

Intel® Speed Select Technology – Performance Profile (Intel® SST-PP) Overview

Authors

Reshma Pattan

David Hunt

1 Introduction

Intel® Speed Select Technology – Performance Profile (Intel® SST-PP) provides customers with multiple processors configurations within a single device, allowing the customer to select between an optimal balance of core count, base frequency, and Thermal Design Point (TDP). Intel SST-PP is available on all 3rd Generation Intel® Xeon® Scalable processors and can be easily integrated across the network. This allows for improved CPU SKU consolidation by using one SKU for multiple workload and application types. As the workload requirements change over time, a more appropriate Intel SST-PP configuration can be selected without replacing the CPU, leading to improved server utilization.

This document describes how to configure the Intel® SST-PP, which is available from 3rd Generation Intel® Xeon® Scalable processor (formerly codenamed Ice Lake) and later processors. Intel® SST-PP offers dynamic performance profile selection of CPU, with each profile offering different performance vectors such as number online cores, base frequency, TDP, and so on. The document also describes the usage of a Linux* kernel tool that was developed to configure Intel® SST-PP technology on a platform. The tool provides a convenient, easy-to-use interface that aids configuration and hides complexity. For more information about Intel® SST-PP, see the References section of this document.

This document is part of the Network Transformation Experience Kit, which is available at <https://networkbuilders.intel.com/network-technologies/network-transformation-exp-kits>.

Table of Contents

1	Introduction.....	1
1.1	Terminology.....	3
1.2	Reference Documentation.....	3
2	Overview	3
3	Deployment.....	4
4	BIOS Settings	4
4.1	Static Intel SST-PP Configuration.....	4
4.2	Dynamic Intel SST-PP Configuration	5
5	Operating System Configuration	5
5.1	Static Configuration.....	5
5.2	Dynamic Configuration.....	5
6	Install and Use the intel-speed-select Tool.....	6
6.1	Check Availability of Intel SST-PP.....	6
6.2	Display Help Text.....	6
6.3	Display System Status and Configuration.....	6
6.4	Display Current Performance Profile	7
6.5	Configure the Performance Profile	8
6.6	Verify the Performance Profile	8
6.7	Example of Intel SST-PP Profile.....	8
6.8	Intel SST-PP Profile Differences.....	9
7	Summary	10

Figures

Figure 1.	One server with multiple configurations	3
Figure 2.	Intel® SSP-PP configurations.....	3

Tables

Table 1.	Terminology.....	3
Table 2.	Reference Documents	3

Document Revision History

REVISION	DATE	DESCRIPTION
001	April 2022	Initial release.

1.1 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
Intel® SST-PP	Intel® Speed Select Technology – Performance Profile
SKU	Stock-keeping Unit
TDP	Thermal Design Point

1.2 Reference Documentation

Table 2. Reference Documents

REFERENCE	SOURCE
Intel® Speed Select Technology (Intel® SST) User Guide	https://www.kernel.org/doc/html/latest/admin-guide/pm/intel-speed-select.html
CPU hotplug in the Kernel	https://www.kernel.org/doc/html/latest/core-api/cpu_hotplug.html
Power Management - Technology Overview	https://networkbuilders.intel.com/solutionslibrary/power-management-technology-overview-technology-guide

2 Overview

Intel SST-PP profiles configure the CPU with different performance vectors like the number of online cores, Thermal Design Point (TDP), base frequency and others. Intel SST-PP can have several different profiles, depending on the CPU model number.

