

PRODUCT BRIEF

High-Performance Computing
Intel® VTune™ Amplifier



Optimize your Software for Modern Hardware

Discover performance bottlenecks faster with accurate data plus advanced analysis using Intel® VTune™ Amplifier

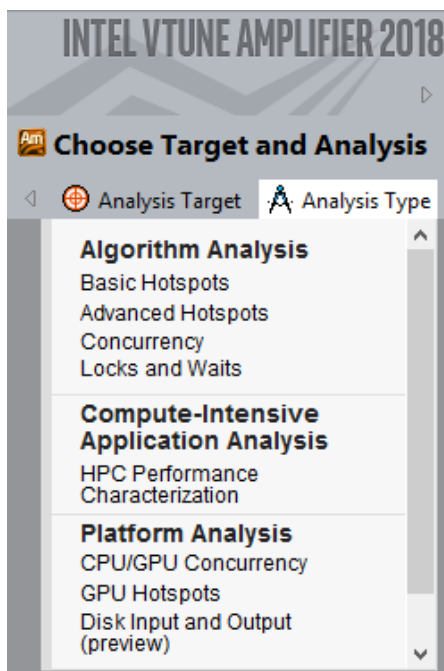


Figure 1. Collect and analyze a rich set of performance data.

Collect a Wide Range of Performance Data

Whether you're tuning a simple application for the first time—or doing advanced performance optimization on a threaded MPI* application—you get the data you need with Intel® VTune™ Amplifier. Collect a rich set of performance data for hotspots, threading, locks and waits, DirectX*, OpenCL*, OpenMP*, Intel® Threading Building Blocks, bandwidth, cache, memory access, non-uniform memory, storage latency, and more (Figure 1). Profile C, C++, C#, Fortran, Python*, Go*, Java*, and OpenCL—or any mix. Unlike single-language profilers, Intel VTune Amplifier analyzes mixed code. You can:

- **See more data.** CPU, FPU, GPU, threading, memory access, and more.
- **Get fast answers.** Easy analysis turns data into insight.
- **Create faster code.** Tune with accurate data and low overhead.
- **Improve your workflow** with both local and remote collection and a command line/graphical interface.

You also get priority support. Connect directly to Intel engineers for confidential answers to technical questions, access older versions of the products, and receive free updates for a year. (Paid license required.)

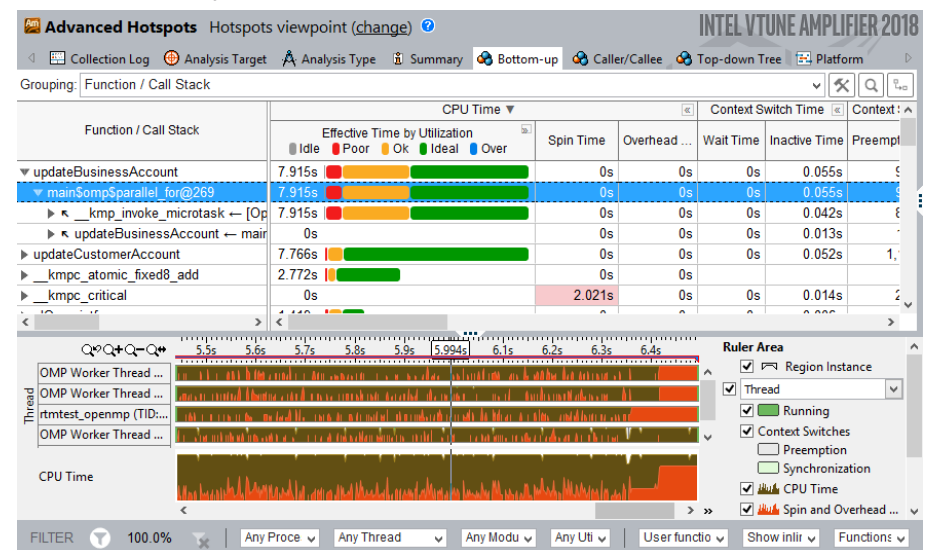


Figure 2. Hotspot analysis shows where your application's time is spent. It can also give a detailed analysis of threading performance showing the potential performance gain and common causes of poor performance such as imbalance, lock contention, forking, scheduling, and reduction.

