

# PRODUCT BRIEF

Performance Profiling  
Intel® VTune™ Profiler



# Optimize Software Performance for Modern Hardware

## Quickly Discover Performance Bottlenecks with Intel® VTune™ Profiler



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, get the data you need with Intel® VTune™ Profiler. Collect a rich set of performance data for hotspots, threading, locks and waits, DirectX\*, OpenMP\*, Threading Building Blocks, bandwidth, cache, memory access, non-uniform memory, storage latency, OpenCL™ applications, and more.

Profile C, C++, C#, Fortran\*, Python\*, Go\*, Java\* and OpenCL\*—or any mix—running on Intel® processors. Unlike single-language profilers, Intel VTune Profiler analyzes mixed code.

- **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.
- **Choose your workflow.** Select from local/remote collection or command line/graphical interface.

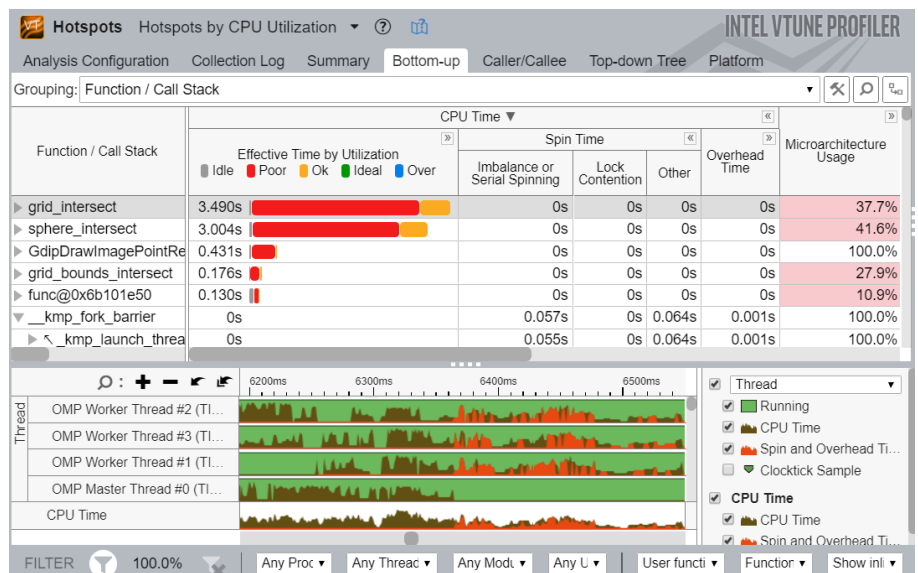


Figure 2. Hotspot analysis shows where your application's time is spent. It can also give a detailed analysis of threading performance showing the potential performance gain and common causes of poor performance such as imbalance, lock contention, forking, scheduling, and reduction.

## Powerful Data Analysis Saves You Time

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

### New for 2020

- **Memory Analysis.** Design and optimize for Intel® Optane™ DC persistent memory.
- **Platform Profiler.** Get a faster, better display of system metrics.
- **I/O Analysis.** PCIe\* device metrics
- **HPC Analysis.** Vectorization metrics, process and thread affinity, and Lustre\* parallel file I/O metrics (preview feature)
- **Snapshot.** Communication pattern diagnosis and Open MPI\*
- **GPU and OpenCL™ Application Analysis.** Inline filer and instruction count
- **Linux\*.** More analysis types enabled without adding drivers.
- **Performance Analysis Cookbooks.** User-friendly recipes.
- **Improved user interface.** New user? Take an overlay tour.
- **More descriptive name** (formerly Intel® VTune™ Amplifier).

## Flexible Options: Get It Standalone or Integrated in Comprehensive Development Tool Suites

Intel VTune Profiler is available as a standalone product and as part of other comprehensive software development tool suites.

- **Intel® Parallel Studio XE** is a development suite to help developers build high-performance, scalable, reliable parallel code for enterprise to cloud, and HPC to AI applications. The suite also includes compilers, libraries, and other analysis tools.
- **Intel® System Studio** Professional and Ultimate editions are used for system and IoT development of smart, connected devices. The suite also includes a compiler, libraries, analysis tools, and cloud connectors and provides access to 400+ sensors.

Get more analysis tools to complement Intel VTune Profiler:

- **Intel® Advisor** helps optimize vectorization, threading, and flow graphs. Roofline analysis finds loops with the most headroom for improvement.

## Product Details

### Quickly Locate Code Taking a Lot of CPU Time

Hotspots analysis gives you a sorted list of the functions 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.

Function / Call Stack	CPU Time				Spin Time	Overhead Time
	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

- **Intel® Trace Analyzer and Collector** examines MPI applications and tells Intel VTune Profiler which loops will benefit most from threading optimization.

