Intel® Parallel Studio XE 2018 Beta Program: What's New

Contents

1. Introduction 3
2. What's New in Intel® Parallel Studio XE 2018 Beta 3
   2.1. Intel® Application Snapshot 3
   2.2. Intel® C++ and Fortran Compilers 18.0 Beta 4
      2.2.1. Common Features 5
      2.2.2. C and C++ Features 5
      2.2.3. Fortran Features 5
   2.3. Intel® Cluster Checker 2018 Beta 5
   2.4. Intel® Data Analytics Acceleration Library (Intel® DAAL) 2018 Beta 6
   2.5. Intel® Math Kernel Library (Intel® MKL) 2018 Beta 6
   2.6. Intel® Integrated Performance Primitives 2018 Beta 6
   2.7. Intel® Advisor 2018 Beta 6
      2.7.1. Cache-Aware Roofline* 7
      2.7.2. Dynamic Instruction Mix Summary 8
      2.7.3 New MPI Command Line feature 9
      2.7.4 New "Joined" Command Line report to provide merged data from all collections in the same report 9
      2.7.5 Track refinement analysis 9
      2.7.6 Improved Trip Counts and FLOPS* 10
      2.7.7 Filtering by Module* 10
   2.8. Intel® Inspector 2018 Beta 11
   2.9. Intel® VTune™ Amplifier 2018 Beta 12
      2.9.1. Enhanced MPI metrics for HPC Performance Characterization analysis 12
      2.8.2. Memory Consumption analysis for native and Python* Linux targets 13
      2.8.3. Support for Intel® Xeon Phi™ coprocessor codenamed Knights Landing 14

2.8.4. Profiling Docker* Container Targets 14
2.8.5. Profile-Guided Optimization report 15
2.10. Intel® MPI Library 2018 Beta 16
2.11. Intel® Trace Analyzer and Collector 2018 Beta 16
2.12. Intel® Distribution for Python* 2018 Beta 16
2.12.1. What's new 16
3. Additional Information 17
4. Legal Information 17

1. Introduction

Intel® Parallel Studio XE 2018 Beta combines Intel's industry-leading C/C++ and Fortran compilers, performance and parallel libraries, error checking and code robustness tools, and performance profiling tools into a single suite offering. This helps boost application performance and increase the code quality, security, and reliability needed by high-performance computing and enterprise applications. At the same time, the single Intel® Parallel Studio XE suite for Windows* or Linux* eases the procurement of all the necessary tools and simplifies the transition from multicore to manycore processors in the future. Visit our website to learn more.

2. What's New in Intel® Parallel Studio XE 2018 Beta

2.1. Intel® Application Snapshot

Take a quick look at your application's performance to see if it is well optimized for modern hardware. **Application Performance Snapshot** now **merges MPI** and **Application** data to give a unified performance snapshot view which is easy, convenient, faster to view and analyze. Use it to understand where your application will benefit from tuning. Explore richer metrics to understand computation efficiency: **MPI** and **OpenMP* parallelism**, **Memory Stalls**, **FPU utilization**, and **I/O efficiency** with recommendations on further in-depth analysis. Supports Intel® MPI, MPICH and Cray MPI.

See the *Get Started with Application Performance Snapshot* help article for more information.

