APPLICATION NOTE

Intel Corporation

Intel® Speed Select Technology – Base Frequency - Enhancing Performance

Authors

John DiGiglio

David Hunt

Ai Bee Lim

Chris MacNamara

Timothy Miskell

1 Introduction

This document provides an overview of a CPU capability called Intel® Speed Select Technology – Base Frequency (Intel® SST-BF), which is available on select SKUs of 2nd generation Intel® Xeon® Scalable processor (formerly codenamed Cascade Lake). The document also includes benchmarking data and instructions on how to enable the capability.

About this capability:

- This capability is offered on selected SKUs of 2nd generation Intel® Xeon® Scalable processor (5218N, 6230N, and 6252N).
- Intel® SST-BF allows the CPU to be deployed with a guaranteed asymmetric core frequency configuration.
- The placement of key workloads on higher frequency Intel® SST-BF enabled cores can result in an overall system workload performance increase and potential overall energy savings when compared to deploying the CPU with symmetric core frequencies.

This document is part of the Network Transformation Experience Kit, which is available at: https://networkbuilders.intel.com/

1.1 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
DPDK	Data Plane Development Kit
NFV	Network Functions Virtualization
SST-BF	Intel® Speed Select Technology – Base Frequency
VM	Virtual Machine

Table of Contents

Figures

Tables

1.2 Intel® SST-BF Overview

Figure 1 shows both symmetric and asymmetric core frequency deployment. In a symmetric core frequency deployment (default), all applications on a processor operate at the same core frequency. In an asymmetric core frequency deployment model, certain cores run at higher frequency and the remaining cores run at lower frequency.

When Intel® SST-BF is enabled, it allows the CPU to be dynamically switched between asymmetric and asymmetric configuration. This enables users to boost performance of targeted applications at runtime.

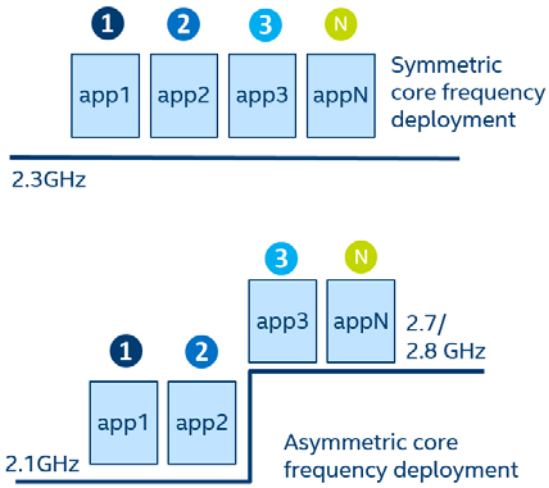

Figure 1. Core Frequency Deployment Methods

Table 2 lists the Intel® Xeon® Scalable processor SKUs that support Intel® SST-BF and compares their corresponding high priority and standard priority frequency values.

Table 2. Intel® SST-BF Enabled CPU SKUs

Intel® Xeon® Scalable processor 6252N (24C @ 2.3GHz @ 150W)			
High Priority		Standard Priority	
Cores	Base Frequency	Cores	Base Frequency
8	2.8 GHz	16	2.1 GHz
Intel® Xeon® Scalable processor 6230N (20C @ 2.3GHz @ 125W)			
High Priority		Standard Priority	
Cores	Base Frequency	Cores	Base Frequency
6	2.7 GHz	14	2.1 GHz
Intel® Xeon® Scalable processor 5218N (16C @ 2.3GHz @ 105W)			
High Priority		Standard Priority	
Cores	Base Frequency	Cores	Base Frequency
4	2.7 GHz	12	2.1 GHz

2 Benchmark Results

To demonstrate the benefit of deploying a system with Intel® SST-BF enabled (asymmetric core frequencies), we simulated a representative NFV deployment scenario where the system node is hosting high priority workloads and low priority workloads.

The High Priority workload hosted in a Virtual Machine is representative of a workload reliant on deterministic compute cycles and benefits from frequency scaling. In this case, we chose the DPDK *Testpmd* packet forwarding sample application. We also targeted the vSwitch component of the system to the higher frequency cores.

We used the *stress* application to generate background work representative of a low priority workload, which is not real-time or performant crucial.

