



Intel[®] VTune[™] Profiler 2021

Release Notes

25 August 2021

Customer Support

For technical support, including answers to questions not addressed in this product, visit the technical support forum, FAQs, and other support information at:

- <https://software.intel.com/en-us/vtune/support/get-help>
- <http://www.intel.com/software/products/support/>
- <https://software.intel.com/en-us/vtune>

Register your product at <https://registrationcenter.intel.com/> by providing your email address. Registration entitles you to free technical support, product updates and upgrades for the duration of the support term. It also helps Intel recognize you as a valued customer in the support forum.

NOTE: If your distributor provides technical support for this product, contact them directly for support issues.

Contents

| | | |
|-------------------|---|--------------------|
| 1 | Introduction | 2 |
| 2 | What's New | 3 |
| 3 | Where to Find the Release | 11 |
| 4 | Installation Notes | 11 |
| 5 | Known Issues | 12 |
| 6 | Attributions | 27 |
| 7 | Legal Information | 27 |

1 Introduction

Intel® VTune™ Profiler provides an integrated performance analysis and tuning environment with a graphical user interface that helps you to analyze code performance on systems with Intel® 64 architectures.

This document provides issues, limitations, and legal information pertaining to Intel VTune™ Profiler.

Intel VTune™ Profiler for Windows and Linux OS has a standalone graphical user interface (GUI) as well as a command-line interface (CLI). VTune™ Profiler for macOS* supports viewing of results collected on other platforms. Native collection on a macOS* system is currently not available.

To learn more about Intel VTune™ Profiler, see:

- New features listed in the [What's New](#) section below.
- [System requirements](#).
- [Help Document](#)
- Installation Guide for:
 - [Linux*](#)
 - [macOS*](#)
 - [Windows*](#)

2 What's New

Detailed [new features](#).

Intel® VTune™ Profiler 2021.7.0

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- **Algorithm Group**
 - **Flame Graph View in Hotspots Analysis**

This version of VTune Profiler introduces support for Flame Graphs in the Hotspots analysis type. The [Hotspots by CPU Utilization viewpoint](#) has been enhanced with a [Flame Graph window](#) that displays a graphical view of hot code paths. Use flame graphs to analyze the time spent on each function and its callee functions.
- **Debug Formats**
 - **Support for DWARF 5 Debug Format**

VTune Profiler now supports version 5 of the DWARF debug format. You can now use debug information in DWARF 5 format to resolve function names and source locations for binaries.
- **GPU Accelerators**
 - **CPU Context for GPU Execution in GPU Offload Analysis**

The [GPU Offload Analysis](#) now presents a richer set of information about execution on the GPU by including context from the CPU.

This includes stack information on:

- Execution
- Data transfer from host to device
- Data transfer from device to host

The viewpoint for the [GPU Offload Analysis](#) now includes the [Call Stack pane](#) with a new grouping by GPU Computing Task/Host Call Stack. Navigate through transfer data contained in this pane to identify inefficient code paths in your application.

- **Analysis of Multiple GPUs**

When you have multiple GPUs connected to your system, you can now analyze all of the GPUs collectively with the [GPU Offload](#) and [GPU Compute/Media Hotspots](#) analyses. Previously, you could analyze a single GPU at a time after VTune Profiler identified all the GPUs connected to the system. When you run these analyses on all connected GPUs, see analysis information about each GPU in the Summary window. Full compute set in **Characterization** mode is not available in multi-adapter and multi-tile analysis.

Intel® VTune™ Profiler 2021.6.0

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- **Command Line Analysis**
 - **Perf Tool Parameters for All Analysis Types**

You can now use the [target-system](#) command to get parameters on the command line for the native perf tool for all CPU hardware-based analysis types, including custom analyses. Use the **get-perf-cmd** argument for this purpose. You can collect the perf trace on a target with the Linux Perf tool and then import the trace to the VTune Profiler UI.
- **GPU Accelerators**
 - **Occupancy Report in GPU Hotspots Analysis**

The [GPU Compute/Media Hotspots](#) analysis has been enhanced to display [occupancy information](#) in the Summary section. Use this data to understand the architectural limitations of the GPU that affect occupancy.
- **Microarchitecture Analyses**
 - **Platform Diagram in Memory Usage View**

This release introduces the Platform diagram in the [Memory Usage viewpoint](#) of the Memory Access analysis type. Use this diagram to understand:

 - System topology
 - Utilization metrics for DRAM
 - Intel® UPI links
 - Physical cores

The platform diagram is available for:

 - All client platforms
 - Server platforms based on Intel® microarchitecture code name Skylake, with up to four sockets.

Intel® VTune™ Profiler 2021.5.0

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- **VTune Profiler – Platform Profiler as Analysis Type**

Intel VTune Profiler – Platform Profiler has been completely integrated into VTune Profiler as an analysis type. Platform Profiler is now fully available as an analysis from the GUI or command line of VTune Profiler. For more information, see [Platform Profiler Analysis](#).
- **Application Performance Snapshot**
 - **Outlier Detection**

This release introduces a mechanism for the detection of outliers, or individual metric values that contribute to an average metric that differ significantly from the overall distribution or break a certain threshold. Outliers can cause imbalance and distort average metric values. You can now see outliers in both HTML and CLI reports, with attribution to specific rank or node where an outlier occurred.
 - **Metric Tooltip Enhancements**

Metric tooltips now visualize ranges of average metrics, with their minimum, maximum, and average contributing values.
- **MPI Applications Support**
 - **Support for MPI applications in GPU and IO analyses**

The [GPU Offload](#), [GPU Compute/Media Hotspots](#), and [Input and Output](#) analyses now support profiling of MPI applications, as described in the [MPI Code Analysis](#) topic.

Intel® VTune™ Profiler 2021.4.0

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- **Improvements to Vectorization Information**

The Vectorization sections of Performance Snapshot and HPC Performance Characterization analyses have been enriched to provide a clearer picture of the state of vectorization in your application. Quickly see if your code is not vectorized at all, if your code does not use the latest vector instruction set extension, or if your code has too many scalar instructions. This version of VTune Profiler also features improved recommendations to resolve vectorization issues.

- **Support for the 3rd Gen Intel® Xeon® Scalable processors (code name Ice Lake Server)**

This release introduces full support for the Ice Lake Server architecture in the Input and Output analysis.

- **GPU Analyses improvements**

- **Advanced Data Transfer Information in GPU Offload**

The following additions to the Graphics window clarify better the data transfer between CPU host and GPU device when you run GPU profiling analyses.

- Allocation time information displays as part of total time by device operation.
- Data Transferred table has been renamed as Transfer Size table. Columns under Transfer Size feature new names for data transferred between host and device.
- Highlights and tooltips for workloads with sub-optimal offload schemes direct your attention to improve offload schema where necessary.

- **Advanced analysis of Low occupancy problems for GPU computing tasks**

The GPU Compute/Media Hotspots analysis is now able to automatically detect what limits peak achievable occupancy for the hottest computing tasks that make EU Array Idle while waiting for the scheduler. Improved tooltips for Occupancy Metrics provide information about peak occupancy and bounding reasons for existing computing task launch configuration.

- **User experience**

- **The Project Navigator User Experience**

- The Project Navigator pane now features menu options to open a new or existing project to better facilitate your VTune experience.

- **GPU Analysis Coverage for Self-Check**

Coverage of checks by the self-check functionality in VTune Profiler now includes GPU analyses as well. Run [vtune-self-checker.sh script](#) on Windows and Linux systems to check for the GPU Compute/Media Hotspots Analysis in source analysis and characterization modes for DPC++ application running on GPU (requires oneAPI Base toolkit installation).

Intel® VTune™ Profiler 2021.3.0

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- **Application Performance Snapshot**

- **Rich metric tooltips in HTML reports**

Metric tooltips in APS HTML reports have been overhauled to present a holistic view of metrics and their properties. The new tooltips present a compact yet comprehensive overview of a metric, which

helps you to better understand the importance of metrics in performance analysis.