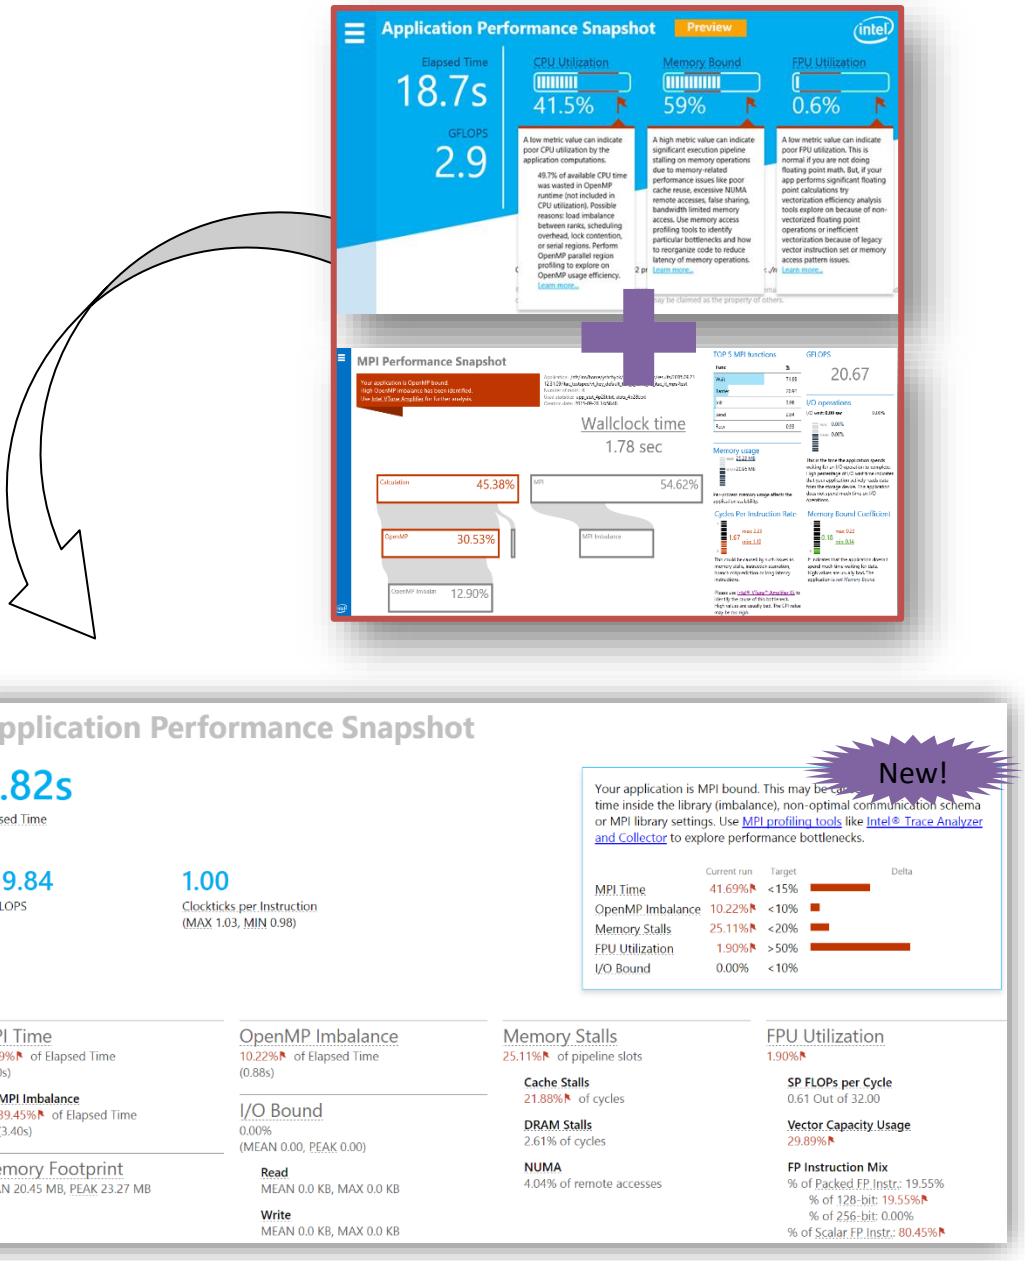

2.2. Intel® C++ and Fortran Compilers 18.0 Beta

The Intel® Parallel Studio XE 2018 Beta includes Intel® C++ Compiler 18.0 Beta and Intel® Fortran Compiler 18.0 Beta. Windows*, Linux*, and macOS* are supported. This section describes the new features we are seeking feedback on. To see other changes, please refer to the product release notes.

2.2.1. Common Features

Intel® C++ and Fortran Compilers 18.0 Beta adds the support of hardware based PGO and extends OpenMP* 5.0 feature support with added support of the following:

- `task_reduction` and `in_reduction` clauses
- `monotonic` and `overlap` keywords for `ordered` block in `simd` context
- `simd early_exit` clause (C/C++)
- conditional scalar assignment for `simd` construct

2.2.2. C and C++ Features

The Intel® C++ Compiler 18.0 beta includes expanded language Standards support. For C++ customers, new C++17 features support includes: nested namespaces, terse static_assert, relaxed range based for loops and Parallel STL. Additional features include C11 _Atomic keyword.

2.2.3. Fortran Features

The Intel® Fortran Compiler 18.0 beta includes full Fortran 2008 support. Additional features include MEMKIND attribute and directive, Visual Studio 2015 shell support and additional "init" options for more memory initialization choices.

2.3. Intel® Cluster Checker 2018 Beta

Intel® Cluster Checker 2018 Beta simplifies diagnosis of issues to improve cluster functionality and performance. It provides rapid reassurance to execute on a fully working cluster environment, which can be easily integrated into other software (using the Intel® Cluster Checker API). Its comprehensive cluster environment checking can be easily extended with custom tests. New features continue to improve usability and checking capabilities:

- Adding support for new Intel silicon and platform elements (processors, fabric, memory, cluster provisioning, HPC platforms)
- Introducing simplified grouping of checks for extensibility
- Improved diagnostic output
- Validation of Intel® Scalable System Framework Classic HPC Cluster Reference Architectures
- Validation with Intel® HPC Orchestrator on installation, compliance
- In-depth Intel® Omni-Path checking, OFI support
- Analysis from multiple database sources
- Bug fixes

2.4. Intel® Data Analytics Acceleration Library (Intel® DAAL) 2018 Beta

- A new documentation page that better organizes a new API References Manual and a new Developer Guide.
- Introduced API modifications to streamline library usage and enable consistency across functionality.
- Introduced support for Decision Tree and Decision Forest for both classification and regression.

2.5. Intel® Math Kernel Library (Intel® MKL) 2018 Beta

- DNN: Average pooling has an option to include padding into mean values computation.

- BLAS Features: Introduced optimized integer matrix-matrix multiplication routines (GEMM_S16S16S16 and GEMM_S16S16S32) to work with quantized matrices for all architectures. Introduced ?TRSM_BATCH to complement the batched BLAS for all architectures.

- Library Engineering: Introduced support for the next generation Intel® Xeon Phi™ processor. Optimizations are not dispatched unless explicitly enabled with mkl_enable_instructions function call or MKL_ENABLE_INSTRUCTIONS environment variable.

2.6. Intel® Integrated Performance Primitives 2018 Beta

- Introduced the patch files for the GraphicsMagick source to provide drop-in optimization by the Intel® IPP functions.

- The cryptography functions are now provided as the standalone packages without the need of installation of the main Intel® IPP packages.

- Improved performance of LZO data compression functions on Intel® Advanced Vector Extensions 2 (Intel® AVX2) and Intel® Streaming SIMD Extensions 4.2 (Intel® SSE4.2).

2.7. Intel® Advisor 2018 Beta

Intel® Advisor provides two tools to help ensure your Fortran and native/managed C++ applications take full performance advantage of today's processors:

- **Vectorization Advisor** is a vectorization analysis tool that lets you identify loops that will benefit most from vectorization, identify what is blocking effective vectorization, explore the benefit of alternative data reorganizations, and increase the confidence that vectorization is safe.

- **Threading Advisor** is a threading design and prototyping tool that lets you analyze, design, tune, and check threading design options without disrupting your normal development.

This section describes the new features we are seeking feedback on. To see other changes, please refer to the product release notes.