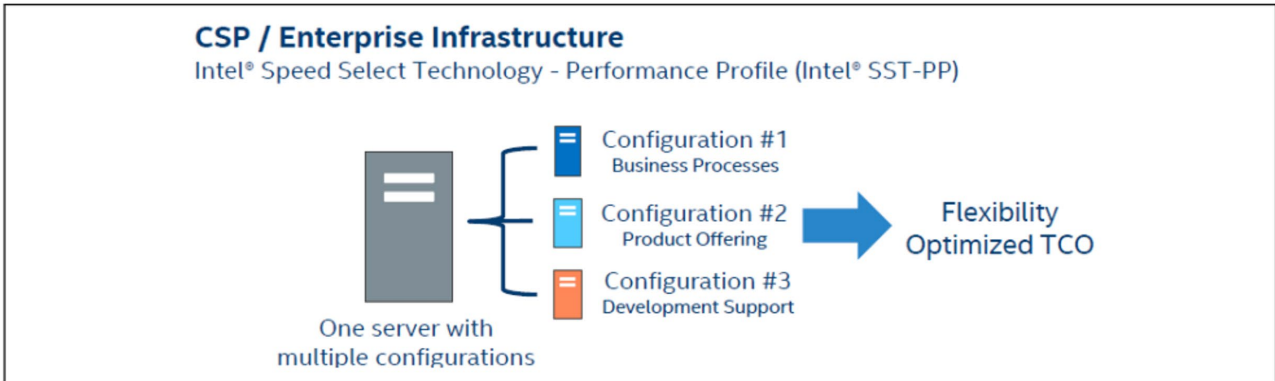


Figure 1. One server with multiple configurations

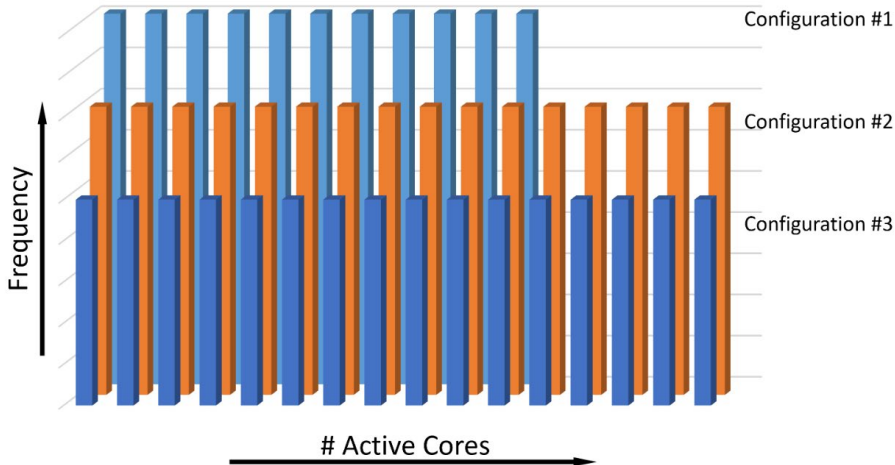


Figure 2. Intel® SSP-PP configurations

The principle use case for Intel SST-PP is:

- CPU SKU consolidation by using one SKU for multiple workload and application types
- Improved server utilization

3 Deployment

In terms of deployment, there are two possible options:

- Static - BIOS discovery and configuration. BIOS may lock the configuration to a selected level such that the OS (operating system) level software cannot switch the configuration at runtime. In this case, cores that are not used by a particular Intel SST-PP profile are offline and not visible to the Linux kernel.
- Dynamic - Performance levels (also called performance profile) can be switched dynamically. This relies on support from OS/software to discover online/offline cores according to the core mask in the selected profile and can be switched using the `intel-speed-select` tool.

