

# Case Study of Scaled-Up SKT 5G MEC Reference Architecture

## Authors

**Dongkee, Lee**  
SK Telecom

**Moonyoung, Chung**  
SK Telecom

**Woohyun, Nam**  
SK Telecom

**Keunhyun, Kim**  
SK Telecom

**Wooram (Alex) Kim**  
Intel Corporation

**Jongup (Jayden) Lee**  
Intel Corporation

**Sy jong, Choi**  
Intel Corporation

**Huisuk (Kevin) Hong**  
Intel Corporation

**Evstratios (Steve) Vandrís**  
Intel Corporation

## 1. Introduction

SK Telecom (SKT) and Intel collaborated on testing of the next generation of the SKT 5G MEC reference architecture. This platform is designed to be highly scalable, flexible, reliable, and to meet service level agreements (SLAs). SKT has designed this MEC platform to be deployed in a variety of scenarios, including at a central office, at the network edge, or on a customer premises. The platform is based on network functions virtualization (NFV) infrastructure (NFVI) for flexibility. The tests described in this paper show performance impact on the SKT 5G MEC from applying Intel® Speed Select Technology – Base Frequency (Intel® SST-BF), a feature that is available on select models of 2nd generation Intel® Xeon® Scalable processors, and Intel® Optane™ DC persistent memory. Intel Optane DC persistent memory is a new technology designed to fill the capacity, cost, and performance gaps between traditional dynamic random access memory (DRAM) and solid state drives (SSDs).

### Key Object of Test

This document details testing of the SKT 5G MEC software running on Intel® architecture server boards, and it includes benchmarking data and instructions on how to replicate the tests. Telecom equipment manufacturers (TEMs) and independent software vendors (ISVs) can use the implementation guidelines from this work to optimize and further develop solutions for high-performance production-level implementations.

Test objectives were:

- Demonstrate packet forwarding performance improvement with Intel Xeon Gold 6252N processor.
- Show performance improvement of Intel Optane DC persistent memory on multiple in-memory database (IMDB) workloads (IMDB from Redis Labs).
- Validate performance improvement with Intel SST-BF by changing asymmetric core frequency configuration on SKT 5G MEC Edge Routing Function.

## 1.1 Multi-access Edge Computing Overview

Multi-access Edge Computing (MEC) is defined by the European Telecommunications Standards Institute (ETSI) MEC Industry Specification Group (ISG). MEC enables placement of data center-grade compute, storage, and network resources in locations at the edge of the network—closer to the users and endpoint devices. MEC servers are designed to serve consumer and enterprise applications that require low latency and high bandwidth. Locations at the edge of the network include on premises at an enterprise or locations in the mobile network operator (MNO) network including wireless access points of presence (POPs), wireless network aggregation sites including base station or cloud remote access network (C-RAN) sites, or in communication service provider (CommSP) wireline network access aggregation sites and central offices.

MEC provides to telecom service providers the ability to deliver new, real-time services with lower latency. This lower latency is a key requirement for new revenue-generating services to enterprises or consumers. Since edge computing

## Table of Contents

1. Introduction .....	1
1.1 Multi-access Edge Computing Overview .....	1
1.2 Overview of SKT 5G MEC Software Platform.....	2
1.3 Key Hardware Technology ..	3
2. Test Setup .....	4
2.1 Device Under Test Configuration .....	4
2.2 Test Procedures .....	7
3. Test Results .....	8
3.1 Packet Forwarding Performance Results .....	8
3.2 VM Scalability Test Result ..	10
4. Summary .....	11
5. Resources .....	11
6. About the Authors.....	11
7. Glossary of Terms .....	12

minimizes the amount of data to be sent to the centralized cloud, it uses network bandwidth and resources more efficiently, reducing the cost for both enterprises and operators.

In addition, the CommSP can offer to consumers, enterprises, and government new revenue-generating services; for example, the CommSP can offer consumer services such as 4K/360 video streaming with content delivery network (CDN), augmented reality/virtual reality (AR/VR), or cloud gaming.

Numerous services become possible such as data analytics, cloud storage services, video analytics for surveillance and computer vision in warehousing, manufacturing, industrial IT/OT integration for industrial control in manufacturing, smart retail with AR/VR, and computer vision/video analytics coupled with in-store payment systems, or medical imaging analytics for healthcare. For government and smart cities, services such as video surveillance, traffic and parking optimization, and smart connected vehicles using vehicle to everything (V2X) technology are possible.

MEC must allow for the dynamic placement and execution of applications across different edge locations. This helps enable optimal placement of workloads at the edge location where they will deliver the most value and enable the flexibility for applications. Good network planning and software portability to different server platforms enable this flexibility.

At the architectural level, a common, scalable hardware and software platform based on Intel architecture-based servers with Intel Xeon Scalable processors and Intel Optane DC persistent memory can serve all edge locations. The ETSI MEC ISG has defined the standard MEC software architecture.<sup>1</sup>

The 3rd Generation Partnership Project (3GPP) has defined edge computing in the mobile network as part of the 5G standard in Technical Specification TS 23.501. The ETSI MEC ISG has defined how the ETSI MEC software architecture can be deployed in 4G and 5G networks defined by 3GPP.<sup>2</sup>

The ETSI MEC standard has also defined standard APIs for software developers to use when developing MEC software applications optimized for the edge using MEC application services.<sup>3</sup>

## 1.2 Overview of SKT 5G MEC Software Platform

The 5G mobile network specification provides for radio delay or latency of up to 8 ms based on Enhanced Mobile Broadband (eMBB). This is a 60% reduction in the radio delay compared to the 4G standard, enabling the 5G network to provide ultra-low latency services. But there's still the delay associated with the wired line backhaul transport to the cloud that must be factored into the total round trip latency. If the backhaul network is unmanaged with no content delivery network (CDN) services, it is likely that latency will be too long for many real-time services.

MEC servers solve this latency issue by bringing the cloud server closer to the user. Properly located MEC servers minimize transport latency. Figure 1 shows the SK Telecom 5G-based MEC architecture, which was developed based

on 3GPP and ETSI MEC standards. The platform has been designed and developed to enable a variety of 5G services that require low latency, security features, or a need to keep edge traffic local to reduce backbone data traffic. To enable this, SK Telecom implemented its own MEC platform including edge routing function, a MEC service enabler composed of an API gateway and MEC service, and MEC platform management system.