2.7.1. Cache-Aware Roofline*

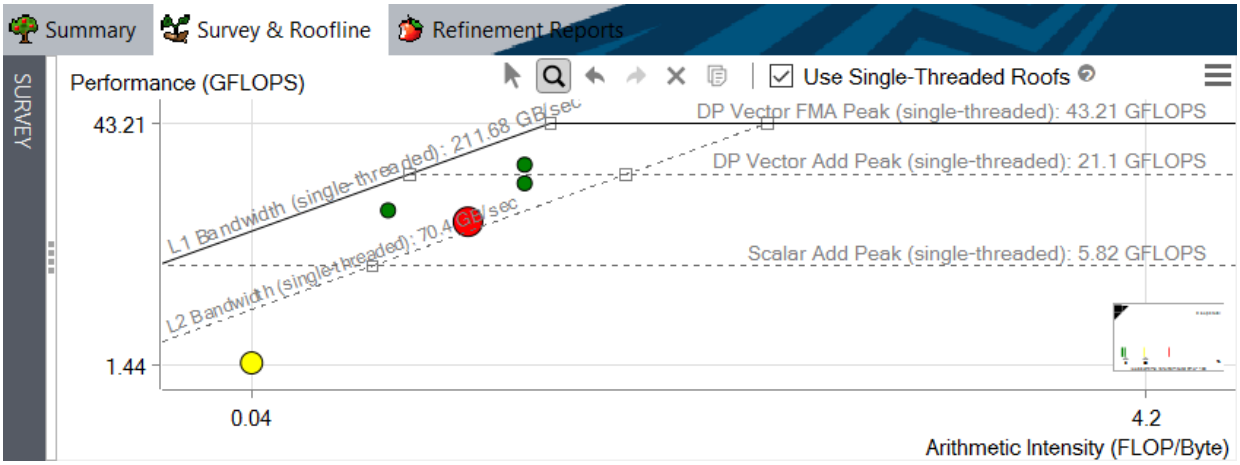

Cache-Aware Roofline, introduced as a preview in the previous update, is now an official feature of Intel® Advisor, and no longer requires setting an environment variable to access. This visual performance model provides insight into the source of performance limitations and reveals how much potential performance your application is currently leaving on the table, allowing you to spend your time efficiently in optimizing where it counts most.

More information about the Roofline feature can be found in the Intel® Advisor Roofline main article.

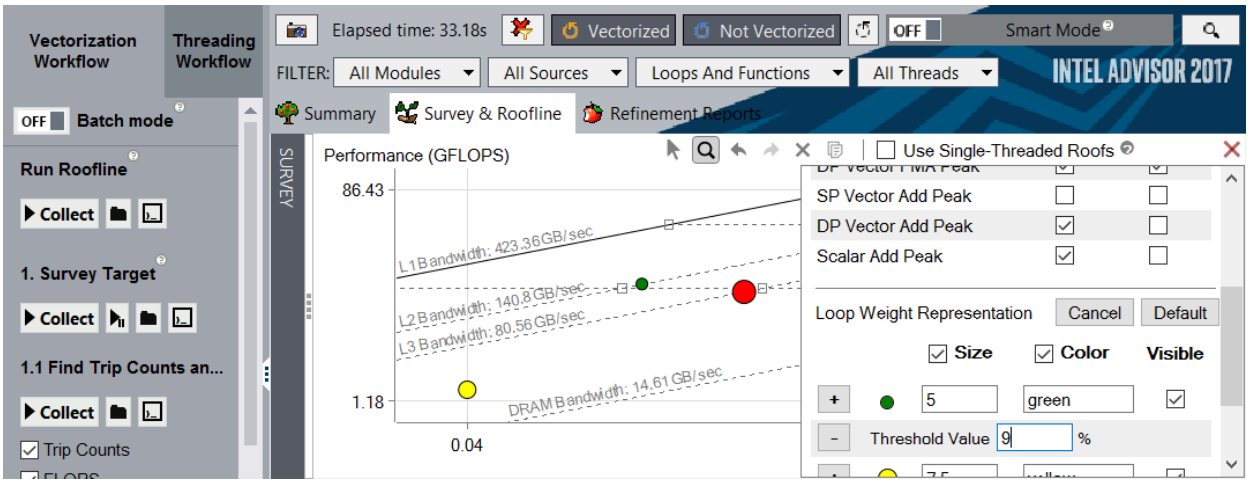

Roofline modeling was first proposed in a paper by Berkeley researchers in 2009. The model has been improved since then, particularly in a 2013 paper from the Technical University of Lisbon that introduced the Cache-Aware model, which is now available in Advisor as an automatically-generated chart when both Survey and Trip Counts with FLOPS have been collected. The Roofline chart and Survey report can either be viewed on their own using the toggle bar on the left to switch between them, or side by side by clicking the white ribbon with four gray dots (initially located by the toggle bar).

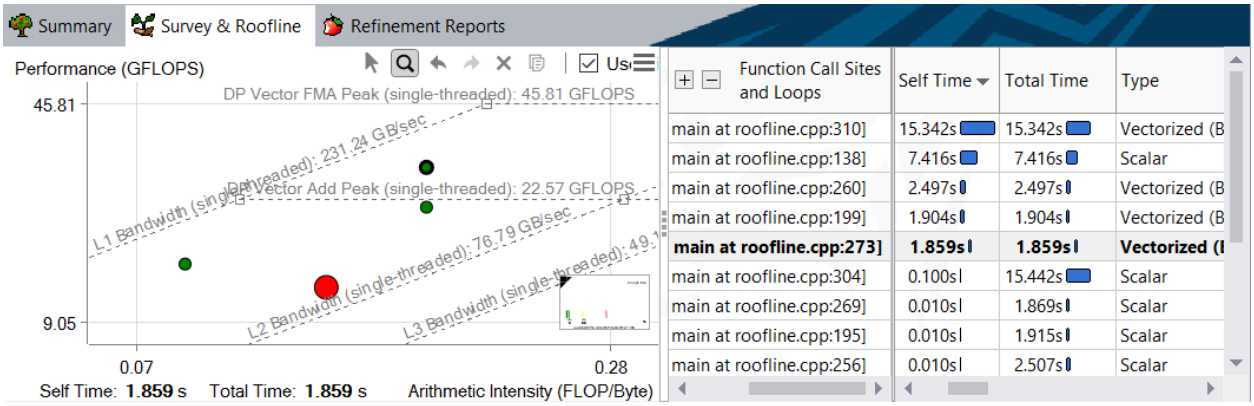

2.7.2. Dynamic Instruction Mix Summary

In addition to providing a Static instruction mix summary we also dynamic count the number of each instruction executed. We provide this as shown in the table below.