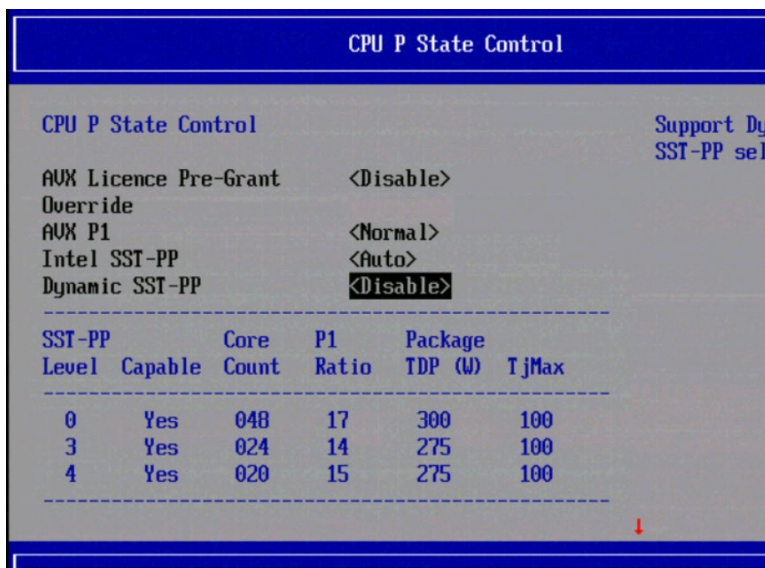
Note: The preferred method of dynamic configuration is to put cores online or offline as per the selected Intel SST-PP profile. This is achieved by using the '-o' option of the `intel-speed-select` tool when switching profile.

4 BIOS Settings

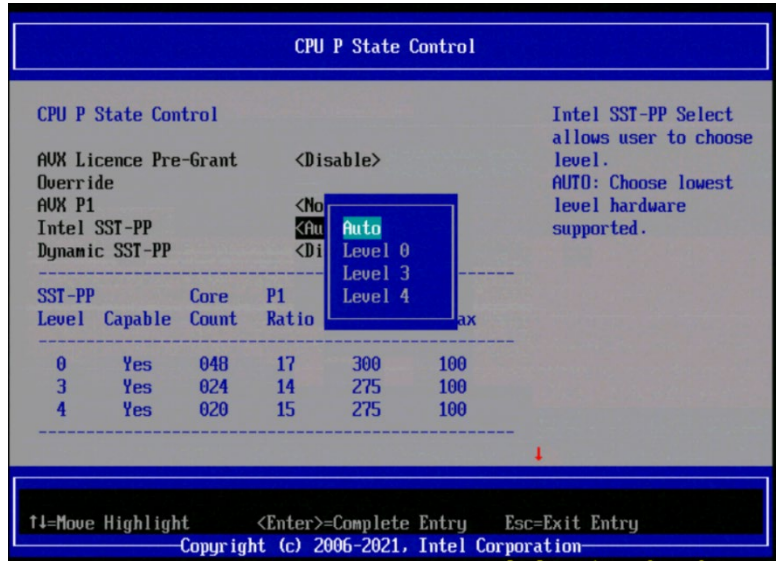
The following screenshots show the different BIOS settings that enable dynamic or static Intel SST-PP on the system.

4.1 Static Intel SST-PP Configuration

To enable static Intel SST-PP, disable the **Dynamic SST-PP** option in the BIOS.



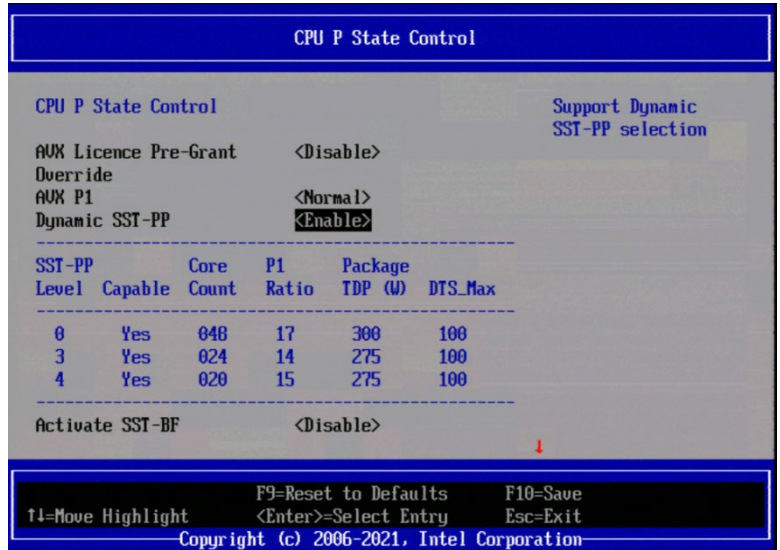
The user can select the SST-PP level to be used when the system is booted. By selecting **Auto**, the profile with the lowest number will be selected.