2.1 Test Setup

Figure 2 shows the representative topology of our test setup. The traffic generator was built using the open source TRex traffic generator built on DPDK 19.05. (DPDK downloads are available at: https://git.dpdk.org/dpdk-stable/)

The representative system was connected to the traffic generator system using Intel® XL710 40 GbE Ethernet Adapters.

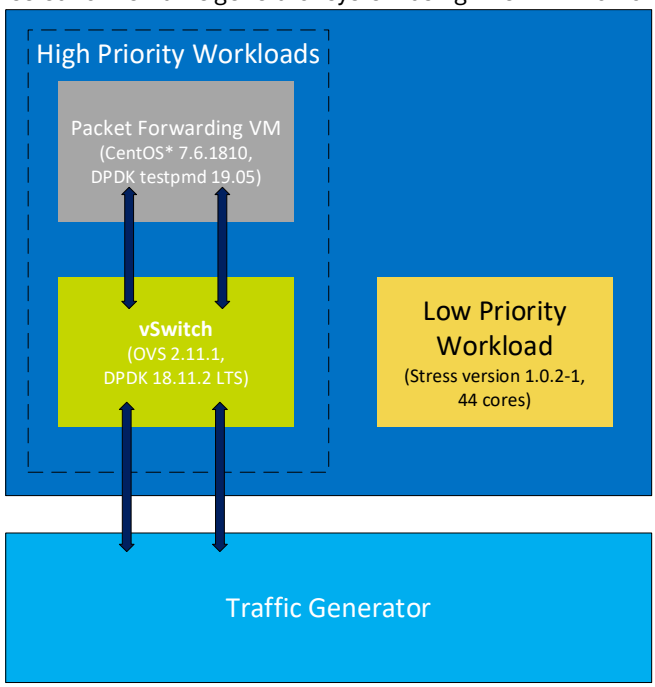

Figure 2. Test Setup

Our representative system had the following characteristics:

- Dual Intel® Xeon® Gold 6252N @2.30 GHz processor
- CentOS* 7.6.1810 (Linux* centos76sstbf 5.1.3-1.el7.elrepo.x86_64 #1).
- BIOS/FW is WW14 BKC SE5C620.86B.02.01.1008.041420190659. Date: 04/14/2019
- OVS 2.11.1 with DPDK 18.11.2
- Virtual Machine uses CentOS* 7.6.1810 as Guest OS.
- Stress application v1.0.2-1

We generated traffic and measured both traffic throughput and consumed wall power with 2 scenarios:

1. Intel® SST-BF feature disabled and the stress application idle (not running in the background).
2. Intel® SST-BF feature disabled and the stress application maximizing activity.

2.2 Test Results

The results from the above 2 scenarios are captured in Figure 3, which uses bars to show throughput (in MPPs) and lines to show wall power (in Watts).

The results indicate that we are able to achieve higher traffic throughput for our high priority workload while staying within the same power consumption. The results were measured with multi-queue enabled for OVS-DPDK with up to 4 queues per port.

