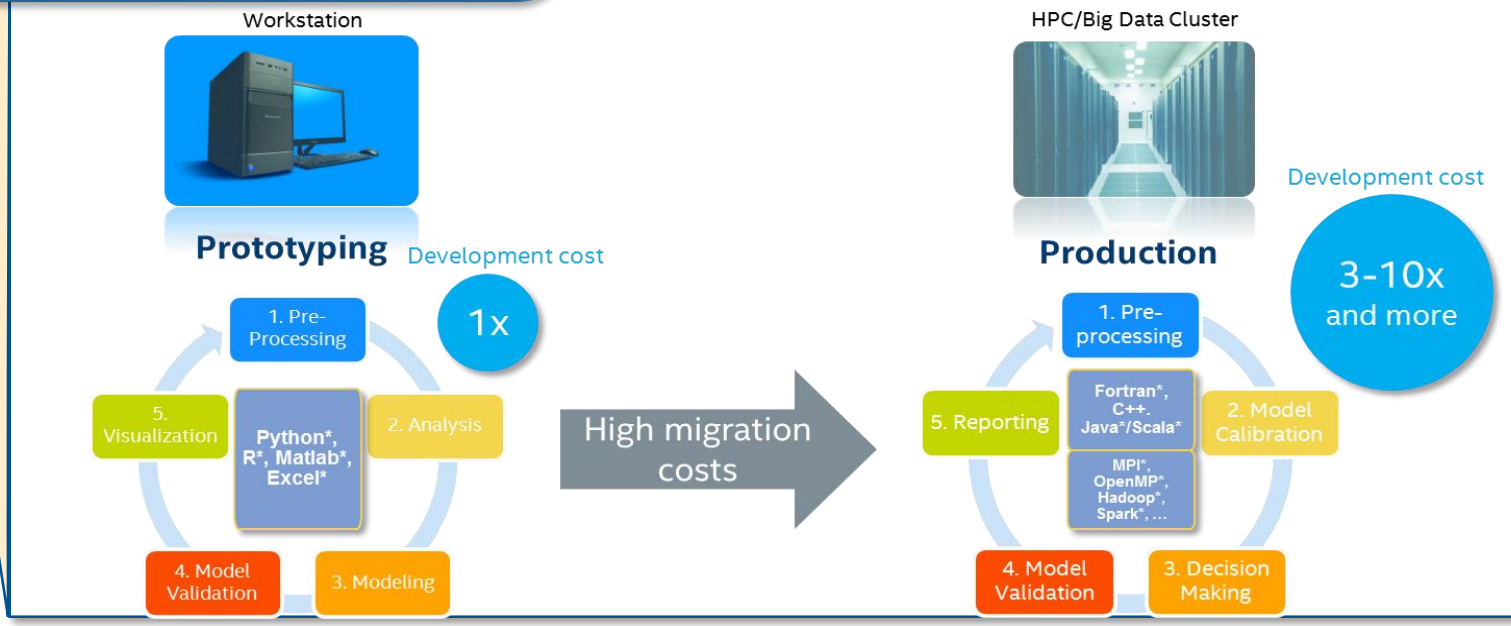

intel Software

What's Inside Intel® Distribution for Python

High Performance Python* for Scientific Computing, Data Analytics, Machine & Deep Learning

FASTER PERFORMANCE	GREATER PRODUCTIVITY	ECOSYSTEM COMPATIBILITY
Performance Libraries, Parallelism, Multithreading, Language Extensions	**Prebuilt & Accelerated Packages**	**Supports Conda & PIP**
Accelerated NumPy/SciPy/scikit-learn with Intel® MKL[1] & Intel® DAAL[2] Data analytics, machine learning & deep learning with scikit-learn, pyDAAL, Caffe*, Theano* Scale with Numba* & Cython* Includes optimized mpi4py, works with Dask* & PySpark* Optimized for latest Intel® architecture	Prebuilt & optimized packages for numerical computing, machine/deep learning, HPC, & data analytics Drop in replacement for existing Python - No code changes required Jupyter* notebooks, Matplotlib included Free download & free for all uses including commercial deployment	Compatible & powered by Anaconda*, supports conda & pip Distribution & individual optimized packages also available at conda & Anaconda.org, YUM/APT, Docker image on DockerHub Optimizations upstreamed to main Python trunk Priority Support through Intel® Parallel Studio XE
Intel® Architecture Platforms		
Operating System: Windows*, Linux*, MacOS[1]*		

[1]Intel® Math Kernel Library
[2]Intel® Data Analytics Acceleration Library

[1]Available only in Intel® Parallel Studio Composer Edition.

Faster Python* with Intel® Distribution for Python*

- Advance Performance Closer to Native Code
  - Accelerated NumPy, SciPy, scikit-learn for scientific computing, machine learning & data analytics
  - Drop-in replacement for existing Python - no code changes required
  - Highly optimized for the latest Intel processors
- What's New in the 2018 edition
  - Updated to support Python 3.6
  - Optimized scikit-learn for machine learning speedups
  - Conda build recipes for custom infrastructure

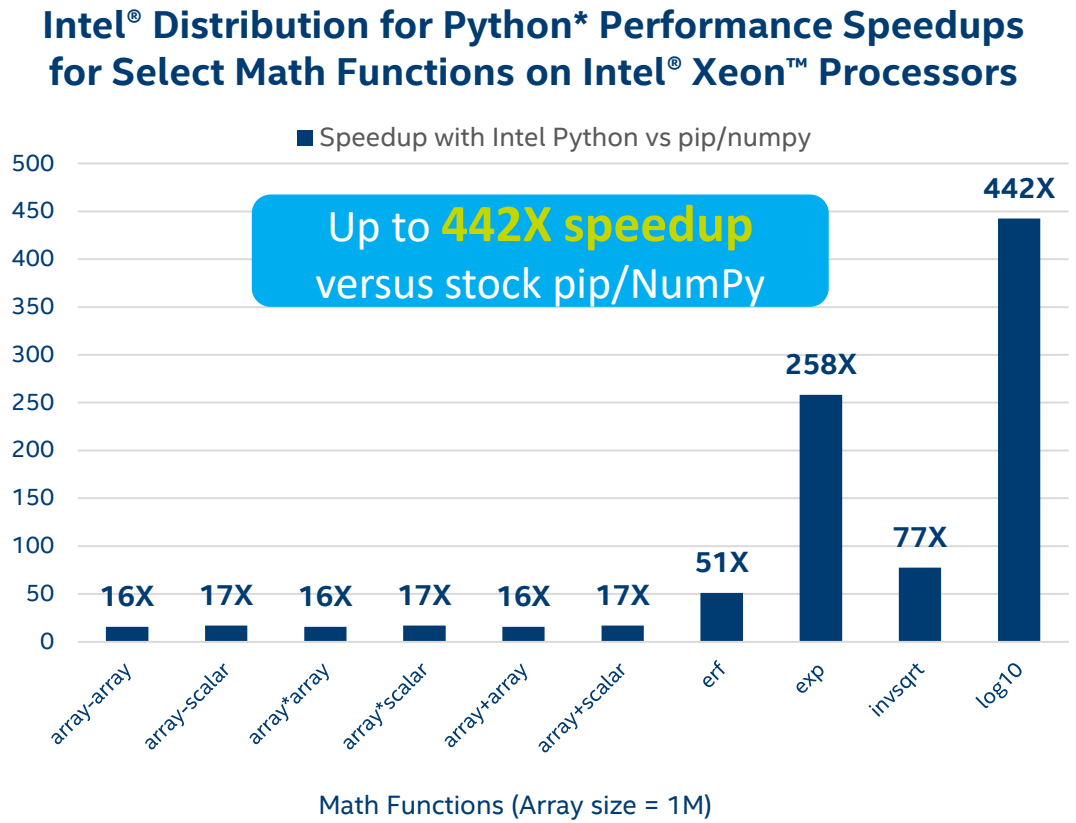

Learn More: software.intel.com/distribution-for-python

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: mkl 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