This data is computed as part of the "Trip Count/FLOPS" Collection.

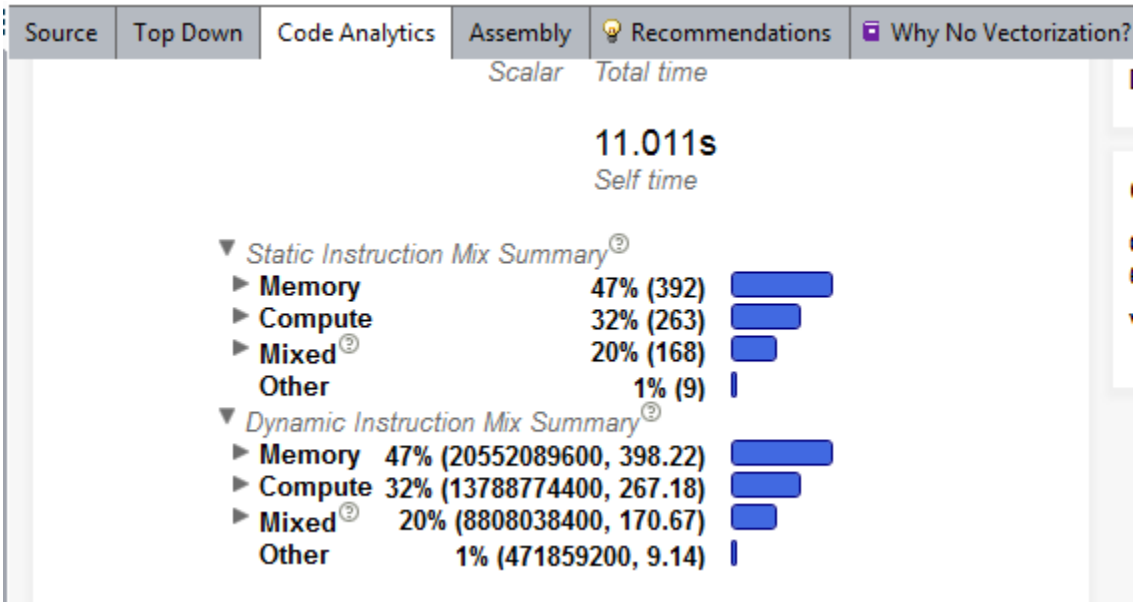

2.7.3 New MPI Command Line feature

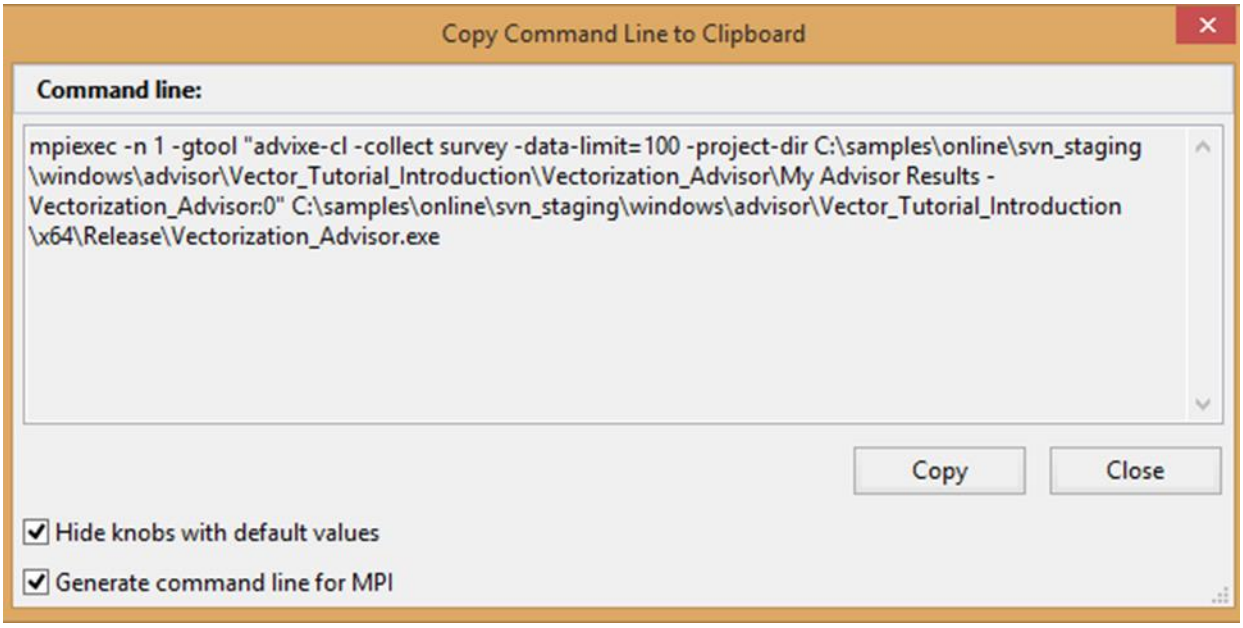

2.7.4 New "Joined" Command Line report to provide merged data from all collections in the same report

```
advixe-cl –report joined –proj-dir ./advixe
```

2.7.5 Track refinement analysis

During our refinement analysis we now report the loops that have been selected for additional analysis.