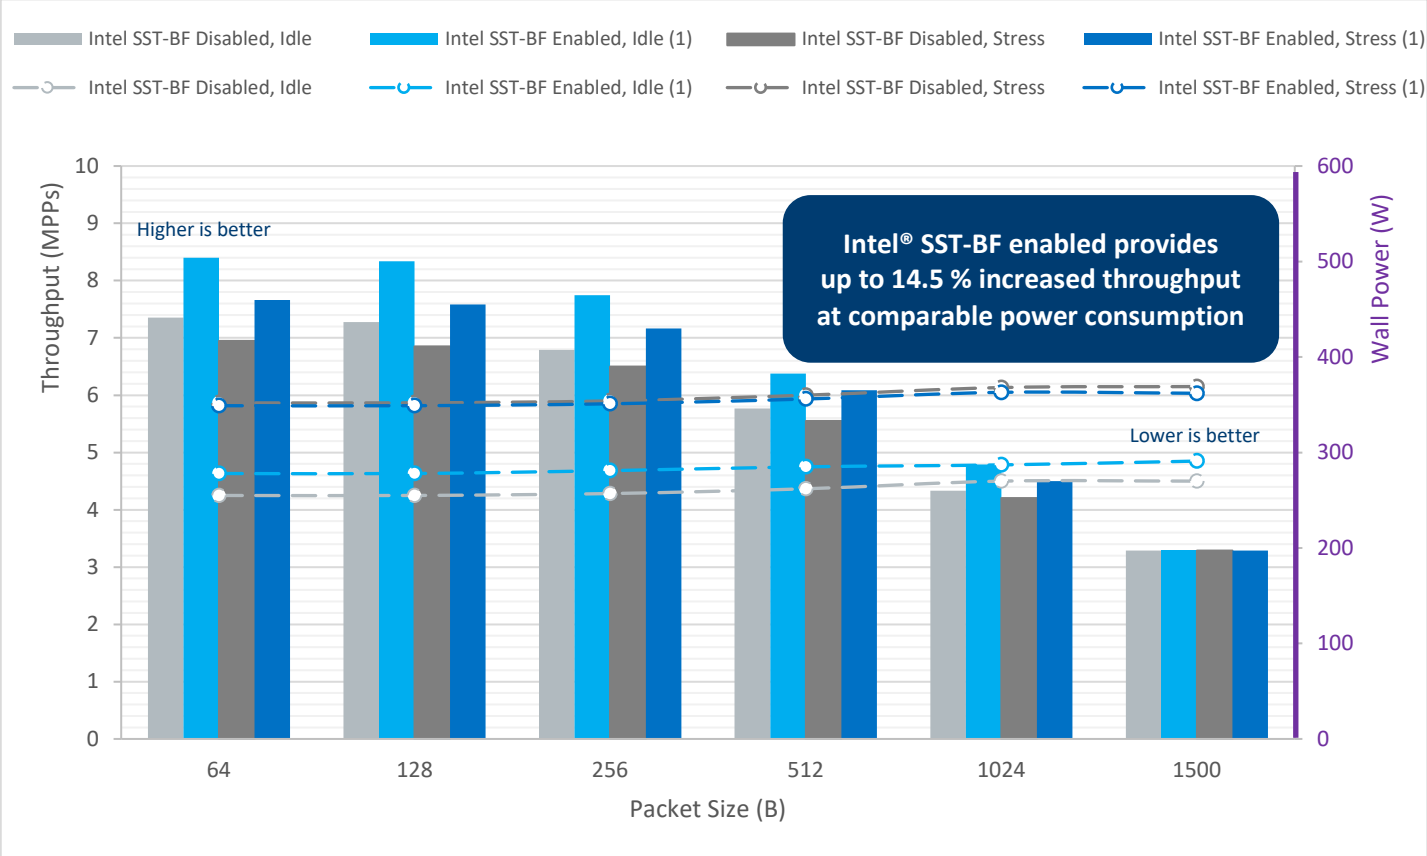

SST-BF Disabled

Setting	Value
Turbo	Disabled
P-state	Disabled
C-state	Disabled
SST-BF	Disabled

SST-BF Enabled

Setting	Value
Turbo	Enabled
P-state	Enabled
C-state	Disabled
SST-BF	Enabled

Figure 3. Test Results

Note:
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

For more information go to www.intel.com/benchmarks.

Performance results are based on testing as of September 25th 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

§ Configurations: Host OS CentOS 7.6.1810, Linux centos76sstbf 5.1.3-1.el7.elrepo.x86_64 #1 SMP Thu May 16 17:49:50 EDT 2019 x86_64 x86_64 x86_64 GNU/Linux, VM CentOS Linux release 7.6.1810, Stress version 1.0.2-1, BIOS/FW: WW14 BKC SE5C620.86B.02.01.1008.041420190659 Date: 04/14/2019, Power measurement: https://help.raritan.com/px3-5000/v3.5.0/en/#46385.htm

3 Setting up Intel® SST-BF Functionality

3.1 Prerequisites

To use the Intel® SST-BF functionality on a Red Hat* Enterprise Linux* based platform you need:

- Linux* Kernel version 5.1 or later. Check with Linux vendors for feature integration.
- Enable the Intel® SST-BF feature in the BIOS.
- Activate the feature using a sample node configuration script, sst_bf.py available at: https://github.com/intel/CommsPowerManagement.

3.2 BIOS Settings

Table 3 shows the required BIOS settings for deterministic performance. For details on how to implement these settings, see Section 3.3.

At the time of writing this document, the latest version of the BIOS is 02.01.0009.

