TECHNOLOGY GUIDE

Intel Corporation

Intel Technologies for Resource Tuning Energy Efficient Network Slices

Authors

Chris MacNamara,
Intel Corporation

John Browne,
Intel Corporation

Paul Veitch,
BT Applied Research

1 Introduction

In a modern network infrastructure, maximizing performance and realizing energy efficiency for end-to-end network slices is paramount. This paper describes the Intel® technologies and resource tuning techniques available to:

- Improve workload performance
- Eliminate noisy neighbor interference
- Provide energy-efficient network slices that meet Service Level Agreements (SLAs)

The technologies in question are:

- Intel® Speed Select Technology – Base Frequency (Intel® SST-BF)
- Intel® Resource Director Technology (Intel® RDT)

The paper provides an overview of how these technologies, when combined and applied locally in a multi-core Intel® Xeon® Scalable processor, contribute to meet the diverse needs of end-to-end energy efficient network slices.

These technologies and techniques are particularly suited to use cases such as:

- Edge compute in 5G and beyond 5G
- Virtual Customer Premises Equipment (vCPE)
- Power-aware slicing in 5G and beyond 5G

This paper references work done by Intel in cooperation with our partner BT Applied Research to validate the benefits of using these technologies in a fully virtualized testbed that provided real, measurable results. Please refer to the Resource Tuning for Energy Efficient Slicing white paper on the IEEE website for details.

This technology guide is intended for communication service providers who are planning and deploying virtualized mobile core infrastructure running on the latest Intel® Xeon® Scalable processors.

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

Figures

Tables

Document Revision History

REVISION	DATE	DESCRIPTION
001	July 2021	Initial release.

1.1 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
CAT	Cache Allocation Technology
CMT	Cache Monitoring Technology
CPU	Central Processing Unit
DPDK	Data Plane Development Kit
HQ	Headquarters
IEEE	Institute of Electrical and Electronics Engineers
KPI	Key Performance Indicator
KVM	Kernel-based Virtual Machine
LLC	Last Level Cache
MBA	Memory Bandwidth Allocation
MBM	Memory Bandwidth Monitoring
NFV	Network Function Virtualization
OPNVF	Open Platform for NFV Project
OVS	Open vSwitch
RDT	Resource Director Technology
SLA	Service Level Agreement
SST-BF	Speed Select Technology – Base Frequency
TCP	Transmission Control Protocol
vCPE	Virtual Customer Premises Equipment
VM	Virtual Machine
VTA	Virtual Test Agent
WAN	Wide Area Network

1.2 Reference Documentation

Table 2. Reference Documents

REFERENCE	SOURCE
Resource Tuning for Energy Efficient Slicing	https://ieeexplore.ieee.org/document/9385531
OPNFV Barometer Project	https://wiki.opnfv.org/display/fastpath
OPNFV Project One-Click-Install of Barometer Containers	https://wiki.opnfv.org/display/fastpath/One+Click+Install+of+Barometer+Containers
Intel® Speed Select Technology (Intel® SST) web page	https://www.intel.com/content/www/us/en/architecture-and-technology/speed-select-technology-article.html
Intel® Speed Select Technology User Guide	https://www.kernel.org/doc/html/latest/admin-guide/pm/intel-speed-select.html
Intel® Resource Director Technology (Intel® RDT) web page	https://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html

2 Overview

A key challenge in modern networks is the allocation and management of resources to fulfill throughput and performance demands. As technology evolves, new services (such as holographic services, extended reality, and autonomous services) place ever more stringent performance demands on network infrastructure. Achieving predictable performance for such services running on top of the network infrastructure is essential.

For multi-tenant workloads competing for shared compute resources on Intel servers, it is critically important to understand the available fine-tuning capabilities that can be leveraged to achieve the required performance levels, while optimizing infrastructure resources. The cost of compute resources and power consumption considerations must be considered when balancing diverse performance and SLA targets.

Research continues in this area, but the focus of this paper is on a relatively new suite of features on specific Intel processors (beginning with the 2nd Gen Intel® Xeon® Scalable processor), most notably Intel® Speed Select Technology – Base Frequency (Intel® SST-BF) and how it can be used to achieve notable performance gains while balancing power consumption within a constrained power envelope.

Any multi-tenant solution must consider the possible effects of a noisy neighbor, where one of the co-tenants of a compute host monopolizes bandwidth, disk I/O, CPU, or other resources, and negatively impacts the other tenants. Intel® Resource Director Technology (Intel® RDT) component features can be used to effectively detect and eliminate the destructive effects of a noisy neighbor.