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 http://www.intel.com/performance. Benchmark Source: Intel Corporation.

intel Software

Fast, Scalable Code with Intel® Math Kernel Library (Intel® MKL)

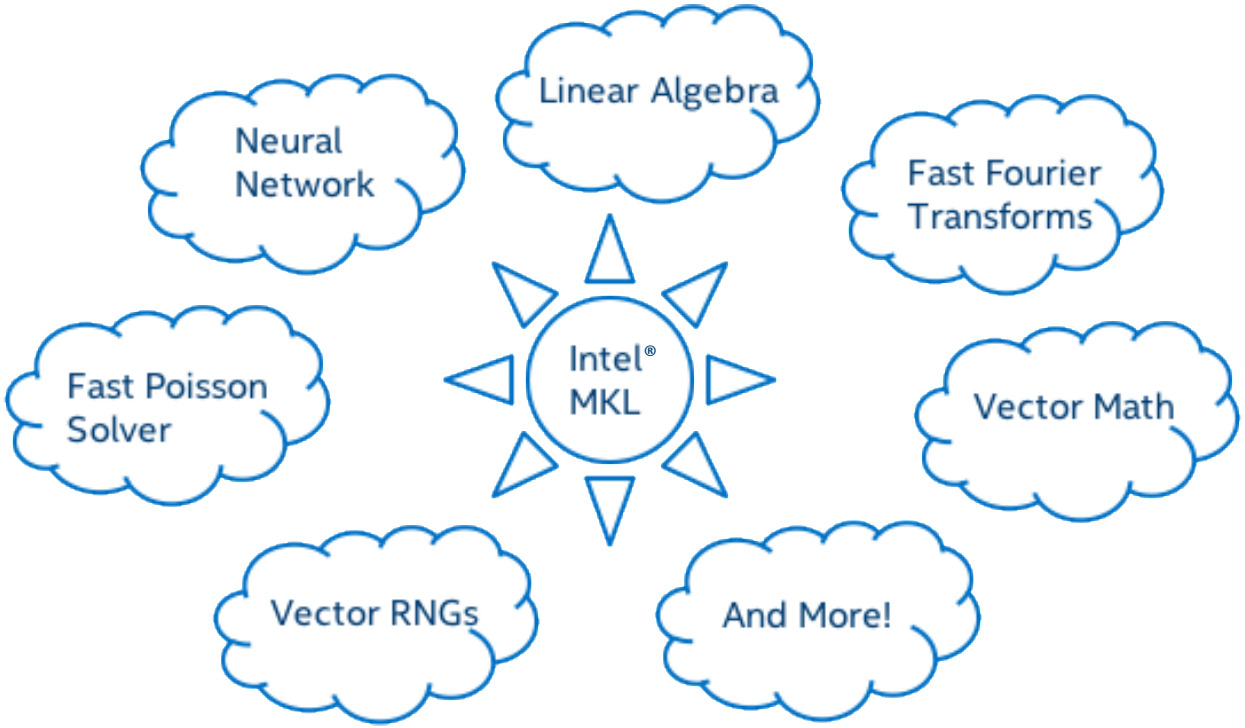

- Highly optimized, threaded, & vectorized math functions that maximize performance on each processor family
- Utilizes industry-standard C and Fortran APIs for compatibility with popular BLAS, LAPACK, and FFTW functions—no code changes required
- Dispatches optimized code for each processor automatically without the need to branch code

- What's New in the 2018 edition
- Improved small matrix multiplication performance in GEMM & LAPACK
- Improved ScaLAPACK performance for distributed computation
- 24 new vector math functions
- Simplified license for easier adoption & redistribution
- Additional distributions via YUM, APT-GET, & Conda

Learn More: software.intel.com/mkl

intel Software

What's Inside Intel® Math Kernel Library

Accelerate HPC, Enterprise, Cloud & **IoT** Applications

Linear Algebra

- BLAS
- LAPACK
- ScaLAPACK
- Sparse BLAS
- Iterative sparse solvers
- PARDISO*
- Cluster Sparse Solver

FFTs

- Multidimensional
- FFTW interfaces
- Cluster FFT

Neural Networks

- Convolution
- Pooling
- Normalization
- ReLU
- Inner Product

Vector RNGs

- Congruential
- Wichmann-Hill
- Mersenne Twister
- Sobol
- Neiderreiter
- Non-deterministic

Summary Statistics

- Kurtosis
- Variation coefficient
- Order statistics
- Min/max
- Variance-covariance

Vector Math

- Trigonometric
- Hyperbolic
- Exponential
- Log
- Power
- Root

& More

- Splines
- Interpolation
- Trust Region
- Fast Poisson Solver

Intel® Architecture Platforms

Operating System: Windows*, Linux*, macOS[1]*

[1]Available only in Intel® Parallel Studio Composer Edition.

intel Software

Get the Benefits of Advanced Threading with Intel® Threading Building Blocks

Use Threading Techniques to fully Leverage Multicore Performance & Heterogeneous Computing

- Parallelize computationally intensive work across CPUs, GPUs & FPGAs,—deliver higher-level & simpler solutions using C++
- Most feature-rich & comprehensive solution for parallel application development
- Highly portable, composable, affordable, & approachable—future-proof scalability

What's New in 2018 edition

- New capabilities in Flow Graph improve concurrency and heterogeneity
- Improves insight into parallelism inefficiencies for Intel® VTune Amplifier 2018
- Support for Cmake file

