Priority Core Turbo Technology (PCT Technology)

Technical Article

December 2025

Revision 1.5.1

Document Number: 846906

Notice: This document contains information on products in the design phase of development. The information here is subject to change without notice. Do not finalize a design with this information.

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

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

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

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.

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.

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site.

Your costs and results may vary.

"Conflict-free" refers to products, suppliers, supply chains, smelters, and refiners that, based on our due diligence, do not contain or source tantalum, tin, tungsten or gold (referred to as "conflict minerals" by the U.S. Securities and Exchange Commission) that directly or indirectly finance or benefit armed groups in the Democratic Republic of the Congo or adjoining countries.

All product plans and roadmaps are subject to change without notice.

Code names are used by Intel to identify products, technologies, or services that are in development and not publicly available. These are not "commercial" names and not intended to function as trademarks.

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.

Altering clock frequency or voltage may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Check with system and component manufacturers for details.

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software, and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology.

Results have been estimated or simulated.

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

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visiting the Resource and Documentation Center.

Contents

Figures

Tables

Revision History

Revision Number	Description	Date
1.5.1	• Scrubbed document for publication as Public. • Fixed typo.	December 2025
1.5	• Added a chapter to document runtime expectations with legacy Linux* OS CPU frequency attributes and Intel® Speed Select Technology (Intel® SST) commands for querying Priority Core Turbo Technology (PCT Technology) attributes.	November 2025
1.4	• Transitioned from single-PCT-core to multi-PCT-core partitions.	September 2025
1.3	• Updated Table 4-1 with "8 PCT Cores Examples". • Added a virtualization implications chapter. • Corrected "p0half" with "p0max" in the Background chapter.	June 2025
1.2	• Differentiated P0max and P0half. • Updated frequencies in the sample table to better match the actual PCT SKUs. • Used 8-core PCT in the illustration.	March 2025
1.1	• Updated values in Table 3-2 to reflect the PCT SKU behavior as compared to a hypothetical example.	March 2025
1.0	• Initial release of the document.	February 2025

1 Introduction

This document describes a Priority Core Turbo Technology (PCT Technology) feature that allows four or eight cores on Intel® Xeon® 6900/6700-series processors with P-cores (Granite Rapids) to achieve elevated Turbo frequencies. Intel leverages existing technology, Intel® Speed Select Technology - Turbo Frequency (Intel® SST-TF), to implement the PCT Technology.

Note: For readability purposes, PCT Technology may be hereinafter referred to as PCT.

2 Background

On a conventional Intel® Xeon® CPU, the maximum frequency at which all cores may simultaneously operate is all-core turbo (P0n, that is, n-core turbo), as shown in Figure 3-1(A). The maximum frequency at which any core may operate is half-core turbo (P0half). Half-core turbo requires that at least half of the cores be asleep in C6 to make the power-budget available to the active cores, as shown in Figure 3-1(C). Turbo frequencies at specific-use conditions are determined by the core and uncore power utilization (that is, idle cores allow increased power budget to active cores) and reliability constraints (that is, time-in-state at temperature and voltage).

Table 3-1 defines the relevant frequencies from low (top row) to high (bottom row). Intel SST-TF allows software to identify priority and non-priority cores. When all cores are active with Intel SST-TF, higher-priority cores may achieve frequencies higher than P0n (and up to p0max) at the expense of non-priority cores running below P0n (usually around P1). Intel SST-TF is shown in Figure 3-1(B). Software identifies Intel SST-TF priority cores, and any physical core may be chosen as a priority core.

For an Intel Xeon CPU (host) for GPU/accelerated systems usage, a few GPUs (usually four or eight) are linked to one or two Intel Xeon CPUs. The CPU cores are divided into multiple “logical partitions”; each partition has one GPU. One or multiple cores in each partition are dedicated to service one GPU. The CPU-to-GPU binding is fixed to help enforce localization and to avoid context switches and page migration. The desire is to run these critical GPU-serving cores at the highest possible turbo frequency, while the remaining cores in the partition can run at a lower frequency. There is a concern about core C6 exit latency, so the software does not force cores into C6, making the legacy P0half feature less desirable.