3 Intel® SST-BF Synopsis

Intel® Speed Select Technology (Intel® SST) is a suite of technologies delivering more granular control over CPU performance, essentially applying power/frequency where and when it is needed. Intel® SST-BF (Base Frequency) permits a two-tier asymmetric system of core frequency deployment within the same multi-core processor.

Figure 1 shows a configuration without Intel® SST-BF enabled. All CPU cores operate at the same guaranteed base frequency (for example, 2.3 GHz) in a symmetric fashion.

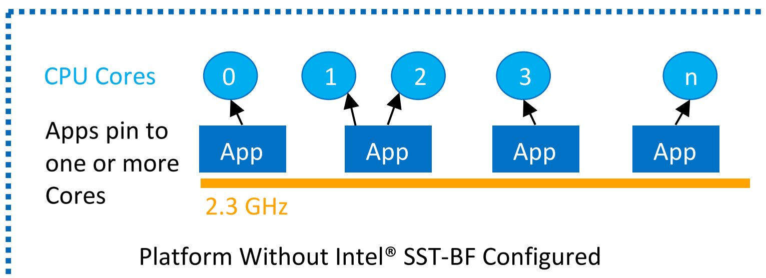

Figure 1. Symmetric Core Frequency Deployment

Figure 2 shows a configuration with Intel® SST-BF enabled. Several designated higher priority cores run at a higher guaranteed base frequency (for example, 2.8 GHz), while other cores run at a lower base frequency (for example, 2.1 GHz). The ability to designate CPU cores as higher or lower priority enables a more flexible method of achieving the balance between required performance, compute costs, and energy consumption.

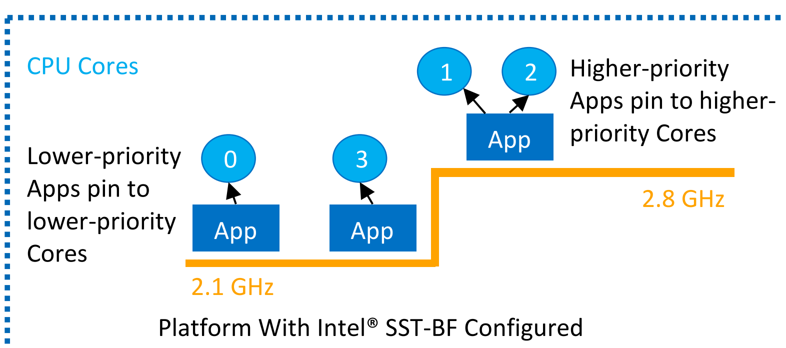

Figure 2. Asymmetric Core Frequency Deployment

For more information, refer to Intel® Speed Select Technology, More Control, More Optimal Performance. See also the Intel® Speed Select Technology User Guide for tools to configure Intel SST in Linux*.

4 Combining Intel® SST-BF with Other Technologies to Realize Energy-Efficient Network Slices

By combining Intel® SST-BF with Intel® Resource Director Technology (Intel® RDT), coordinated resource tuning and optimization of energy-efficient network slices can be realized, while eliminating noisy neighbor interference between multi-tenanted workloads on the same x86 server.

Intel® RDT is a suite of technologies that bring control over how shared resources such as last-level cache (LLC) and memory bandwidth are used by applications, virtual machines (VMs), and containers. In the context of this paper, the important Intel® RDT technologies are:

- Cache monitoring technology (CMT) and cache allocation technology (CAT) to help monitor and address shared resource concerns by providing software control of where data is allocated into the LLC, enabling isolation and prioritization of key applications.
- Memory bandwidth monitoring (MBM) and memory bandwidth allocation (MBA) help to monitor memory bandwidth and provide approximate and indirect per-core control over memory bandwidth.

Intel® RDT provides prioritized access to cache and memory bandwidth, while Intel® SST provides prioritized frequency ranges. For more information, refer to Intel® Resource Director Technology (Intel® RDT).

5 Mapping Network Slice KPIs to Local Platform Resources

Using Intel® SST-BF and Intel® RDT, it is possible to map end-to-end network slice Key Performance Indicator (KPI) targets to local (platform) slice mapping and resource management capabilities as described in Table 3.

