



Intel® SoC Watch 2.7 for Linux* OS

Release Notes

15 August 2018

Version History

These are the main releases of Intel® SoC Watch for Linux* OS.

Date	Revision	Description
July 2015	2.0	Initial release for 2.0 Product
October 2015	2.1	Update to 2.1 Product release
June 2016	2.2	Update to 2.2 Product release
October 2016	2.3	Update to 2.3 Product release
May 2017	2.3.1	Update to 2.3.1 Product release
November 2017	2.4	Update to 2.4 Product release. First release that aligns all command line parameters and output formats across all supported operating systems.
February 2018	2.5	Update to 2.5 Product Release. Includes support for Intel platform code named Gemini Lake and other fixes.
May 2018	2.6.1	Enhancements include new hw-cpu-hwp metric, --log option, and improved support in gfx metrics.
August 2018	2.7	Added average frequency report, new options (program-delay, disable-alts), new metric (pkg-pwr), new group names, fixed issues in ddr-bw, automation summary, and multiple pkg handling.

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: [Intel System Studio forum](#) discuss issues with your peers, ask questions of the development team, etc.

To submit an issue, go to Intel® Premier Support:

(<https://employeeportal.intel.com/irj/portal/IntelPremierSupportUser>)

Please remember to 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.

Contents:

1	Introduction	4
2	New in This Release	5
3	System Requirements	6
4	Where to Find the Release	7
5	Installation Notes	8
6	Fixed Issues	11
7	Known Issues	12
8	Related Documentation	15
9	Legal Information	16

1 Introduction

Intel® SoC Watch is the data collector for power-related data that helps identify issues on a platform that are preventing entry to power-saving states. The metrics it captures include system sleep states, CPU and GPU sleep states, processor frequencies, temperature data and device sleep states among others. The collected data can be correlated over time using Intel® VTune™ Amplifier.

This document provides system requirements, installation instructions, issues and limitations, and legal information.

To learn more about this product, see:

- New features listed in the [New in this Release](#) section below, or in the help.
- Reference documentation listed in the [Related Documentation](#) section below
- Installation instructions can be found in the [Installation Notes](#) section below.
- For a detailed quick start guide to running the tool, see the Intel SoC Watch User's Guide in your installed documentation.

2 New in This Release

Release v2.7 includes the following significant changes compared to v2.6.1:

- New group names supported to make it easier to find and select features:
 - `cpu-hw`: Collects most CPU-specific metrics that are obtained from hardware data sources (`hw-cpu-cstate`, `hw-cpu-hwp`, `hw-cpu-pstate`, `ia-throt-rsn`, `cpu-gpu-concurrency`, `core-temp`, `soc-temp`)
 - `gfx-hw`: Collects most GPU-specific metrics that are obtained from hardware data sources (`hw-gfx-cstate`, `hw-gfx-pstate`, `gfx-throt-rsn`, `cpu-gpu-concurrency`)
- Command line changes:
 - The `-f sys` group name includes most metrics, including the metrics in the expanded groups above plus bandwidth and package power. (`sstate + device + temp`, `cpu-hw`, `gfx-hw`, `ddr-bw`, `pkg-pwr`).
 - A new option for generating help output in JSON format has been added: `--export-help`. This option may be used to generate a JSON format help output file, useful in automating discovery of features available on a platform.
 - Added a new switch to delay start of program execution: `--program-delay`. If the `-p` option is used to launch a workload before collection, you can use `--program-delay <seconds>` to delay the starting of that program for the specified number of seconds. Collection will begin after the program starts unless the `-s` option is also given. When both are specified, Intel SoC Watch first waits for the number of seconds specified by the `--program-delay` option then starts the program and waits for the amount of time specified by the `-s` option before starting data collection.
- The `-f sstate` feature is now consistent across all OSes. It is a group name that includes all metrics related to system idle states.
- The `-f hw-cpu-pstate` summary report includes a new Core P-state Average Frequency table. The table reports the Core P-state *Approximate Avg* and *Actual Avg* frequency. Note that the *Core P-State/Frequency Summary* report shows sampled frequency Residency as a % of total collection time, but the *Core P-State Average Frequency* table is calculated as a % of Active (C0) time, excluding idle time. This is required in order to allow comparison of Approximate Avg to Actual Avg frequency as a measure of accuracy for the sampled data. (The Actual Avg frequency comes from registers Aperf/Mperf that accumulate average frequency during active time, providing the most accurate average frequency value.)