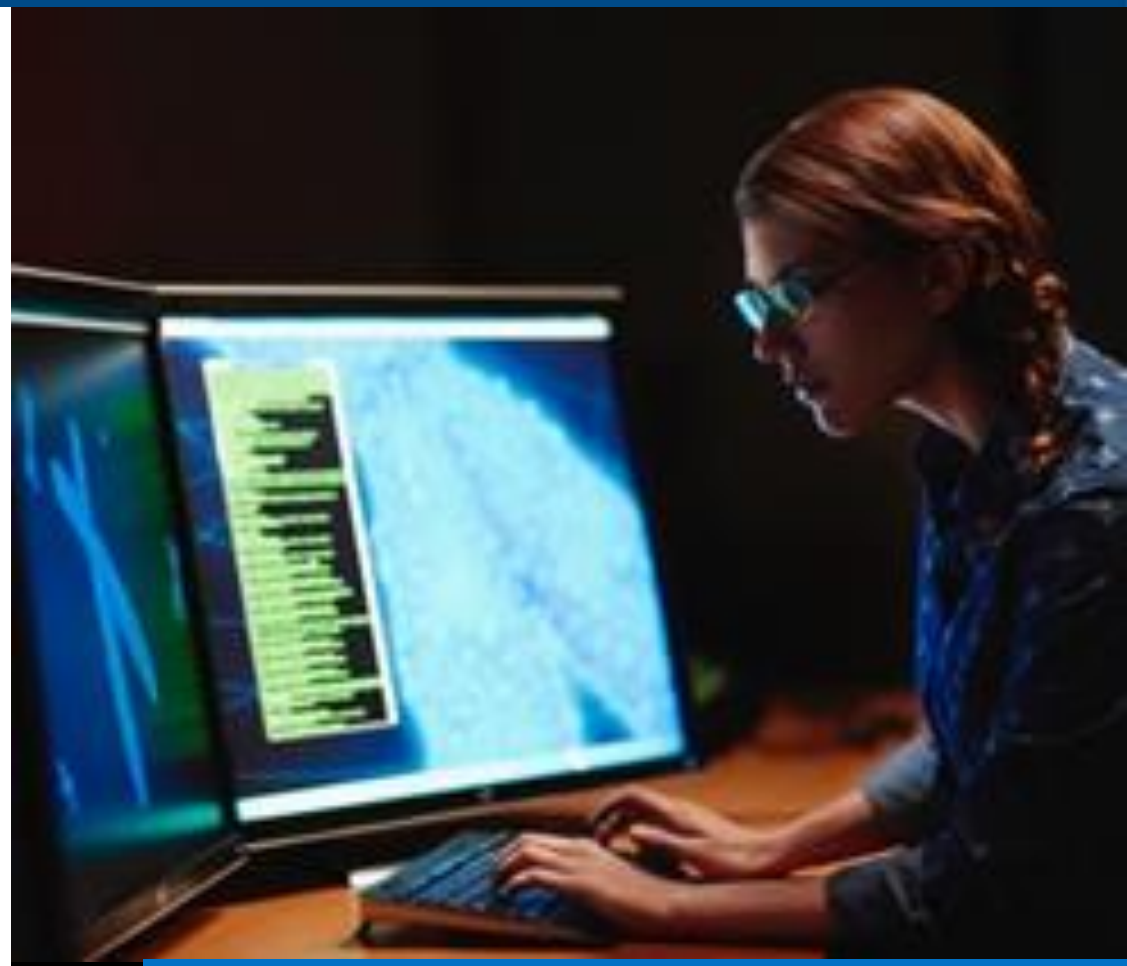

Learn More: software.intel.com/intel-tbb

Speedup Analytics & Machine Learning with Intel® Data Analytics Acceleration Library (Intel® DAAL)

- Highly tuned functions for classical machine learning and analytics performance across a spectrum of Intel® architecture devices
- Optimizes data ingestion together with algorithmic computation for highest analytics throughput
- Includes Python*, C++, Java* APIs, and connectors to popular data sources including Spark* and Hadoop*

Learn More: software.intel.com/daal

What's New in the 2018 Edition

- New Algorithms
  - Classification & Regression Decision Tree and Forest
  - k-NN
  - Ridge Regression
- Spark* MLlib-compatible API wrappers for easy substitution of faster Intel® DAAL functions
- Improved APIs for ease of use
- Repository distribution via YUM, APT-GET, and Conda

Pre-processing	Transformation	Analysis	Modeling	Validation	Decision Making
	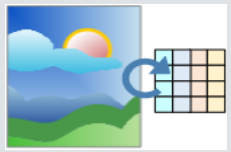				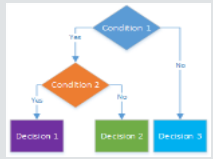
Decompression, Filtering, Normalization	Aggregation, Dimension Reduction	Summary Statistics Clustering, etc.	Machine Learning (Training) Parameter Estimation Simulation	Hypothesis Testing Model Errors	Forecasting Decision Trees, etc.

INTEL® PARALLEL STUDIO XE COMPONENT TOOLS

BUILD

Intel® C++ Compiler
Intel® Fortran Compiler
Intel® Distribution for Python*
Intel® Math Kernel Library
Intel® Integrated Performance Primitives
Intel® Threading Building Blocks
Intel® Data Analytics Acceleration Library
Included in Composer Edition

ANALYZE

Intel® VTune™ Amplifier XE
Intel® Advisor
Intel® Inspector

Part of the Professional Edition

SCALE

Intel® MPI Library
Intel® Trace Analyzer & Collector
Intel® Cluster Checker

Part of the Cluster Edition

Application Performance Snapshot Adds MPI

Data in One Place: MPI+OpenMP+Memory Floating Point—Intel® VTune™ Amplifier

Quick & Easy Performance Overview

- Does the app need performance tuning?

MPI & non-MPI Apps†

- Distributed MPI with or without threading
- Shared memory applications

Popular MPI Implementations Supported

- Intel® MPI Library
- MPICH & Cray MPI

Richer Metrics on Computation Efficiency

- CPU (processor stalls, memory access)
- FPU (vectorization metrics)