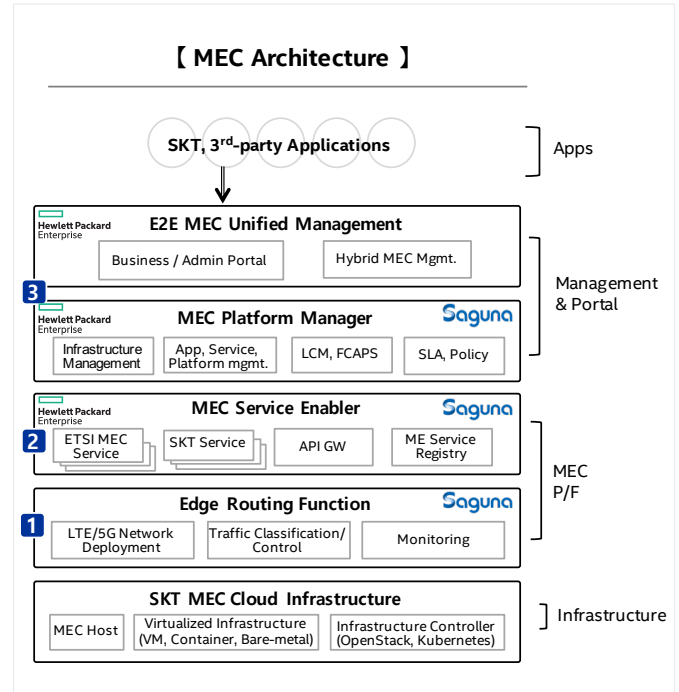


Figure 1. SK Telecom's 5G-based MEC platform architecture

SK Telecom's Edge Routing Function, which resides in the MEC Host, enables the deployment of the MEC platform using 3GPP S1-U/N3 or SGI/N6 interfaces in an NSA/SA 5G core network. In the case of the S1-U/N3 deployment scenario, the Edge Routing Function acts as a "bump-in-the-wire" system by bringing this routing functionality to the legacy network. SK Telecom's Edge Routing Function provides ETSI-standard MEC services such as radio network information service (RNIS), location service, and the others. The main role of Edge Routing Function is classifying user traffic and forwarding it to the corresponding MEC application instance.

Edge Routing Function supports three types of MEC applications. In breakout mode, the session's connection to a service on an internet server is redirected to the same service hosted in a MEC application that is hosted on the SKT MEC platform. In in-line mode, the session's connectivity is maintained to the original internet server, while all traffic passes through the MEC application. In tap mode, specified data traffic is duplicated and forwarded to the tap MEC application, for example, when deploying virtual network probes to measure network quality.

ME Service Enabler is also a key element of SK Telecom's MEC platform. The main functions of the MEC Service Enabler include MEC service discovery, routing, authentication, authorization, throttling control, and MEC service level agreement (SLA) management. MEC Service

Enabler also provides MEC services that can be utilized in the development of MEC applications. These include application enablement, RNIS, bandwidth management (BWM), UE identity, and location services as defined by ETSI MEC ISG.

In addition the standards-based MEC services, the MEC Service Enabler also enables SK Telecom's differentiated MEC services, such as Precise Positioning Service, location spoofing protection, software developer support services, and service analytics. In the near future, SKT will enhance its location-based service to provide a location spoofing prevention capability that is necessary for location-based gaming services that are subject to the possibility of GPS spoofing.

The MEC Platform Manager (MEPM), which resides in the MEC Host, enables MEC Platform software and MEC applications to operate on the MEC host. MEPM features include MEC Orchestrator, MEC Platform Manager, and key elements of the virtualization infrastructure manager defined in the ETSI MEC standard.

1.3 Key Hardware Technology

1.3.1 Intel® Speed Select Technology - Base Frequency (Intel® SST-BF)

Intel Speed Select Technology - Base Frequency (Intel SST-BF) is available on select models of the 2nd generation Intel Xeon Scalable processors. It allows a part of the CPU to process data asynchronously to improve performance. The default CPU operation mode is a symmetric core frequency deployment, where all applications on the processor operate at the same core frequency. When Intel SST-BF is enabled, it allows the CPU clock frequency to be dynamically distributed across cores in an asymmetric configuration. This enables users to boost performance of targeted applications at runtime.

One option for performance improvement is Intel® Turbo Boost Technology, but it is non-deterministic (opportunistic) and can introduce jitter and latency issues even as it boosts performance. Intel SST-BF provides prioritization across cores, supporting NFV workloads that result in different tasks for different cores. Intel SST-

BF addresses those cores that can become a bottleneck when operating symmetrically. The placement of key workloads on higher frequency Intel SST-BF-enabled cores can result in an overall system workload increase and potential overall energy savings when compared to deploying the CPU with symmetric core frequencies.

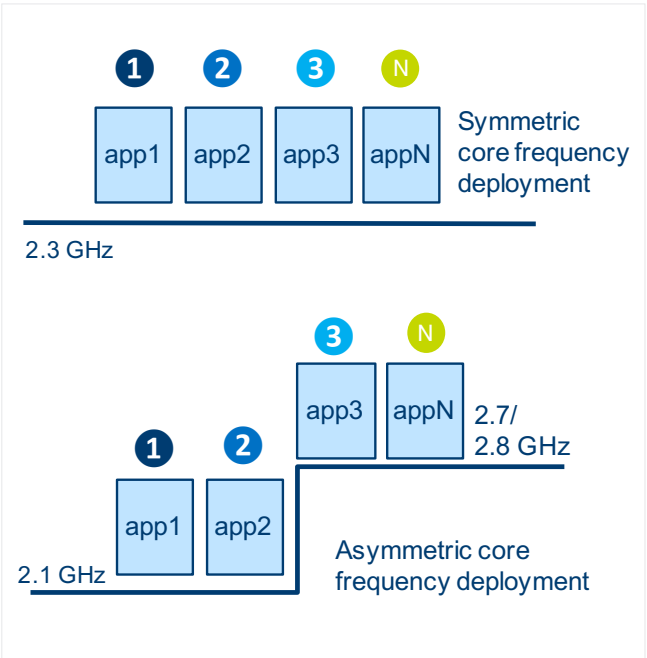


Figure 2. Intel® SST-BF

Figure 2 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.