Download the latest BIOS at: https://downloadcenter.intel.com/download/29105/Intel-Server-Board-S2600WF-Family-BIOS-and-Firmware-Update-Package-for-UEFI

Table 3. BIOS Settings

Menu (Advanced)	Path to BIOS Setting	BIOS Settings	Required Setting for Deterministic Performance
Advanced	Processor Configuration	Hyperthreading	Enabled
Power Configuration	Power and Performance	CPU Power and Performance Policy	Performance
		Workload Configuration	I/O Sensitive
	Power and Performance → CPU P-State Control	Enhanced Intel® SpeedStep Technology	Disabled (See **Note**.)
	Power and Performance → Hardware P States	Hardware P States	Disabled (See **Note**.)
	Power and Performance -→ CPU C State Control	Package C-State	C0/C1 state
		C1E	Disabled
		Processor C6	Disabled
	Power and Performance →Uncore Power Management	Uncore Frequency Scaling	Disabled
		Performance P-limit	Disabled
Memory Configuration	Advanced → Memory Configuration	IMC Interleaving	2-Way Interleaving
Virtualization Configuration	Processor Configuration	Intel® Virtualization Technology (VT)	Enabled
	Processor Configuration	Intel® VT for Directed I/O	Enabled
Thermal Configuration	Advanced → System Acoustic and Performance Configuration	Set Fan Profile	Performance

NOTE: Enabled in the case where Intel® SST-BF is enabled to allow for configuration of individual core speeds.

3.3 Configuring BIOS

The following steps, specifically for enabling SST-BF from the BIOS, were performed on an Intel® Server Board S2600WF (formerly codenamed Wolf Pass).

In the BIOS, perform the following steps.

1. Select **Advanced**.

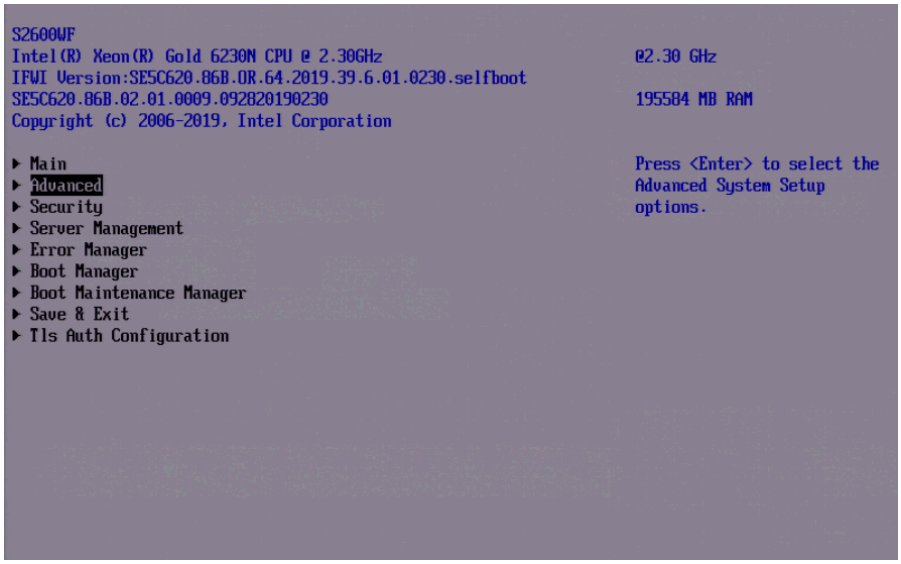

2. Select **Power & Performance**.
   In the **Power & Performance** screen, there are two sections we must check:
   - Hardware P States
   - CPU P State Control