3 PCT Technology Design

PCT Technology leverages Intel SST-TF on an Intel Xeon 6900/6700-series processors with P-cores SKU so that a small number of High Priority (HP) cores are capable of operating at an elevated turbo frequency within acceptable reliability constraints. Figure 3-1(B) shows eight cores identified as PCT HP cores and running at the elevated F3 frequency; Figure 3-1(C) shows that half of the cores can operate at high turbo frequency when another half of the cores are in core C6 sleep state.

Table 3-2 provides examples of potential frequencies. The cores beyond the eight fast cores are not specified to run at the elevated frequency and will operate alongside the priority cores at an Intel selected frequency somewhere between Pn and P0n, which can be supported by the power budget. These non-priority cores run at a frequency below P0n to allocate the maximum possible power budget to the fast priority cores. To reiterate, when PCT is enabled, Low Priority (LP) cores are capped at ~P1.

The OS maximizes performance by assigning priority work to the priority cores, and non-priority work (or no work) to the remaining cores. The system will automatically boot up with PCT Technology enabled if the part supports the capability. The default PCT Technology setting is to divide cores into four partitions and assign the first two cores of each partition as the PCT HP cores, assuming the SKU has eight PCT cores. There are BIOS setup knobs available to support advanced PCT configurability, including:

1. Enable/disable the PCT Technology.
2. Select the number of PCT partitions.
3. Select the PCT HP cores within the partition.

When PCT Technology is enabled, the BIOS programs the Intel SST-TF CLOS registers for priority cores and non-priority cores and kicks off the Intel Primecode flow to enforce HP versus LP frequencies. If PCT Technology is not enabled, the CPU will operate in conventional turbo behavior with a maximum core frequency between P0half and P0n based on the number of active cores and package power budget.

Table 3-1. Frequencies

Frequency for SSE Instructions (Low to High)	Description
0	C6 (sleep)
Pn	Minimum core frequency
P1	All-core guaranteed frequency
P0n	All-core turbo frequency
P0half	Half-core turbo frequency
P0max	Maximum turbo frequency