Intel SST-BF is available on the Intel Xeon Gold 6252N processor and the Intel Xeon Gold 6230N processor (see Table 1). When in Intel SST-BF mode, the Intel Xeon Gold 6252N processor is available with up to 16 cores operating at 2.1 GHz and eight cores operating at 2.8 GHz for priority workloads. The Intel Xeon Gold 6230N processor features up to 14 cores operating at 2.1 GHz and six cores operating at 2.7 GHz for priority workloads.


INTEL® XEON® GOLD 6252N PROCESSOR (24C @ 2.3 GHZ @ 150W)

HIGH PRIORITY		STANDARD PRIORITY	
Cores	Base Frequency	Cores	Base Frequency
8	2.8 GHz	16	2.1 GHz

INTEL® XEON® GOLD 6230N PROCESSOR (20C @ 2.3 GHZ @ 125W)

HIGH PRIORITY		STANDARD PRIORITY	
Cores	Base Frequency	Cores	Base Frequency
6	2.7 GHz	14	2.1 GHz

Table 1. Intel® SST-BF enabled CPU SKUs



### 1.3.2 Intel® Optane™ DC Persistent Memory

Intel Optane DC persistent memory is a new technology that offers the capacity, cost, and performance benefits of both traditional DRAM and solid state device (SSD) storage. DRAM is popular for storage because it offers good performance, but it is relatively expensive and is volatile, which means data stored in the memory disappears when the server is powered off or rebooted. SSDs are not nearly as fast as DRAM, but cost less on a per-Gigabyte basis and they are non-volatile, so data is retained when the server is turned off or rebooted.

Intel Optane DC persistent memory is made up of non-volatile media placed onto a dual in-line memory module (DIMM), which is the traditional packaging for DRAM. Intel Optane DC persistent memory is installed on the memory bus alongside DRAM DIMMs. Intel Optane DC persistent memory is designed for use with 2nd generation Intel Xeon Scalable processors and is electrically and physically compliant to standard DDR4 sockets.

Intel Optane DC persistent memory provides an option between DRAM and SSD, bringing the economics and capacity of non-volatile SSDs to the form factor of standard memory. This helps lower the cost of memory while potentially increasing the size of server memory and/or providing the fastest persistent storage of data. Up to half of a server's DIMM slots can be used for Intel Optane DC persistent memory.

For example, there are now two scenarios to provide per-socket memory of 3 Terabytes:

- 12x 256 GB DDR4 DIMMs, or
- 6x 64 GB DDR4 DIMMs + 6x 512 GB Intel Optane DC persistent memory modules

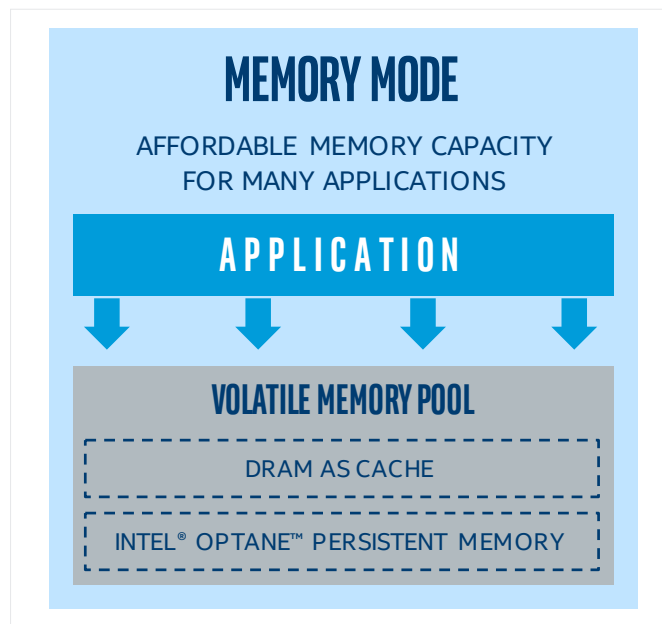
In this Memory Mode configuration (see Figure 3), the DDR4 DIMMs are “invisible” to the operating system and instead act as a fast cache for the memory modules.

## 2. Test Setup

### 2.1 Device Under Test Configuration

Figures 4 and 5 show the DUT setup and the systems used to exercise various traffic profiles for each test:

- Edge Routing Function: The packet forwarding performance of the Kernel-based Virtual Machine



**Figure 3.** Intel® Optane™ DC persistent memory configured in Memory Mode

(KVM)/ Quick Emulator (QEMU) virtual machine (VM) infrastructure was tested by enabling or disabling Intel SST-BF on Intel Xeon Gold 6252N CPU. In this configuration, eight high-priority CPUs (2.8 GHz) in Socket0 were pinned to the Edge Routing Function VM, then SR-IOV pass-through was configured. TRex traffic generators were connected to Edge Routing Function VM via an Intel® Ethernet Converged Network Adapter XL710 40G Ethernet connection.

- Memory Bandwidth: The Redis Labs server VM was used to measure GET/SET operations per second (OPS/sec), and memory bandwidth using the memtier-benchmark tool. The configuration includes one 2.1 GHz vCPU in socket1 that was pinned to each Redis VM. Then an SR-IOV VF as configured and various memory configurations were applied to compare DRAM-only mode (one-layer memory mode) with Intel Optane DC persistent memory (two-layer memory mode). The DUT system ran KVM/QEMU to instantiate multiple VMs. Each 20 to 40 VMs were interconnected with memtier-benchmark tools via a 25G Ethernet Intel® Ethernet Network Adapter XXV710 NIC.