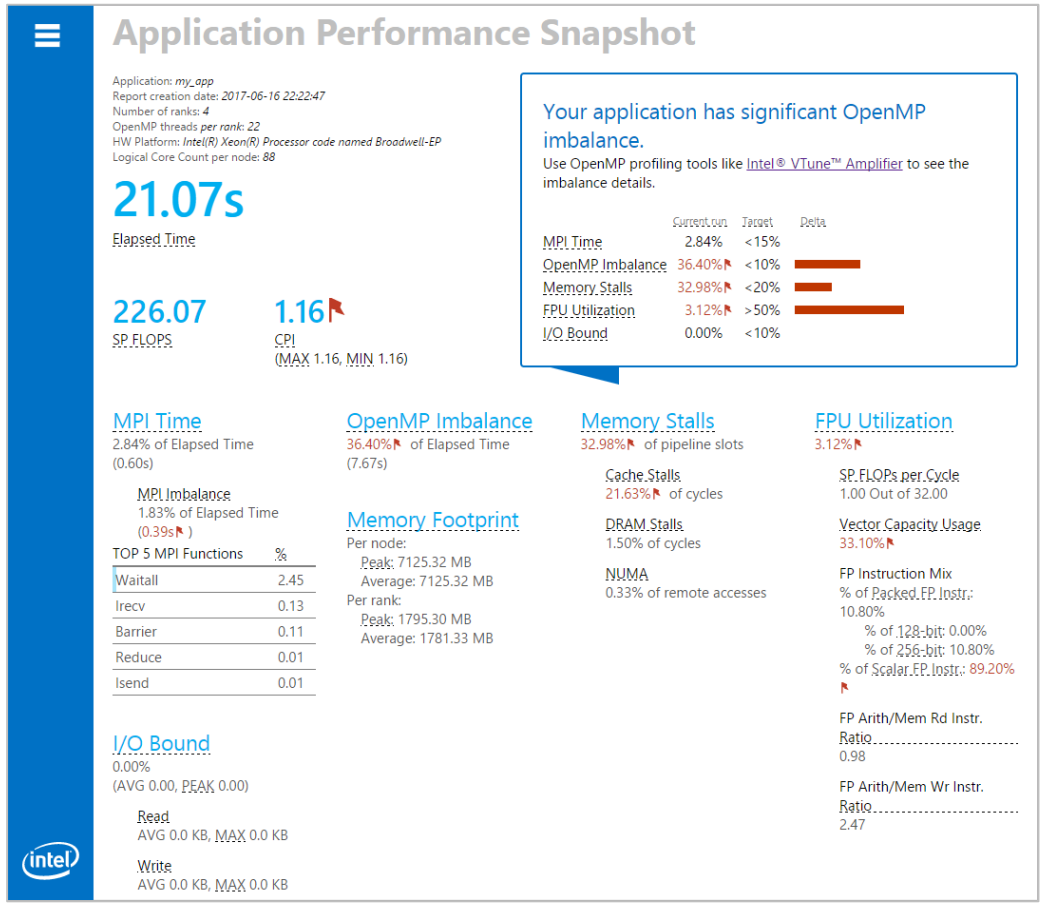

†MPI supported only on Linux*

Analyze & Tune Application Performance & Scalability with Intel® VTune™ Amplifier—Performance Profiler

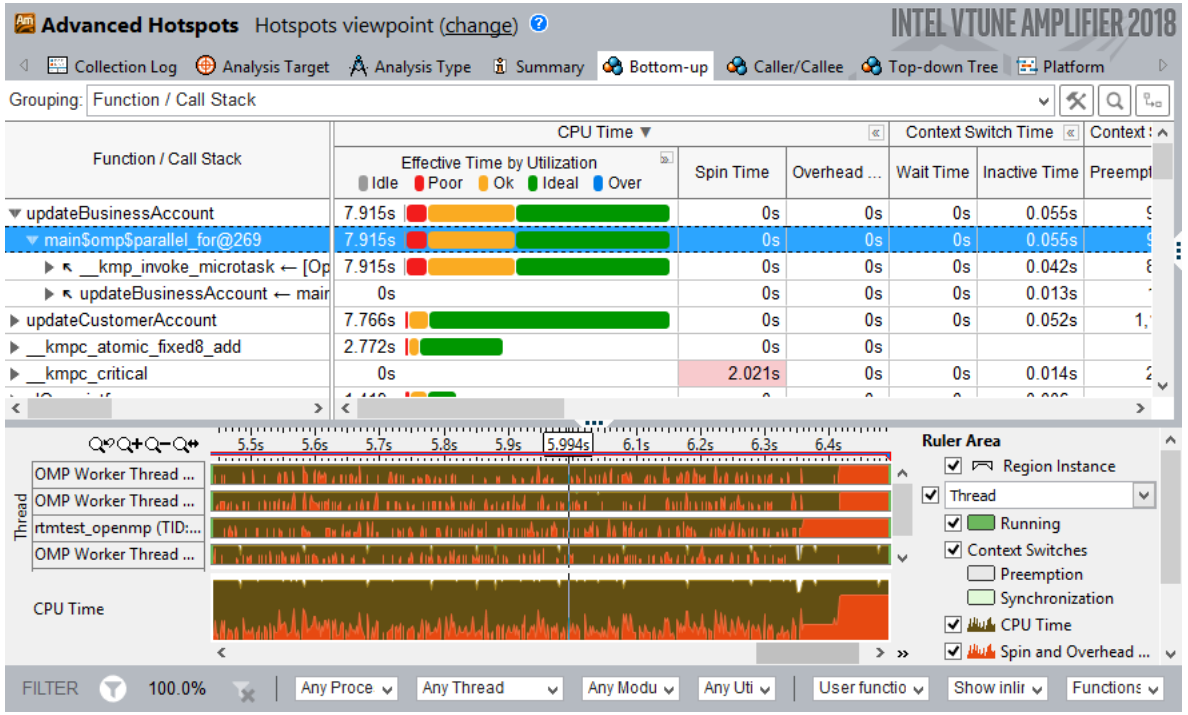

Save Time Optimizing Code

- Accurately profile C, C++, Fortran*, Python*, Go*, Java*, or any mix
- Optimize CPU, threading, memory, cache, storage & more
- Save time: rich analysis leads to insight

New for 2018 edition (partial list)

- Quick metrics for shared & distributed memory apps
- Cross-OS analysis – e.g. analyze Linux* from Windows* or macOS*
- Profile inside containers

Learn More: software.intel.com/intel-vtune-amplifier-xe

Rich Set of Profiling Features for Multiple Markets

Intel® VTune™ Amplifier—Performance Profiler

Basic Profiling
- Hotspots

Threading Analysis
- Concurrency, Locks & Waits
- OpenMP, Intel® Threading Building Blocks

Micro Architecture Analysis
- Cache, branch prediction, …

Vectorization + Intel® Advisor
- FLOPS estimates

MPI + Intel® Trace Analyzer & Collector
- Scalability, imbalance, overhead

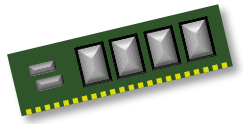

Use Memory Efficiently
- Tune data structures & NUMA

Optimize for High Speed Storage
- I/O and compute imbalance

Intel® Media SDK Integration
- Meaningful media stack metrics

Low Overhead Java*, Python*, Go*
- Managed + native code

Containers
- Docker*, Mesos*, LXC*

‘Automatic’ Vectorization is Often Not Enough

A good compiler can still benefit greatly from vectorization optimization—Intel® Advisor

Compiler will not always vectorize

- Check for Loop Carried Dependencies using Intel® Advisor
- All clear? Force vectorization. C++ use: pragma simd, Fortran use: SIMD directive

Compiler directives are used to force vectorization & get more performance.

Not all vectorization is efficient vectorization

- Stride of 1 is more cache efficient than stride of 2 & greater. Analyze with Intel® Advisor
- Consider data layout changes Intel® SIMD Data Layout Templates can help

Arrays of structures are great for intuitively organizing data, but are less efficient than structures of arrays. Use Intel® SIMD Data Layout Templates to map data into a more efficient layout for vectorization.