This change includes a visual bar that indicates where the metric value stands in terms of current performance and tuning potential.

- **PCIe bandwidth info in CLI reports**

APS command line reports now include PCIe bandwidth metrics. This data is only available on server platforms when using the Sampling Driver.

- **New reports and filters**

APS now features the following new types of reports and filters:

- Node topology report: view relations between ranks, nodes, and PCIe devices.
- Metrics report: get a configurable table that displays any collected metric for each rank, node, or device.
- Ability to filter data by node.

- **Support for Code Targets**

- This release introduces support for running the Hotspots analysis on .NET 5 targets in Launch Application mode when using hardware event-based sampling.

- **Hardware Support**

- This version of VTune Profiler includes support for Intel processor architecture code named Ice Lake Server including Microarchitecture Exploration Analysis.

Intel® VTune™ Profiler 2021.2.0

This version of Intel® VTune™ Profiler contains improvements and additions in the following:

- **User Interface**

This release introduces a new [main vertical toolbar](#) to enhance your user experience. All controls previously located in the main horizontal toolbar are now located on this toolbar.

- **Hardware Support**

This version includes support for Intel Atom® Processor P Series code named Snow Ridge, including Hotspots, Microarchitecture Exploration, Memory Access, and Input and Output analyses.

- **GPU Accelerators**

- **Source-level analysis for DPC++ and OpenMP applications running on GPU over Level Zero**

The following modes in GPU Compute/Media Hotspots analysis are now available when profiling Level Zero applications:

- [Dynamic Instruction Count](#)
- [Basic Block Latency](#)
- [Memory Latency](#)

Support also includes full-scale analysis of the kernel source per code line, including Source/Assembly mapping.

- **Input and Output Analysis**

- **New major features in Input and Output analysis**

- This release introduces the **Platform Diagram**, a new starting point for the [Input and Output analysis](#). It reveals system topology and high-level utilization metrics for hardware resources including PCIe devices, Intel® Ultra Path Interconnect, and memory. It enables you to examine the utilization of your hardware at a glance.

This feature is enabled for 1st and 2nd Generation Intel® Xeon® Scalable Processors in up to four-socket configurations, excluding the Intel® Xeon® Platinum 9200 series processors code named Cascade Lake AP. This feature is also supported on Intel Atom® Processors P Series code named Snow Ridge.

- **Intel® Data Direct I/O (Intel DDIO) utilization efficiency metrics** are extended with average Inbound PCIe read/write latency and core/IO contention indicator.

- It is now possible to perform **Linux perf-based data collection** without root access on 1st and 2nd Generation Intel Xeon® Scalable Processors on Linux kernel versions 5.10 and newer.

Intel® VTune™ Profiler 2021.1.2

This version of Intel® VTune™ Profiler contains improvements and additions in the following:

Fixed a problem of command line analysis based on User-Mode Sampling that didn't work under a non-root account.

Details: If VTune Profiler was installed by a root/sudo user, some executable files were requiring only root permission to run an analysis based on User-Mode Sampling collector, e.g. Hotspots.

Intel® VTune™ Profiler 2021

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

GPU Accelerators:

- **GPU Adapter Selection for Profiling Analyses in Multi-GPU Systems**
When you have multiple Intel GPUs connected to your system, you can now select a specific GPU adapter directly in the user interface for your GPU Offload Analysis or GPU Compute/Media Hotspots Analysis. The **Target GPU** pulldown menu appears in the HOW pane of the analysis configuration when Intel VTune Profiler detects multiple Intel GPUs on your system. The menu lists available GPU adapters with their Bus/Device/Function (BDF) values.
- **Energy Consumption Metrics in GPU Compute/Media Hotspots Analysis**
When you run the GPU Compute/Media Hotspots Analysis on an Intel® Iris® Xe MAX graphics discrete GPU in a Linux environment, you can now use the Analyze power usage option to collect information about energy consumed by the GPU. The analysis results display energy consumption metrics over time and per discrete GPU kernel. Use this data to better monitor power usage with processing time and optimize for either purpose.

FPGA Accelerators:

- **Multiple enhancements to CPU/FPGA Interaction Analysis**
The CPU/FPGA Interaction analysis type features several new additions to enhance your FPGA profiling experience.
 - Analysis results now display **Activity percentage** and **Idle percentage** metrics to describe the proportion of cycles when a channel instruction was enabled or absent.
 - The analysis type can now profile loops and display occupancy information.
 - You can now adjust the depth of channels using **Average depth** and **Maximum depth** information that displays in the analysis results.

Energy Analysis:

- **Rootless Data Collection on Linux Systems**
You do not require root privileges to run energy analysis using Intel VTune Profiler in a Linux environment. You can run this analysis without root privileges once your system administrator installs sampling drivers for VTune Profiler and configures relevant permissions for the drivers. Administrator privileges are required to collect energy data in Windows machines.

Intel® VTune™ Profiler 2020 Update 3

Hardware Support:

- Added support for 11th Gen Intel® Core™ processors codenamed Tiger Lake, including Hotspots, Microarchitecture Exploration, Memory Access, and GPU analyses.

Input and Output Analysis:

- Source-level Memory Mapped I/O (MMIO) analysis now supports InfiniBand devices.

Profiling Applications Annotated with ITT API:

- Profiling of applications annotated with ITT API has been enhanced by the introduction of additional **Average Task Time** and **Average Frame Time** metrics.

Profiling Remote Amazon Web Services* Instances:

- Added support for remote profiling of applications running in Amazon Web Services* (AWS) EC2 instances.

Intel® VTune™ Profiler 2020 Update 2

- Intel® VTune™ Profiler has been updated to include more recent versions of 3rd party components, which include functional and security updates. Users should update to the latest version.

Performance Snapshot analysis type for quick summary:

- This release introduces the [Performance Snapshot](#) analysis type. Start with this analysis and get a quick overview of issues that affect your application performance. Performance Snapshot characterizes the workload on the system. It also provides recommendations for next steps to help you select other analyses for deeper profiling.

Platform Analysis:

- Platform I/O metrics can now be attributed to individual devices managed by [Intel® VMD technology](#)
- I/O Analysis has been enhanced for Sky Lake and Cascade Lake servers by highlighting code potentially performing MMIO reads

New hardware/operating systems/IDEs support:

- Support for Intel's processors code named Cooper Lake and Comet Lake
- Ubuntu* 20.04, Fedora* 32
- Microsoft* Windows* 10, May 2020 Update

Deprecation:

- We recommend that Storage Snapshot users switch to using the Platform Profiler feature of Intel® VTune™ Profiler. It provides a more informative set of data with a similar low overhead. We are deprecating Storage Snapshot and will discontinue it in our next major release.

- Preview of [Input and Output](#) analysis on Windows* is deprecated and will be removed in a future release. This analysis continues to be supported on Linux* OS.

Intel® VTune™ Profiler 2020 Update 1

This version of Intel® VTune™ Profiler contains improvements and additions in these areas:

- Microarchitecture Exploration analysis is now supported on Intel processors codenamed Icelake.
- GPU accelerators support:
 - [GPU Compute/Media Hotspots](#) analysis in the Dynamic Instruction Count mode has been extended to include SIMD utilization metrics at the kernel and instruction level. These metrics help identify instructions in the OpenCL™ kernel that utilize SIMD poorly.
 - A deeper GPU utilization analysis has been introduced in Application Performance Snapshot (APS) and the HPC Performance Characterization analysis. The GPU utilization analysis now includes these GPU computation metrics:
 - GPU Time
 - GPU IPC
 - GPU Utilization
 - % of Stalled and Idle EUs.
 - There is now a simplified dependency on the Intel® Metrics Discovery Application Programming Interface library to collect GPU hardware statistics on Linux* systems. VTune Profiler now automatically selects the latest libstdc++ available in runtime to satisfy the GPU analysis requirements. For older versions of the product, follow procedures to enable [manual configuration](#).
- Platform analysis improvements:
 - [CPU/FPGA Interaction analysis](#) has been extended to process data sources collected either with AOCL Profiler (new mode) and via OpenCL Profiling API (legacy mode). Specify the name of your application target and its parameters directly in the WHAT pane.
 - The [Hardware Tracing mode](#) in the System Overview analysis has been extended to include new metrics to make the analysis more kernel-aware:
 - OS Kernel Activity and OS Scheduling metrics identify anomaly issues caused by unexpected kernel activity or preemptions.
 - The CPU Time metric has been split to User Time and Kernel Time metrics to show the number of kernel mode switches and their frequency (switches per second).
 - A new **Module Entry Point** grouping level has been added to the System Overview viewpoint to display data collected in the Hardware Tracing mode. The grouping shifts the focus to precise CPU time spent within sys calls, interrupts or within particular API of the runtime library.