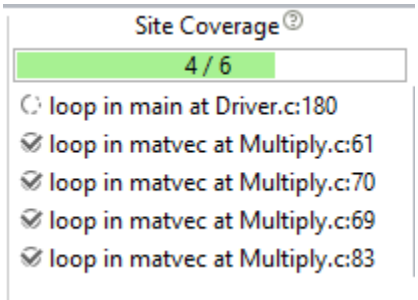

2.7.6 Improved Trip Counts and FLOPS*

The general quality and coverage of the Trip Counts collection has been improved.

Collecting Trip Counts data also collects a new **Call Count** metric for functions.

Trip Counts and FLOPS can now be **collected independently** of each other. In the GUI, there are checkboxes in the work flow. On the command line, FLOPS can be enabled by adding the -flops-and-masks option to a tripcounts collection, while Trip Counts can be disabled by adding -no-trip-counts.

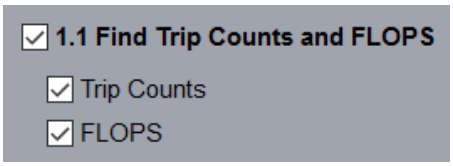

2.7.7 Filtering by Module*

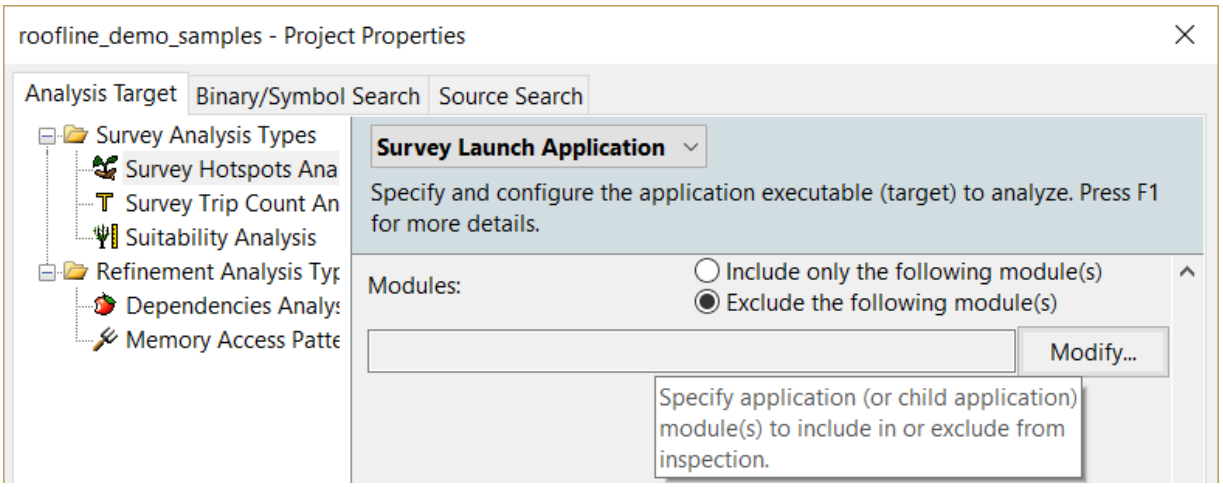

Application analysis scope can now be narrowed down before collection begins to cut out extraneous data and reduce overhead for Survey and Trip Counts analysis types, by **restricting collection to specific modules**. Modules can be filtered by excluding unwanted files or by restricting collection to a specific list of desired modules, using either the command line or the GUI project properties.

*Also available in Intel® Advisor 2017 update 2

2.8. Intel® Inspector 2018 Beta

Intel® Inspector is an easy-to-use memory and threading error debugger for C, C++ and Fortran* applications that run on Windows* and Linux*. No special compilers or builds are required. Just use a normal debug or production build. Use the graphical user interface or automate regression testing with the command line. The user interface can be used standalone on both Windows and Linux or integrated with Microsoft Visual Studio*. Enhance productivity, cut costs, and speed time-to-market. This section describes the new features we are seeking feedback on. To see other changes, please refer to the product release notes.

- Support for C++17 std::shared_mutex that enables threading error analysis for applications with read/write synchronization primitives (also available in Intel® Inspector 2017 update 2)