3 System Requirements

Supported Architectures

Intel SoC Watch for Linux* OS supports the following Intel microarchitecture or platform code names:

- Apollo Lake (Broxon-P)
- Skylake
- Kaby Lake
- Denverton
- Coffee Lake
- Broadwell
- Cherryview
- Anniedale
- Haswell
- Skylake-Xeon

Dependencies

Intel SoC Watch depends on specific OS configurations and hardware capabilities. If these are not present on the target system, Intel SoC Watch may fail to work properly.

- Linux Kernel version needs to be $\geq 2.6.32$
- GNU C Library must be version GLIBC_2.17 or later
- KERNEL_CONFIG_TRACEPOINTS must be enabled
- Kernel must have been compiled with "CONFIG_MODULES" enabled.
- P States
 - Kernel config CONFIG_X86_SFI_CPUFREQ or CONFIG_X86_ACPI_CPUFREQ must be enabled (i.e. set to 'y' or 'm').
 - One of the following pstate drivers must be utilized: sfi-cpufreq, acpi-cpufreq, or intel_pstate. To determine which driver is loaded, check the sysfs `/sys/devices/system/cpu/cpu0/cpufreq/scaling_driver` file.
 - If one of these pstate drivers is not loaded, the kernel needs to be reconfigured and recompiled.
- C States
 - Kernel config CONFIG_TIMER_STATS must be enabled.
 - Kernel config CONFIG_INTEL_IDLE must be enabled and the intel_idle kernel module has to support the target platform's core.
 - To determine if the intel_idle kernel module is loaded, check the sysfs `/sys/devices/system/cpu/cpuidle/current_driver` file. It must equal intel_idle. If it equals acpi_idle, only C0 and C1 will be used by the core.

4 Where to Find the Release

Go to the [Intel® System Studio website](#) to get either an Evaluation (30-day trial release) license or a commercial license, and download the package from [Intel Registration Center](#).

5 Installation Notes

Intel SoC Watch for Linux* OS is available as part of Intel® System Studio. Use the steps below to install Intel SoC Watch on a target Linux system.

Extracting the Intel SoC Watch package

Users will need to extract the Intel SoC Watch package to the system containing the target device's kernel (this may be the host system if the target device is running an Android kernel, or the target system if the device is running a Linux kernel). By default the package can be found in the following locations:

- On Linux systems:
/opt/intel/system_studio_<version>/energy_profiler_and_socwatch/socwatch_for_target
- On Windows* systems: C:\Program Files
(x86)\IntelSWTools\energy_profiler_and_socwatch\socwatch_for_target

Use the `find . -name Module.symvers` command on the device containing the target system's kernel to determine the kernel build directory.

Build the Kernel Modules

If the Intel SoC Watch kernel modules (i.e. device drivers) are not present in the target system's OS image, you will need to build and possibly sign them. Building and signing device drivers requires access to the kernel build directory for the OS image running on your target device. A kernel build directory is generated while building the target system's OS image.

When building the kernel modules, the Intel SoC Watch package (i.e. ZIP file) should not be opened (i.e. unzipped) on a Windows* based system and then copied to a Linux system. The package should be unzipped on the Linux build system using the `unzip` command to make sure the build scripts and make files are unmodified.

If a kernel is built with the `CONFIG_MODULE_SIG` kernel config enabled, any device driver loaded into that kernel must be signed with the same keys used to build the kernel. In general, drivers built for Linux targets do not need to be signed and the following description assumes the drivers do not need to be signed. But, if an end user tries to load an unsigned driver into a kernel that requires signed drivers, the `insmod` command will fail with the error *Required key not available*. If a signed driver is loaded into a kernel that does not require signed drivers, the load will succeed.

The `build_drivers` script is provided to simplify building all of the drivers. The script supports multiple switches including;

`-n` // do not build the socperf driver; used for Intel® Core™ processor based systems

`-l` // build the kernel for a Linux target

`-s <full path to sign-file>` // signs the drivers; the path is normally `.../kernel/*/scripts/sign-file`

Building Linux* Kernel Modules

Linux kernel modules may only be built after the Intel SoC Watch package and kernel headers are copied to and installed on the target. See the section [Intel SoC Watch for Linux Installation](#) below for instructions on how to build the kernel modules for a target device running Linux.