Get Breakthrough Vectorization Performance

Intel® Advisor—Vectorization Advisor

Faster Vectorization Optimization

- Vectorize where it will pay off most
- Quickly ID what is blocking vectorization
- Tips for effective vectorization
- Safely force compiler vectorization
- Optimize memory stride

Data & Guidance You Need

- Compiler diagnostics + Performance Data + SIMD efficiency
- Detect problems & recommend fixes
- Loop-Carried Dependency Analysis
- Memory Access Patterns Analysis

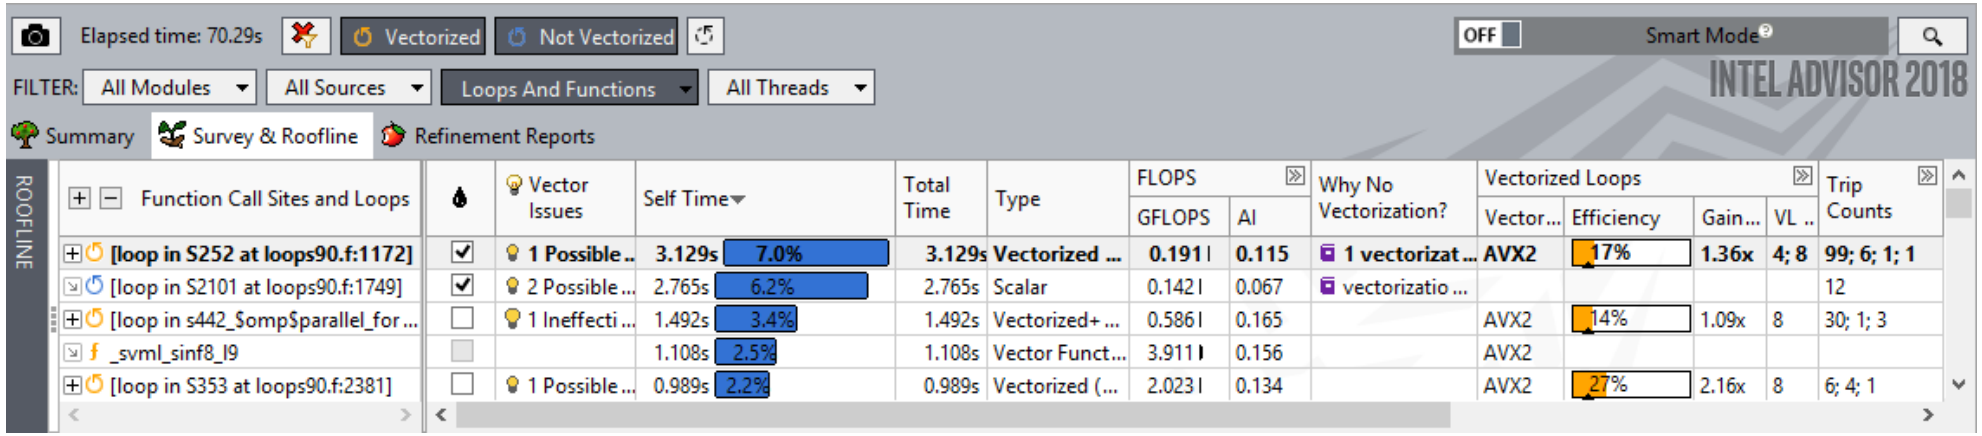

Optimize for Intel® AVX-512 with or without access to AVX-512 hardware

Find Effective Optimization Strategies

Cache-aware Roofline Analysis—Intel® Advisor

Roofline Performance Insights

- Highlights poor performing loops
- Shows performance 'headroom' for each loop
  - Which can be improved
  - Which are worth improving
- Shows likely causes of bottlenecks
- Suggests next optimization steps

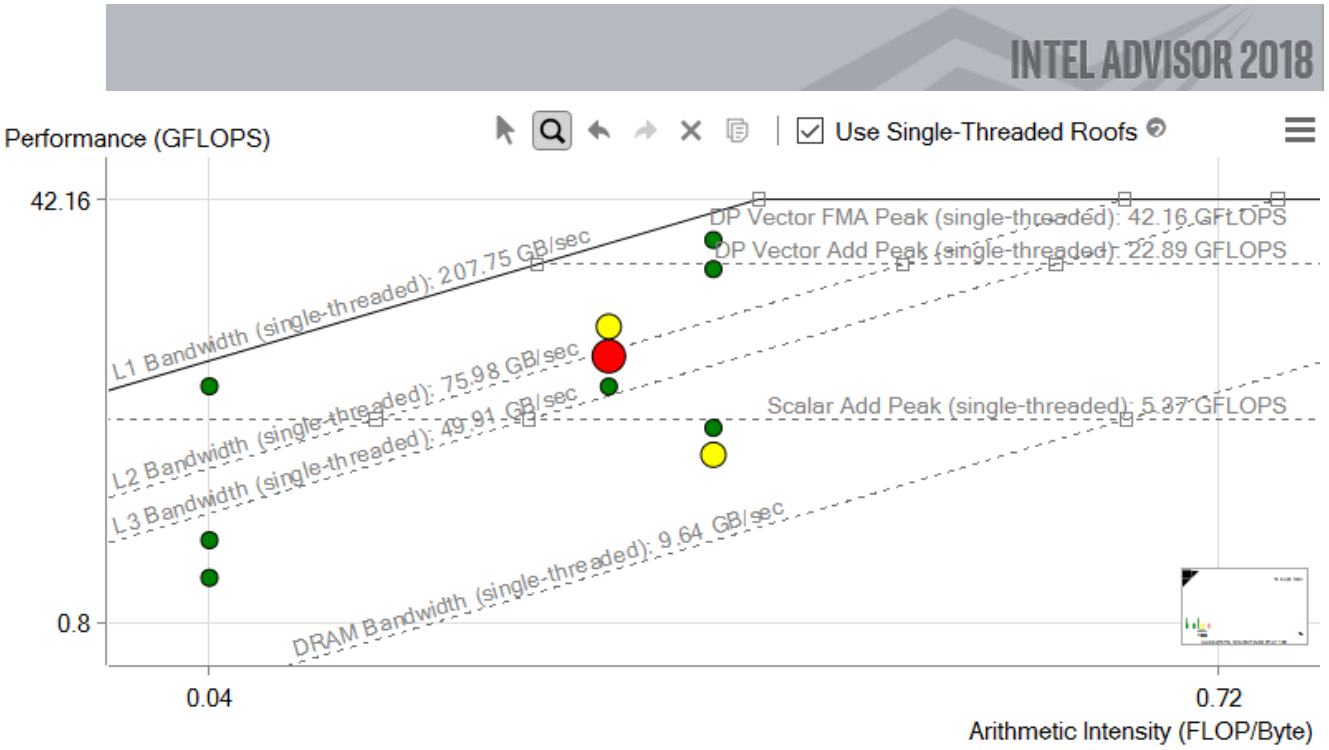

Debug Memory & Threading with Intel® Inspector

Find & Debug Memory Leaks, Corruption, Data Races, Deadlocks

Debugger Breakpoints

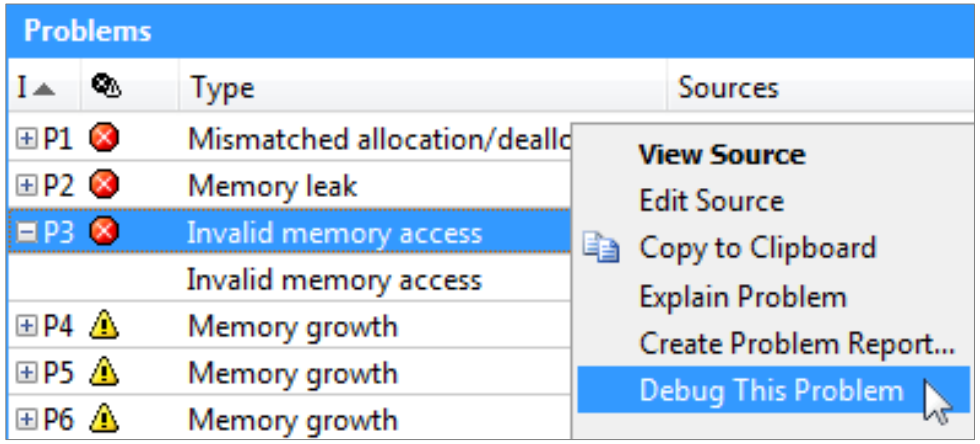

Diagnose in hours instead of months

Learn More: intel.ly/inspector-xe

Correctness Tools Increase ROI by 12%-21%[1]