Powerful Data Analysis Saves You Time

Good data isn't enough. You need to mine the data for insights. Save time with high-level summaries and powerful analysis to sort, filter, and visualize results on the timeline and on your source.

What's New for 2018

- **More MPI metrics** including imbalance and critical path rank.
- **Application Snapshot** merges MPI and OpenMP data and adds MPICH* support.
- **Profile inside containers** and attach to running Java services and daemons.
- **Python profiling** adds threading and memory consumption analyses.
- **OpenCL kernel hotspot analysis** finds kernel bottlenecks.
- **Multi-OS.** Collect and analyze data on all supported operating systems.

Integrated With Intel® Parallel Studio XE Analysis Tools, Runtimes, and Compilers

Intel VTune Amplifier is available standalone or as part of Intel® Parallel Studio XE, a comprehensive software development suite for high-performance software. Additional analysis tools complement Intel VTune Amplifier when purchased as an Intel Parallel Studio bundle, including:

- **Intel® Advisor** to optimize vectorization and prototype threading.
- **Intel® Trace Analyzer and Collector** to examine MPI applications and tell Intel VTune Amplifier which loops will benefit most from threading optimization.

Get the Data You Need

- Hotspot (statistical call tree), call counts (statistical)
- Thread profiling with locks and waits analysis

- Memory access, cache miss, bandwidth, NUMA analysis
- FLOPS and FPU utilization
- Storage accesses mapped to source, latency histogram, and I/O wait
- OpenCL program kernel tracing and GPU offload

Easy to Use

- No special compiles: C, C++, C#, Fortran, Java, Python, Go, ASM*
- Microsoft Visual Studio* IDE integration
- Graphical interface and command line
- Local and remote data collection, multi-rank setup for MPI applications
- Collect on: Linux*, Windows*, FreeBSD*, Android* and select embedded operating systems.
- Analyze results on: Linux, Windows, and macOS* hosts.

Find Answers Fast

- View results on the source/assembly.
- OpenMP scalability analysis and graphical frame analysis
- Memory analysis: Tune data structures and optimize NUMA latency.
- Filter out extraneous data with the timeline and viewpoints.
- Visualize thread and task activity on the timeline.

Low-Overhead/High-Resolution Hardware Profiling

In addition to basic analysis that works on both Intel® and compatible processors, Intel VTune Amplifier has advanced analysis that uses the on-chip Performance Monitoring Unit (PMU) on Intel processors to collect data with very low overhead. This also finds important performance issues like cache misses, branch mispredictions, bandwidth, and more.

Product Details

Quickly locate code taking a lot of CPU time. Hotspots analysis gives you a sorted list of the functions that are using a lot of CPU time. This is where tuning gives you the biggest benefit. Click > for the call stacks. Double-click to see the source.

Grouping: Function / Call Stack		CPU Time			Spin Time	Overhead Time		
Function / Call Stack	Effective Time by Utilization	Idle	Poor	Ok	Ideal	Over		
▼ FireObject::checkCollision	3.348s						0s	0s
▶ [Loop at line 1453 in FireObject::P	2.771s						0s	0s
▶ [Loop at line 1491 in FireObject::P	0.578s						0s	0s
▶ [Loop at line 1453 in FireObject::Proces	1.052s						0s	0s
▶ rand	0.696s						0s	0s
▶ ParticleEmitter::FirePatch::initParticle	0.520s						0s	0s

Product Details (Continued)

See the results on your source. A double-click from the function list takes you to the hottest spot in the function: C, C++, Fortran, assembly, Java, Python, Go, and now OpenCL kernels. See line-level profiling details on the source.