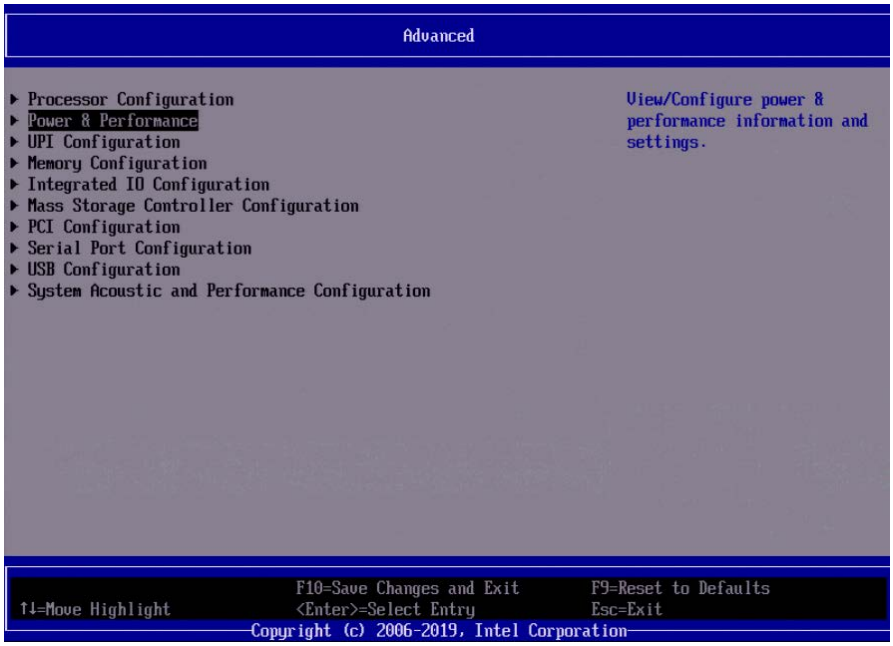

3. In the **Power & Performance screen**, select **Hardware P States**.

4. In the **Hardware PM State Control** screen, ensure that the following options are set:
   - **Hardware P-States**: Native Mode with No Legacy Support
   - **EPP Enable**: Enabled
   - **RAPL Prioritization**: Enabled