Intel® VTune™ Profiler 2020

- There is a new, more descriptive name: “Intel® VTune™ Profiler” (formerly “Intel® VTune™ Amplifier”).
 - Command line interface `amplxe-cl` and GUI interface `amplxe-gui` were re-named to `vtune` and `vtune-gui` respectively
- GPU accelerators support:

- New [GPU Offload](#) analysis added to explore and correlate code execution across CPUs and GPUs. You can identify a kernel of interest for GPU-bound applications and explore further with GPU Compute/Media Hotspots analysis.
- [GPU Compute/Media Hotspots](#) analysis updated with GPU in-kernel analysis for OpenCL™ code and an option to filter by a kernel of interest.
- Command line [hotspots](#) report now supports GPU analysis types. You can apply the `computing-task` and `computing-instance` groupings to your collected data to focus on time-intensive computing tasks.
- [Dynamic instruction count](#) collection (available as part of the GPU Compute/Media Hotspots Analysis) improved to provide better accuracy for basic block Assembly analysis.
- Support for Intel® Processor Graphics Gen11.
- Platform analysis support:
 - [System Overview](#) analysis updated to serve as an entry point to platform analysis. Use this analysis to assess system (IO, accelerators and CPU) performance and review guidance for next steps.
 - New [Hardware Tracing mode](#) in the System Overview analysis enables application analysis on the micro-second level and identification of causes for latency issues.
- HPC analysis improvements:
 - Max and Bound Bandwidth metrics added to Application Performance Snapshot to better estimate the efficiency of the DRAM, MCDRAM, Persistent Memory and Intel® Omni-Path usage.
- Platform Profiler new features and improvements:
 - Overview and Memory views extended with new metrics to analyze Non-Uniform Memory Access (NUMA) behavior.
 - User authentication and authorization implemented to enable access control to user data.
 - Added a new option for users to choose or modify the location of Platform Profiler data files.
- Energy analysis improvements:
 - New [Throttling analysis](#) added to identify causes for system throttling, including violation of safe thermal or power limits.
 - Options for [Energy analysis](#), based on the Intel SoC Watch data collector, extended to monitor processor package energy consumption over time and identify how it correlates with CPU throttling.
- Cloud and containerization support:
 - [Containerization support](#) extended with an option to install and run VTune™ Profiler in a Docker* container and profile targets inside and outside the same container.
 - Added support to profile applications running in Amazon Web Services* (AWS) EC2 Instances based on Intel microarchitecture code name Cascade Lake X.
- New [Fabric Profiler](#) performance tool added to VTune™ Profiler in Preview mode. Use Fabric Profiler to identify detailed characteristics of the runtime behavior for an OpenSHMEM application.
- Quality and usability improvements:
 - Symbol resolution for effective source-level analysis enabled for crossgen (Ahead-of-JIT compilation) functions on Linux* systems.
 - Interactive **Help Tour** (available on the Welcome page) guides you through the product interface using a sample project.
- New hardware/operating systems/IDEs support:
 - 10th Gen Intel® Core™ processors

- Ubuntu* 19.10
- Microsoft* Windows* 10, November 2019 Update

3 Where to Find the Release

If you have a qualifying license for this product (i.e. an unexpired free license or a paid license with active support at the time of the build date for this product release), you can download the installer by logging in to <https://registrationcenter.intel.com/> and selecting the appropriate product.

You may need to create an account and/or register your product. For additional information please see:

- <https://software.intel.com/en-us/faq/registration>
- <https://software.intel.com/en-us/faq/downloading>

4 Installation Notes

For instructions on installing, updating, or removing Intel® VTune™ Profiler, see the appropriate installation guide for your operating system:

- Linux*: <https://software.intel.com/en-us/VTune-Install-Guide-Linux>
- macOS*: <https://software.intel.com/en-us/VTune-Install-Guide-macOS>
- Windows*: <https://software.intel.com/en-us/VTune-Install-Guide-Windows>

System requirements

Detailed [system requirements](#).

Intel® Software Manager

Intel® Software Manager installs with Intel Software Development Products on Windows* operating systems. Intel Software Manager is a utility that allows you to perform these tasks:

- Download and install updates for your Intel® Software Development Products.
- Manage subscription status of installed software.
- Activate serial numbers.
- Find out about the latest news for Intel Software Development Products.

Intel Software Manager requires an Internet connection to connect to a remote server for information and updates.

For more information about the Intel Software Manager, see <https://registrationcenter-ssl.intel.com/Docs/ism.htm> .

Default Installation Folders

The default top-level installation folders for this product are:

- Windows: C:\Program Files (x86)\Intel\oneAPI\vtune\latest

- If you are installing on a system with a non-English language version of Windows* OS, the name of the `Program Files` folder may be different.
- Linux: `/opt/intel/oneAPI/vtune/latest`

In both cases, there is also a generic soft link to the most recently installed version.

Setting Up the Command Line

Before using Intel® VTune™ Profiler from the command line, you must set up the environment. Run the appropriate script for your OS

- Windows: `<installation_dir>\vtune-vars.bat`
- Linux: `source <installation_dir>/vtune-vars.sh`

NOTE: Commands to invoke the command line interface (`amplxe-cl`) and GUI interface (`amplxe-gui`) have been renamed to `vtune` and `vtune-gui` respectively.

5 Known Issues

Use these links to see OS-specific limitations:

- [Windows* Specific Limitations](#)
- [Linux* Specific Limitations](#)
- [macOS* Specific Limitations](#)

Common Limitations

RUNNING TIME IS ATTRIBUTED TO THE NEXT INSTRUCTION (VASP-3854)

To collect data about time-intensive running regions of the target, Intel® VTune™ Profiler interrupts the execution of target threads and attributes the time to the context IP address.

Due to the collection mechanism, the captured IP address points to the instruction *after* the most time-intensive one. This causes the running time to get attributed to the next instruction (or sometimes to subsequent instructions) in the Assembly view. In rare cases, this can also lead to incorrect attribution of running time in the source. The time can be erroneously attributed to the source line *after* the actual hot line.

If the inline mode is ON and the program has small functions inlined at the hotspots, this can cause the running time to be attributed to the wrong function since the next instruction can belong to a different function in tightly inlined code.

VTUNE CRASHES ON FREEBSD 11/12 AFTER PRESSING CTRL-C DURING COLLECTION (VASP-17676)

Upon termination of collection on FreeBSD with a Ctrl-C operation, the next VTune hardware event-based collection fails with this error- PMU resource(s) currently being used by another profiling tool or process: another PMU collection session may be running.

As a workaround, use `-command stop` to stop the collection.

COLLECTION MISSING DATA FOR CORE_0 ON SKYLAKE-X SYSTEMS (VASP-11480)

On some systems based on Intel microarchitecture code name Skylake-X with enabled SMT mode, VTune™ Profiler hardware sampling analysis can collect data against less logical CPUs than available on the system.

For example, the problem was reproduced on Intel® Core™ i9-7900X 10C20T CPU installed on an Asus* PRIME X299-DELUXE motherboard.

As a workaround, upgrade the BIOS and the firmware of the board to 0802 or a newer version.