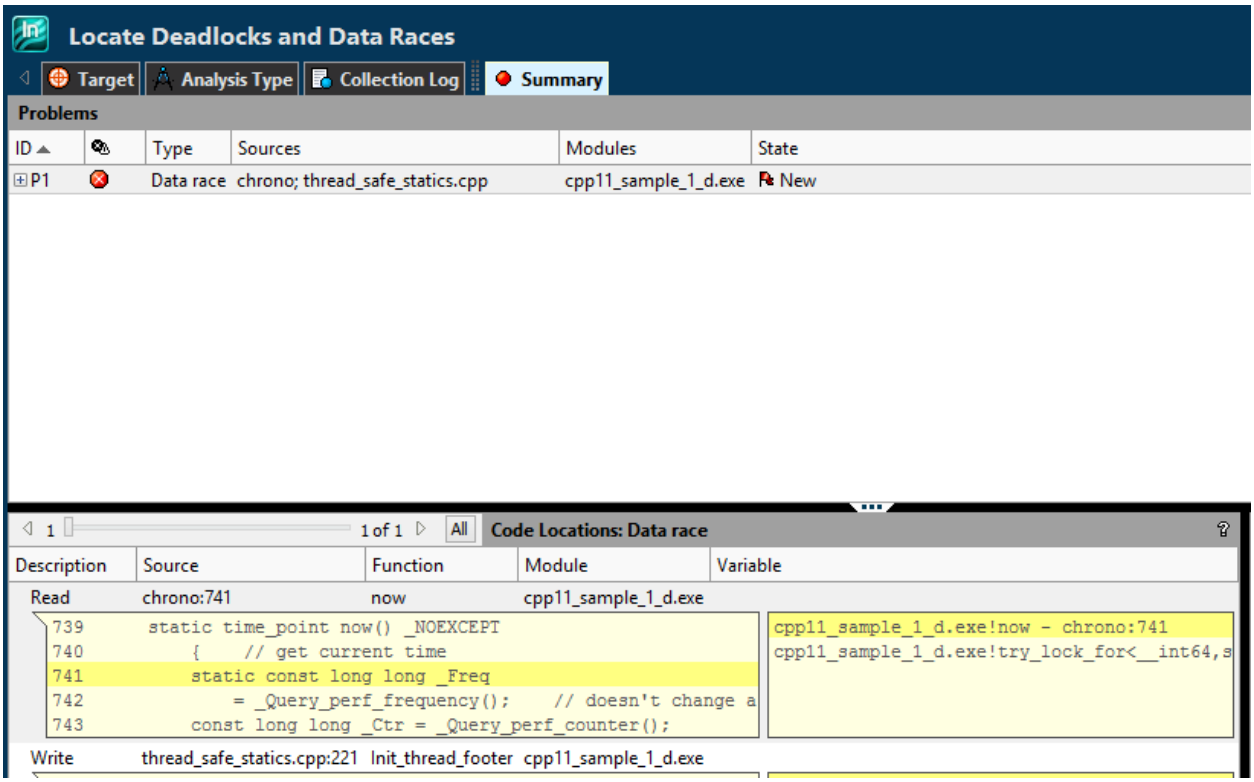

- Added support for Fedora 25 and Ubuntu 16.10 (also available in Intel® Inspector 2017 update 2)
- Support for cross-OS analysis to all license types. The installation packages for additional operating systems can be downloaded from registrationcenter.intel.com (also available in Intel® Inspector 2017 update 2)

2.9. Intel® VTune™ Amplifier 2018 Beta

2.9.1. Enhanced MPI metrics for HPC Performance Characterization analysis

Use the **HPC Performance Characterization analysis** as a starting point to analyze the performance of hybrid OpenMP* and MPI applications.

For MPI applications, review the **MPI Imbalance metric** that shows the CPU time spent by ranks spinning in waits on communication operations, normalized by number of ranks on the profiling node.

The metric issue detection description generation is based on minimal MPI Busy Wait time by ranks. If the minimal MPI Busy wait time by ranks is not significant, then the rank on with the minimal time most likely lies on the critical path of application execution. In this case, review the CPU utilization metrics by this rank.

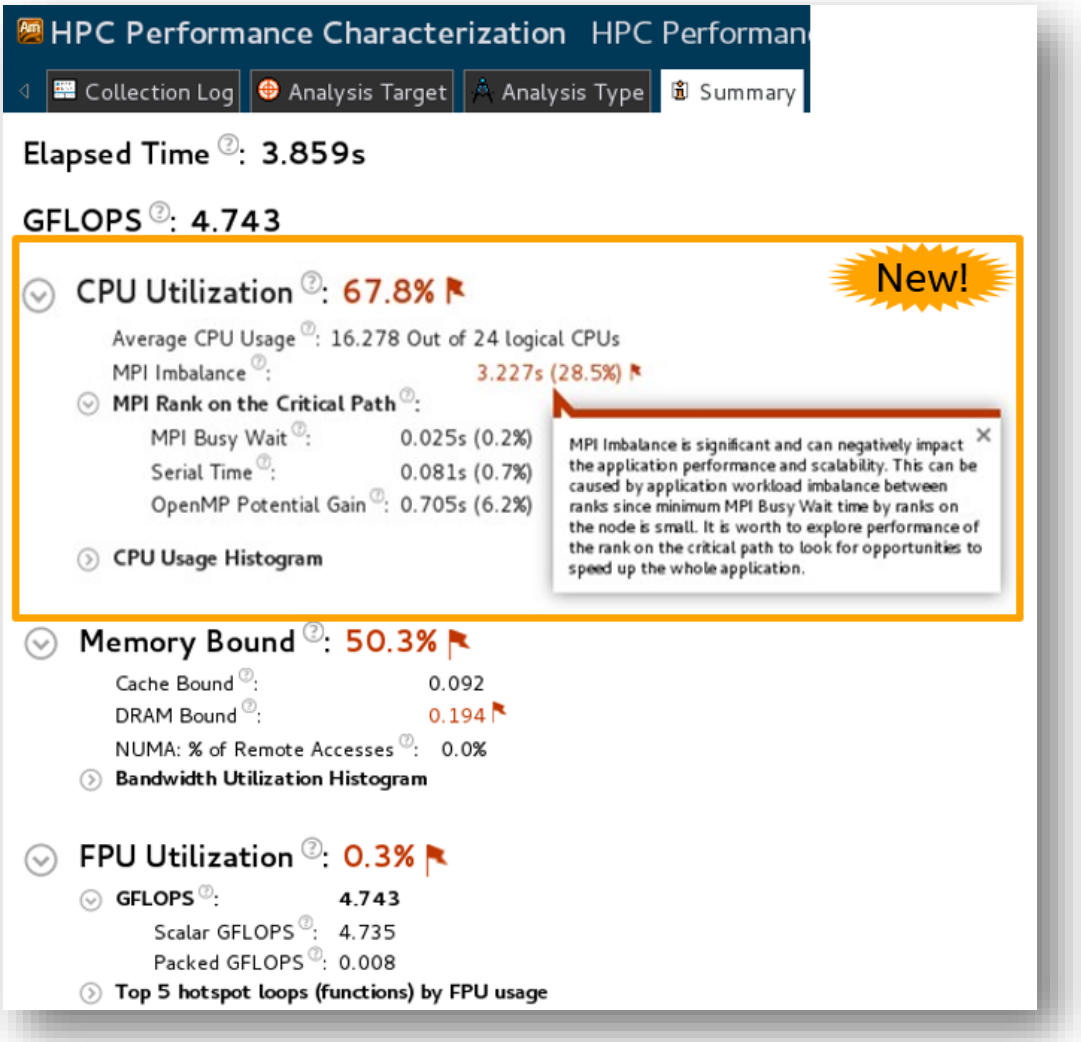

See the help topic *Interpreting HPC Performance Characterization Data* for more information.