- Errors found earlier are less expensive to fix
- Races & deadlocks not easily reproduced
- Memory errors are hard to find without a tool

Debugger Integration Speeds Diagnosis

- Breakpoint set just before the problem
- Examine variables and threads with the debugger

What's New in 2018 edition

- Fewer false positives
- C++ 17 std::shared_mutex added
- Windows SRW Locks added

[1]**Cost Factors – Square Project Analysis -** *CERT: U.S. Computer Emergency Readiness Team, and Carnegie Mellon CyLab NIST: National Institute of Standards & Technology: Square Project Results*

INTEL® PARALLEL STUDIO XE COMPONENT TOOLS

BUILD

Intel® C++ Compiler
Intel® Fortran Compiler
Intel® Distribution for Python*
Intel® Math Kernel Library
Intel® Integrated Performance Primitives
Intel® Threading Building Blocks
Intel® Data Analytics Acceleration Library
Included in Composer Edition

ANALYZE

Intel® VTune™ Amplifier XE
Intel® Advisor
Intel® Inspector

Part of the Professional Edition

SCALE

Intel® MPI Library
Intel® Trace Analyzer & Collector
Intel® Cluster Checker

Part of the Cluster Edition

Boost Distributed Application Performance with Intel® MPI Library

Performance, Scalability & Fabric Flexibility

Standards Based Optimized MPI Library for Distributed Computing

- Built on open source MPICH Implementation
- Tuned for low latency, high bandwidth & scalability
- Multi fabric support for flexibility in deployment

What's New in 2018 edition[1]

- Up to **11x** faster in job start-up time
- Up to **25%** reduction in job finalization time
- Supports the latest Intel® Xeon® Scalable processor

Learn More: software.intel.com/intel-mpi-library

[1]See following benchmarks slide for more details

Intel® MPI Library Features

Optimized MPI Application Performance

- Application-specific tuning
- Automatic tuning
- Support for latest Intel® Xeon® & Intel® Xeon Phi™ Processors
- Support for Intel® Omni-Path Architecture Fabric

Multi-vendor Interoperability & Lower Latency

- Performance optimized support for the fabric capabilities through OpenFabrics* (OFI)
- Industry leading latency

Faster MPI Communication - Optimized collectives

Sustainable Scalability

Native InfiniBand* interface support allows for lower latencies, higher bandwidth, and reduced memory requirements

More Robust MPI Applications

Seamless interoperability with Intel® Trace Analyzer & Collector

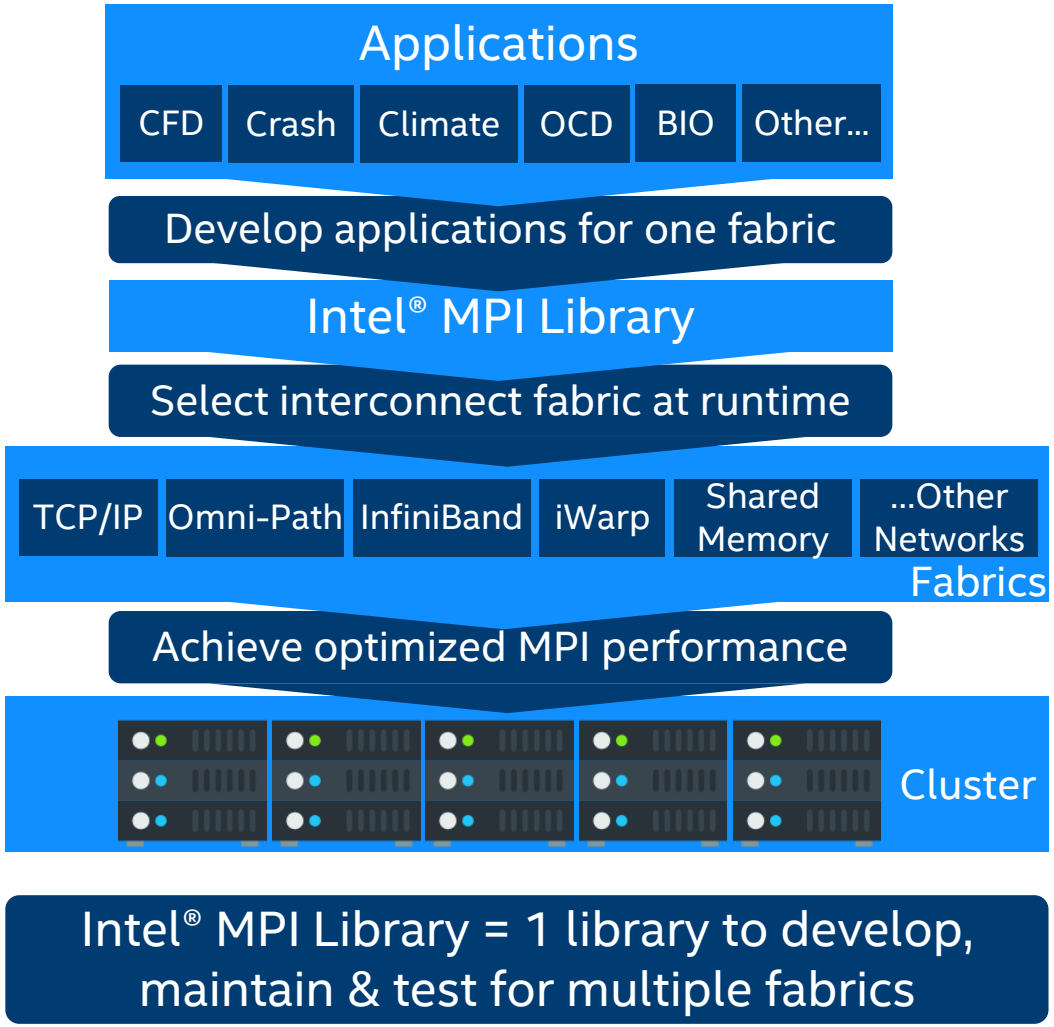

Profile & Analyze High Performance MPI Applications

Intel® Trace Analyzer & Collector

Powerful Profiler, Analysis & Visualization Tool for MPI Applications

- Low overhead for accurate profiling, analysis & correctness checking
- Easily visualize process interactions, hotspots & load balancing for tuning & optimization
- Workflow flexibility: Compile, Link or Run