AN APPLICATION THAT ALLOCATES MASSIVE CHUNKS OF MEMORY MAY FAIL TO WORK UNDER VTUNE™ PROFILER

If a 32-bit application allocates massive chunks of memory (close to 2 GB) in the heap, it can fail to launch with VTune™ Profiler but execute independently without error. This happens because VTune™ Profiler requires additional memory when profiling an application. As a workaround, use a larger address space. For example, convert the project to 64-bit.

HARDWARE EVENT-BASED ANALYSIS MAY CRASH CERTAIN INTEL® CORE™ I7 PROCESSOR-BASED SYSTEMS WHEN DEEP SLEEP STATES ARE ENABLED

On some Intel® Core™ i7 processor-based (based on Intel microarchitecture code named Nehalem) systems with C-states enabled, sampling may cause system hanging due to a known hardware issue. To avoid this, disable the “Cn (ACPI Cn) report to OS” BIOS option before sampling with the VTune™ Profiler on Intel Core i7 processor-based systems.

INSTRUCTION SET REFERENCE DOCUMENT IS NOT POSITIONED ON DESCRIPTION OF PROPER INSTRUCTION

The reference information for assembly instructions can be opened in any PDF viewer, but only Adobe* Acrobat Reader* supports the positioning of the instruction reference document on the required page. To ensure correct functionality of this feature, install the latest version of Adobe Acrobat Reader.

VTUNE™ PROFILER DOES NOT SUPPORT PROFILING AN APPLICATION LAUNCHED UNDER DEBUGGER

The Hotspots, Threading, Memory Access with object tracking, and Input and Output analysis types provide incorrect results when analyzing an application launched under a debugger. VTune™ Profiler does not detect whether a debugger is attached to a profiled application. Make sure that no debugging tools are attached to the application that is profiled with the VTune™ Profiler.

ONLY ONE STACK FRAME IS SHOWN IN JAVA* CODE IF IBM* J9* JVM IS USED (VASP-3535)

Currently, IBM* J9* JVM does not support Java* stack unwinding.

RESULTS NOT COLLECTED FOR PROCESSES TERMINATED WITH `TERMINATEPROCESS ()`

No results are displayed for a profiled process that is terminated by a different process calling `TerminateProcess ()`. Use an alternate method to terminate the process.

DO NOT USE `-IPO (LINUX*)//QIPO (WINDOWS*)` OPTION WITH INTEL COMPILER VERSION EARLIER THAN 13.0. (VASP-3534)

If you are using a version of Intel® compiler that is 13.0 or older, use the additional option `-inline-debug-info (Linux*)//debug:inline-debug-info (Windows*)` to get performance data on inline functions. However, avoid using the `/Qipo` option. Currently this option disables the generation of the inline debug information in the compiler. Note that the Intel compiler integrated into the Microsoft Visual Studio* IDE uses the `-ipo//Qipo` by default in the Release configuration.

JAVA* SOURCE LINE MAY BE SHOWN INCORRECTLY FOR ORACLE JDK* 6U25 X64

Drilling down to Java* source code from VTune™ Profiler results may show incorrect source line. The issue occurs with Java* applications built with Oracle JDK* 6u25 x64. This is a JDK* problem. Refer to issues JDK-7051767 and JDK-7047624 in the Oracle bug database.

RESULT FINALIZATION MAY BE SLOW IF RESULTS ARE LOCATED IN A REMOTE NETWORK DIRECTORY (VASP-995)

The problem is caused by slow network access. To improve performance, use local directory for result storage.

CALL STACK CANNOT BE UNWOUND VIA NO RETURN FUNCTIONS (VASP-3350)

If the analyzed application contains functions without `ret` instruction (for example, calling `exit()`), unwinding the call stack to its caller and higher may fail (no stack frames shown).

ATTACHING THE VTUNE™ PROFILER TO A PROCESS MAY TAKE SOME TIME (VASP-2289)

Sometimes, your target application may complete execution before data collection begins. In this case, VTune™ Profiler may report an error. Increase the duration of your test application if necessary.

Command line option "`vtune -command status`" is currently supported only for user-mode sampling and tracing based analysis types, and not for EBS analysis types.

HARDWARE EVENT-BASED ANALYSIS MAY CAUSE UNPREDICTABLE SYSTEM BEHAVIOR ON PROCESSORS BASED ON INTEL MICROARCHITECTURE CODE NAMED SANDY BRIDGE

On processors based on Intel microarchitecture code name Sandy Bridge, hardware event-based sampling may cause unpredictable system behavior due to a known hardware issue (see erratum BK105 in <http://www.intel.com/content/dam/www/public/us/en/documents/specification-updates/2nd-gen-core-family-mobile-specification-update.pdf>). To avoid this, avoid running Microarchitecture Exploration (formerly, General Exploration) or Custom Hardware event-based analysis with precise events on these systems unless a BIOS workaround for the erratum is provided.

VTUNE™ PROFILER MAY DETECT WHICH TIMER SOURCE TO USE INCORRECTLY ON INTEL® XEON® PROCESSOR E5-XXXX PROCESSORS

In User-mode Sampling and Tracing based analysis (Hotspots, Threading), the following message may be displayed in command line or Debug Window:

Warning: Cannot load data file '<path>\data.0\18385-18389.0.trace' (SampleCallback: timestamps aren't ascended!).

As a workaround, collect results from the command line using the `-run-pass-thru` option:

```
vtune -collect hotspots -run-pass-thru=-timestamp=sys -- <application>
```

HARDWARE EVENT-BASED ANALYSIS DOES NOT WORK IF MORE THAN 128 EVENTS ARE GATHERED SIMULTANEOUSLY (VASP-1992)

Decrease number of PMU events in analysis settings to resolve it.

ITT API TASK OR FRAME IS NOT SHOWN IN RESULTS IF ITS END OCCURS WHEN COLLECTOR IS INACTIVE (VASP-3354)

When ITT task or frame end notification occurs after collection was paused, or detach or stop command issued, the frame or task is not displayed in the result.

AN APPLICATION MAY EXPERIENCE THE STACK OVERFLOW EXCEPTION WHEN RUNNING UNDER THE VTUNE™ PROFILER (DOC-200)

An application allocating massive chunks of memory on a thread stack may experience the stack overflow exception and fail when running under the VTune™ Profiler, while running flawlessly on its own. This happens because VTune™ Profiler requires additional space on an application thread stack for profiling needs. To work around this problem, consider using larger thread stack space.

VTUNE™ PROFILER CANNOT LOAD DRIVERS INTO KERNEL

Most Android* kernels only allow signed drivers. If you try and install a driver (`sep3_10.ko`, `pax.ko`, `vtsspp.ko`, `SOCWATCH1_3.ko` or `apwr3_1.ko`), which was built after the kernel was built, you will get an error message stating that no signing key is found. To work around this problem, you must build drivers with the exact signing key created when the Android* OS (`boot.bin`) was built. Use drivers which have been integrated into the kernel.

STARTING PLATFORM PROFILER SERVER CAUSES SERVER CREDENTIALS ERROR: "DATABASE SERVER CREDENTIALS ARE INVALID OR CORRUPTED" DXET-1990)

Platform Profiler server is unable to connect to the backend of the database due to an unexpected change in database credentials.

To work around this problem:

- Verify that you have read/write access to the directory where Platform Profiler data is stored.
- Verify that the Platform Profiler server is not running.
- Open a Terminal window (Linux) or a Command Prompt window (Windows*) and type the following command:

```
$ <vtune install dir>/vpp/server/vpp-server reset-db-credentials
```
- Start the Platform Profiler server

DATA COLLECTED WITH PREVIOUS VERSIONS OF PLATFORM PROFILER IS NOT AVAILABLE AFTER UPGRADE TO VTUNE™ PROFILER 2020 (DXET-2014)

Platform Profiler 2020 allows users to choose the location for data storage. Previous versions of Platform Profiler stored the data in these locations:

- Linux: `/opt/intel/vtune-Profiler-platform-profiler/vpp-data`
- Windows: `C:\ProgramData\Intel\VTune Profiler Platform Profiler\data`

To view the data collected with previous versions of Platform Profiler:

- **Linux:** Start the Platform Profiler server with the `--data-dir` option:

```
<vtune install dir>/vpp/server/vpp-server start --data-dir /opt/intel/vtune-Profiler-platform-profiler/vpp-data/
```
- **Windows:** In the Platform Profiler server startup dialog, click the "Browse Directory" button. Select "C:\ProgramData\Intel\VTune Profiler Platform Profiler\data" and then click "Start" to start the sever.