Figure 3-1. P0n, P0half, Intel SST-TF, and PCT Technology

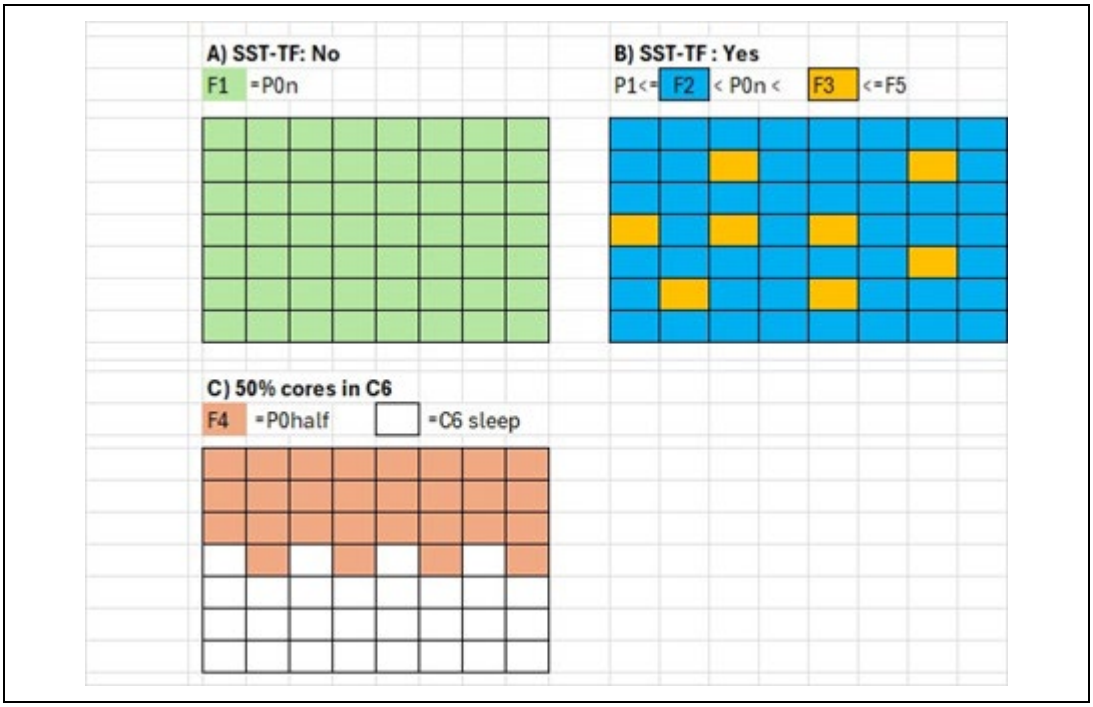

Table 3-2. Example Frequency Ranges

Description	Name	Example Frequency
All-Core Turbo (P0n)	F1	3.6 GHz
Low-Priority Frequency	F2	2.3 GHz
PCT High-Priority Frequency	F3	4.6 GHz
Half-Core Turbo (P0half)	F4	3.9 GHz
Maximum Turbo Frequency (P0max)	F5	4.6 GHz

4 Configuration and Usage

The solution assumes that software will divide all processors into N partitions and will schedule the "important" tasks onto one or multiple consecutive processors of each partition. The term "processor" refers to the CPU thread as seen in Linux/proc/cpuinfo. All processors within a partition are associated with one of the GPUs.

- The software divides all available "processors" into N partitions. Default N=4.
  - The BIOS knob is available for the user to select the number of partitions. The default is 4.
- One or more consecutive cores in each partition will be configured as PCT HP cores. By default, selection starts from the first core.
- The software should schedule "important" tasks to these HP cores.
- The BIOS automatically activates the PCT Technology if the part supports the capability.
  - The BIOS knob is available to disable the PCT Technology. The default setting is enabled.
- The HP cores will run at maximum turbo frequency (P0max) while LP cores will be clipped to ~P1 frequency.
- The OS kernel's processor number ordering must follow the MADT structure convention as defined by the Intel reference BIOS.

No software programing or configuration is required if the CPU-to-GPU mapping is aligned with the default assumptions because the PCT Technology is activated automatically during the BIOS boot time.

Users can use the Intel SST-TF tool to disable or enable the PCT Technology and/or reconfigure HP versus LP cores at OS runtime dynamically, as the PCT Technology is implemented through Intel SST-TF.

Figure 4-1 shows the HP processor selection for a CPU with n cores on a two-socket system with Intel® Hyper-Threading Technology (Intel® HT Technology) enabled. The yellow highlighted processors are the HP ones, specifically the first two cores of each partition. There are four partitions and eight PCT cores for each CPU socket in this example.

Figure 4-1. PCT Technology Core Selection

Processor(p)	Socket	Thread	Processor(p)	Socket	Thread
0, 1	0	0	2n, 2n+1	0	1
...	0	0	...	0	1
n/4, n/4+1	0	0	2n+n/4, 2n+n/4+1	0	1
..	0	0	..	0	1
2n/4, 2n/4+1	0	0	2n+2n/4, 2n+2n/4+1	0	1
...	0	0	...	0	1
3n/4, 3n/4+1	0	0	2n+3n/4, 2n+3n/4+1	0	1
...	0	0	...	0	1
n-1	0	0	3n-1	0	1
n, n+1	1	0	3n, 3n+1	1	1
...	1	0	...	1	1
n+n/4, n+n/4+1	1	0	3n+n/4, 3n+n/4+1	1	1
...	1	0	...	1	1
n+2n/4, n+2n/4+1	1	0	3n+2n/4, 3n+2n/4+1	1	1
...	1	0	...	1	1
n+3n/4, n+3n/4+1	1	0	3n+3n/4, 3n+3n/4+1	1	1
...	1	0	...	1	1
2n-1	1	0	4n-1	1	1

Figure 4-2 provides examples of the PCT HP processor numbers in various system configurations.