What's New in 2018 edition

- Support of OpenSHMEM* applications
- Supports the latest Intel® Xeon® Scalable and Intel® Xeon Phi™ processors

Learn More: software.intel.com/intel-trace-analyzer

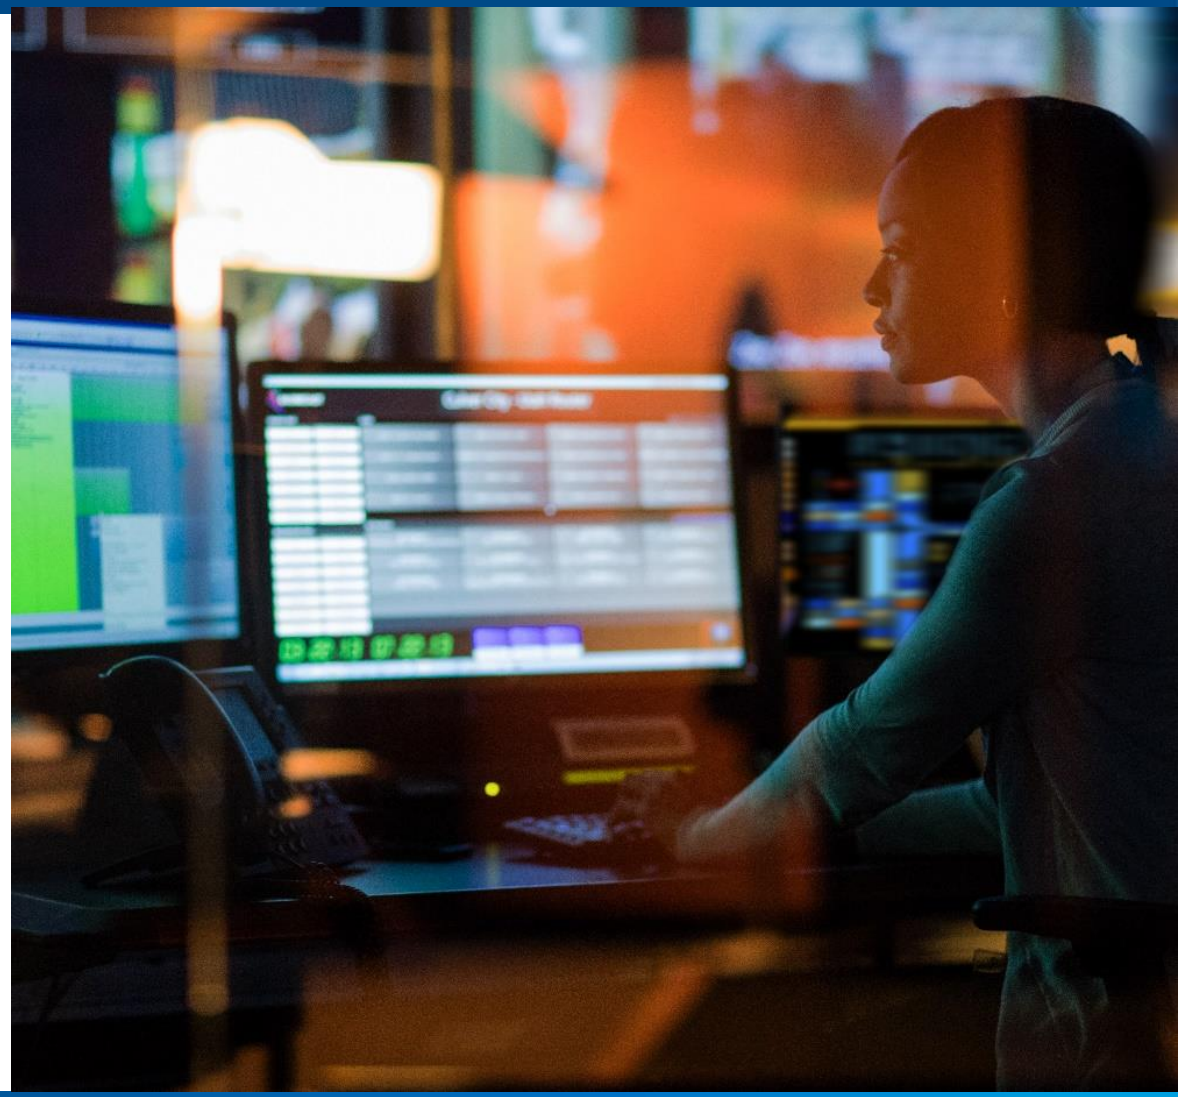

Efficiently Profile MPI Applications

Intel® Trace Analyzer & Collector

Helps Developers

- Visualize & understand parallel application behavior
- Evaluate profiling statistics & load balancing
- Identify communication hotspots

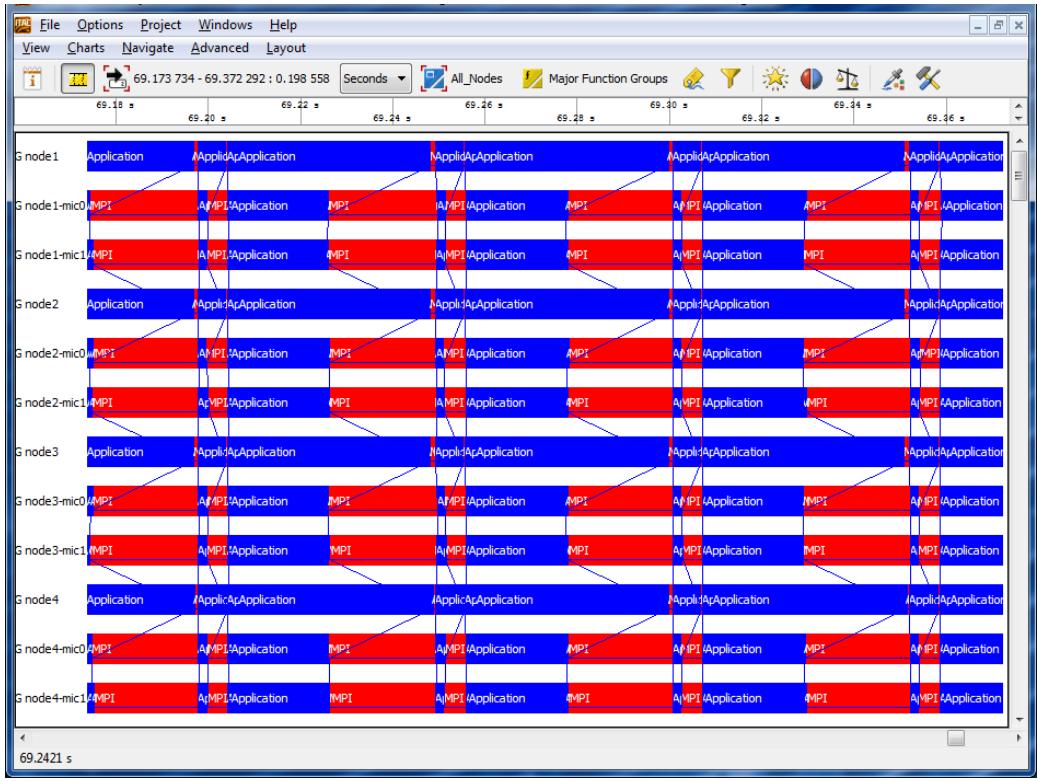

Features

- Event-based approach
- Low overhead
- Excellent scalability
- Powerful aggregation & filtering functions
- Idealizer
- Scalable