PLATFORM PROFILER IS UNABLE TO UPLOAD RESULTS OR VIEW RESULT DIRECTORIES AND FILES WITH GOOGLE CHROME* VERSION 80 OR LATER (DXET-2045)

Users of Google Chrome* version 80 or later are unable to browse Platform Profiler collected results or upload new result files. Clicking the "View Results" on the Platform Profiler home page displays an empty page.

To work around this problem:

- In Google Chrome*,
 - Navigate to the "New Results" page.
 - Press F12 to open the Chrome Developer tools.
 - Click and hold the "refresh" toolbar button.
 - Select "Empty Cache and Hard Reload" from the menu.
- Alternatively, you can also use Mozilla Firefox* or Microsoft* Edge. Platform Profiler is fully functional with Mozilla Firefox* and Microsoft* Edge. However, some some degradation in user experience (e.g. topology diagrams are not fully rendered) is likely.

SIZE LIMITATIONS WHEN UPLOADING DATA COLLECTION RESULTS TO THE PLATFORM PROFILER SERVER (DXET-1200)

The Platform Profiler server currently supports data collection results up to 1GB in size (compressed).

CPU TIME IN EXCESS OF 100% SHOWN FOR INDIVIDUAL THREADS (VASP-21841)

For data collected using the sampling driver on both Windows OS and Linux OS, the timeline may show CPU Time larger than 100% for individual threads, typically for the duration of one sampling interval at a time. This issue is caused by the thread migrating between logical CPUs.

As a workaround, you can switch to hardware sampling with stacks on both Windows OS and Linux OS or use driverless per-process sampling on Linux OS.

PYTHON CODE PROFILING LIMITATION: WHEN COLLECTING DATA IN ATTACH TO PROCESS MODE, PYTHON USER CODE IS NOT SHOWN IN VTUNE SOURCE VIEW FOR PYTHON* 3.6 VERSION AND HIGHER.

The limitation is valid starting from Python* 3.6.

APS TOOL FROM VTUNE PROFILER 2021 PACKAGE LOSES APPLICATION ARGUMENTS

You might observe a problem that Application Performance Snapshot (APS) tool loses application arguments while starting a data collection.

The issue affects the following versions: Intel® VTune™ Profiler 2020 Update 3, Intel® VTune™ Profiler 2021 Beta Update 10, and future Intel® VTune™ Profiler 2021.

As a workaround you can create a small shell script where to place a command line of the application with parameters and then launch the script from APS tool.

VTUNE PROFILER'S GPU OFFLOAD OR GPU COMPUTE/MEDIA HOTSPOTS ANALYSIS MAY FAIL ON COLLECTING DATA AT ANDROID TARGET DEVICE (VASP-22213)

With GPU Offload or GPU Compute/Media Hotspots analysis types for a target Android device you need to disable the 'Trace GPU programming APIs' option. Otherwise the collection will fail with error: 'Unknown option: --gpu-kernels-to-profile=*#1#1#4294967295'. With command line interface use the 'collect-programming-api' knob to disable this option, for example: `vtune -collect gpu-offload -knob collect-programming-api=false ...`

Windows* Specific Limitations

VTUNE 2020 U1 SHOWS EMPTY WINDOW WHEN A NEW RESULT IS OPENED AFTER SYSTEM AWAKE (VASP-19132)

Open VTune 2020 U1 GUI, open a result, put system on sleep and login back. In VTune GUI, if a new result is opened then it shows an empty window. Workaround for this issue is to restart VTune GUI.

JIT WRITER API CANNOT HANDLE TOO LONG FILE NAMES ON WINDOWS* (VASP-484)

This issue happens when file name of JIT trace file is longer than 260 symbols. VTune™ Profiler is unable to generate a proper result when the result directory has such a long name.

THREAD STACK SIZE LIMITATION (VASP-1808)

VTune™ Profiler may crash with the following error message:

Error: failed to create a sampling thread: not enough storage is available to process this command.

This happens when the profiled application has big reserved and committed thread stack size (/STACK:reserve[,commit] command line switch of `link.exe`). The suggested workaround is reducing the reserved/committed thread stack size to profile the target.

INCORRECT TIMING RESULTS WHEN RUNNING ON A 32-BIT VIRTUAL MACHINE

VTune™ Profiler may fail to collect correct timing data when running on a virtual machine with problematic virtualization of time stamp counters. In this case the VTune™ Profiler throws a warning message:

Warning: Cannot load data file '<path_to_a_trace_file>.trace' (syncAcquiredHandler: timestamps aren't ascended!)

UNINSTALLING LIMITATION: PIN.EXE STAYS RUNNING AFTER DETACHING.

Once attached to a target, VTune™ Profiler cannot be uninstalled until execution on the target is over. The cause is that `pin.exe` keeps working after detaching from the target and exits only after the profiled application/process execution finishes.

SECOND ATTACH TO THE SAME APPLICATION SHOULD PRINT AN ERROR AND EXIT IMMEDIATELY

VTune™ Profiler allows you to run an analysis while the previous one is in progress. However, it does not store any data from the second analysis run.

TRUNCATED .NET* MODULE NAMES MAY BE DISPLAYED IN RESULTS VIEW (VASP-3772)

When viewing results collected for a .NET application, you may observe truncated .NET module names. Ensure that the system was rebooted after installing the .NET application and before profiling with the VTune™ Profiler.

.NET MODULES LOADED DYNAMICALLY VIA REFLECTION API ARE DISPLAYED AS "UNKNOWN" HOTSPOTS (VASP-1945)

If dynamic .NET modules (e.g. loaded via `Assembly.LoadFrom`) are used in target application, VTune™ Profiler will show them as "unknown" functions and modules in hotspots list.

ATTACHMENT AS ADMINISTRATOR TO A PROCESS THAT LAUNCHED BY THE 'SYSTEM' USER FAILS (VASP-476)

You may use a utility from <http://technet.microsoft.com/en-us/sysinternals/bb897553> to profile a system service (for example, `w3wp.exe`-based code) from the command line using `vtune.exe` that is located in `<product_install_dir>/bin64`. Do the following:

1. Configure the `w3wp` service to run with the permissions you use to log in: Open IIS Manager, right-click an application pool you are using and set "Process model:Identity" to the account under which `w3wp` needs to be run.
2. Run the `w3wp` service and make sure you run it with proper credentials and remember its PID.
3. Start data collection:


```
psexec -i 0 /path/to/vtune.exe -c=hotspots -r /path/for/your/data_dir -- target-pid=PID
```
4. Run your workload.
5. Stop data collection:


```
psexec -i 0 /path/to/vtune.exe -command detach -r /path/for/your/data_dir
```
6. Open `/path/for/your/data_dir` in GUI:


```
<product_install_dir>/bin64/vtune-gui /path/to/your/data_dir
```

LIMITED SUPPORT OF WINDOWS STORE* APPLICATIONS (VASP-531)

VTune™ Profiler supports attaching to Windows Store* applications, but not launching them. Only hardware event-based analysis types are supported.

VTUNE™ PROFILER ANALYSIS MAY FAIL ON MACHINE WITH VERDASYS DIGITAL GUARDIAN* SOFTWARE INSTALLED. (VASP-2701)

User mode sampling and tracing based analysis types (Hotspots, Threading) run from GUI may fail if Verdasys Digital Guardian* software is installed on the machine. However, it is still possible to run the collection from the command line.

You can also use EBS collection with stacks, either from the GUI or command line.