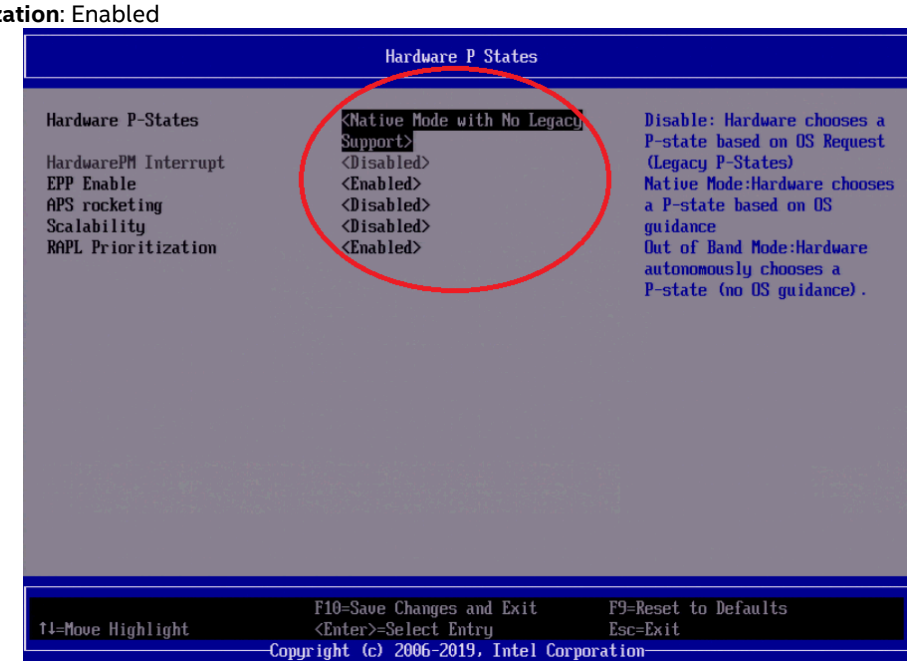

5. Choose **Exit** to return to the **Power & Performance** screen, then select the **CPU P State Control** option.

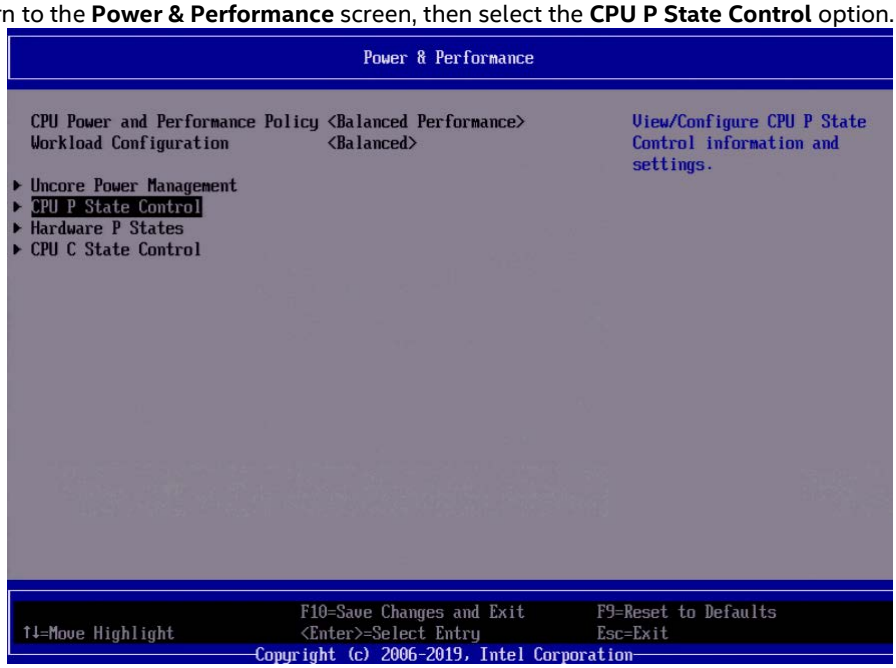

6. In the **CPU P State Control** screen, configure Intel® SST-BF and Intel® Turbo Boost Technology. The following options should be set.

   Note: The **PBF** options in the screenshot below are used for **Intel(R) SST-BF**.

   - **Enhanced Intel SpeedStep(R) Tech**: Enabled
   - **Activate PBF**: Enabled
   - **Configure PBF**: Disabled
   - **Intel Configurable TDP**: Disabled
   - **Intel(R) Turbo Boost Technology**: Enabled
   - **Energy Efficient Turbo**: Disabled

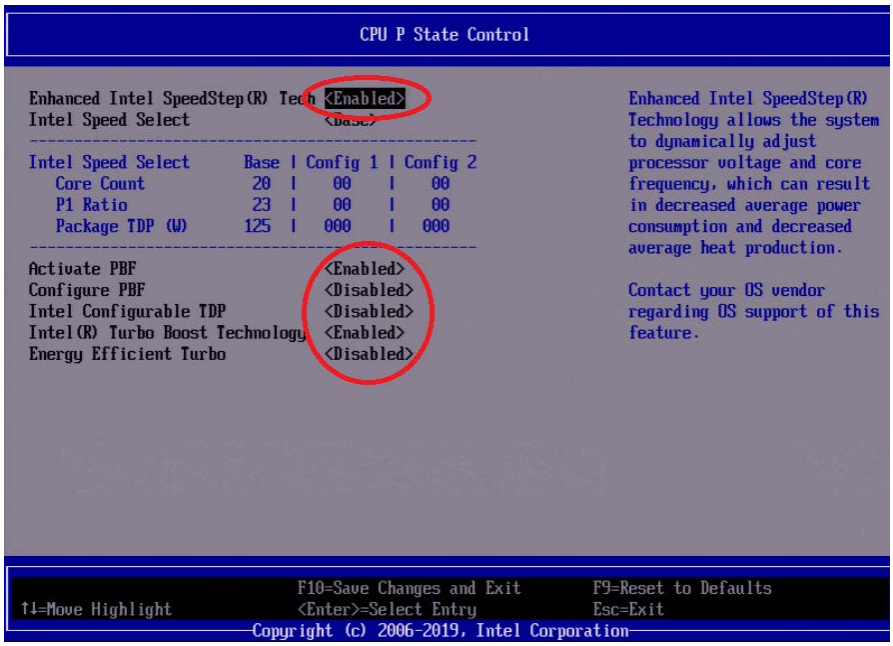

Note: Intel® SST-BF can be used with hyper-threading on or off, but you must be sure that both hyper-thread siblings have the same configuration to get the expected results.

Save your changes and exit BIOS setup.

3.4 Operating System Settings

On a Red Hat* Enterprise Linux* based platform, Intel® SST-BF requires kernel 5.1 and later. In addition, you must boot with the `intel_pstate` driver active. Although the `intel_pstate` driver is typically disabled in order to provide deterministic performance, the results shown in Figure 3 indicate that we are still able to achieve deterministic performance with the `intel_pstate` driver enabled and Intel® SST-BF enabled.