2.8.2. Memory Consumption analysis for native and Python* Linux targets

Use the **Memory Consumption analysis** for your Linux* native or Python* targets to explore memory consumption (RAM) over time and identify the most memory-consuming objects allocated and released during the analysis run.

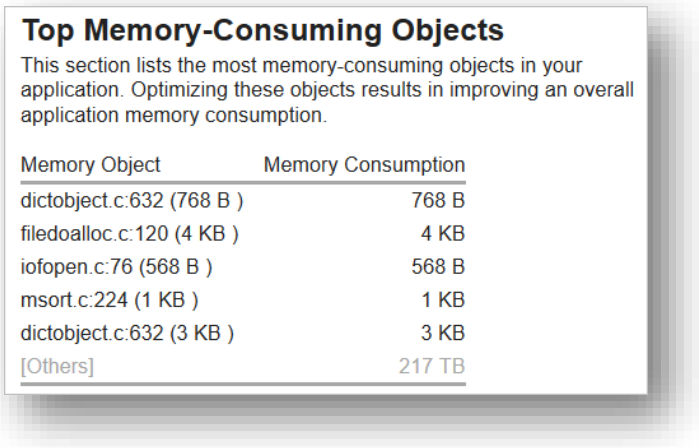

Explore the memory consumption distribution over time. Focus on the peak values on the Timeline pane, select a time range of interest, right click and use the Filter In by Selection context menu option to filter in the program units (functions, modules, processes, and so on) executed during this range:

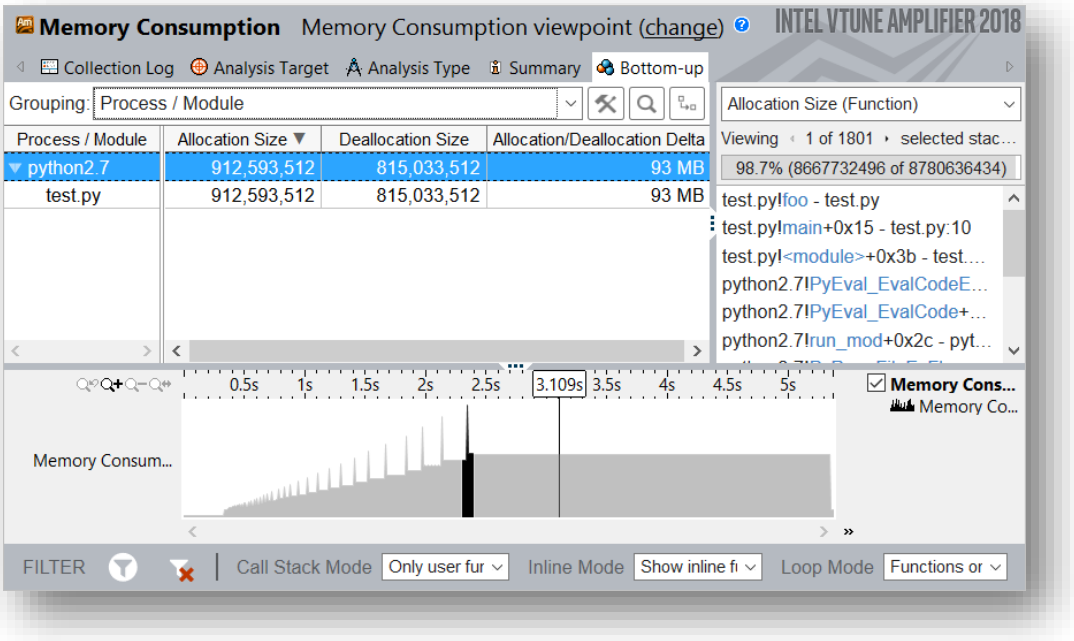

See the help topic Memory Consumption Analysis for more information.

2.8.3. Support for Intel® Xeon Phi™ coprocessor codenamed Knights Landing

Intel® Xeon Phi™ coprocessor codenamed Knights Landing is supported by VTune Amplifier from Linux* OS host with hardware event based sampling analysis including General Exploration, Memory Access and HPC Performance Characterization. Support is enabled for both native and offload applications. Installation automatically configures SEP drivers and collector on the card.

For more details on analysis configuration see the *Intel® Xeon Phi™ Coprocessor Analysis Workflow* product help article.

2.8.4. Profiling Docker* Container Targets

Profile native or Java* applications running in a Docker container with help of **Advanced Hotspots** analysis on a Linux system. VTune Amplifier automatically detects an application running in the container and doesn't require container configuration specific.