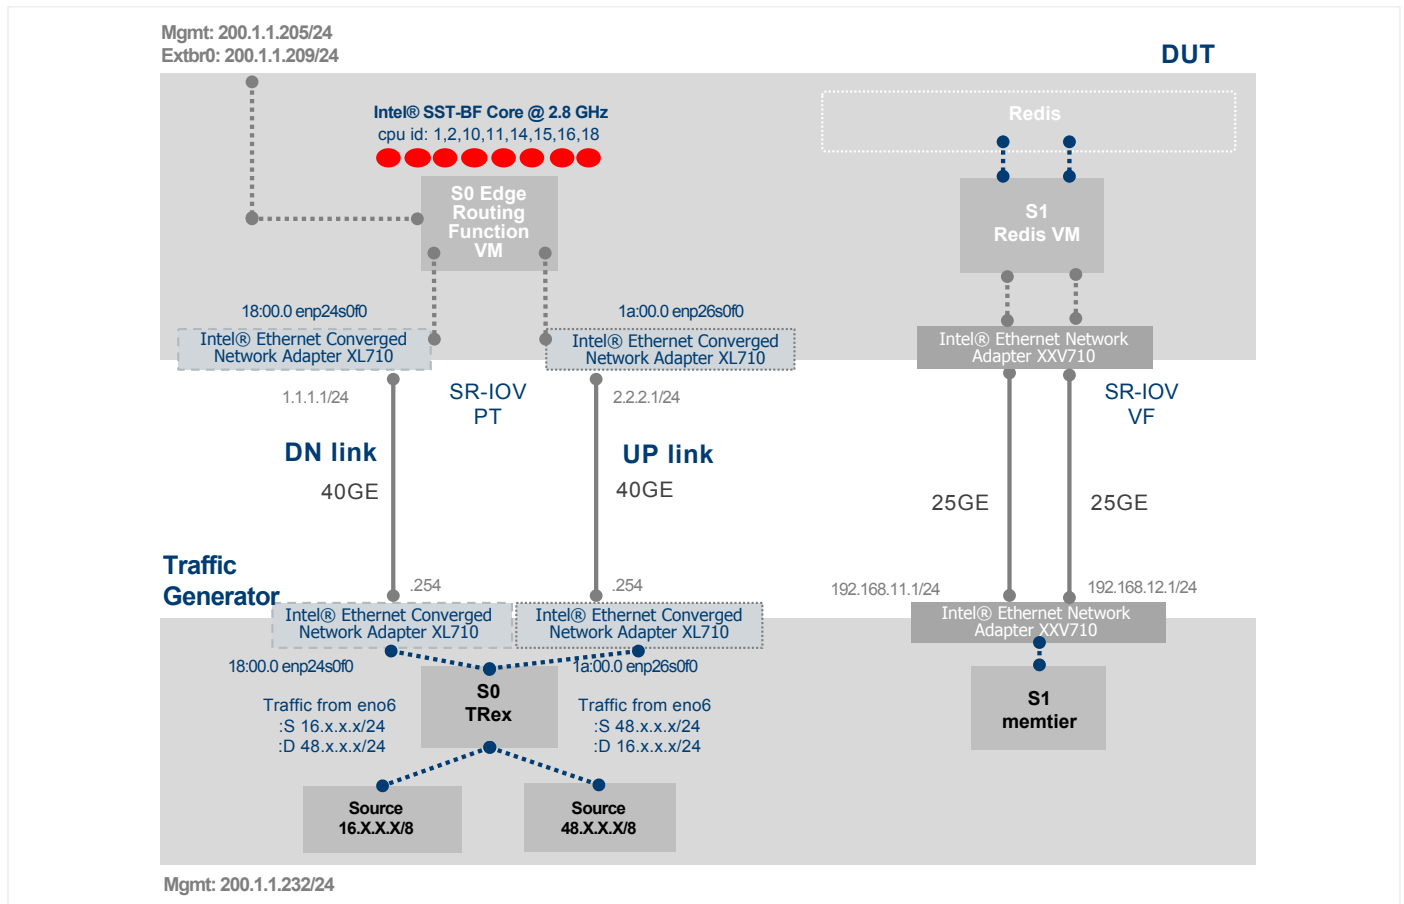


Figure 4. Test system overview (complete DUT details listed in Table 2)



Figure 5. Photos of the Intel® Server Board S2600WFT 2RU commercial off-the-shelf DUT

Tables 2 and 3 show the hardware and software composition of the DUT and traffic generator. Each of the tests were conducted using a fixed CPU configuration, utilizing socket0 for Intel SST-BF and socket1 for Intel

Optane DC persistent memory performance testing, respectively. This configuration helped to validate the scaled-up MEC architecture and also to isolate any inconsistency that occurred in each of the test results.

CATEGORY	DESCRIPTION	
PROCESSOR	Product	Intel® Xeon® Gold 6252N processor
	Microcode	0x500001c
	Frequency	2.3 to 2.8 GHz
	Cores per Processor	24 Cores/48 Hyper threads
MEMORY	DIMM Slots per Processor	6 Channels per Processor
	Capacity	192 GB DRAM (12x16 GB), 1.5 TB Intel® Optane™ DC persistent memory (12x 128 GB), configured as Memory Mode (FW version 01.02.00.5346)
	Memory speed	2666 MHz, DDR4
NETWORK	Number of ports	2 port Intel® Ethernet Converged Network Adapter XL710 40 GbE QSFP for Intel SST-BF test 2 port Intel® Ethernet Network Adapter XXV710 25 GbE SFP28 for Redis benchmark test
2RU SERVER	Vendor	Intel® Server Board S2600WFT
HOST OS	Vendor/version	CentOS Linux Release 7.6.1810 4.20.4-1.el7.elrepo.x86_64
KVM	Vendor/version	CentOS 2.12.0 (qemu-kvm-ev-2.12.0-18.el7_6.3.1)
BIOS	Vendor/version	Intel Corporation /SE5C620.86B.0D.01.0387.021220191105 Release Date: 02/12/2019
INTEL® SST-BF SCRIPT	Vendor/version	Intel Corporation /v1.2h
EDGE ROUTING FUNCTION	Vendor/version	SKT Custom Edge Routing Function version: 18.1-2101
REDIS SERVER	Vendor/version	Redis 5.0.3

Table 2. DUT Configuration

CATEGORY	DESCRIPTION	
PROCESSOR	Product	Intel® Xeon® Platinum 8280 processors
	Microcode	0x400000a
	Frequency	2.7 GHz
	Cores per Processor	28 Cores
MEMORY	DIMM Slots per Processor	6 Channels per Processor
	Capacity	192 GB DRAM
	Memory speed	2666 MHz, DDR4
NETWORK	Number of ports	2 port Intel® Ethernet Converged Network Adapter XL710 40 GbE QSFP for Intel SST Test 2 port Ethernet Network Adapter XXV710 25 GbE SFP28 for Redis Test
2RU SERVER	Vendor	Intel® Server Board S2600WFT
HOST OS	Vendor/version	CentOS Linux release 7.6.1810 3.10.0-957.12.1.el7.x86_64
BIOS	Vendor/version	Intel Corporation /SE5C620.86B.0D.01.0134.100420181737 Release Date: 10/04/2018
TREX	Vendor/version	Cisco TRex v2.53 Cisco TRex stateless gui v4.5.0-gb80ac59
REDIS MEMENTIER	Vendor/version	Redis Labs memtier_benchmark 1.2.15

Table 3. Traffic Generator Configuration

## 2.2 Test Procedures

### 2.2.1 Packet Forwarding Performance Enhancement by Intel SST-BF

To demonstrate the asymmetric core frequency benefit of deploying a system with Intel SST-BF enabled, Intel and SKT tested the SKT Edge Routing Function's packet forwarding performance with a configuration that has the MEC hosting a high priority VM that must have deterministic performance and would benefit from frequency scaling. During the test, Intel and SKT also ran the Redis memtier benchmark tool as a low priority workload that does not need real-time performance.