ANALYZED APPLICATION MAY CRASH DURING ANALYSIS IF "SYMANTEC ENDPOINT PROTECTION APPLICATION AND DEVICE CONTROL" SOFTWARE IS INSTALLED (VASP-2286)

The issue may occur with user mode sampling and tracing based analysis types (Hotspots, Threading). Ways to overcome the problem:

- If user mode sampling and tracing based analysis is needed exclude the analyzed application from the list of monitored applications of "Symantec Endpoint Protection Application and Device Control" software.
- Use Hotspots analysis with Hardware Event based Sampling and **Collect stacks** option turned on.

DOCUMENTATION CANNOT BE OPENED IN MICROSOFT INTERNET EXPLORER* 10 ON WINDOWS SERVER* 2012 (DOC-204)

To fix the error, modify security settings in Microsoft Internet Explorer* 10. From **Tools > Internet Options > Security**, add "about:internet" to the list of trusted sites.

CHILD PROCESS ANALYSIS DOES NOT WORK FOR EBS ANALYSIS TYPES IF PROCESS IS RUN FROM SCRIPT USING START COMMAND (DOC-231)

If a process is started from a batch file by the `start` command, it does not have original command interpreter as a parent process. So the new process cannot be identified as a child and not profiled. This is true for all Hardware event-based analyses. Possible workarounds:

- Use user-mode sampling and tracing based analysis types instead
- Use `call` command in a batch file instead of `start`
- Use "attach to process" collection mode to connect directly to the process of interest

RESULTS MAY CONTAIN PROCESS NAME CUT ON THE 14TH CHARACTER (VASP-2702)

VTune™ Profiler may truncate a lengthy process name at the 14th character in results collected with EBS analysis types with stacks. For example, a result should contain `find_hotspots.exe`, but contains only `find_hotspots..` This is an operating system limitation.

WARNING "UNSIGNED DRIVER" DURING INSTALLATION FOR SEPDRV AND VTSS DRIVERS ON WINDOWS* 7 AND WINDOWS SERVER* 2008 R2 (VASP-234)

VTune™ Profiler hardware event-based sampling drivers (`sepdrv.sys` and `vtss.sys`) are signed with digital SHA-2 certificate key for compliance with Windows 10 requirements.

To install the drivers on Windows* 7 and Windows Server* 2008 R2 operating systems, add functionality for the SHA-2 hashing algorithm to the systems by applying Microsoft Security Update 2033929:

<https://technet.microsoft.com/en-us/library/security/3033929>.

STARTING PLATFORM PROFILER CAUSES ERROR: "THE SERVER WAS STARTED FROM AN UNEXPECTED LOCATION" (DXET-2063)

When starting the Platform Profiler server, the server may fail to start and report the following error:

"The server was started from an unexpected location. Please stop this server instance before trying to start the server again".

This error can happen when the server was previously started from the Windows Command Prompt using the "vpp-server start" command.

To work around this problem:

- Stop these processes:
 - influx
 - <vtune install dir>/vpp/server/dist/3rdparty/vpp-python/python.exe
- Restart the Platform Profiler server.

Note that starting the Platform Profiler server from the VTune user interface does not result in this error.

LONG FINALIZATION TIME IF WINDOWS BINARIES ARE BUILT WITH THE "/OPT:ICF" LINKER OPTION (VASP-21607)

Debugging information generated for large (Several hundreds of MBs) binaries built with the "/OPT:ICF" option can contain many overlapping symbols. That can cause two following problems:

- Long finalization of VTune results (because VTune is not optimized to handle such symbols);
- Impossibility to distinguish certain symbols because they have the same address range in debug info.

A workaround is to remove the "/OPT:ICF" from linker options.

VTUNE™ PROFILER 2021 TOOLBAR BY DEFAULT IS HIDDEN IN MICROSOFT* VISUAL STUDIO 2019 IDE (VASP-22133)

To enable the toolbar right-click on the Visual Studio toolbar area and check "Intel VTune Profiler". The VTune Profiler toolbar will appear, but unfortunately, the extension toolbar needs to be enabled on each launch of the IDE.

Alternatively, you can launch VTune Profiler using Menu->Tools->Intel VTune Profiler or from a project context menu.

VTUNE™ PROFILER 2021 IS NOT SUPPORTING A REMOTE COLLECTION WITH VIA PSEXEC, SSH, POWERSHELL UTILITIES ARE NOT SUPPORTED ON WINDOWS PLATFORMS (VASP-22269)

As a workaround, you can use a Windows Remote Desktop session and collect VTune profiles as on a local machine.

VTUNE™ PROFILER IS NOT ABLE TO RESOLVE DEBUG INFORMATION WHEN AN APPLICATION IS COMPILED WITH THE /DEBUG:FASTLINK OPTION. (VASP-3093)

SOMETIMES VTUNE PROFILER HANGS ON PROCESSING .PDB FILES BUILT WITH /DEBUG:FASTLINK OPTION

To fix this issue, please use the /DEBUG:FULL option instead. For more information, see

<https://docs.microsoft.com/en-us/cpp/build/reference/debug-generate-debug-info>

VTUNE™ PROFILER IS REPORTING AN ERROR WHEN CONFIGURING AN ANALYSIS WITH A PROJECT CREATED OR RESULTS STORED IN A BOX DRIVE FOLDER (E.G. C:\USERS\<USER NAME>\BOX ON WINDOWS) (VASP-22392)

A discussion of the problem with Box Drive can be found here: <https://support.box.com/hc/en-us/community/posts/1500001146222-Directory-to-Box-drive-is-an-empty-symlink->

A workaround would be to place a project and result files in local folders or other shared drives.

A DPC++ APPLICATION COMPILED WITH OPTIONS THAT PRESERVE DEBUG INFORMATION MAY BE SIGNIFICANTLY SLOWED DOWN WHILE RUNNING ON CPU

To work around this issue, explicitly set an environment variable `CL_CONFIG_CPU_NO_DBG_CBK=1` before running the application or profiling it with VTune Profiler

VTUNE PROFILER DOES NOT SUPPORT DEBUG INFORMATION FOR INLINE FUNCTIONS ON WINDOWS (VASP-24541)

VTune Profiler is unable to read debug information for inline functions when it contains in a PDB file.

Consequently, when you profile applications on Windows systems, you may not see source line information for inline functions.

To work around this issue, before you profile an application with VTune Profiler, set the `CL_CONFIG_USE_NATIVE_DEBUGGER` environment variable to 0 (`CL_CONFIG_USE_NATIVE_DEBUGGER=0`)

Linux* Specific Limitations

SOME FEATURES REQUIRE INSTALLING AS ROOT

Most of Intel® VTune™ Profiler profiling features work with a non-root install. Many features work on either a genuine Intel processor or a compatible processor.

Some advanced features that use event-based sampling require the latest OS kernel or sampling driver to be installed. Intel Atom® processors also require this driver for analysis.

To install the driver on a system with a genuine Intel processor, launch the installer as root or ask your system administrator to install the driver later. For information on building and setting up the drivers, see https://software.intel.com/en-us/sep_driver.

SLES 15 MACHINE REBOOTS WHEN STARTING INPUT/OUTPUT ANALYSIS WITH BOTH PERF/SEP (VASP-14205)

When running Input and Output analysis on SUSE* Enterprise Linux (SLES) 15, consider switching to **Custom Analysis** and turn off page faults collection. This helps prevent a system crash that may happen when kernel trace points for page faults are turned on.

