

Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF) Overview

Authors

Reshma Pattan
Anatoly Burakov

1 Introduction

Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF) offers an ability to assign prioritization on turbo frequency to specific cores. By using this feature, some cores can be configured to get higher turbo frequency by designating them as high priority at the cost of low or no turbo frequency on the low priority cores. Intel® SST-TF is one of the technologies that falls under the Intel® Speed Select Technology umbrella. For basic overview of Intel SST-TF and other technologies that comes under the Intel Speed Select Technology umbrella, refer to the *Power Management – Technology Overview guide* provided in the References section of this guide. This document shows how to configure Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF), which is available on selected SKUs of Intel® Xeon® Scalable processors. The document also describes the usage of a Linux* kernel tool that was developed to configure Intel SST-TF technology on a platform. The tool provides a convenient, easy-to-use interface that aids configuration and hides complexity. For more information about Intel Speed Select Technology – Turbo Frequency, see the References section.

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	System Requirements.....	4
3.1	Hardware	4
3.2	BIOS.....	4
3.3	Operating System Configuration	4
4	Install the intel-speed-select Tool.....	4
4.1	Check Availability of Intel SST-TF	4
4.2	Display Help Text.....	4
4.3	Intel SST-TF dependency on Intel® Speed Select Technology - Performance Profile	5
4.4	Get Base Turbo Capability	5
4.5	Get Intel SST-TF Capability.....	5
4.6	Intel SST-TF Buckets.....	6
4.6.1	Understanding the Intel SST-TF Bucket Information.....	6
4.6.2	Enabling Intel SST-TF on Single Socket	7
4.6.3	Enabling Intel SST-TF on Multiple Socket	7
4.7	Different Intel SST-TF Configurations and their Behavior	7
5	Summary	8

Figures

Figure 1.	Intel® SST-TF configuration example.....	3
-----------	--	---

Tables

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

Document Revision History

REVISION	DATE	DESCRIPTION
001	July 2022	Initial release.

1.1 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
CLOS	Class of Service
TDP	Thermal Design Point
Intel® AVX2	Intel® Advanced Vector Extensions 2
Intel® AVX-512	Intel® Advanced Vector Extensions 512
Intel® SSE	Intel® Streaming SIMD Extensions
Intel® SST-CP	Intel® Speed Select Technology - Core Power
Intel® SST-TF	Intel® Speed Select Technology - Turbo Frequency
Intel® SST-PP	Intel® Speed Select Technology - Performance Profile

1.2 Reference Documentation

Table 2. Reference Documents

REFERENCE	SOURCE
Intel® Speed Select Technology (Intel® SST) Tool Guide	https://www.kernel.org/doc/html/latest/admin-guide/pm/intel-speed-select.html
Power Management – Technology Overview	https://networkbuilders.intel.com/solutionslibrary/power-management-technology-overview-technology-guide

2 Overview

Intel SST-TF allows the maximum turbo frequency of certain number of CPU cores to be increased (high priority cores) in exchange for lowering the maximum turbo frequency on the rest of the CPU cores (lower priority cores).

In certain workloads, there are small number of cores that block the overall system performance. By assigning higher frequencies to those cores with Intel SST-TF, significant opportunistic frequency improvements can be realized with no change in software. [Figure 1](#) gives an example of the differences between Legacy All-Core Turbo mode and Intel SST-TF mode. With a low number of high priority turbo cores, a higher turbo frequency can be achieved by these cores. As the number of high priority cores increases, the turbo frequency reduces until it is at the normal Legacy All-Core Turbo Limit.