Throughput the tests, the performance of the Intel Xeon Gold 6252N CPU-based server board (Intel® Server Board S2600WFT) and the SKT Edge Routing Function was measured both with Intel SST-BF enabled and disabled. In both cases, the configuration included:

- SR-IOV pass-through
- 6x Intel SST-BF cores pinned to DPDK forwarding engine
- 1 GB Huge Pages
- 8 GB memory
- L3 static route was applied to TRex route-sourced networks

The TRex traffic generators were connected to Edge Routing Function via an Intel Ethernet Converged Network Adapter XL710 40G Ethernet controller. TRex generated the traffic and measured max throughput.

The BIOS configurations of the test setup were important and needed to be set to configure Intel SST-BF and Intel Turbo Boost Technology. The following BIOS options were used in the tests:

- Active Intel SST-BF: Enabled
- Configure Intel SST-BF: Disabled
- Intel Configurable TDP: Disabled
- Intel Turbo Boost Technology: Enabled
- Energy Efficient Turbo: Disabled

Next comes configuring the Intel SST-BF functionality by issuing `./sst-bf.py` with a configuration that is shown in Table 4. The Intel SST-BF was enabled on cores 1,2,10,11,14,15,16,18. On socket0, the frequency rates guaranteed minimum 2.8 GHz and a non-guaranteed maximum of 3.6 GHz; all the other cores were configured at 2.1 GHz.

An operating system kernel patch to enable Intel SST-BF was applied. This patch is available in the upstream Linux kernel version 4.20 and later. In addition, the server board must be booted with the `Intel_pstate` driver active.

The traffic profile of the packet generator reflects the characteristics of mobile traffic on a production SKT telecommunications network with the following IMIX packet size distributions:

- 64B (20%)
- 128B (10%)
- 256B (20%)
- 512B (10%)
- 1024B (20%)
- 1500B (20%)

These differing packet sizes were generated as an interleaving of multiple streams. Each TRex port generated full-duplex traffic processing conditions and the traffic rate was changed by L1 bandwidth rates until the non-drop rate (NDR) was zero as determined by a zero packet loss report from the TRex stream counters. The TRex generated traffic for one hour, which included 1,000 packets with adjustments made to the source and destination IP addresses, the source and destination port, with both TCP and UDP transport protocols. Packet sizes used ranged from 64 bytes to 1500 bytes but also included random packet sizes generated by IMIX (see IMIX profile above). RX errors also needed to be zero, and that was determined using the `softQFull` RX error counter that is built into the Edge Routing Function.

```
[root@mec test]# ./sst_bf.py -n 0
-----
SST-BF SETUP SCRIPT v1.2h
-----
Option: l
1,2,10,11,14,15,16,18,32,33,36,37,38,39,46,47,49,50,58,59,62,
63,64,66, 80,81,84,85,86,87,94,95
0xc0f30005cc06c0f30005cc06

# Enabling SST-BF
[root@mec test]# ./sst_bf.py -n 0
Option: i
Name = 6252N
CPUs = 96
Base = 2300
|-----sysfs-----|
Core | base  max  min |
-----|-----|
1 | 2800  2800  2800 |
2 | 2800  2800  2800 |
...
10 | 2800  2800  2800 |
11 | 2800  2800  2800 |
14 | 2800  2800  2800 |
15 | 2800  2800  2800 |
16 | 2800  2800  2800 |
18 | 2800  2800  2800 |
-----|-----|
```

**Table 4.** Intel® SST-BF configuration on DUT

## 2.2.2 VM Scalability on Single MEC Hosts by Intel® Optane™ DC Persistent Memory

To demonstrate the benefit of Intel Optane DC persistent memory in memory mode for MEC performance, Intel and SKT tested both Redis VM scalability and Redis VM performance. The tests used Redis-memtier benchmark tool with automation provided by Intel-developed Redis BKM WW10 scripts that automatically start and create VM/Redis memtier benchmark tools. The following memory configurations were tested to compare DRAM vs Intel Optane DC persistent memory performance:

- Average OPS/sec from the instantiated Redis VM
- Average memory bandwidth (KB/sec) from the instantiated Redis VM

**Table 5** shows Redis VM test configuration utilizing SR-IOV VF for each VM, 16 GB memory was allocated for VM scalability, 16 GB for Redis VM performance tests. The Redis-memtier benchmark tool configuration had a set packet size of 1024B. Per thread, the tests featured a 4:1 GET/SET ratio with Gaussian key-pattern, and three connections. All of the results were measured as a sub-sec SLA condition as shown below

```
memtier_benchmark --ratio=1:4 -d 1024 -n 210000000
--key-pattern=G:G --key-minimum=1 --key-
maximum=210000001 --threads=3 --pipeline=64 -c 3
--hide-histogram -s <server ip> -p <port>
```