Intel® Cluster Checker 2018
For Linux* High Performance Compute Clusters

- Clusters are Complex Systems!
- Challenge is to reduce this complexity barrier for
  - Application developers
  - Cluster architects
  - Cluster users
  - System administrators

- Intel® Cluster Checker is an expert system approach that provides cluster systems expertise
  - Verifies system health
  - Offers suggested actions
  - Provides extensible framework
  - API for integrated support

What's New in Intel® Cluster Checker 2018
Ensure Your HPC Cluster Components Work Together

- New Features Improve Usability & Checking Capabilities
  - Adds support for new Intel silicon & platform elements (processors, fabric, memory, storage, cluster provisioning, HPC platforms)
  - Introduces simplified grouping of checks for extensibility
  - Improves diagnostic output
  - Validates Intel® Scalable System Framework Classic HPC Cluster Reference Architectures
  - Check Intel® Omni-Path in-depth
  - Analyze data from multiple database sources

CODE THAT PERFORMS AND OUTPERFORMS

Download a *free*, 30-day trial of
Intel® Parallel Studio XE 2018 today

software.intel.com/en-us/parallel-studio-xe

AND DON'T FORGET...

To fill out the evaluation survey via a URL that will be provided at the end of the day

OR

Watch your email for a link to the survey

P.S.

Everyone who fills out the survey will receive a personalized certificate indicating completion of the training!

Legal Disclaimer and Optimization Notice

- 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.
- 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.
- OpenCL and the OpenCL logo are trademarks of Apple Inc. used by permission by Khronos.
-

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


intel
Software

Backup: APS for a scientific application

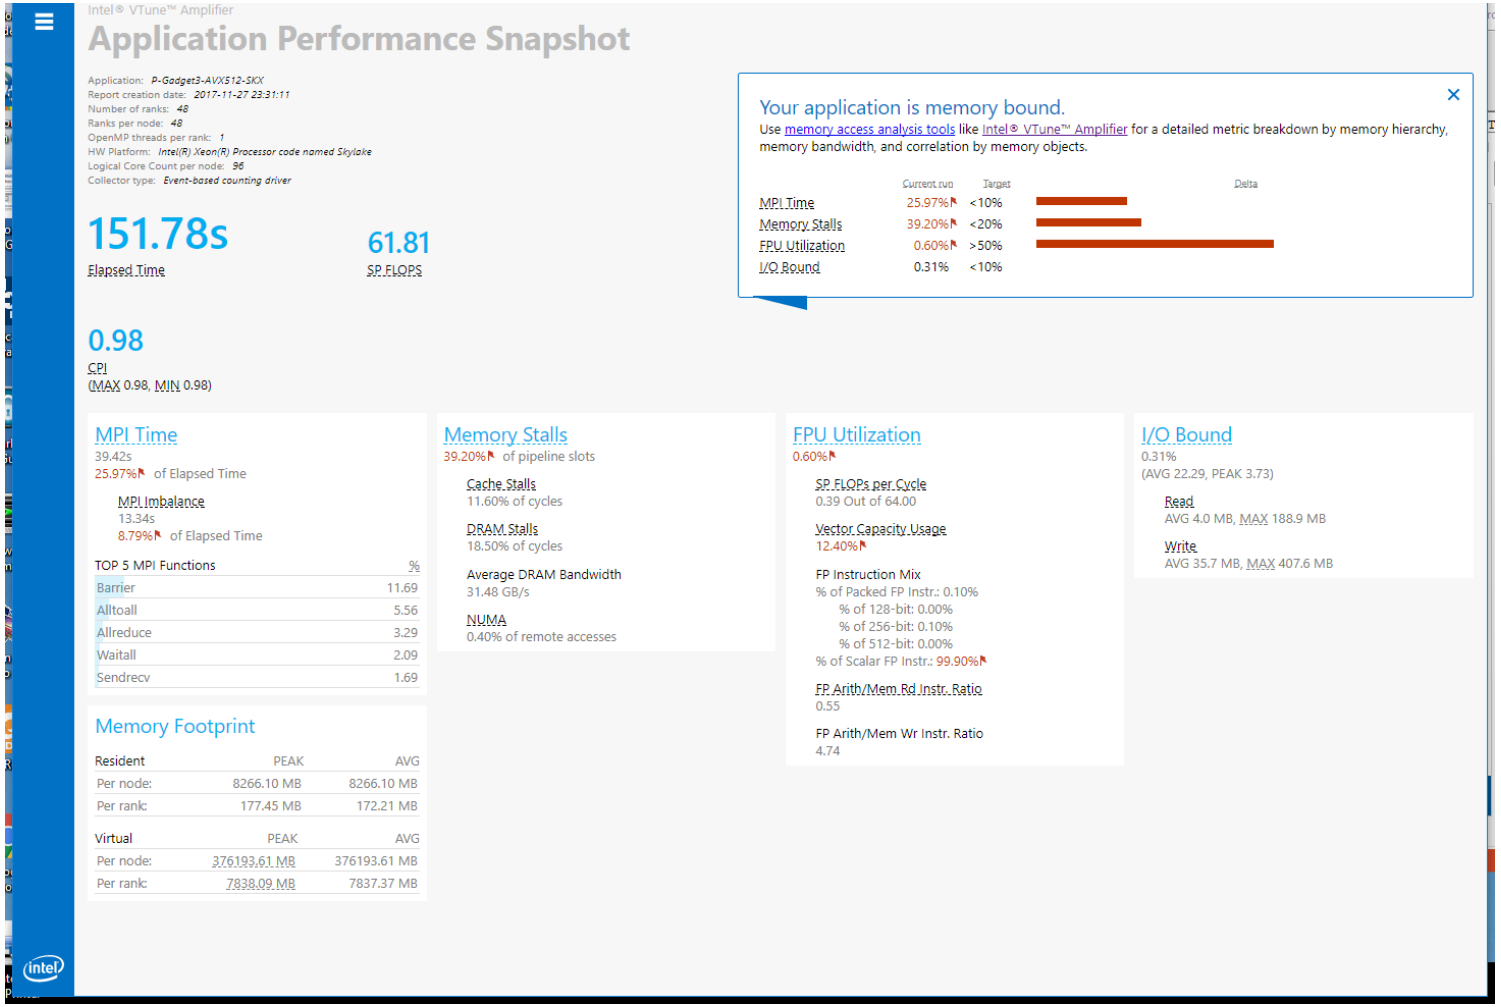

†MPI supported only on Linux*