When static SST-PP is selected, the dynamic profile selection from the OS is not available.

4.2 Dynamic Intel SST-PP Configuration

To configure dynamic Intel SST-PP using `intel-speed-select` tool, BIOS should have the **Dynamic SST-PP** option enabled.



Note: When Dynamic SST-PP is enabled, the BIOS menu option to select an **SST-PP profile** is removed (as seen in static configuration above), and the selection of Intel SST-PP profiles is expected to be done with the `intel-speed-select` tool in the OS.

5 Operating System Configuration

5.1 Static Configuration

Static Intel SST-PP profile selection is done in BIOS. The selected profile comes up after the system is booted up.

5.2 Dynamic Configuration

In order to use Intel SST-PP, the Linux OS must use a supported kernel version. The minimum supported kernel version is 5.3.

There are two important kernel configuration options that need to be enabled for the recommended configuration of Intel SST-PP.

1. CONFIG_HOTPLUG_CPU option is set to 'y': This allows the `intel-speed-select` tool to put cores online or offline as required by the various profiles.
2. "cpu0_hotplug" must be added to the kernel boot parameters: This allows the Linux kernel to be migrated from core 0 in the event that core 0 is put offline by a profile change.

See below for the proper usage of the `intel-speed-select` tool, including the recommended `-o` option.

6 Install and Use the `intel-speed-select` Tool

This guide references a command-line `intel-speed-select` tool. It requires a Linux kernel version 5.3 or higher, and usually is installed by default. For this guide, version 1.8 was used.

To find version of the tool, add `-v` as a command-line parameter:

```
# intel-speed-select -v
Intel(R) Speed Select Technology
---snip---
Version v1.8
Build date Jul 28 2021 time 08:07:56
```

If the version that comes with the distribution is older than 1.4, it is advised to either upgrade the kernel, or compile the tool from source. The source code can be downloaded either from Linux kernel source tree, or from the following page:

<https://github.com/spandrudava/intel-speed-select-utility-src-packages>. The `intel-speed-select` tool is intended to be run from the root account.

6.1 Check Availability of Intel SST-PP

Platforms may or may not support the Intel SST-PP, so ensure to check its availability by using the following command. Also look for “Intel® SST-PP (feature perf-profile) is supported”.

```
# intel-speed-select --info
Intel(R) Speed Select Technology
--snip--
Intel(R) SST-PP (feature perf-profile) is supported
TDP level change control is locked
--snip--
```

6.2 Display Help Text

To display the help text that is specific to Intel SST-PP technology, the “perf-profile” option must be specified, along with the `-h` option to display help message. For example, the following command will list all available Intel SST-PP related commands.

```
# intel-speed-select perf-profile -h
Intel(R) Speed Select Technology
---snip---
Caching topology information
perf-profile:  An architectural mechanism that allows multiple optimized
                performance profiles per system via static and/or dynamic
                adjustment of core count, workload, Tjmax, and
                TDP, etc.
Commands : For feature=perf-profile
            info
            get-lock-status
            get-config-levels
            get-config-version
            get-config-enabled
            get-config-current-level
            set-config-level
```

To list help for a specific command from the above list, command from the above list must be added as a command-line parameter, followed by a `-h` option.

```
# intel-speed-select perf-profile get-lock-status -h
Intel(R) Speed Select Technology
---snip---
Print TDP lock status [No command arguments are required]
```

6.3 Display System Status and Configuration

To display the general package status and which Intel SST-PP technology profiles are supported, specify the `info` command.

```
# intel-speed-select perf-profile info
Intel(R) Speed Select Technology
---snip---
```