CONFIGURATION	MEMORY CAPACITY	MEMTIER (GET/SET)	PACKET SIZE (BYTE)	ASSIGNED VM MEM (GB)	USED VM MEM (GB)	CPU/ VM	NO. OF REDIS- VMS
Redis VM Scalability	384 GB DRAM	4:1	1024	16GB	12GB	1	20
	768 GB Intel Optane DC Persistent Memory (6x 128 GB + 96 GB Cache)	4:1	1024	16GB	12GB	1	40
	768 GB Intel Optane DC Persistent Memory (6x 128 GB + 192 GB Cache)	4:1	1024	16GB	12GB	1	40
Redis VM Performance	384 GB DRAM	4:1	1024	18GB	13GB	1	20
	384 GB Intel Optane DC Persistent Memory (3x 128 GB + 48 GB Cache)	4:1	1024	18GB	13GB	1	20
	384 GB Intel Optane DC Persistent Memory (3x 128 GB + 96 GB Cache)	4:1	1024	18GB	13GB	1	20
	512 GB Intel Optane DC Persistent Memory+ 64 GB DRAM (4x 128 GB+2x 32 GB+ 64 GB Cache)	4:1	1024	18GB	13GB	1	20
	512 GB Intel Optane DC Persistent Memory+ 64 GB DRAM (4x 128 GB+2x 32 GB+ 128 GB Cache)	4:1	1024	18GB	13GB	1	20

**Table 5.** Redis memtier test configuration

## 3. Test Results<sup>4</sup>

### 3.1 Packet Forwarding Performance Results

Intel and SKT measured packet throughput using overall bit rate with zero packet loss. Figure 6 shows the throughput improvement from 2.92 Gbps to 3.47, Gbps which is an 18.65% increase using the same IMIX traffic profile but using Intel SST-BF to change the core frequency from 2.3 GHz to 2.8 GHz. The tests also showed a 28%

improvement in throughput (measured in both Mpps and Gbps) when the core frequency was set to 3.6 GHz. Figures 7 and 8 show the test results of each traffic profile. Intel SST-BF achieved a higher traffic throughput for Edge Routing Function workload and adjusted the frequency of other cores to stay within the chip's thermal design power (TDP) consumption. Overall, the tests showed a 15% to 28% improvement of layer-three packet forwarding performance in six core-pinned test cases with 100% utilization of the pinned core.





Figure 6. Throughput improvement from 2.3 GHz, 2.8 GHz, and 3.3 GHz

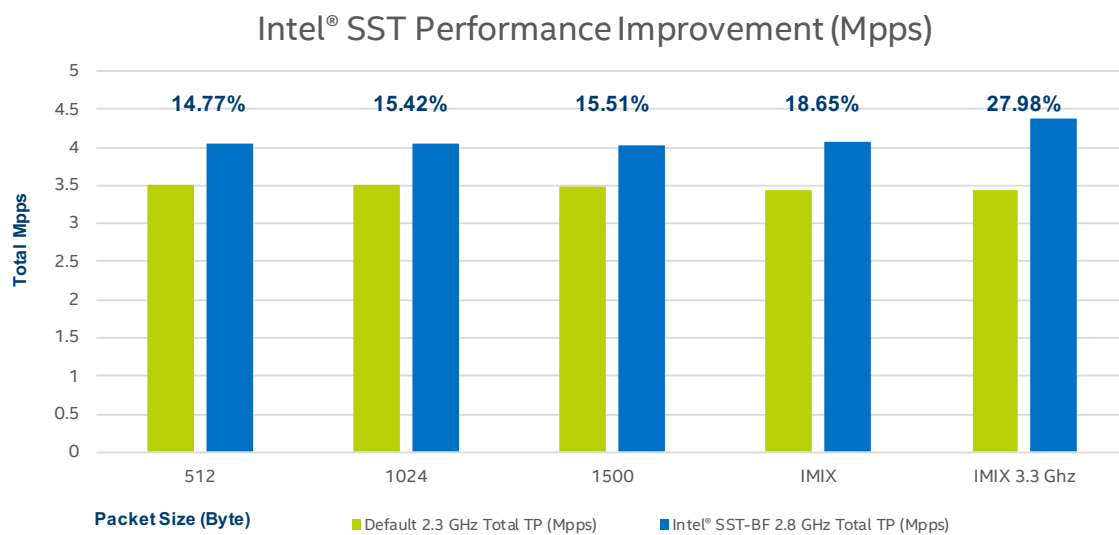


Figure 7. Layer 3 packet forwarding performance-millions of packets per second (Mpps)


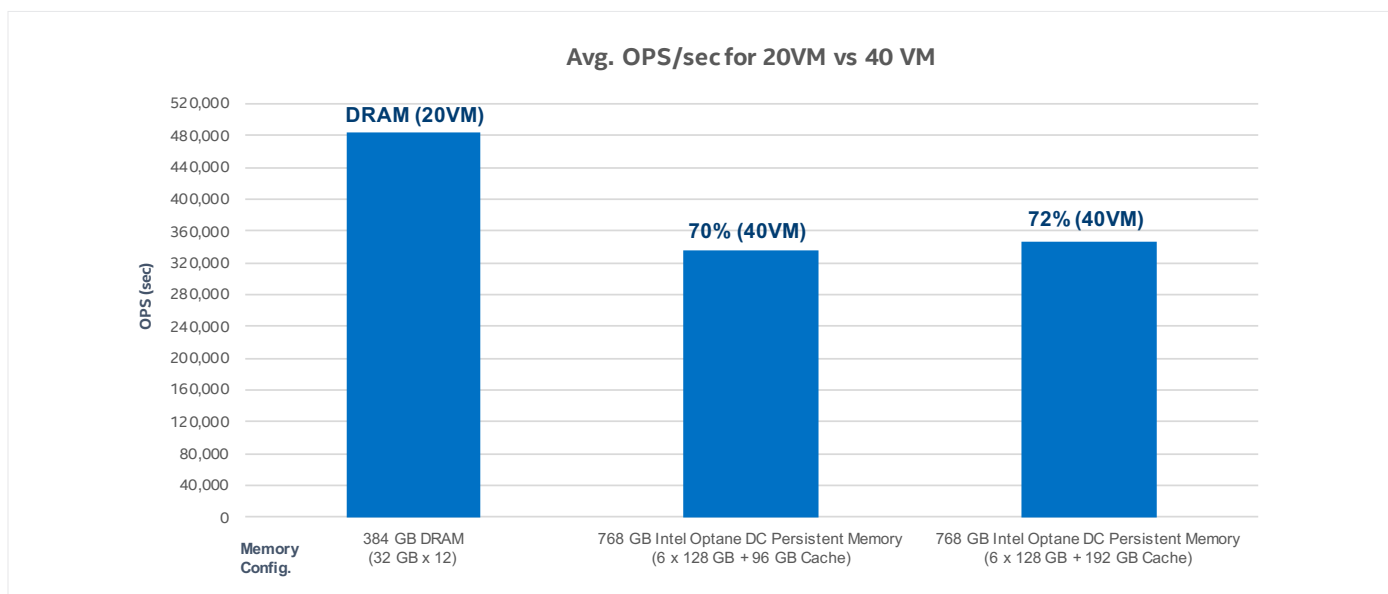


Figure 8. Layer 3 packet forwarding performance-Gbps

### 3.2 VM Scalability Test Result

For the Redis performance comparison, the tests measured overall average Redis throughput in OPS/sec of every VM that utilized the 2.3 GHz fixed frequency CPU in socket1. That performance was compared to that of the DRAM. Figure 9 provides Redis VM scalability test results for each memory configuration with a varying number of VMs. All of the tests meet SLA of less than 1 ms. Test highlights include:

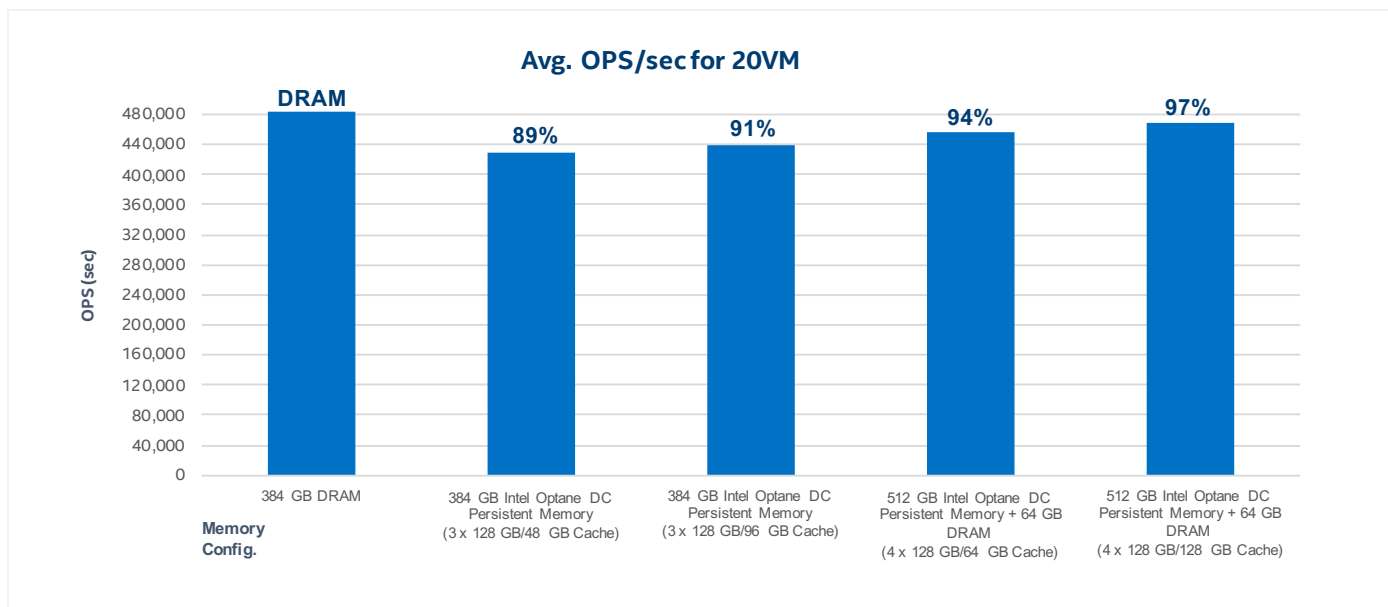
- 2x increase of Redis-server (40 VM per KVM) instances by increasing Intel Optane DC persistent memory capacity to 768 GB.
- 2x0.7 increase of total system throughput when using 768 GB of Intel Optane DC persistent memory vs 384 GB DRAM
- Reduced per-VM throughput with an increase in number of instances while achieving net higher system throughput.



**Figure 9.** Redis VM scalability test results. Charts show doubling of number of Redis servers per KVM and an almost doubling of total system throughput when Intel Optane DC persistent memory is increased to 768 GB.

Figure 10 provides the throughput performance results of Redis VM with the same configuration. All of the tests meet SLAs of less than 1 ms, demonstrating DRAM-like performance when used with the MEC software. Highlights include:

- Provided between 89% and 97% of DRAM throughput performance.
- Memory bandwidth increases contributed to the throughput cache memory increases.



**Figure 10.** Redis VM performance results. Intel Optane DC persistent memory provides up to 97% of performance of DRAM (which is measured in left bar).

## 4. Summary

These tests present some platform considerations that match the architecture of SKT MEC Platform with new Intel technologies to improve throughput. When Edge Routing Function software is run on a server using Intel Xeon Gold 6252N processors, then CommSPs can also utilize Intel SST-BF and Intel Optane DC persistent memory to boost performance. Some of the results of the tests described in this paper include:

- Increase in throughput of between 15% and 18% when Intel SST-BF is used to increase core frequency to 2.8 GHz and a ~28% increase at a base frequency of 3.6 GHz. Both of these results obtain these increases while remaining within the CPU's TDP because of a feature within Intel SST-BF that reduces frequencies in other cores to maintain chip TDP. Under abnormal traffic congestion, CommSPs can improve the quality of experience by adopting a mixed configuration that is to set the guaranteed minimum at 2.8 GHz and the non-guaranteed maximum at 3.6 GHz.<sup>4</sup>
- Improvement in deterministic forwarding performance for Edge Routing Function workloads can be achieved by pinning the application to an Intel SST-BF core (operating at a higher frequency). This proved to result in an overall system workload increase and stabilized production services by delivering constant Layer 3 forwarding throughput of about ~3.4 Mpps regardless of packet size.<sup>4</sup>

The tests also showed that Intel Optane DC persistent memory provides a viable alternative to DRAM-only implementations. Performance of this memory in these tests was close to that of DRAM and proved quite suitable for virtualized in-memory databases, such as Redis, and for virtual desktop infrastructure. Intel Optane DC persistent memory offered very good scale-out capability, near physical memory throughput, increased memory size and a balance between memory and CPU utilization. Specifically, the testing showed:

- Low cost of Intel Optane DC persistent memory could improve TCO for workloads that need large memory on single commercial off-the-shelf servers.<sup>4</sup>
- Up to 2X increase of VM instantiation, can reduce system cost per VM with Intel Optane DC persistent memory without sacrificing performance.<sup>4</sup>

Overall, the test results show the performance increases using the SKT 5G MEC platform that TEMs and ISVs can anticipate when they utilize the Intel Optane DC persistent memory and the Intel SST-BF functionality that is available on select models of 2nd generation Intel Xeon Scalable processors.

