



Intel[®] VTune[™] Profiler 2022

Release Notes

27 September 2022

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	7
4	Installation Notes	7
5	Known Issues	8
6	Attributions	23
7	Legal Information	23

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](#)

2 What's New

Detailed [new features](#).

Intel® VTune™ Profiler 2022.4.0

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

- **Several critical bugs were fixed**

Intel® VTune™ Profiler 2022.4 may not include all the latest functional and security updates. New versions of Intel® VTune™ Profiler are targeted to be released in December 2022 and will include additional functional and security updates. Customers should update to the latest version as it becomes available.

As announced previously, Windows driver support of integrated graphics processors included with 6th - 10th Gen Intel Core Processor and related Intel Atom®, Pentium®, and Celeron® processors is deprecated and has moved to maintenance mode. Only security and critical bug fixes will be updated.

oneAPI tools using existing integrated graphics processor functionality in the aforementioned processors may continue to work, but will no longer be supported. Note that CPU functionality for these processors remains fully supported and unaffected.

Intel® VTune™ Profiler 2022.3.0

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

- **GPU Accelerators**
 - **Support for Unified Shared Memory extension of OpenCL™ API**

When you use the [GPU Offload analysis type](#) to profile OpenCL applications, you can now profile the CPU-side stacks for GPU computing tasks and identify bottlenecks related to Unified Shared Memory (USM) for the OpenCL API.
 - **Support for DirectML API**

This release also extends profiling support in the GPU Offload and GPU Compute/Media Hotspots analysis types for Microsoft® DirectX* applications to include support for the DirectML API.
- **Platform Support**
 - Support for Legacy Processors**

VTune Profiler now supports the following generations of processors in client and server platforms:

 - **Server CPUs:** Intel® Xeon® Processor v3 Family and newer.
 - **Client CPUs:** Intel® Core™ 4th generation processors and newer.

Starting with this release, VTune Profiler does not support processors older than the versions listed above. To analyze performance on older processors, use an older version of VTune Profiler.

Intel® VTune™ Profiler 2022.2.0

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

- **HPC Performance Characterization Analysis**
 - **Better Hardware Observability**

This release adds the **Platform Diagram** to the **Summary** tab of the [HPC Performance Characterization analysis](#) result. The **Platform Diagram** reveals system topology, utilization metrics for physical cores, DRAM, and Intel® Ultra Path Interconnect (Intel® UPI) links. Available for server platforms based on Intel® microarchitecture code named Skylake and newer.
- **VTune Profiler Server**
 - **New Command-Line Options for Convenience**

The `vtune-backend` binary that launches VTune Profiler server features new command-line options to make setup in certain environments more convenient. You can now specify a base URL that VTune Profiler server will use as a basis for URL generation. Additionally, new options were added to suppress automatic help tours on startup and to provide/decline consent to collect usage information right from the command line. These new options can be especially useful if you are running VTune Profiler server inside a container.

Intel® VTune™ Profiler 2022.1.0

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

- **Managed Code Targets**
 - **.NET 6 Support**

This release introduces support for analyzing .NET 6 targets using User-Mode Sampling. You can analyze .NET 6 workloads in **Launch Application** and **Attach to Process** modes on both Windows* and Linux* hosts.
- **Input and Output Analysis**
 - **Intel® VT-d Observability**

Intel® Virtualization Technology for Directed I/O ([Intel® VT-d](#)) observability is introduced in the [Input and Output](#) analysis for server platforms based on 3rd Gen Intel® Xeon® Scalable processors (code named “Ice Lake”), the Intel Atom® P5900 Processor Family (code named “Snow Ridge”), and newer. New performance metrics reveal efficiency of hardware-driven DMA addresses remapping and penalties for sub-optimal Intel VT-d utilization.
- **Application Performance Snapshot**
 - **Histograms in Metric Tooltips**

The metric tooltips in Application Performance Snapshot HTML reports were enhanced with histograms that clearly visualize the distribution of metric values that were observed during analysis.
- **Operating System Support**
 - **New Host Operating Systems Supported**

This release introduces support for these operating system hosts:

 - Microsoft Windows* 11
 - Ubuntu 21.10
- **Hardware Support**
 - **Support for First Generation of Intel® Arc™ High-performance Discrete GPUs**

This release of VTune Profiler supports the first generation of Intel® Arc™ high-performance discrete GPUs code named Alchemist, and previously known as DG2. The support includes:

- Explicit support for DPC++, DirectX, Intel® Media SDK, OpenCL™, and OpenMP-offload software technologies
- Support for multi-GPU systems. You can now profile all Intel GPU devices, including integrated and discrete GPUs.
- Support for GPU Offload and GPU Hotspots analyses, including source level in-kernel profiling

Intel® VTune™ Profiler 2022.0.0

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

• User Interface

○ Rich Metric Tooltips in Multiple Analyses

This release introduces rich metric tooltips in [Performance Snapshot](#), [Hotspots](#), [HPC Performance Characterization](#), and [Microarchitecture Exploration](#) analyses. The new tooltips aim to make metrics more intuitive by providing visualizations for thresholds, desired direction (more/less is better), and tuning potential. Hover over a metric to get this tooltip.

○ Detection of Compilation with Low Optimization Level in Hotspots Analysis

When debug information is available, VTune Profiler now detects and flags modules that may have been compiled using non-optimal compiler optimization flags in the **Top Hotspots** section of the [Hotspots analysis](#) result. This can help detect underutilization of compiler optimization capabilities and correct the build system setup.