Figure 4-2. HP Processors Example

SKU Core Count	Sockets	HT	PCT Core Count	Partitions	HP Cores Per Parition	HP Processor Numbers
64	1	Disable	8	4	2	S0: 0,1,16,17,32,33,48,49
64	1	Enable	8	4	2	S0: 0,1,16,17,32,33,48,49, 64,65,80,81,96,97,112,113
64	2	Enable	8	4	2	S0: 0,1,16,17,32,33,48,49,128,129,144,145,160,161,176,177 S1: 64,65,80,81,96,97,112,113,192,193,208,209,224,225,240,241
64	1	Enable	8	1	8	S0: 0,1,2,3,4,5,6,7,64,65,66,67,68,69,70,71
64	1	Enable	8	8	1	S0: 0,8,16,24,32,40,48,56,64,72,80,88,96,104,112,120
64	1	Enable	8	3	2	S0: 0,1,21,22,42,43,64,65,85,86,106,107
48	1	Enable	8	4	2	S0: 0,1,12,13,24,25,36,37,48,49,60,61,72,73,84,85
64	1	Disable	4	4	1	S0: 0,16,32,48
64	1	Enable	4	4	1	S0: 0,16,32,48,64,80,96,112
64	2	Enable	4	4	1	S0: 0,16,32,48, 128,144,160,176 S1: 64,80,96,112,192,208,224,240
48	1	Enable	4	4	1	S0: 0,12,24,36,48,60,72,84
86	2	Enable	4	4	1	S0: 0,21,42,63,172,193,214,235 S1: 86,107,128,149,258,279,300,321
72	1	Disable	8	4	2	S0: 0,1, 18,19,36,37,54,55
72	1	Enable	8	4	2	S0: 0,1, 18,19,36,37,54,55, 72,73,90,91,108,109,126,127
72	2	Enable	8	4	2	S0: 0,1, 18,19,36,37,54,55, 144,145,162,163,180,181,198,199 S1: 72,73,90,91,108,109,126,127, 216,217,234,235,252,253,270,271

This solution is designed to support any SKU, any number of Cdies, and any number of core count. If the core count is not divisible by the number of partitions, the algorithm will round down the core count per partition and pick one or consecutive multiple cores of each partition as the PCT HP cores. If the SKU PCT core count cannot be evenly distributed across partitions, the surplus PCT cores will operate as non-PCT cores.

For example, the 86-core SKU is not divisible by 4, so it would have 21 cores in each partition, and the last two cores will be part of the PCT LP cores. For a SKU that has eight PCT cores, where BIOS configures it to have three partitions, the algorithm will assign two HP cores per partition, with a total of six PCT HP cores, and the remaining two surplus PCT cores will operate as non-PCT cores.

Note: The BIOS-based PCT Technology solution may not work if the system is built with a different MADT convention.

OS Interaction with PCT Technology Operating Modes:

- **PCT Technology Enabled by the BIOS**:
  - The BIOS builds an ACPI MADT table to pass the "processor" numbering information to the OS using the standard ACPI specification (not unique to the PCT Technology). The term "processor" refers to the CPU thread as seen in Linux/proc/cpuinfo.
  - The BIOS divides all available processors into N partitions. Default N=4; N is configurable in the BIOS menu.

- The BIOS configures one or multiple processors in each partition as the PCT HP cores.
- It is the OS/SW responsibility to schedule and affinitize "important" tasks to these PCT HP cores.
- This mode is a good fit for users with fixed pinning between the CPU core and the GPU, not wishing to deploy the Intel® Speed Select Technology (Intel® SST) tool.

- **PCT Technology Enabled by the Intel® SST Tool**:
  - Users use the Intel SST tool to choose HP cores in lieu of BIOS actions in the preceding first two steps.
  - The OS/SW is responsible for scheduling and affinitizing "important" tasks to these PCT HP cores (same as the previous point).
  - This mode is a good fit for Intel SST-TF users that require the flexibility of runtime pinning and/or runtime enable/disable of PCT Technology.
  - The Intel SST tool can override the BIOS setting at runtime without requiring a reboot.
  - If the user configures more PCT cores than what the SKU supports, the frequency that the user will get will not be the PCT frequency. It will be based on the next Intel SST-TF Turbo Ratio Limits (TRL) bucket. By default, bucket-0 specifies the number of PCT cores and PCT frequency in the TRL table.
  - The user can find out the Intel SST-TF TRL buckets using the following command:

    ```
    sudo intel-speed-select turbo-freq info -l 1
    ```

- **The following commands configure the two-socket Intel Xeon 6900/6700-series processors with P-cores system with 8 PCT cores (16 threads) at 4.6 GHz. It caps the non-PCT cores to 2.5 Ghz.**
  - The Intel SST tool needs to be run with "sudo".

```
# Set the performance level to 0
intel-speed-select perf-profile set-config-level -l 0 -o
# Disable SST Base Frequency
intel-speed-select base-freq disable -a
# Enable SST Turbo Frequency
intel-speed-select turbo-freq enable
# Set the min/max frequency for CLOS0 (HP)
intel-speed-select core-power config -c 0 --min 4600 --max 4600
# Set the min/max frequency for CLOS3 (LP)
intel-speed-select core-power config -c 3 --min 800 --max 2500
# Enable Ordered Priority for CLOS based prioritization
intel-speed-select -c 0-255 core-power enable --priority 1
# First put all the cores in CLOS3
intel-speed-select --cpu 0-255 core-power assoc -c 3
# Select cores in Socket 0 and associate them with CLOS0 (4
partitions/socket)
```

```
intel-speed-select --cpu
0,1,16,17,32,33,48,49,128,129,144,145,160,161,176,177 core-power
assoc -c 0
# Select cores in Socket 1 and associate them with CLOS0 (4
partitions/socket)
intel-speed-select --cpu
64,65,80,81,96,97,112,113,192,193,208,209,224,225,240,241 core-power
assoc -c 0
```

5 Usage Expectations

Intel is not providing or asking for time-in-state enforcement.

6 PCT Technology and Virtualization

If a Virtual Machine Manager (VMM) requires exposing PCT Technology to virtual machines, the manager can specifically assign one or more priority logical host cores (as determined by the BIOS knobs or Intel SST-TF settings) to specific virtual machines. All other non-priority host cores will be limited to the Intel SST-TF clip frequency.

For example, in a 64-core CPU with eight priority cores, divided into four partitions - logical cores 0, 1, 16, 17, 32, 33, 48, and 49, the VMM could assign logical cores to virtual machines as follows: 0-15, 16-31, 32-47, and 48-63. Each 16-core virtual machine would have two priority cores that could reach the PCT frequency, and 14 cores limited to the Intel SST-TF clip frequency.

In Intel Xeon 6900/6700-series processors with P-cores, it is not possible to enable PCT Technology (or Intel SST-TF) at any granularity other than the entire SoC. If Intel SST-TF is available to some virtual machines, all non-priority cores will be limited to the Intel SST-TF clip frequency.

Figure 6-1. PCT Technology Example with Four Virtual Machines

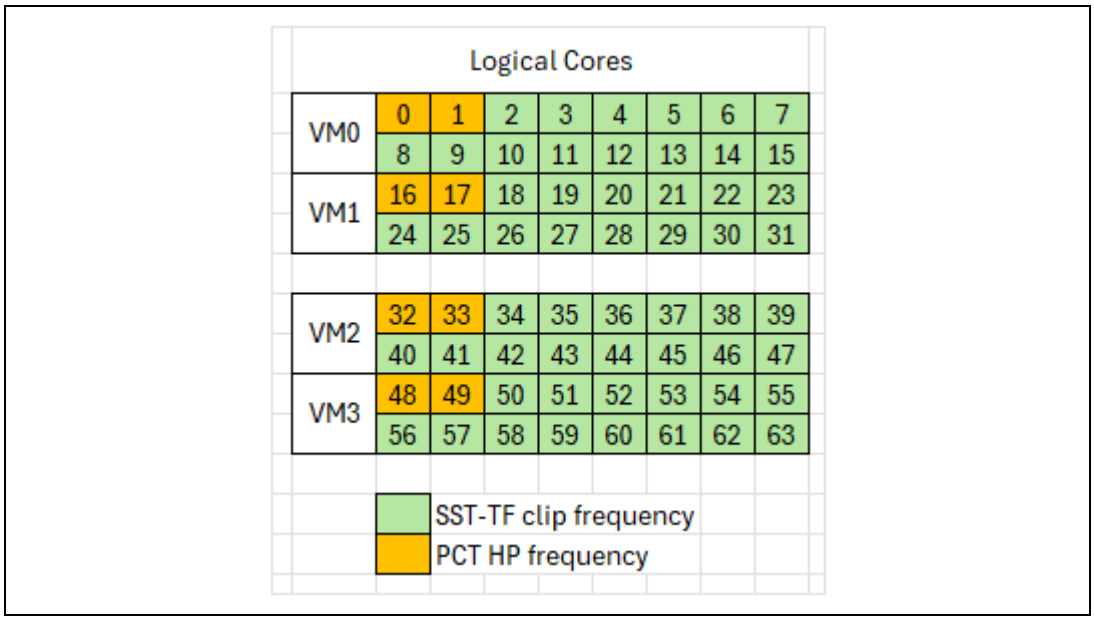

7 Runtime Expectations

When the OS boots, the OS pstate driver does not have information about the selection of PCT HP cores and their frequencies, so the Linux* OS sysfs entries (`/sys/devices/system/cpu/cpu*/cpufreq/`) do not show differences for the CPU frequency attributes between PCT cores and non-PCT cores. If the user desires to get information about the PCT and non-PCT cores in the OS, Intel's recommendation is to use the Intel SST tool.