## Get the Data You Need

- **Hotspot** (statistical call tree)
- **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, 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 OSs.
- **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.
- **Storage analysis:** Find I/O bottlenecks.
- **Filter out extraneous data** with the timeline and view-points.
- **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 Profiler 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 (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	Creation	Scheduling	Reduction
1,456	for( u32 j = rangeBegin; j < range	0.5%	0.0%	0.0%	0.0%	0.0%
1,457	{	0.0%	0.0%	0.0%	0.0%	0.0%
1,458	FireObject *pfo = m_pFireObj	0.4%	0.0%	0.0%	0.0%	0.0%
1,459	if( checkCollision(ttp, ttp,	5.4%	0.0%	0.0%	0.0%	0.0%
1,460	{	0.0%	0.0%	0.0%	0.0%	0.0%
1,461	// if it passes this test	0.0%	0.0%	0.0%	0.0%	0.0%

**Tuning OpenMP and 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**

Sort results by impact of improved OpenMP performance.

Process / OpenMP Region / Function / Thread / Call Stack	Effective Time by Utilization	Spin Time	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	
▶ make_rk_step\$omp\$parallel:	21.239s	1.320s	0.003s	0.350s	0.461s	

**Identify Which Memory Objects Are Bottlenecks**

A typical hotspot analysis shows code that is taking the most time. Memory Access analysis offers a different perspective. It shows which memory objects cause performance issues and consume bandwidth. This can yield new insight on how to improve performance.

Memory Object	Total Latency	Loads	Stores	LLC Miss Count
alloc_test.cpp:157 ( 30 MB )	65.6%	4,239,327,176	4,475,334,256	0
alloc_test.cpp:135 ( 305 MB )	6.8%	411,212,336	441,613,248	0
alloc_test.cpp:109 ( 305 MB )	6.3%	439,213,176	449,613,488	0
alloc_test!_data_init.436.0.6 ( 576 B )	5.2%	742,422,272	676,820,304	0
[vmlinux]	4.6%	173,605,208	116,003,480	0
[Others]	11.5%	1,533,646,008	1,674,450,232	0

*\*N/A is applied to non-summable metrics.*

**Optimize Data Structures**

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

**Optimize NUMA Latency and Scalability**

- Tune true and false sharing
- Inter-socket bandwidth

**Design and Optimize for Intel® Optane™ DC Persistent Memory**

Bandwidth Utilization Type / Function / Thread	CPU Time	L1 Bound	DRAM Miss
▼ Low	10.631s	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%

Bandwidth Domain: DRAM, GB/sec

Elapsed Time: 0s to 2s

Bandwidth Utilization: Low, Medium, High

## Get Priority Support for Confidential Technical Consultation

Every paid license of Intel® Software Development Products automatically gives you access to Priority Support via the Online Service Center for at least a year from the date of purchase. You can extend/renew it at a significantly reduced rate. Benefits include:

- **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
- **Quick responsiveness** to your product needs and technical questions (old and new versions)
- **Community product forums** covering all Intel Software Development Products
- **Access to a vast library of self-help documents** that build off decades of experience creating high performance code

## Specifications at a Glance

<b>Processors</b>	Intel® processors and compatible processors, Intel® Arria® 10
<b>Languages</b>	C, C++, C#, Fortran, Java, Python, Go, ASM, OpenCL, and more
<b>Compilers</b>	Works with compilers from Microsoft, GCC, Intel, and others that follow OS standards
<b>Development Environment</b>	Integrated with Microsoft Visual Studio or runs standalone
<b>Host Operating Systems</b>	Windows, Linux, macOS
<b>Target Operating Systems</b>	Linux, Windows, FreeBSD, Android, and select embedded operating systems
<b>Threading Analysis</b>	Intel® OpenMP, Intel Threading Building Blocks, and native threads
<b>MPI Parallelism</b>	Integration with Intel® Trace Analyzer and Collector MPI Profiler
<b>GPU</b>	OpenCL program and media application tuning on the latest Intel processors. Hotspot analysis for OpenCL kernels.

“We achieved a significant improvement (almost 2x) even on one core by optimizing the code based on the information provided by Intel® VTune™ Profiler. Good scalability is a result of using a 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 Profiler >**

**Get a Free 30-Day Evaluation >**

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 <http://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.

For more information regarding performance and optimization choices in Intel® Software Development Products, see our Optimization Notice: <https://software.intel.com/articles/optimization-notice#opt>

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 © 2019 Intel Corporation. All rights reserved. Intel, Xeon, Xeon Phi, VTune, Optane, 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

1119/SS

Please Recycle