User Guide | Intel® Speed Select Technology – Performance Profile (Intel® SST-PP) Overview

```
package-0
  die-0
    cpu-0
      perf-profile-level-0
        cpu-count:96
        enable-cpu-count:96
      enable-cpu-list:0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,
        16,17,18,19,20,21,22,23,24,25,26,27,28,29,30,31,32,33,34,
        35,36,37,38,39,40,41,42,43,44,45,46,47,96,97,98,99,100,101,
        102,103,104,105,106,107,108,109,110,111,112,113,114,115,116,
        117,118,119,120,121,122,123,124,125,126,127,128,129,130,131,
        132,133,134,135,136,137,138,139,140,141,142,143
        base-frequency(MHz):1700
        thermal-design-power(W):300
      --snip--
      perf-profile-level-3
        cpu-count:96
      enable-cpu-count:48
      enable-cpu-list:1,2,3,5,7,10,15,16,19,21,22,23,25,27,28,29,
        34,35,36,40,41,42,43,46,97,98,99,101,103,106,111,112,115,
        117,118,119,121,123,124,125,130,131,132,136,137,138,139,142
        thermal-design-power-ratio:14
        base-frequency(MHz):1400
        thermal-design-power(W):275
      --snip--
      perf-profile-level-4
        cpu-count:96
        enable-cpu-count:40
        enable-cpu-list:1,3,7,10,11,13,15,20,21,23,24,25,27,28,34,36,42,
          44,46,47,97,99,103,106,107,109,111,116,117,119,120,121,123,
          124,130,132,138,140,142,143
        base-frequency(MHz):1500
        thermal-design-power(W):275
      --snip--
```

You can fetch the info about each performance profile by using the `info` command as shown below:

```
# intel-speed-select perf-profile info -l 3
Intel(R) Speed Select Technology
---snip---
package-0
  die-0
    cpu-0
      perf-profile-level-3
        cpu-count:96
        enable-cpu-count:48
        enable-cpu-list:1,2,3,5,7,10,15,16,19,21,22,23,25,27,28,29,
          34,35,36,40,41,42,43,46,97,98,99,101,103,106,111,112,115,
          117,118,119,121,123,124,125,130,131,132,136,137,138,139,142
        base-frequency(MHz):1400
        thermal-design-power(W):275
      --snip--
```

6.4 Display Current Performance Profile

By default, performance profile 0 is enabled. To view the current performance profile, issue the sub-command `get-config-current-level` as shown below.

```
# intel-speed-select perf-profile get-config-current-level
Intel(R) Speed Select Technology
---snip---
package-0
  die-0
    cpu-0
      get-config-current_level:0
```

```
package-1
die-0
cpu-48
get-config-current_level:0
```

6.5 Configure the Performance Profile

To configure the new profile, use the `set-config-level` sub-command as shown below.

```
# intel-speed-select perf-profile set-config-level -l 3 -o
Intel(R) Speed Select Technology
---snip---
package-0
die-0
cpu-0
perf-profile
set_tdp_level:success
```

Option `-l`: `-l` represents the SST-PP level

Option `-o`: `-o` indicates that the cores should be put online or offline, depending on the core count of the profile.

Note: Using the `-o` option is highly recommended when changing profiles. This will put cores online or offline as appropriate for a particular profile. In the event that the `-o` option is not used the unused cores will still be available to the Linux scheduler to schedule the work on to those cores however, the cores will run at a significantly reduced frequency.

6.6 Verify the Performance Profile

To verify that the correct Intel SST-PP profile is set, run the `get-config-current-level` command.

The user should also verify the following to confirm the settings.

- Check that the offline cores are listed under `/sys/devices/system/cpu/offline`.
- Check that the enabled cores are listed under `/sys/devices/system/cpu/online`.
- If Hyper Threading is enabled, all sibling threads that belong to the physical core are also enabled for that profile.
- After profile is set, check whether the correct base frequency is set for the cores listed under “`enable_cpu_mask/enable-cpu-list`” of the profile by checking the `/sys/devices/system/cpu/cpu<N>/cpufreq/base_frequency`. The “N” represents the core number from the “`enable_cpu_mask/enable-cpu-list`”.
- The cores running the workload should have Intel SST-PP profile base frequency reached during the load. To check if the average frequency of the cores is equal to the base frequency, disable the turbo as shown below:

```
echo 1 > /sys/devices/system/cpu/intel_pstate/no_turbo
```