PROXY-RELATED ERRORS WHEN STARTING THE PLATFORM PROFILER SERVER (DXET-1338)

On some systems, starting the server generates the following error:

Malformed Proxy variables. Please set your proxy variables with "http" scheme and restart the PPE server :[Errno socket error] [Errno 10060] A connection attempt failed because the connected party did not properly respond after a period of time, or established connection failed because connected host has failed to respond

To work around this problem:

1. Remove the following proxy-related environment variables and start the server again: `http_proxy`, `https_proxy`, `noproxy`
2. Start the server using the following commands:

```
$ cd <vtune install directory>/vpp/server  
$ dist/3rdparty/vpp-python/python -m scripts.start --skip-proxy-check
```

SECURITY-ENHANCED LINUX* IS NOT SUPPORTED

Security-enhanced Linux settings (SELinux) are currently not supported by the VTune™ Profiler and need to be either disabled or set to permissive for a successful tool suite installation. If your Linux distribution has SELinux enabled the following error message will be issued by the installer:

Your system is protected with Security-enhanced Linux (SELinux). We currently support only "Permissive" mode, which is not found on the system. To rectify this issue, you may either disable SELinux by - setting the line "SELINUX=disabled" in your /etc/sysconfig/selinux file - adding "selinux=0" kernel argument in lilo.conf or grub.conf files or make SELinux mode adjustment by - setting the line "SELINUX=permissive" in your /etc/sysconfig/selinux file or ask your system administrator to make SELinux mode adjustment. You may need to reboot your system after changing the system parameters. More information about SELinux can be found at <http://www.nsa.gov/selinux/>

VTUNE™ PROFILER USER MODE COLLECTION MAY FAIL TO RUN ON UBUNTU* (VASP-3352)

VTune™ Profiler may fail to collect data for Hotspots and Threading analysis types on the Ubuntu operating system. Once a collection is started, this message appears in the output:

Failed to start profiling because the scope of ptrace() system call application is limited. To enable profiling, please set /proc/sys/kernel/yama/ptrace_scope to 0. See the Release Notes for instructions on enabling it permanently.

To work around this problem for the current session, set contents of the `/proc/sys/kernel/yama/ptrace_scope` sysctl file to 0.

To make this change permanent, set `kernel.yama.ptrace_scope` value to 0 at `/etc/sysctl.d/10-ptrace.conf` file using root permissions and reboot the machine.

THREADING ANALYSIS MAY PROVIDE INCORRECT WAIT TIME DATA IF LINUX FUTEXES ARE USED

In certain cases, custom synchronization constructions based on Linux futexes are not detected by the VTune™ Profiler. System libraries may also use such constructions. To ensure Wait time correctness in the Threading analysis, use `syscall(SYS_futex, ...)` API in your code.

VTUNE™ PROFILER MAY CAUSE KVM GUEST VIRTUAL MACHINE TO CRASH (VASP-3806)

In some cases, when you run Event-Based Sampling analysis on a host machine, a guest KVM virtual machine may experience kernel panic. This happens only if precise events are used, for example, with Microarchitecture Exploration (formerly, General Exploration) analysis type. Upgrade KVM to a newer version; there is a KVM patch that fixes the problem.

SYMBOLS ARE NOT RESOLVED FOR BINARIES REMAPPED WITH HUGTLBFS UTILITY (VASP-707)

`hugetlbfs` remaps segments from executable file to its temporary file system or to anonymous memory regions. VTune™ Profiler can't map those segments to the original executable.

The workaround for this problem is to rename the original executable to the one that `hugetlb` exposes, i.e. the module name that you see in VTune™ Profiler results.

`my_exe` (original executable) → `my_exe_32_2` (file name on `hugetlb` file system in case of sharing between processes)

`my_exe` (original executable) → `libhugetlbfs.tmp.<random hash>` (memory region name in case of anonymous mappings)

RUNNING OTHER PMU-BASED TOOLS SIMULTANEOUSLY WITH HARDWARE EVENT-BASED SAMPLING COLLECTION IS NOT SUPPORTED

The Sampling Driver requires exclusive use of the PMU resources. Running multiple PMU based tools will cause incorrect results and/or system instability.

HOTSPOTS AND THREADING ANALYSIS TYPES MAY NOT WORK ON EXECUTABLES THAT DO NOT DEPEND ON THE `libpthread.so.0` LIBRARY

There is currently a limitation in the product regarding profiling application targets where the executable does not depend on the `libpthread.so.0` library. The message "*Link libpthread.so to the application statically and restart profiling*" appears when profiling an application where program image does not depend on `libpthread.so.0` but then it `dlopen()`s a shared library which does depend on `libpthread.so.0`. The collector is not able to follow the program execution and module load/unload so the collection results are likely to be misleading. A workaround is to set "`LD_PRELOAD=libpthread.so.0`" before running the collection.

OPENING RESULTS SIGNIFICANTLY DELAYED ON LINUX SYSTEMS CONFIGURED WITH "LARGE PAGES"

To work around this issue make sure that you have `hugectl` utility and `libhugetlbfs` library on the system. VTune™ Profiler can be started this way:

```
hugectl --heap <vtune install dir>/bin64/vtune-gui
```

```
LD_PRELOAD=libhugetlbfs.so <vtune install dir>/bin64/vtune-gui
```

HELP TOC IS TRUNCATED IN SOME LINUX BROWSERS

When using certain HTML browsers to view the VTune™ Profiler offline documentation, if the Contents tab width appears truncated and hides some of the topic titles:

- Click the **Index** (or **Search**) button
- Click the **Contents** button

COLLECTION FAILS WITH ERROR "NO CONFIG_PERF_EVENTS=Y KERNEL SUPPORT CONFIGURED" (DOC-241)

Driverless EBS collection on Linux can fail with the error: "*No CONFIG_PERF_EVENTS=y kernel support configured*" even if the kernel config has been set. This is a known Linux kernel issue that is resolved in kernel version 4.1. There is a simple patch that provides a workaround: <http://halobates.de/ivb-allow-mem-load-uops>

USER TASKS, EVENTS AND FRAMES MARKED WITH ITT APIs MAY NOT BE SHOWN IF TARGET APPLICATION IS FINISHED BY `CTRL+C` (VASP-1367)

Use another method to stop application, or VTune™ Profiler `stop` command.

WHEN INSTALLING ON A XEN* VIRTUAL MACHINE, IT MAY HAPPEN THAT INSTALLATION SCRIPT EXITS CLAIMING CPU IS UNSUPPORTED

To skip the check for CPU model, add `--ignore-cpu` parameter to install script command line:

```
./install.sh --ignore-cpu
```

VTUNE™ PROFILER HELP DOES NOT OPEN FROM GUI ON SOME LINUX SYSTEMS (VASP-3322)

On some Linux* systems, like Wind River Linux, default ways to open web browser may not work. You may export the environment variable `BROWSER` (before starting VTune™ Profiler) with the path to the web browser to get VTune™ Profiler help working.

For example:

```
$ export BROWSER=/usr/bin/firefox
```

VTUNE™ PROFILER USER APIS ARE NOT SUPPORTED FOR STATIC BINARIES (VASP-2287)

To use ITT and JIT APIs with software-based collectors (Hotspots, Threading), you must link their application dynamically. Otherwise ITT and JIT API notifications will be missing in the result.

VTUNE™ PROFILER TARGET EXECUTABLES DO NOT WORK ON YOCTO PROJECT X64 DUE TO DIFFERENT PATHS TO LD (VASP-8531)

Error message: `"-sh: ./amplxe-perrunss: No such file or directory"`. Usually `ld` is located in `/lib64/ld-linux-x86-64.so.2` but on Yocto x64 it is located in `/lib/ld-linux-x86-64.so.2`. The workaround is to create `/lib64/ld-linux-x86-64.so.2` as a symlink to `/lib/ld-linux-x86-64.so.2`.

EVENT-BASED SAMPLING DRIVER FAILS TO BUILD ON YOCTO PROJECT 1.6 INTEL X86_64 BSPTS DUE TO MISSING CONFIG_COMPAT #DEFINE IN LINUX KERNEL HEADERS (VASP-1558)

Sampling driver fails to build on Yocto x86_64 BSPTS images that do not have `"CONFIG_COMPAT=y"` setting (https://bugzilla.yoctoproject.org/show_bug.cgi?id=6777)

To resolve the issue, you need to re-build the image with the option manually.