○ Platform Diagram Extended with Persistent Memory Block

For [Input and Output](#) and [Memory Access](#) analyses, the **Platform Diagram** shown in **Summary** windows now features a dedicated block for Persistent Memory devices, together with average per-socket bandwidth.

This data is available on server platforms based on microarchitectures code named Cascade Lake and Ice Lake.

○ Changes to Viewpoint Selection

The Viewpoint selection was adjusted with respect to each analysis type. Now, the viewpoint selection is disabled for certain analysis types, and only features a managed set of most helpful viewpoints for other analysis types. You can re-enable the display of all applicable viewpoints in the [Options pane](#).

• Analysis Targets

○ Extended FreeBSD* OS Support

- The [Instrumentation and Tracing Technology API \(ITT API\)](#) is now fully supported on FreeBSD OS. The appropriate header and library files are provided as part of the FreeBSD target package. You can use ITT API to annotate your code and collect arbitrary statistics with little to no overhead.
- On FreeBSD OS, The [Input and Output](#) analysis now supports Storage Performance Development Kit (SPDK) analysis. You can now get [SPDK-specific performance data](#) on FreeBSD OS.

• Platform Analyses

○ CPU Throttling Data in System Overview Analysis

The [System Overview analysis](#) now displays information about factors that can cause the CPU

to throttle. Use this information to examine if your system is overheated or consumes significant power, both of which could result in frequency drops that affect system performance.

- **Architecture Support**
 - **Analysis Support for Intel® microarchitecture code named Alder Lake**
This version of Intel® VTune Profiler introduces support for Intel® microarchitecture code named Alder Lake in these analysis types:
 - [Microarchitecture Exploration analysis](#)
 - [Memory Access analysis](#)
- **VTune Profiler Server**
 - **New Capabilities for Account and Privilege Handling**
 - VTune Profiler Server now supports [profiling of workloads that require sudo elevation](#).
 - Introduced support for a collector [wrapper script](#) to [elevate privileges](#) before launching or attaching to a workload.
- **Accelerators**
 - **Host-GPU Bandwidth Information in GPU Offload Analysis**
Previously, you checked the **Analyze memory bandwidth** option in the GPU Offload Analysis to see data required for this computation. Starting with this release of VTune Profiler, you can now use the **Analyze host-GPU bandwidth** option instead. Depending on your hardware configuration, this selection displays DRAM bandwidth, PCIe bandwidth, or both sets of data on the timeline.
 - **PCIe Bandwidth Information in Custom and Command Line Runs of GPU Offload Analysis**
Use new options to collect information about PCIe bandwidth (between the host and GPU sides) when you run custom and command line runs of the GPU Offload analysis:
 - Use the `collect-host-gpu-pci-bandwidth` switch for both custom and command line runs.
 - In the UI, check the **Analyze host-GPU PCIe bandwidth** option for custom analysis.
 - **Improvements to Peak Occupancy Metric**
The **GPU Peak Occupancy** metric for a computing task now flags the factors that limit peak occupancy in the order of priority. Start tuning your application by addressing the most restricting factor. VTune Profiler customizes recommendations for potential improvements based on the launch parameters of the compute kernel (work size, SLM and barriers usage).
 - **Support to Trace DirectX API on CPU Host**
This release of VTune Profiler introduces support to profile DirectX applications on the CPU host. These versions of the DirectX API can be traced:
 - DXGI
 - Direct3D 11
 - Direct3D 12
 - Direct3D-11-On-12 (D3D11On12)
 - **Data Collection of CPU Host Stacks**
When you collect information about host stacks in the GPU Offload and GPU Compute/Media Hotspots analyses, you can now filter the data by selecting a call stack mode from the filter bar.
 - **Enhancements to GPU Offload Summary**
The Summary window of the GPU Offload Analysis contains these enhancements for an improved user experience.
 - Locate hotspots in your function when the GPU is not busy. See the new **“Top hotspots when GPU was idle”** table in the **GPU Time, % of Elapsed Time** (formerly **GPU Utilization**) section.
 - The **Hottest Computing Functions** section now includes occupancy information.

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:

- <https://www.intel.com/content/www/us/en/develop/documentation/vtune-install-guide/top.html>

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 ...`

CREATING OF MANY ITT API INSTANCES CAN LEAD TO ERROR REPORT THROUGH __ITT_REPORT_ERROR() AND/OR ABSENCE OF ITT API DATA.

To avoid this issue consider using latest ITT API from open source repository: <https://github.com/intel-innersource/applications.analyzers.components.userapi>

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. Howeverm 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:

Intel® VTune™ Profiler 2022

- o influx
- o <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-pttrace.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 `libhugectlbfs` library on the system. VTune™ Profiler can be started this way:

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

```
LD_PRELOAD=libhugectlbfs.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: `./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 2022

```
<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
```

DRIVERLESS PERF BASED COLLECTIONS ON ADL IS SUPPORTED BY LINUX KERNEL SINCE 5.15.5. (VASP-27934)

Driverless Perf based collections on ADL is supported by Linux kernel since 5.15.5. To check if your kernel supports driverless collection on ADL architecture, check for the following device files:

Intel® VTune™ Profiler 2022
/sys/bus/event_source/devices/cpu_atom

/sys/bus/event_source/devices/cpu_core

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 - 2022 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.