Install Intel SoC Watch

Host: laptop, desktop, or server used to communicate with target device.

Target: device to be analyzed with Intel SoC Watch.

Intel SoC Watch for Linux* Installation

If Intel SoC Watch was previously installed on the target, delete the `socwatch _linux_*` directory before installing a new version. Then, perform the following steps on the target device.

1. Login to the target as root:

```
> ssh root@<your_target_IP>
```

2. Extract the Intel SoC Watch package to the target Linux system. For more information, see [Extracting the Intel SoC Watch package](#). The `<extract_dir>/system_studio_target/socwatch_linux_v<version>_x<architecture>` directory will be created.

3. Copy the Intel SoC Watch package to a working directory:

```
> mkdir -p /home/socwatch
> cp
<extract_dir>/system_studio_target/socwatch_linux_v<version>_x<architecture>
/home/socwatch/.
```

4. Navigate to the Intel SoC Watch directory:

```
> cd /home/socwatch/socwatch_linux_v<version>_x<architecture>
```

5. Use one of the following commands:

- a. If the target system has an Intel Atom® processor, use the following command to build the `socwatch2_6.ko` and `socperf2_0.ko` files.

```
> sh ./buid_drivers.sh -l -k <kernel-build-dir> -s <full-path-to-sign-
file>
```

- b. If the target system has an Intel Core processor, use the following command to build the `socwatch2_6.ko` file.

```
> sh ./build_drivers.sh -l -k <kernel-build-dir> -n -s <full-path-to-
sign-file>
```

Key Files

The following table describes the key files.

File	Description
<code>build_drivers.sh</code>	The build script used to build all of the device drivers utilized by Intel SoC Watch.
<code>socperf2_0.ko</code>	The socperf kernel module used to measure bandwidth and DRAM self refresh on systems with an Intel Atom processor.
<code>socwatch2_x.ko</code>	The Intel SoC Watch kernel module used to collect both hardware and kernel data at runtime.
<code>setup_socwatch_env.sh</code>	The script used to setup the Intel SoC Watch runtime environment.
<code>socwatch</code>	The Intel SoC Watch executable built as a native application. Use this file to collect data and generate additional results from a raw SW2 file.
<code>SOCWatchConfig.txt</code>	The Intel SoC Watch configuration file. The configuration file is read by Intel SoC Watch immediately before each collection. It contains hardware addresses utilized by the device driver during the collection.
<code>EULA.txt</code>	End User License Agreement file.
<code>third-party-programs.txt</code>	List of third party programs included in the package.
<code>plugins/libSWCore.so</code>	A library providing Intel SoC Watch functionality.

Remove the Intel SoC Watch Drivers

After using Intel SoC Watch, remove the drivers using the `rmmmod` command (e.g. `rmmmod socwatch2_6`). The `socwatch` driver must be unloaded before the `socperf` driver is unloaded.

6 Fixed Issues

- Corrected issue where some tables from the summary report were not included in the Automation_Summary file. Also changed the column headers for all data to make them unique by prepending the table name, as seen in the summary file, to each column header.
- Error messages are now reported on the console for command line syntax errors such as including an argument with `-m` or using invalid value for `-r` option.
- Fixed an issue where pathnames for output files were incorrectly displayed on the console when relative pathnames were specified with the `-o` option.
- Improved labeling of cores on systems with multiple packages. Previously, on systems with more than one CPU core, metrics that reported data per core would have column headers with multiple occurrences of each core number (ex: Core_0 would appear twice). Now the package number will be prepended to the core number (ex: Package_0/Core_0) for multi-package systems. If the system has a single package, the column header does not change, and gives just the core number (ex: Core_0).
- The `-f dram-srr` reports have been modified for consistency with other metrics, transposing states into rows and adding Total Samples Received table.
- The handling of the filename argument for the `--log` option has been improved so that it behaves like the `--output` option. Folders in the specified pathname will be created if they don't exist, use of "/" in pathname is allowed, and an error is reported if the file location is not writable rather than the tool crashing.