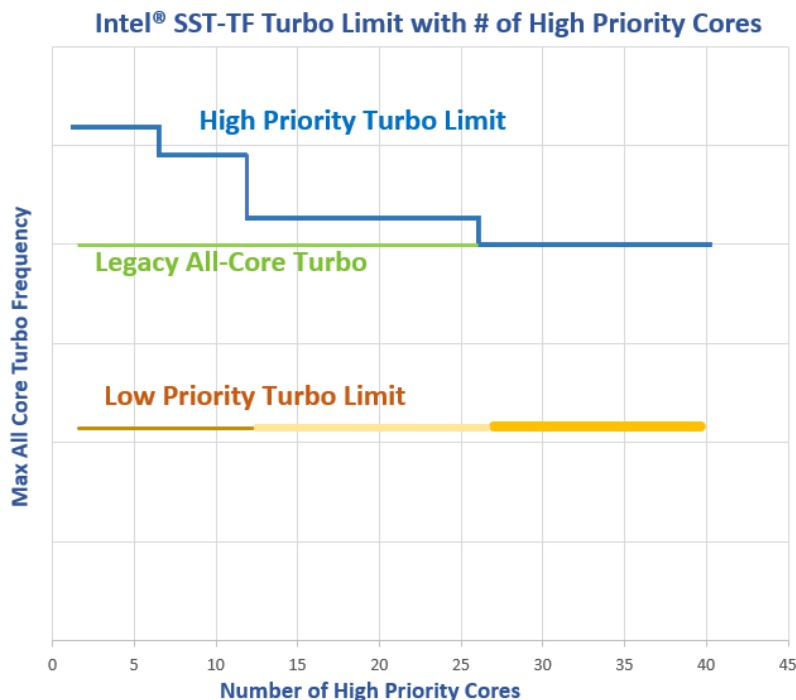


Figure 1. Intel® SST-TF configuration example

3 System Requirements

3.1 Hardware

Requires Intel® CPU enabled with Intel SST-TF and Intel® Speed Select Technology - Core Power (Intel® SST-CP)

3.2 BIOS

The following BIOS settings are needed to configure and use Intel SST-TF on the system:

- Enable the Intel SST-CP bios option
- Disable Hyper-Threading to simplify the usage demonstration for the document

3.3 Operating System Configuration

The Intel® Speed Select config tool allows the configuration of Intel SST-TF via Intel SST OS mailbox interface . Linux kernel v5.8 or later is required for full functionality. The Intel SST mailbox interface is enabled in kernel config using the config option CONFIG_INTEL_SPEED_SELECT_INTERFACE=m. To check if CONFIG_INTEL_SPEED_SELECT_INTERFACE is set on the system, check the boot config file as below.

```
cat /boot/config-$(uname -r) | grep SPEED_SELECT
```

For custom kernels, refer to your Linux distribution documentation on how to configure and build the kernel with option CONFIG_INTEL_SPEED_SELECT_INTERFACE=m enabled.

4 Install the intel-speed-select Tool

This paper references a command-line `intel-speed-select` tool. It requires a Linux kernel version 5.8 or higher, and usually is installed by default. For this paper, version 1.10 of tool 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 v1.10
```

If the tool 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 from the below Linux kernel source tree:

<https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git/tree/tools/power/x86/intel-speed-select>

To understand and install the tool using the source code please Refer to the [Intel® Speed Select Technology \(Intel® SST\) Tool Guide](#).

The `intel-speed-select` tool is intended to be run from the root account.

4.1 Check Availability of Intel SST-TF

Platforms may or may not support Intel SST-TF, so ensure to check the availability of Intel® SST-TF by using the following `intel-speed-select` command. Also look for Intel(R) SST-TF (feature `turbo-freq`) is supported.

```
# intel-speed-select --info
Intel(R) Speed Select Technology
--snip--
Intel(R) SST-TF (feature turbo-freq) is supported
--snip--
```

4.2 Display Help Text

To display the help text that is specific to Intel SST-TF technology, the `turbo-freq` option must be specified along with the `-h` option to `intel-speed-select` command. For example, the following command will list all available Intel SST-TF related commands.

```
#intel-speed-select turbo-freq -h
Intel(R) Speed Select Technology
---snip---
turbo-freq:      Enables the ability to set different turbo ratio
                  limits to cores based on priority.
Command: For feature=turbo-freq
            command : info
            command : enable
            command : disable
```

User Guide | Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF) Overview

To view help for a specific command from the above list, add them as a command-line parameter, followed by a -h option.