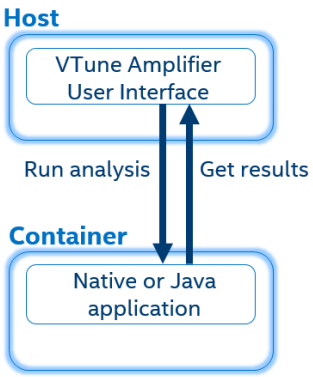

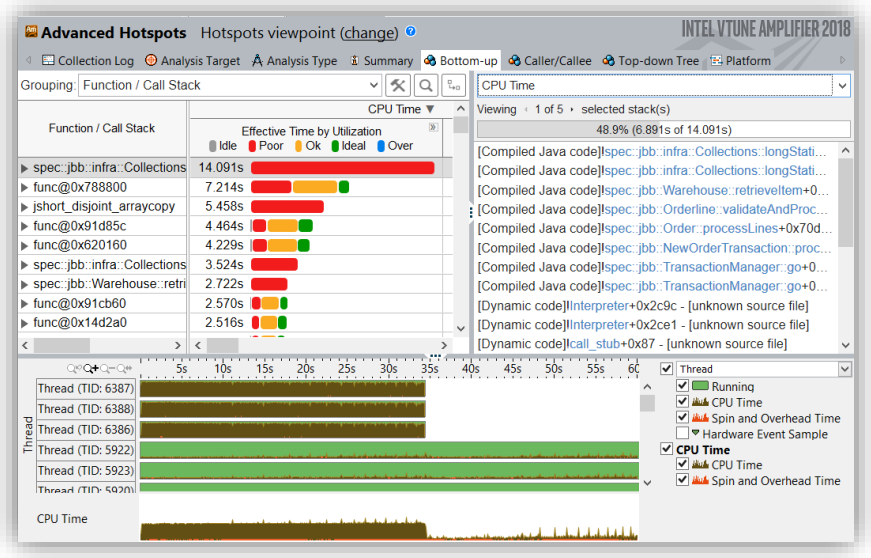

See the help topic Profiling Docker* Container Targets for more information.

2.8.5. Profile-Guided Optimization report

VTune Amplifier provides an option to generate a Profile-Guided Optimization (PGO) report for Intel® Compiler (Linux* only), Clang* compiler, and GCC* compiler. Using this information the compilers can provide a better performance optimization when compiling your code. To generate a PGO report, use the *amplxe-pgo-report.sh* utility.

See the help topic Profile-Guided Optimization Report for more information.

2.10. Intel® MPI Library 2018 Beta

Intel® MPI Library 2018 Beta includes performance improvements to the Hydra launcher to increase stability and reduce job start time. Hard finalization is now the default for jobs using OFI or TMI fabrics. Hard finalization is a less graceful but faster finalization process. Support is removed for the Intel® Xeon Phi™ Coprocessor (code named Knights Corner) in order to focus on improvements for newer architectures.

2.11. Intel® Trace Analyzer and Collector 2018 Beta

Intel® Trace Analyzer and Collector 2018 Beta introduces experimental support for tracing SHMEM over MPI. This allows a user to trace SHMEM functions using the -trace functionality in mpirun and display SHMEM functions in Intel® Trace Analyzer. The help system has been modified for this version. Support is removed for the Intel® Xeon Phi™ Coprocessor (code named Knights Corner).

2.12. Intel® Distribution for Python* 2018 Beta

The Intel® Distribution for Python* 2018 Beta includes additional performance packages and optimizations to help accelerate industry standard software. Updated 2018 Beta versions of Intel® performance libraries have also been included with the distribution. OpenCV, a computer vision package, has been included to the list of available packages in the distribution.

2.12.1. What's new

The Beta adds additional accelerations for NumPy, SciPy, and scikit-learn through the following optimization work:

- Arithmetic & transcendental functions in NumPy have been accelerated with additional 60X speedups
- Further FFT optimizations in NumPy & SciPy with ~400X speedups
- Additional scikit-learn optimizations over version 2017 Update 2 have been made with additional integration of Intel® DAAL (Intel® Data Analytics Acceleration Library) components.

New packages targeting computer vision have also been optimized and included:

- OpenCV is now accelerated with Intel® IPP (Intel® Integrated Performance Primitives), and is included with the distribution.

3. Additional Information

Please visit our product website to learn more about the Intel® Software Development Tools. If you have any questions regarding the Beta program, please refer to the Intel Parallel Studio XE 2018 Beta page. If you have any requests, issues, or feedback contact us at Online Service Center.

4. Legal Information

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to: http://www.intel.com/design/literature.htm

This document contains information on products in the design phase of development.

MPEG-1, MPEG-2, MPEG-4, H.261, H.263, H.264, MP3, DV, VC-1, MJPEG, AC3, AAC, G.711, G.722, G.722.1, G.722.2, AMRWB, Extended AMRWB (AMRWB+), G.167, G.168, G.169, G.723.1, G.726, G.728, G.729, G.729.1, GSM AMR, GSM FR are international standards promoted by ISO, IEC, ITU, ETSI, 3GPP and other organizations. Implementations of these standards, or the standard enabled platforms may require licenses from various entities, including Intel Corporation.

BlueMoon, BunnyPeople, Celeron, Celeron Inside, Centrino, Centrino Inside, Cilk, Core Inside, E-GOLD, Flexpipe, i960, Intel, the Intel logo, Intel AppUp, Intel Atom, Intel Atom Inside, Intel CoFluent, Intel Core, Intel Inside, Intel Insider, the Intel Inside logo, Intel NetBurst, Intel NetMerge, Intel NetStructure, Intel SingleDriver, Intel SpeedStep, Intel Sponsors of Tomorrow., the Intel Sponsors of Tomorrow. logo, Intel StrataFlash, Intel vPro, Intel Xeon Phi, Intel XScale, InTru, the InTru logo, the InTru Inside logo, InTru soundmark, Itanium, Itanium Inside, MCS, MMX, Pentium, Pentium Inside, Puma, skoool, the skoool logo, SMARTi, Sound Mark, Stay With It, The Creators Project, The Journey Inside, Thunderbolt, Ultrabook, vPro Inside, VTune, Xeon, Xeon Inside, X-GOLD, XMM, X-PMU and XPOSYS 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.

Microsoft, Windows, and the Windows logo are trademarks, or registered trademarks of Microsoft Corporation in the United States and/or other countries.

Java is a registered trademark of Oracle and/or its affiliates.

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