Source Line	Source	Effective Time by Utilization	Spin Time	Overhead Time
480	maxP = FMin(neg? param1 : p	1.7%	0.0%	0.0%
479	minP = FMax(neg? param2 : p	2.4%	0.0%	0.0%
478	bool neg = (rz < 0.f);	0.2%	0.0%	0.0%
477	float param2 = (AABB.zMax -	3.0%	0.0%	0.0%
476	float param1 = (AABB.zMin -	4.4%	0.0%	0.0%

Tune threading with locks and waits analysis. Quickly find a common cause of slow performance in parallel programs waiting too long on a lock while the cores are underutilized during the wait.

Tuning OpenMP and Intel Threading Building Blocks is easier with the right data. See the cause of threading inefficiencies sorted by potential impact for accurate data and low overhead.

Optimize multi-rank hybrid MPI/OpenMP. Profile multiple MPI ranks selected in Intel Trace Analyzer and Collector. Sort by impact of improved OpenMP performance.

Process / OpenMP Region / Function / Thread / Call Stack	Effective Time by Utilization	Imbalance or Serial Spinning	Lock Contention	Comm. (MPI)	Other
heart_demo (rank 15)	99.641s	21.705s	0.009s	8.027s	6.780s
heart_demo (rank 17)	99.569s	21.650s	0.017s	8.012s	6.864s
[Serial - outside any region]	10.336s	15.719s	0.004s	6.602s	4.992s
make_rk_step\$omp\$parallel:	21.290s	1.418s	0.003s	0.286s	0.445s
make_rk_step\$omp\$parallel:	21.183s	1.431s	0.004s	0.341s	0.419s

Optimize data structures:

- Attribute cache misses to data structures (not just code lines).

Optimize NUMA latency and scalability:

- Tune true and false sharing.
- Tune inter-socket bandwidth.
- Tune Intel® Xeon Phi™ processor MCDRAM (high bandwidth memory).

Bandwidth Utilization Type / Function / Thread	CPU Time	L1 Bound	DRAM Bound	LLC Miss
Low	10.531s	0.133	3.3%	
grid_intersect	5.795s	0.154	3.3%	
sphere_intersect	3.282s	0.102	1.2%	
shader	0.135s	0.109	0.0%	
tri_intersect	0.059s	0.301	0.0%	
func@0x1401513f0	0.040s	0.000	58.0%	
func@0x10009c00	0.037s	0.147	0.0%	

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. 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

3

Specifications at a Glance

Processors	Intel processors, compatible processors, and Intel Xeon Phi processors and coprocessors
Languages	C, C++, C#, Fortran, Java, Python, Go, ASM, OpenCL, and more.
Compilers	Works with compilers from Microsoft, GCC, Intel, and others that follow standards
Development Environment	Integrated with Microsoft Visual Studio or runs standalone
Host Operating Systems	Windows, Linux, macOS
Target Operating Systems	Linux, Windows, FreeBSD, Android, and select embedded operating systems.
Basic Threading Analysis	Intel® OpenMP, Intel Threading Building Blocks, and native threads
Extended Threading Performance Analysis	Intel OpenMP, Intel Threading Building Blocks
MPI Parallelism	Application Performance Snapshot and integration with Intel Trace Analyzer and Collector MPI Profiler

“We achieved a significant improvement (almost 2x) even on one core by optimizing the code based on the information provided by Intel® VTune™ Amplifier. Good scalability is a result of usage of combination of Intel® TBB and OpenMP* parallelization techniques. We achieved over 8x the performance of the previous version on eight cores and almost 11x the performance on 16 cores.”

Alexey Andrianov
R&D Deputy Director
Mechanical Analysis Division
Mentor Graphics Corporation



[Learn more about Intel VTune Amplifier >](#)

[Get a Free 30-Day Evaluation >](#)

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'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

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