Table 3. Mapping End-to-End Slice KPIs to Local Resources

END-TO-END SLICE KPI TARGETS	LOCAL PLATFORM SLICE MAPPING AND RESOURCE MANAGEMENT CAPABILITIES
High Throughput Low Latency Energy Efficient	Intel® SST-BF: - Use high-frequency cores for priority workloads Intel® RDT (CMT, CAT and MBM, MBA): - Reserve adequate LLC resources - Do not limit memory bandwidth
Tolerate Lower Throughput Require Low Latency Energy Efficient	Intel® SST-BF: - Use low-priority cores Intel® RDT (CMT, CAT and MBM, MBA): - Reserve adequate LLC resources - Do not limit memory bandwidth
Best Effort Throughput No Latency or Jitter Requirements Energy Efficient	Intel® SST-BF: - Use low-priority cores Intel® RDT (CMT, CAT and MBM, MBA): - Use minimal LLC resources - Rate-limit memory bandwidth This preempts Noisy Neighbor effects

This capability gives rise to some compelling customer-identified use cases as described in Table 4.

Table 4. Customer-Identified Use Cases for Intel® SST-BF

END-TO-END SLICE KPI TARGETS	LOCAL PLATFORM SLICE MAPPING AND RESOURCE MANAGEMENT CAPABILITIES
Edge Compute in 5G and Beyond 5G	• In distributed edge and fog compute scenarios, singleton edge servers are likely, with a mixed profile of workload types and functions, some of which need more throughput, lower latency, and so on. • Intel® SST-BF can be used to map different workloads to different cores using asymmetric frequency settings.
Virtual Customer Premises Equipment (vCPE)	• The rationale for vCPE is a single box hosting multiple functions (a router, firewall, WAN accelerators, and so on). • Intel® SST-BF can be leveraged to map different workloads to different cores using asymmetric frequency settings. This is more likely to be applicable in larger branch/HQ locations for vCPE.
Power-aware Slicing in 5G and Beyond 5G (general purpose scenarios where intra-server power/performance control is required).	• When considering 5G and beyond 5G slice-aware core networks, with high throughput/power sensitive slices, Intel® SST-BF can be leveraged to map different workloads to different cores using asymmetric frequency settings.

6 Validation in a Test Environment with Platform Telemetry

The Resource Tuning for Energy Efficient Slicing white paper, available on the IEEE website, provides details of a testbed used to validate just how effective Intel technologies and tuning techniques can be in a real-world scenario. The following is a summary of the setup to give a flavor for what is required:

- An Intel server board with Intel® Xeon® CPUs
- A Kernel-based Virtual Machine (KVM) hypervisor that creates and runs the following VMs, all of which are pinned to specific CPU cores:
  - Virtual test agents (VTAs) to generate synthetic traffic
  - A high-performance router
  - A noisy neighbor workload
- The KVM hypervisor employs Open vSwitch (OVS) to bridge the traffic between VMs, and the Data Plane Development Kit (DPDK) for fast packet processing
- OPNFV Barometer components are used to collect and display platform telemetry.

The collection and visualization of telemetry associated with power consumption and temperature, as well as the CPU core frequency settings, is integrated using OPNFV Barometer containers running: Collectd*, InfluxDB*, and Grafana* (for more information, see OPNFV Barometer Project). It is possible to use OPNFV Barometer "one-click install" to get up and running with these components (see OPNFV Project One-Click-Install of Barometer Containers).

This is an important part of the test setup because it enables greater insights into platform and network behavior to complement the TCP performance results acquired from the VTAs.

7 Summary

Intel provides technologies and tuning techniques that enable local multi-core processor platform tuning to meet the diverse needs of end-to-end energy-efficient network slices and achieve real, measurable benefits.

A testbed developed in cooperation with our research partner demonstrated just how much can be achieved by applying these technologies and tuning techniques. The following summarizes what was achieved.

- Improved workload performance – increasing throughput by 26%
- Elimination of noisy neighbor interference – using resource allocation techniques
- Energy-efficient network slices that meet SLAs – demonstrating an overall improvement of 34% for a power consumption increase of just 5%

From these results, we conclude that by applying the Intel technologies and tuning techniques described in this guide, you can significantly improve the energy efficiency of network slices, while eliminating noisy neighbor effects.

While the telemetry used in this scenario provides good visibility into the behavior of various components on the platform, a topic of further research is the use of telemetry results to dynamically modify the behavior of those platform components to provide even greater opportunities for energy savings.

0721/DN/WIPRO/PDF 644893-001US