7 Known Issues

Intel® VTune™ Amplifier Visualization

- Intel VTune Amplifier 2017 for Systems Update 1 or later is required for visualizing and analyzing Intel SoC Watch v2.6.1 PWR files. We recommend using the latest version of Intel VTune Amplifier, currently Intel VTune Amplifier 2018.
- If the bandwidth is 0 Mb throughout the collection for a particular bandwidth type, Intel VTune Amplifier will not show a timeline entry for it. The timeline is shown only if there is at least one non-zero value.
- Sometimes the summary CSV results produced by Intel SoC Watch do not match exactly the summary results shown by Intel VTune Amplifier even though they represent the same collection. For example, the summary CSV file may report a specific `cpu-pstate` residency of 50.78% and Intel VTune Amplifier may report the same `cpu-pstate` residency as 50.8%.
- Intel VTune Amplifier currently does not support bandwidth ranges used for `ReadPartial` and `WritePartial`. In order to keep the visualization consistent with Intel SoC Watch v1.x, Intel VTune Amplifier uses the upper bound of the range to visualize the bandwidth.
- The Intel SoC Watch v2.6.1 `ddr-bw` feature measures both bandwidth and bandwidth requests on Intel platforms code named Haswell, Broadwell, and Skylake. When this data is imported and visualized with older versions of Intel VTune Amplifier, only actual bandwidth is shown.
- The minimum and average calculations displayed in the grid for `Sampled Value` metrics don't take 0 values into consideration in older versions of Intel VTune Amplifier. For example, `Sampled Graphics P-States` minimum values may show a value higher than 0 Mhz even when some samples have 0 Mhz values. This in turn affects the average value calculation.
- VTune Amplifier reports the following error if the `.pwr` file being imported includes too many metrics: *Error 0x4000002a (Database interface error) - Precompute error.*

Bandwidth and DRAM Self Refresh

- Intel SoC Watch v2.6.1 requires loading the `socperf v2.0` driver to measure bandwidths or DRAM self-refresh on Intel platforms code named Cherry View, Broxton, and Apollo Lake. Use the `-v` switch to determine which version of the `socperf` driver is *loaded*. If a mismatch occurs (`socperf v1.2.0` used with `socwatch >=v2.6.1`), Intel SoC Watch will report `-1, SOCPERF ERROR configuring SOCPERF interface...`
- On a very small number of systems, results from the `all-approx-bw` feature may be one half of the correct result. During testing, this issue was only experienced on the first collection after the system was booted. All subsequent collections correctly measured the systems bandwidth as expected.
- If Intel SoC Watch crashes while collecting a bandwidth feature (e.g. `-f ddr-bw`) or the DRAM self-refresh feature (i.e. `-f dram-srr`) AND a subsequent collection prints the following error

ERROR: ERROR configuring SOCPERF interface!

then both the `socperf2_0.ko` and `socwatch2_6.ko` kernel modules must be unloaded with the `rmmod` command and reloaded with the `insmod` command before Intel SoC Watch can be used to collect additional data.

- When measuring DRAM self refresh using the `-f dram-srr` feature on cost reduced systems (e.g. BayTrail cost reduced or CherryTrail cost reduced), Intel SoC Watch *may* report 100% self refresh residency on Channel 1. These systems are single channel systems and therefore, the result should be 0%.
- Only one bandwidth (`ddr-bw`, `cpu-ddr-mod0-bw`, `cpu-ddr-mod1-bw`, `io-bw`, `disp-ddr-bw`, `gfx-ddr-bw`, `isp-ddr-bw`, `all-approx-bw`) or DRAM self refresh (`dram-srr`, `dram-srr-ch0`, `dram-srr-ch1`) can be measured during a collection due to hardware limitations. For example, the following Intel SoC Watch command line will fail. `./socwatch -f ddr-bw -f cpu-ddr-mod0-bw ...`
- Bandwidth and DRAM self refresh metrics that are read from hardware signals are not supported on Intel platforms code named Apollo Lake.