Change kernel configuration according to <http://www.yoctoproject.org/docs/1.6.1/kernel-dev/kernel-dev.html#changing-the-configuration> :

1. Create the `compat.cfg` file in `meta-yocto-bsp/recipes-kernel/linux/files/` directory with content:


```
CONFIG_IA32_EMULATION=y
CONFIG_COMPAT=y
```
2. Add the following strings to `meta-yocto-bsp/recipes-kernel/linux/linux-yocto_3.10.bbappend`:


```
FILESEXTRAPATHS_prepend := "${THISDIR}/files:"
SRC_URI += file://compat.cfg
```
3. Clean and build a modified kernel according to <http://www.yoctoproject.org/docs/1.6.1/dev-manual/dev-manual.html#patching-the-kernel> and <http://www.yoctoproject.org/docs/1.6.1/dev-manual/dev-manual.html#build-the-modified-qemu-kernel-image> :


```
$ bitbake linux-yocto -c cleansstate
$ bitbake -k linux-yocto
```
4. Build image:


```
$ bitbake core-image-sato
```


COLLECTING PERFORMANCE DATA FOR INTEL® OPTANE™ PERSISTENT MEMORY (DXET-1844)

If you are using an account with root privileges and you try to collect data on a system enabled with Intel® Optane™ Persistent Memory, this warning appears:

User needs sudo privilege to collect Persistent Memory metrics. Persistent Memory metrics will not be collected.

Use the following command to start data collection with elevated privileges:

```
sudo <vtune install dir>/vpp/collector/vpp-collect start
```

Use the following command to stop data collection:

```
sudo <vtune install dir>/vpp/collector/vpp-collect stop
```

NOTE: elevated privileges are required to stop a data collection session that was started with elevated privileges.

VTUNE COLLECTION CAN SOMETIMES FAIL REPORTING LIBSTDC++.SO.6 LIBRARY NOT FOUND (VASP-21417)

Due to a known Linux loader bug, VTune collection can sometimes fail when a user had libraries in LD_PRELOAD environment prior to launching VTune collection. The example of how the bug can manifest when having libzet_intercept.so in LD_PRELOAD:

```
/opt/intel/vtune_profiler_2021.1.0/bin64/pin:
/opt/intel/vtune_profiler_2021.1.0/bin64/./lib64/libstdc++.so.6: version
`GLIBCXX_3.4.26' not found (required by /lib/x86_64-linux-gnu/libze_loader.so.1)
```

There are several workarounds for this problem:

1. Clear LD_PRELOAD before launching VTune and configure LD_PRELOAD as an environment for the target application if it is required.
2. Add libstdc++.so.6 either from VTune installation or from system location to LD_PRELOAD (see <https://software.intel.com/content/www/us/en/develop/documentation/vtune-help/top/troubleshooting/error-message-cannot-collect-gpu-hardware-metrics.html> that has a similar workaround)
3. Do 'export LD_PROFILE=none' before launching collection (this could cause an insignificant performance hit for the loader, does not work in security hardened environment)

Any one of these workarounds should solve the problem.

USING A WRAPPER SCRIPT IN VTUNE PROFILER COMMAND MIGHT NOT WORK PROPERLY IF "SUDO -A SU <USER> -C "\$*" IS USED IN THE SCRIPT FOR LOCAL CONNECTION. (VASP-22952)

Correct "su" commands are the following:

The command "sudo -A su <user> -c "\$*" " works for Remote Linux (SSH) collection.

The command "sudo -A -C 65000 su <user> -c "\$*" " works for local (or VTune Server) collection.

UNABLE TO RUN PLATFORM PROFILER ON FEDORA VERSIONS 30 OR NEWER(DXET-2342)

On Fedora versions 30 or newer, Platform Profiler can fail to start. This error message appears:

Intel® VTune™ Profiler 2021

```
<vtune-install-dir>/vpp/server/dist/3rdparty/vpp-python/bin/python3: error while loading
shared libraries: libcrypt.so.1: cannot open shared object file: No such file or
directory
```

Failed to configure Intel(R) VTune(TM) Profiler - Platform Profiler server.

Intel(R) VTune(TM) Profiler - Platform Profiler is supported on Ubuntu, CentOS, Debian and RedHat based systems only.

To work around this problem, install the libxcrypt-compat package using the Fedora package manager.

PLATFORM PROFILER IS UNABLE TO COLLECT PERSISTENT MEMORY METRICS FROM SYSTEMS WITH INTEL® OPTANE™ PERSISTENT MEMORY 200 SERIES MODULES (FORMERLY BARLOW PASS) (DXET-2330)

When collecting Platform Profiler data from systems with Intel(R) Optane(TM) Persistent Memory 200 Series Modules (formerly Barlow Pass), you may see this warning:

```
WARNING: PMWatch was not found. Persistent Memory metrics will not be collected.
For more details, see ~/vpp-collect.log" displayed
```

The VTune error log file contains more information:

```
WARNING: ipmwatch pre-requisites:
    Firmware version >= 01.00.00.5142 for PMON functionality
                                01.00.00.5127 for rest of the functionality
    ipmctl >= 01.00.00.3296 for PMON functionality
                                01.00.00.3262 for rest of the functionality
    ipmctl found on the system: 02.00.00.3457
ipmwatch is not guaranteed to work with versions less than the ones mentioned above.
```

To work around this issue,

1. Identify the versions of libipmctl installed on the system:

```
$ ls -la /lib64 | grep ipmctl
lrwxrwxrwx 1 root root 18 Dec 18 12:59 libipmctl.so.3 -> libipmctl.so.3.1.0
-rwxr-xr-x 1 root root 2525496 Feb 28 2020 libipmctl.so.3.1.0
lrwxrwxrwx 1 root root 18 May 1 2020 libipmctl.so.4 -> libipmctl.so.4.0.0
-rwxr-xr-x 1 root root 2562712 May 1 2020 libipmctl.so.4.0.0
```

2. Update the link libipmctl.so.3 to point to libipmctl.so.4.0.0

```
$ unlink /lib64/libipmctl.so.3
$ ln -s /lib64/libipmctl.so.4.0.0 /lib64/libipmctl.so.3
$ ls -l /lib64/ | grep ipmctl
lrwxrwxrwx 1 root root 25 Apr 22 15:47 libipmctl.so.3 ->/lib64/libipmctl.so.4.0.0
-rwxr-xr-x 1 root root 2525496 Feb 28 2020 libipmctl.so.3.1.0
lrwxrwxrwx 1 root root 18 May 1 2020 libipmctl.so.4 -> libipmctl.so.4.0.0
-rwxr-xr-x 1 root root 2562712 May 1 2020 libipmctl.so.4.0.0
```

macOS* Specific Limitations

NATIVE COLLECTION IS NOT AVAILABLE ON MACOS*

Intel VTune™ Profiler for macOS* supports viewing of results collected on other OSs. Native collection on macOS is not currently available.

SOME SYSTEMS DO NOT SUPPORT CONTEXT-SENSITIVE HELP

On some macOS* systems, such as macOS 10.13, VTune™ Profiler context-sensitive help is not supported.

6 Attributions

The license information for third party software that was used to develop Intel® VTune™ Profiler can be found in `<install_dir>/licensing/third-party-programs.txt`.

For the avoidance of doubt, the Intel® VTune™ Profiler is solely governed by the terms and conditions of the End User License Agreement for Intel® Software Development Product that accompanies the Intel® VTune™ Profiler.

7 Legal Information

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

The products and services described may contain defects or errors which may cause deviations from published specifications. Current characterized errata are available on request. No product or component can be absolutely secure.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or by visiting www.intel.com/design/literature.html.

Intel, the Intel logo, Intel Core, Intel Xeon Phi, VTune, and Xeon 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.

Copyright 2010 - 2021 Intel Corporation.

This software and the related documents are Intel copyrighted materials, and your use of them is governed by the express license under which they were provided to you (**License**). Unless the License provides otherwise, you may not use, modify, copy, publish, distribute, disclose or transmit this software or the related documents without Intel's prior written permission.

This software and the related documents are provided as is, with no express or implied warranties, other than those that are expressly stated in the License.