```
# intel-speed-select turbo-freq info -h
Intel(R) Speed Select Technology
---snip---
```

Enable Intel Speed Select Technology Turbo frequency feature
Optional: -t|--tr1 : Specify turbo ratio limit

Optional Arguments: -a|--auto: Designate specified target CPUs with-C|--cpu option as high priority using core-power feature

4.3 Intel SST-TF dependency on Intel® Speed Select Technology - Performance Profile

The support of Intel SST-TF depends on the performance level configuration of Intel® Speed Select Technology - Performance Profile (Intel® SST-PP). It is possible that only a certain performance level supports Intel SST-TF. It is also possible that only the base performance level (level = 0) has the support of Intel SST-TF. Hence, first select the desired performance level to enable this feature. In the system under test here, Intel SST-TF is supported at the base performance level 0. However, it is currently disabled.

Note: On systems where Intel® SST-PP feature is not supported, level 0 is still valid. So make sure to check if Intel SST-TF is supported by this level 0.

```
# intel-speed-select -c 0 perf-profile info -l 0
package-0
  die-0
    cpu-0
      perf-profile-level-0
---snip---
```

speed-select-turbo-freq:disabled

4.4 Get Base Turbo Capability

To get the base turbo capability of performance level 0, run the following command:

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

package-0
die-0
cpu-0
perf-profile-level-0
---snip---

turbo-ratio-limits-sse
bucket-0
core-count:8
max-turbo-frequency (MHz) :3400
bucket-1
core-count:12
max-turbo-frequency (MHz) :3400
bucket-2
core-count:16
max-turbo-frequency (MHz) :3400
bucket-3
core-count:20
max-turbo-frequency (MHz) :3300
bucket-4
core-count:24
max-turbo-frequency (MHz) :3100
bucket-5
core-count:26
max-turbo-frequency (MHz) :3000
bucket-6
core-count:30
max-turbo-frequency (MHz) :2800
bucket-7
core-count:32
max-turbo-frequency (MHz) :2800

4.5 Get Intel SST-TF Capability

To display Intel SST-TF capability for performance level 0, use the “turbo-freq info” command:

```
#intel-speed-select turbo-freq info -l 0
Intel(R) Speed Select Technology
package-0
  die-0
```

```
cpu-0
speed-select-turbo-freq-properties
  bucket-0
    high-priority-cores-count:8
    high-priority-max-frequency (MHz) :3400
    high-priority-max-avx2-frequency (MHz) :3400
    high-priority-max-avx512-frequency (MHz) :3300
  bucket-1
    high-priority-cores-count:16
    high-priority-max-frequency (MHz) :3100
    high-priority-max-avx2-frequency (MHz) :3100
    high-priority-max-avx512-frequency (MHz) :3100
  bucket-2
    high-priority-cores-count:24
    high-priority-max-frequency (MHz) :2900
    high-priority-max-avx2-frequency (MHz) :2800
    high-priority-max-avx512-frequency (MHz) :2800
speed-select-turbo-freq-clip-frequencies
  low-priority-max-frequency (MHz) :2200
  low-priority-max-avx2-frequency (MHz) :1700
  low-priority-max-avx512-frequency (MHz) :1400
```

4.6 Intel SST-TF Buckets

Before enabling Intel SST-TF, we need to understand the internals of how Intel SST-TF works.

Intel SST-TF works by setting up two tiers of cores - high priority cores and low priority cores (using Intel SST-CP CLOS groups). The cores can be freely assigned to be high or low priority. However, the number of cores that can be assigned to each group cannot be changed. It is up to the predefined configurations, called "buckets". There are different buckets available for different Config TDP levels. For config TDP level 0, the buckets can be read using the following command:

```
#intel-speed-select turbo-freq info -l 0
Intel(R) Speed Select Technology
package-0
  die-0
    cpu-0
      speed-select-turbo-freq-properties
        bucket-0
          high-priority-cores-count:8
          high-priority-max-frequency (MHz) :3400
          high-priority-max-avx2-frequency (MHz) :3400
          high-priority-max-avx512-frequency (MHz) :3300
        bucket-1
          high-priority-cores-count:16
          high-priority-max-frequency (MHz) :3100
          high-priority-max-avx2-frequency (MHz) :3100
          high-priority-max-avx512-frequency (MHz) :3100
        bucket-2
          high-priority-cores-count:24
          high-priority-max-frequency (MHz) :2900
          high-priority-max-avx2-frequency (MHz) :2800
          high-priority-max-avx512-frequency (MHz) :2800
      speed-select-turbo-freq-clip-frequencies
        low-priority-max-frequency (MHz) :2200
        low-priority-max-avx2-frequency (MHz) :1700
        low-priority-max-avx512-frequency (MHz) :1400
```

4.6.1 Understanding the Intel SST-TF Bucket Information

Intel® SST-TF supports three buckets bucket-0, bucket-1, bucket-2, and a clip frequency. Based on the number of high priority cores the turbo bucket will be selected. Each bucket has Intel® Streaming SIMD Extensions (Intel® SSE), Intel® Advanced Vector Extensions 2 (Intel® AVX2), and Intel® Advanced Vector Extensions 512 (Intel® AVX-512) turbo ratio limits (frequency) points assigned to it, so each bucket has ability to take the high priority core frequencies to high as the given turbo ratio limit.

As an example, if we were to set up 16 high priority cores and run an Intel® AVX2 workload, we should see 3.1 GHz on high priority cores (because having 16 high priority cores puts the cores into bucket-1), and low priority cores should go up to 2.2 GHz (clip frequency) at most. They may be lower due to TDP constraints and the fact that lower priority cores get scaled down first.

4.6.2 Enabling Intel SST-TF on Single Socket

While prioritizing the cores for Intel SST-TF, user can pass different number of cores to the Intel SST tool command line. In the below commands, the cores passed with `-c` arguments are marked as high priority, including its Hyper-Threading siblings.

The option `-auto` sets the specified cores as high priority and rest of the cores as low priority using the Intel SST-CP CLOS groups.

Refer to the [Intel® Speed Select Technology \(Intel® SST\) Tool Guide](#) to know more about the “-auto” option.

Single Core: Configures one core as high priority

```
intel-speed-select -c 5 turbo-freq enable -auto
```

List of cores: Configures comma separated cores as high priority

```
intel-speed-select -c 0,1,3 turbo-freq enable -auto
```

Enabling range of cores: Configures specified range of cores as high priority

```
intel-speed-select -c 0-6 turbo-freq enable -auto
```

All cores: Configures all the cores of the socket as high priority

```
intel-speed-select -c 0-31 turbo-freq enable -auto
```

4.6.3 Enabling Intel SST-TF on Multiple Socket

The below commands show how to configure Intel SST-TF with the cores from the multi-socket system. For example, on a machine that has two CPUs with 32 cores, a command affecting cores 5-8 on each CPU will look like the following:

```
#intel-speed-select -c 5-8,37-40 turbo-freq enable --auto
```

Similarly, to ensure the package wide command affects all cores on the system, add a `-c` parameter specifying first core on each CPU (assuming there are two 32-core CPUs)

```
#intel-speed-select -c 0,32 turbo-freq enable --auto
```

4.7 Different Intel SST-TF Configurations and their Behavior

This section shows how to use Intel SST-TF to prioritize different set of cores and understand the system behavior with respect to core frequency changes with the number of cores configured as high priority cores.

- When Intel SST-TF is configured with the single core (for example, core 5), as shown below, the configured core will have the frequency specified in bucket-0 and all other cores of the CPU will have the frequencies up to those specified in the frequency clip bucket. Check Intel SST-TF information using the command listed in [section 4.6](#) to see the bucket-0 and clip bucket frequencies of the system for various workload types (Intel SSE, Intel AVX2, Intel AVX-512). Assuming that core5 is running with Intel SSE workload, the core5 will run at 3400 Mhz as per Intel SST-TF information in [section 4.6](#) and rest of the cores will run at clip frequency of 2200 Mhz.