6.7 Example of Intel SST-PP Profile

This is an example of how to set an Intel SST-PP profile and check that the attributes after the profile, is set. Firstly, the performance profile level is set, then a workload is run on the cores to ramp them up to their base frequency, and then check to see what base frequency is reached. In this example, Intel SST-PP profile 3 is used.

```
# intel-speed-select perf-profile set-config-level -l 3 -o
```

This results in the following output, which indicates which cores it has put online/offline:

```
set_tdp_level:success
offline cpu 0
online cpu 1
online cpu 2
---snip---
online cpu 46
offline cpu 47
```

Next, we show the information of that profile, making note of the base frequency, in this case, 1400 MHz.

```
# intel-speed-select perf-profile info -l 3
Intel(R) Speed Select Technology
---snip---
package-0
die-0
```


User Guide | Intel® Speed Select Technology – Performance Profile (Intel® SST-PP) Overview

```
cpu-0
perf-profile-level-3
cpu-count:96
enable-cpu-count:48
enable-cpu-list:1,2,3,5,7,10,15,16,19,21,22,23,25,27,28,
29,34,35,36,40,41,42,43,46,97,98,99,101,103,106,111,
112,115,117,118,119,121,123,124,125,130,131,132,136,
137,138,139,142
base-frequency (MHz) :1400
base-frequency-avx2 (MHz) :1300
thermal-design-power (W) :275
```

Disable turbo boost so the cores will not go above the base frequency (that is, the cores will not go into turbo frequency) when under load.

```
# echo 1 > /sys/devices/system/cpu/intel_pstate/no_turbo
```

Run a workload on all the cores on the CPU, for example sysbench.

<https://github.com/akopytov/sysbench>

Run turbostat and observe the output.

```
# turbostat --show Package,Core,CPU,Bzy_MHz -i 1 -n 1
```

The turbostat output shows that the cores enabled for the level3 profiles are running at level3 profile base frequency of 1400 Mhz.

Pkg	Core	CPU	Bzy_MHz
-	-	-	1400
0	1	1	1400
0	1	97	1400
0	2	2	1400
0	2	98	1400
0	3	3	1400
0	3	99	1400
--- snip ---			
0	42	42	1400
0	42	138	1400
0	43	43	1400
0	43	139	1400
0	46	46	1400
0	46	142	1400

6.8 Intel SST-PP Profile Differences

For the performance profiles on a particular CPU model, we may see different numbers of cores, or different base frequencies. This depends on the specifications of that particular CPU model.

As an example, from the `intel-speed-select` tool output we can see the differences between Intel SST-PP profiles on one particular SKU.

```
perf-profile-level-0
enable-cpu-count:20
enable-cpu-list:0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19
base-frequency (MHz) :1700
perf-profile-level-3
enable-cpu-count:14
enable-cpu-list:0,1,2,3,6,8,9,11,12,13,14,15,16,18
base-frequency (MHz) :1400
perf-profile-level-4
enable-cpu-count:20
enable-cpu-list:0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19
base-frequency (MHz) :1500
```

In this case, profile 0 and profile 4 have the same number of cores, but a different base frequency. This is a case where a lower TDP can be configured.

Also, a lower TDP can be achieved by selecting profile 3, in which case the number of cores and the base frequency is reduced.

7 Summary

This document describes the Intel Speed Select Technology – Performance Profile (Intel SST-PP), and how to configure it on the 3rd Generation Intel® Xeon® Scalable processors. Intel SST-PP facilitates multiple CPU configurations, allowing the same SKU to be used for different workload and application types. This is achieved by allowing the CPU to be reconfigured into one of the available Intel SST-PP profiles, which may contain different number of cores or TDP levels.

The document also contains key recommendations on the use of the software when configuring the Intel SST-PP profiles.



Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

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.

Intel technologies may require enabled hardware, software or service activation.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.