The capabilities enable larger data tables and more VM instances with improved throughput and utilization, which reduced SKT 5G MEC's server footprint. Even with higher performance, the tests showed low server power and cooling costs in all MEC deployment locations

MNOs can leverage SKT 5G MEC to meet customer demand and generate new services revenue from consumers and businesses. With the flexibility of Intel memory and CPU technologies, the software can offer very high levels of performance for challenging use cases.

## 5. Resources

Intel® Speed Select Technology – Base Frequency - Enhancing Performance Application Note 338928-001

[https://github.com/intel/CommsPowerManagement/blob/master/sst\\_bf.md](https://github.com/intel/CommsPowerManagement/blob/master/sst_bf.md)

Intel\_pstate driver: <https://www.kernel.org/doc/html/v4.15/admin-guide/pm/cpufreq.html>

<https://www.intel.com/content/www/us/en/architecture-and-technology/speed-select-technology-article.html>

Intel Optane DC Persistent Memory

<https://www.intel.com/content/www/us/en/products/memory-storage/optane-dc-persistent-memory.html>

<https://youtu.be/f9pIXw1ndRI>

<https://www.intel.com/content/www/us/en/architecture-and-technology/optane-dc-persistent-memory.html>

<https://Redislabs.com/press/Redis-labs-delivers-fastest-multi-model-database-intel-optane-dc-persistent-memory-3/>

TRex Packet Generator

<https://TRex-tgn.cisco.com/>

<https://github.com/cisco-system-traffic-generator/TRex-stateless-gui>

Redis Benchmark Tool

[https://github.com/RedisLabs/memtier\\_benchmark](https://github.com/RedisLabs/memtier_benchmark)

<https://www.datadoghq.com/blog/how-to-monitor-Redis-performance-metrics/#performance-metrics>

## 6. About the Authors

This whitepaper is jointly developed by SK Telecom and Intel Corporation.

**SK Telecom—MEC Dev. Team, Cloud Labs and ICT R&D Center**

Dongkee, Lee, [dongkee.lee@sk.com](mailto:dongkee.lee@sk.com)

Moonyoung, Chung, [mychung@sk.com](mailto:mychung@sk.com)

Woohyun, Nam, [namwh@sk.com](mailto:namwh@sk.com)

Keunhyun, Kim, [keunhyun.kim@sk.com](mailto:keunhyun.kim@sk.com)

**Intel Corporation— Data Center Group**

Wooram (Alex) Kim, NCS Solutions Architects (CCIE#9351/ MPC2.0/VCP5-DCV/RHCSA), [alex.kim@intel.com](mailto:alex.kim@intel.com)

Jongup (Jayden) Lee, NPG Platform Application Engineer, [jayde.lee@intel.com](mailto:jayde.lee@intel.com)

Sy jong, Choi, DCG Platform Application Engineer, [sy.jong.choi@intel.com](mailto:sy.jong.choi@intel.com)

Huisuk (Kevin) Hong, NPG Market Development Manager, [huisuk.hong@intel.com](mailto:huisuk.hong@intel.com)

Evstratios (Steve) Vandris, NPG Strategic Planner [steve.vandris@intel.com](mailto:steve.vandris@intel.com)

## 7. Glossary of Terms

TERM	DESCRIPTION
3GPP	Third Generation Partnership Project
5G	Fifth generation mobile wireless networks
AR/VR	Augmented reality/virtual reality
BWM	Bandwidth management
CDN	Content delivery network
COTS	Commercial off the shelf
CPU	Central processing unit
DIMM	Dual-inline memory module
DL	Downlink
DPDK	Data Plane Development Kit
DRAM	Dynamic random access memory
eMBB	Enhanced mobile broadband
EPC	Evolved packet core
ETSI	European Telecommunications Standards Institute
IMDB	In-memory database
Intel SST-BF	Intel Speed Select Technology- Base Frequency
IP	Internet protocol
ISV	Independent software vendor
MEC	Multi-access Edge Computing
MEPM	MEC platform manager

Mpps	Millions of packets per second
NFV	Network function virtualization
NFVI	Network function virtualization infrastructure
NIC	Network interface card
NSA	Non-standalone
OPS/sec	Operations per second
PDN	Packet data network
PGW	PDN Gateway
PGW-C	PDN gateway control plane
POP	Point of presence
RNIS	Radio network information service
SA	Standalone
SGW	Serving gateway
SGW-C	Serving gateway control plane
SLA	Service level agreement
SSD	Solid state drive
TEM	Telecom equipment manufacturer
UE	User equipment. Synonymous with mobile device.
UL	Uplink
UPF	User plane function
V2X	Vehicle to everything
VM	Virtual machine
VNF	Virtual network function



<sup>1</sup> See [https://www.etsi.org/deliver/etsi\\_gs/MEC/001\\_099/003/02.01.01\\_60/gs\\_MEC003v020101p.pdf](https://www.etsi.org/deliver/etsi_gs/MEC/001_099/003/02.01.01_60/gs_MEC003v020101p.pdf) for a diagram of the standard MEC software architecture.

<sup>2</sup> "ETSI MEC Deployments in 4G and Evolution towards 5G" ETSI White Paper No. 24, ETSI MEC ISG, February 2018 and "MEC in 5G Networks", ETSI White Paper No. 28, ETSI MEC ISG, June 2018

<sup>3</sup> "Developing SW for Multi-Access Edge Computing", ETSI White Paper No. 20, February 2019

<sup>4</sup> Testing conducted by SKT and Intel as of June 7, 2019. See the Test Setup and Test Results sections for configurations and full test results.

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 complete information visit [www.intel.com/benchmarks](http://www.intel.com/benchmarks).

Performance results are based on testing as of June 7, 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at [intel.com](http://intel.com).

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice Revision #20110804

Intel, the Intel logo, Intel Optane, and Xeon are trademarks of Intel Corporation or its subsidiaries.

Other names and brands may be claimed as the property of others.

© Intel Corporation

0919/DO/H09/PDF

Please Recycle

340676-002US