Note: Although the OS might disable the intel_pstate driver by default, **do not** set `intel_pstate=disable` or `intel_pstate=no_hwp` within the grub cmdline settings. This is necessary to allow `base_frequency` to appear in the cpufreq directory in sysfs.

3.5 Setup Script

Run the Intel® SST-BF Python script called `sst_bf.py`, which is available at: https://github.com/intel/CommsPowerManagement, and use option `s` to configure Intel® SST-BF. Use option `i` to confirm the settings.

The script sets up the desired minimum and maximum frequencies for operation with Intel® SST-BF. The script provides several options to configure the system and to revert to a state that does not use Intel® SST-BF.

This script can set up multiple configurations, as reflected in the menu options below.

```
[s] Set SST-BF config (set min/max to 2800/2800 and 2100/2100)
[m] Set P1 on all cores (set min/max to 2300/2300)
[r] Revert cores to min/Turbo (set min/max to 1000/3600)
[i] Show current SST-BF frequency information
[l] List High Priority cores
[n] List Normal Priority cores
[v] Show script version
[h] Print additional help on menu options

[q] Exit Script
------------------------------------------------------------
Option:
```

The main setup options are `s` and `m`.

- (Recommended) Option `s` sets the `scaling_max_freq` and the `scaling_min_freq` settings according to whether the core is a P1 High or P1 Normal core. The script automatically queries the system for the frequencies that it sets when this option is selected and displays those frequencies in the menu option.
- Option `m` sets the `scaling_max_freq` and `scaling_min_freq` settings to the SKU P1 frequency. The script automatically queries the system for the frequencies that it sets when this option is selected and display those frequencies in the menu option.

One option is provided to undo Intel® Speed Shift settings, which allows the pstate driver to scale up and down as the pstate driver algorithm requires.

- Option `r` sets the `scaling_max_freq` for all cores to P1.

There is also a command line interface for use with scripts without the need for a menu. Use the `sst_bf.py -h` command to see the available options.

Figure 4 and Figure 5 contain sample output of the `sst_bf.py` script.

4 Summary

Select SKUs of 2nd generation Intel® Xeon® Scalable processor (5218N, 6230N, and 6252N) offer a capability called Intel® Speed Select Technology – Base Frequency (Intel® SST-BF). This document has shown that enabling Intel® SST-BF can result in an overall system performance increase at comparable power consumption.

Note: The test results in this document are for the Intel® Xeon® Scalable processor 6252N SKU. Future generations of Intel® Xeon® Scalable processors will incorporate Intel® SST-BF technology without significant modification to the existing methodology presented above.

Figure 4 presents the output of the `sst_bf.py` script detailing the frequencies of the individual cores, with hyper-threading enabled, for the device under test with Intel® SST-BF disabled.

Figure 5 presents the output of the `sst_bf.py` script detailing the frequencies of the individual cores, with hyper-threading enabled as before, for the device under test with Intel® SST-BF enabled.