Miscellaneous

- In order to graph graphics cstates using the `filename_trace.csv` file generated with the `-r int` switch, the table headers should be manually modified. *Render* and *Media* should be added to the C0, C1, and C6 cells as appropriate in order to properly graph the results.
- The graphics-cstate metric is obtained by frequently polling GPU counters that provide the `gfx-cstate` residencies. On a few devices, we have noticed that for 1 out of every ~3000 samples the residency of one of the c-states (Render C0, C1, C6 or Media C0, C1, C6) as obtained from the GPU counters is greater than the sample duration by 5% or more. When this happens Intel SoC Watch discards that sample and throws the error, *"ERROR: Residency counter for GPU C-state = xxxx TSC ticks, but actual sample duration = yyy TSC ticks. Difference is more than 5.000000 percent."* The error by itself is non-fatal and Intel SoC Watch results give a good idea about the `gfx-cstate` residencies since the bad sample is collected only 1 out of ~3000 samples. This issue is being investigated further.
- Intel SoC Watch reads PMIC and Skin Temperatures from the system's `sysfs`. Rarely, a `sysfs` read may not return before a subsequent `sysfs` read occurs. When this occurs, specific sample results may be missing in the timed trace CSV and raw text files.
- If `socperf` reads occur before the start of collection, a `dmesg` error message is generated: *"socperf2_0: [ERROR] ERROR: RETURNING EARLY from Read_Data"*. This message is benign and can safely be ignored.
- Metrics such as CPU C-state, that report state residency that comes from hardware accumulators will show the *Unknown* state with 0 time and the remaining states will not sum to the total collection duration if the system entered hibernation during the collection and the `-m` option was not specified. When hibernation occurs, a message reporting time spent in hibernation appears at the beginning of the summary report. The *Unknown* state is then included for all appropriate metrics and the time in hibernation is included in that state. In order to find the hibernation time, data must be sampled throughout the collection. Data that comes from hardware accumulators and noted as Snapshot collection mode (in the Intel SoC Watch User's Guide *Options Quick Reference* section) are only collected at the start and end of collection unless the `max-detail` option (`-m`) is specified.
- Total DDR bandwidth does not include EDRAM. On systems using EDRAM, the `ddr-bw` feature report may have a discrepancy between the total data read and writes and the total component requests. The `Data Reads+Data Writes` will be significantly higher than the total

IA+GT+IO requests, because the EDRAM requests are not included. There is no software access to a counter for the traffic between EDRAM and DDR at this time.

- The `-f pkg-pwr` feature will report an incorrect value for energy and power on systems with multiple packages. A fix is being developed and will be included in the next release.
- Permission issues with SELinux will cause Intel SoC Watch collection to fail. Some distributions enable SELinux by default. If you have the following file your system may have SELinux enabled:
`/selinux/enforce`

If that file exists, you can disable by issuing:

```
`echo 0 > /selinux/enforce
```

C-States / P-States

- If all cores in a module request C6FS but actual sleep time is short, hardware's Auto-Demotion logic resolves the module state to module C0. Consequently, you may find module C0 to be greater than the sum of core C0 and C1 on all the cores in a module. On the same lines (auto demotion at the package level), the package C0 may be greater than module C0 residencies of the two modules.
- During the transition time from core C1/C1e/C6 to core C0, a core may run in LFM which will be properly measured by Intel SoC Watch. Therefore, results that include a large number of C1/C1e/C6 residencies may show a lower PState than expected.
- On Intel platforms code named Broxton and Apollo Lake, the `cpu-cstate` metric results do not contain module C-state information.

S States & D States

- On Intel platform code named CherryView based devices, even when the device's screen is off, the NC DState called Display DPIO is reported in the D0i0 state 100% of the time. This result may or may not be correct.
- When collecting NC D0ix states with the `-f nc-dstate` switch, note that the Display Island B (HDMI) IP block will remain in D0i0 when the primary display is enabled even if an HDMI cable is removed.
- When using the `sc-dstate` feature on Intel platforms code named CherryTrail or Braswell based systems, the SEC IP block results are incorrect and should be ignored. This issue is under investigation. Also, the UFS IP block results are incorrect because an internal fuse is disabled.

8 Related Documentation

The below documents are available with this release.

- Intel® SoC Watch for Android* OS and Linux* OS User's Guide
- [Energy Analysis help](#)

9 Legal Information

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

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>).

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 known as errata which may cause deviations from published specifications. Current characterized errata are available on request.

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.htm.

Intel, Intel Core, Intel Atom, the Intel logo, and Intel VTune Amplifier 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.

Copyright 2013 - 2018 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.