```
#intel-speed-select -c 5 turbo-freq enable -auto
```

- When Intel SST-TF is configured with all cores of the CPU (for example, assuming 32 cores on the CPU), all cores of CPU will run at the same frequency as they would run when SST-TF is not enabled.

```
#intel-speed-select -c 0-31 turbo-freq enable -auto
```

- When Intel SST-TF configured with the maximum number of cores allowed for bucket-0 (for example, assuming bucket-0 allows up to 8 cores) as shown below, the configured cores will have frequencies specified in bucket-0, while all other cores will have frequencies up to those specified in the frequency clip bucket.

```
#intel-speed-select -c 1-8 turbo-freq enable --auto
```

- When Intel SST-TF is configured with the number of cores that is one bigger than maximum number of cores for bucket-0 (for example, if bucket-0 allows up to 8 cores, configuring with 9 cores), will let the configured cores to have frequencies specified in bucket-1, while all other cores will have frequencies up to those specified in the frequency clip bucket.
Note: Configuring one more core than bucket-0 cores, all the configured cores will have their frequency transitioned from bucket-0 to bucket-1 frequency. Check Intel SST-TF information using the command listed in [section 4.6](#) to see the bucket-0, bucket-1, and clip bucket frequencies of the system for various workload types (Intel SSE, Intel AVX2, Intel AVX-512). Assuming the cores 1-9 are running with Intel SSE workload, the cores 1-9 will run at 3100 Mhz as per Intel SST-TF information in [section 4.6](#) and rest of the cores will run at clip frequency of 2200 Mhz.

```
#intel-speed-select -c 1-9 turbo-freq enable -auto
```

- When configuring Intel SST-TF with maximum number of cores allowed for bucket-1 (for example, assuming bucket-1 allows up to 16 cores), the configured cores will have frequencies specified in bucket-1, while all other cores will have frequencies up to those specified in the frequency clip bucket.

```
#intel-speed-select -c 1-16 turbo-freq enable -auto
```

User Guide | Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF) Overview

- When configuring Intel SST-TF with the number of cores that's one bigger than maximum number of cores for bucket-1 (for example, if bucket-1 allows for 16 cores, configuring with 17 cores), will let the configured cores have frequencies specified in bucket-2, while all other cores will have frequencies up to those specified in the frequency clip bucket.
Note: Configuring one more core than bucket-1 cores, all the configured cores will have their frequency transitioned from bucket-1 to bucket-2 frequency.

```
#intel-speed-select -c 1-17 turbo-freq enable -auto
```

- When configuring Intel SST-TF with maximum number of cores allowed for bucket-2 (for example, assuming bucket-2 allows up to 24 cores), the configured cores will have frequencies specified in bucket-2, while all other cores will have frequencies up to those specified in the frequency clip bucket.

```
#intel-speed-select -c 1-24 turbo-freq enable -auto
```

- When configuring Intel SST-TF with the number of cores that's one bigger than maximum number of cores for bucket-2 (for example, if bucket-2 allows for 24 cores, configuring 25 cores), the configured cores will have frequencies specified in bucket-2 or lower, while all other cores will have frequencies up to those specified in the frequency clip bucket.

Note: Due to the way Intel SST-TF is intended to function, specifying more cores than bucket-2 will not force all cores into the frequency clip bucket. Instead, the high priority cores may have their max frequency adjusted by one or more bins, subject to TDP constraints.

```
#intel-speed-select -c 1-25 turbo-freq enable --auto
```

5 Summary

This document describes Intel® Speed Select Technology – Turbo Frequency (Intel® SST-TF), and how to configure it on selected Intel® Xeon® Scalable processors. Intel® SST-TF facilitates increased turbo ratio limits for high priority cores helping the cores to perform faster processing.



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.