The Intel SST tool is bundled with several leading Linux distributions. If it is not available, it can be found in the upstream kernel source tree and can be compiled from the source.

Upstream kernel source tree:
`git://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git`

The tool is located at the `linux/tools/power/x86/intel-speed-select` folder.

Note: The BIOS knob is a simple way to enable PCT where logical IDs and maximum frequencies of HP cores are fixed. Using the Intel SST tool to change the configuration could affect the PCT Technology configuration done via the BIOS. There is no need to resolve any conflict with the BIOS as using the Intel SST tool for configuration will override the BIOS settings.

- Downloading and compiling the Intel SST tool:

```
git clone
https://git.kernel.org/pub/scm/linux/kernel/git/torvalds/linux.git
cd linux/tools/power/x86/intel-speed-select
make
make install
```

Note: There is dependency on following packages for compiling the tool. Use the following command to install them:

```
sudo apt-get install libnl-3-dev libnl-genl-3-dev
```

- Intel SST tool commands:
  - Information regarding the CLOS associated with a CPU (for example, cpu 1):

    ```
    sudo intel-speed-select --cpu 1 core-power get-assoc
    ```

  - Information on the CLOS definition (for example, cpu=1 and CLOS=0):

    ```
    sudo intel-speed-select --cpu 1 core-power get-config -c 0
    ```

- To get the maximum frequency associated with a core, get the CLOS information associated with that core and feed that information into the CLOS definition (for example, cpu=10).
  - `cpu=10`
  - `sudo intel-speed-select --cpu ${cpu} core-power get-config -c $(intel-speed-select --cpu ${cpu} core-power get-assoc 2>&1 | awk -F':' '/clos/ {print $2}') 2>&1 | awk -F'[: ]' '/clos-max/ {print $(NF-1)}'`

8 Appendix

Sample Python* programming language script to print the frequency of the PCT and non-PCT cores:

Note: The script must be run as root.

```
#!/usr/bin/python3
import subprocess
pct_cores=dict()
non_pct_cores=dict()

def get_pct_freq():
    command = "lscpu | awk -F'[ .]' '/CPU max MHz/{print $(NF-1)}'"
    result=subprocess.run(command, shell=True, executable="/bin/bash", che
ck=True, capture_output=True, text=True)

    return int(result.stdout.strip())

def get_core_count():
    command = "lscpu | awk '/^CPU\\(s\\):/{print $NF}'"
    result=subprocess.run(command, shell=True, executable="/bin/bash", che
ck=True, capture_output=True, text=True)

    return int(result.stdout.strip())

pct_max_freq=get_pct_freq()
total_cores=get_core_count()

for core in range(total_cores):
    command = ( f"intel-speed-select --cpu {core} core-power get-config -
c " f"$(intel-speed-select --cpu {core} core-power get-assoc 2>&1 | awk -
F':' '/clos/ {{print $2}}') 2>&1 | " f"awk -F'[: ]' '/clos-
max/ {{print $(NF-1)}}'" )
    result=subprocess.run(command, shell=True, executable="/bin/bash", che
ck=True, capture_output=True, text=True)
    max_core_freq = int(result.stdout.strip())
    if max_core_freq == pct_max_freq:
        pct_cores[core]=max_core_freq
    else:
        non_pct_cores[core]=max_core_freq

print("\nPCT Cores")
print(pct_cores)
print("\nNon PCT Cores")
print(non_pct_cores)
```

Sample output of the script:

The output is on a two-socket Intel Xeon 6900/6700-series processors with P-cores system with 64 cores. It has eight PCT cores (16 threads) per socket in four partitions.

```
bash# python print_pct.py
```