	sysfs		
Core	base	max	min
0	2100	2300	1000
1	2100	2300	1000
2	2100	2300	1000
3	2100	2300	1000
4	2100	2300	1000
5	2100	2300	1000
6	2100	2300	1000
7	2100	2300	1000
8	2100	2300	1000
9	2800	2300	1000
10	2800	2300	1000
11	2100	2300	1000
12	2800	2300	1000
13	2800	2300	1000
14	2800	2300	1000
15	2800	2300	1000
16	2800	2300	1000
17	2100	2300	1000
18	2100	2300	1000
19	2100	2300	1000
20	2100	2300	1000
21	2100	2300	1000
22	2800	2300	1000
23	2100	2300	1000
24	2800	2300	1000
25	2800	2300	1000
26	2800	2300	1000
27	2100	2300	1000
28	2100	2300	1000
29	2100	2300	1000
30	2100	2300	1000
31	2100	2300	1000
32	2100	2300	1000
33	2800	2300	1000
34	2800	2300	1000
35	2800	2300	1000
36	2100	2300	1000
37	2100	2300	1000
38	2800	2300	1000
39	2800	2300	1000
40	2100	2300	1000
41	2100	2300	1000
42	2100	2300	1000
43	2100	2300	1000
44	2100	2300	1000
45	2100	2300	1000
46	2100	2300	1000
47	2100	2300	1000
48	2100	2300	1000
49	2100	2300	1000
50	2100	2300	1000
51	2100	2300	1000
52	2100	2300	1000
53	2100	2300	1000
54	2100	2300	1000
55	2100	2300	1000
56	2100	2300	1000
57	2800	2300	1000
58	2800	2300	1000
59	2100	2300	1000
60	2800	2300	1000
61	2800	2300	1000
62	2800	2300	1000
63	2800	2300	1000
64	2800	2300	1000
65	2100	2300	1000
66	2100	2300	1000
67	2100	2300	1000
68	2100	2300	1000
69	2100	2300	1000
70	2800	2300	1000
71	2100	2300	1000
72	2800	2300	1000
73	2800	2300	1000
74	2800	2300	1000
75	2100	2300	1000
76	2100	2300	1000
77	2100	2300	1000
78	2100	2300	1000
79	2100	2300	1000
80	2100	2300	1000
81	2800	2300	1000
82	2800	2300	1000
83	2800	2300	1000
84	2100	2300	1000
85	2100	2300	1000
86	2800	2300	1000
87	2800	2300	1000
88	2100	2300	1000
89	2100	2300	1000
90	2100	2300	1000
91	2100	2300	1000
92	2100	2300	1000
93	2100	2300	1000
94	2100	2300	1000
95	2100	2300	1000

Figure 4. Intel® SST-BF disabled

	sysfs		
Core	base	max	min
0	2100	2100	2100
1	2100	2100	2100
2	2100	2100	2100
3	2100	2100	2100
4	2100	2100	2100
5	2100	2100	2100
6	2100	2100	2100
7	2100	2100	2100
8	2100	2100	2100
9	2800	2800	2800
10	2800	2800	2800
11	2100	2100	2100
12	2800	2800	2800
13	2800	2800	2800
14	2800	2800	2800
15	2800	2800	2800
16	2800	2800	2800
17	2100	2100	2100
18	2100	2100	2100
19	2100	2100	2100
20	2100	2100	2100
21	2100	2100	2100
22	2800	2800	2800
23	2100	2100	2100
24	2800	2800	2800
25	2800	2800	2800
26	2800	2800	2800
27	2100	2100	2100
28	2100	2100	2100
29	2100	2100	2100
30	2100	2100	2100
31	2100	2100	2100
32	2100	2100	2100
33	2800	2800	2800
34	2800	2800	2800
35	2800	2800	2800
36	2100	2100	2100
37	2100	2100	2100
38	2800	2800	2800
39	2800	2800	2800
40	2100	2100	2100
41	2100	2100	2100
42	2100	2100	2100
43	2100	2100	2100
44	2100	2100	2100
45	2100	2100	2100
46	2100	2100	2100
47	2100	2100	2100
48	2100	2100	2100
49	2100	2100	2100
50	2100	2100	2100
51	2100	2100	2100
52	2100	2100	2100
53	2100	2100	2100
54	2100	2100	2100
55	2100	2100	2100
56	2100	2100	2100
57	2800	2800	2800
58	2800	2800	2800
59	2100	2100	2100
60	2800	2800	2800
61	2800	2800	2800
62	2800	2800	2800
63	2800	2800	2800
64	2800	2800	2800
65	2100	2100	2100
66	2100	2100	2100
67	2100	2100	2100
68	2100	2100	2100
69	2100	2100	2100
70	2800	2800	2800
71	2100	2100	2100
72	2800	2800	2800
73	2800	2800	2800
74	2800	2800	2800
75	2100	2100	2100
76	2100	2100	2100
77	2100	2100	2100
78	2100	2100	2100
79	2100	2100	2100
80	2100	2100	2100
81	2800	2800	2800
82	2800	2800	2800
83	2800	2800	2800
84	2100	2100	2100
85	2100	2100	2100
86	2800	2800	2800
87	2800	2800	2800
88	2100	2100	2100
89	2100	2100	2100
90	2100	2100	2100
91	2100	2100	2100
92	2100	2100	2100
93	2100	2100	2100
94	2100	2100	2100
95	2100	2100	2100

Figure 5. Intel® SST-BF enabled