- PCT Cores

```
{0: 4600, 1: 4600, 16: 4600, 17: 4600, 32: 4600, 33: 4600, 48: 4600,
49: 4600, 64: 4600, 65: 4600, 80: 4600, 81: 4600, 96: 4600, 97: 4600,
112: 4600, 113: 4600, 128: 4600, 129: 4600, 144: 4600, 145: 4600, 160:
4600, 161: 4600, 176: 4600, 177: 4600, 192: 4600, 193: 4600, 208:
4600, 209: 4600, 224: 4600, 225: 4600, 240: 4600, 241: 4600}
```

- Non-PCT Cores

```
{2: 2300, 3: 2300, 4: 2300, 5: 2300, 6: 2300, 7: 2300, 8: 2300, 9:
2300, 10: 2300, 11: 2300, 12: 2300, 13: 2300, 14: 2300, 15: 2300, 18:
2300, 19: 2300, 20: 2300, 21: 2300, 22: 2300, 23: 2300, 24: 2300, 25:
2300, 26: 2300, 27: 2300, 28: 2300, 29: 2300, 30: 2300, 31: 2300, 34:
2300, 35: 2300, 36: 2300, 37: 2300, 38: 2300, 39: 2300, 40: 2300, 41:
2300, 42: 2300, 43: 2300, 44: 2300, 45: 2300, 46: 2300, 47: 2300, 50:
2300, 51: 2300, 52: 2300, 53: 2300, 54: 2300, 55: 2300, 56: 2300, 57:
2300, 58: 2300, 59: 2300, 60: 2300, 61: 2300, 62: 2300, 63: 2300, 66:
2300, 67: 2300, 68: 2300, 69: 2300, 70: 2300, 71: 2300, 72: 2300, 73:
2300, 74: 2300, 75: 2300, 76: 2300, 77: 2300, 78: 2300, 79: 2300, 82:
2300, 83: 2300, 84: 2300, 85: 2300, 86: 2300, 87: 2300, 88: 2300, 89:
2300, 90: 2300, 91: 2300, 92: 2300, 93: 2300, 94: 2300, 95: 2300, 98:
2300, 99: 2300, 100: 2300, 101: 2300, 102: 2300, 103: 2300, 104: 2300,
105: 2300, 106: 2300, 107: 2300, 108: 2300, 109: 2300, 110: 2300, 111:
2300, 114: 2300, 115: 2300, 116: 2300, 117: 2300, 118: 2300, 119:
2300, 120: 2300, 121: 2300, 122: 2300, 123: 2300, 124: 2300, 125:
2300, 126: 2300, 127: 2300, 130: 2300, 131: 2300, 132: 2300, 133:
2300, 134: 2300, 135: 2300, 136: 2300, 137: 2300, 138: 2300, 139:
2300, 140: 2300, 141: 2300, 142: 2300, 143: 2300, 146: 2300, 147:
2300, 148: 2300, 149: 2300, 150: 2300, 151: 2300, 152: 2300, 153:
2300, 154: 2300, 155: 2300, 156: 2300, 157: 2300, 158: 2300, 159:
2300, 162: 2300, 163: 2300, 164: 2300, 165: 2300, 166: 2300, 167:
2300, 168: 2300, 169: 2300, 170: 2300, 171: 2300, 172: 2300, 173:
2300, 174: 2300, 175: 2300, 178: 2300, 179: 2300, 180: 2300, 181:
2300, 182: 2300, 183: 2300, 184: 2300, 185: 2300, 186: 2300, 187:
2300, 188: 2300, 189: 2300, 190: 2300, 191: 2300, 194: 2300, 195:
2300, 196: 2300, 197: 2300, 198: 2300, 199: 2300, 200: 2300, 201:
2300, 202: 2300, 203: 2300, 204: 2300, 205: 2300, 206: 2300, 207:
2300, 210: 2300, 211: 2300, 212: 2300, 213: 2300, 214: 2300, 215:
2300, 216: 2300, 217: 2300, 218: 2300, 219: 2300, 220: 2300, 221:
2300, 222: 2300, 223: 2300, 226: 2300, 227: 2300, 228: 2300, 229:
2300, 230: 2300, 231: 2300, 232: 2300, 233: 2300, 234: 2300, 235:
2300, 236: 2300, 237: 2300, 238: 2300, 239: 2300, 242: 2300, 243:
2300, 244: 2300, 245: 2300, 246: 2300, 247: 2300, 248: 2300, 249:
2300, 250: 2300, 251: 2300, 252: 2300, 253: 2300, 254: 2300, 255:
2300}
```