White Paper

intel.

Maximizing vCMTS Data Plane Performance with 4th Gen Intel® Xeon® Scalable Processor Architecture

Intel[®] platform technologies boost virtualized cable modem termination system (vCMTS) data plane performance by 32%, with 41% better performance per watt.

Authors

David Coyle

Senior Software Engineer Intel

Kevin O'Sullivan

Network Software Engineer Intel

Muhammad A. Siddiqui Platform Solutions Architect

Intel

Subhiksha Ravisundar Network Software Engineer

Intel

Introduction

The need for more network bandwidth is growing rapidly to meet the demands of today's broadband services consumers. Internet traffic surged during the COVID-19 pandemic and the rise of streaming video usage for remote working with conferencing applications, cloud computing, virtual-private networking (VPN), and the explosion of video-intensive user-generated content (UGC) seems to be here to stay [1] [2]. A few years ago, it was normal for a few dozen content producers like Netflix to serve billions of video content consumers worldwide. Today, with the explosion of social-media content producers, broadband networks are seeing more bandwidth usage than ever before [3]. The upstream traffic growth has outpaced downstream traffic growth; hence the need for cable multiple systems operators (MSOs) to augment their networks to bolster the upstream network capacity. Consequently, it is imperative for cable MSOs to increase their capacity and efficiency by modernizing their Hybrid Fiber-Coaxial (HFC) network.

Additionally, governments around the world have recently passed laws to encourage private sector investment in broadband network upgrades [4], with the aim in some countries to improve the access for unserved and underserved rural [5], as well as urban, communities. With these government incentives typically available to all network operators, including cable MSOs, the competitive landscape for fixed broadband has heated up considerably in the past year.

As a result of these recent consumer behavior changes and government incentives, the cable industry is aggressively rolling out Distributed Access Architecture (DAA), as defined by CableLabs [6]. DAA isn't a single technology but rather an umbrella term that describes the network architecture cable operators use to future-proof their access networks.

Several options exist to implement DAA. The most widely adopted approach to-date is the Remote-PHY (R-PHY) architecture as illustrated in Figure 1 below. In an R-PHY architecture, the physical layer function is moved to the fiber node and the Data Over Cable Service Interface Specification (DOCSIS) MAC layer, also known as virtualized Cable Modem Termination System (vCMTS), runs as a software component in the MSO's headend facilities on commercial-off-the-shelf industry standard servers. As MSOs transition to a vCMTS-based approach, some of the initial benefits they are realizing are reductions in energy consumption, space, and operational costs of heating, ventilation, and air-conditioning (HVAC) at the headend facilities [7].

This becomes increasingly valuable as consumer demand for network bandwidth continues to increase and the enhancements in traditional headend equipment, for efficiency and capacity, prove challenging within the limited space and power constraints.



Figure 1. Remote-PHY (R-PHY) and virtual-MAC-core architecture.

Table of Contents

Introduction
Running a vCMTS on 4 th Gen Intel® Xeon® Scalable Processors
Intel® vCMTS Reference Dataplane
Crypto and CRC Processing in the vCMTS Data Plane Pipeline
Threading Model Considerations for a vCMTS Data Plane Pipeline
vCMTS Data Plane Performance Analysis7
System Scalability and Server Sizing7
Gen-on-Gen Performance Improvements
Per Service Group Performance Comparison9
Compelling vCMTS Performance and Power Savings on 4 th Gen Intel Xeon Scalable Processors
Appendix A: Intel vCMTS Reference Dataplane Packet-processing Pipeline Stages
Appendix B: Performance Test Environment
Appendix C: Test Environment Configuration Information and Relevant Variables
Appendix D: System Configuration
Appendix E: Acronyms and Definitions

Table of Figures

Figure 1. Remote-PHY (R-PHY) and virtual-MAC-core architecture	
Figure 2. Intel® Xeon® Scalable CPU generational enhancements	
Figure 3. Intel® vCMTS Reference Dataplane	
Figure 4. Crypto/CRC enhancements in 4 th Gen Intel [®] Xeon [®] Scalable Processor architecture	
Figure 5. Core assignment on a 32-core 4 th Gen Intel [®] Xeon [®] Scalable processor7	
Figure 6. Platform throughput scalability with maximum bi-directional bandwidth per service-group	
Figure 7. Platform performance comparison for 3 rd Gen Intel® Xeon® Scalable and 4 th Gen Intel® Xeon® Scalable processors	
Figure 8. Per service group downstream performance comparison for 3 rd Gen Intel® Xeon® Scalable and 4 th Gen Intel® Xeon® Scalable processors	
Figure 9. Performance test environment used to measure vCMTS performance	2

Running a vCMTS on 4th Gen Intel[®] Xeon[®] Scalable Processors

The performance of vCMTS DOCSIS MAC software has been, and continues to be, greatly boosted by technologies such as open source Data Plane Development Kit (DPDK) [8] and the Intel® Multi-Buffer Crypto for IPSec library [9], which provides highly optimized packet processing in software that is tightly coupled to the continuously innovative Intel® architecture (IA). IA provides native instructions and features that specifically accelerate data plane packet processing for access networks.

4th Gen Intel[®] Xeon[®] Scalable processors, however, provide several major enhancements that can further boost access network performance in the most power-efficient manner to date. Multiple workload accelerators are incorporated into the 4th Gen Intel Xeon Scalable processor package, one of which is Intel[®] QuickAssist Technology (Intel[®] QAT).

Designed to help deliver higher performance on workloads requiring data encryption, Intel QAT allows expensive cryptography operations to be offloaded to a dedicated hardware block, thus freeing up CPU core cycles for other packet processing operations. With up to four Intel QAT devices available per CPU on certain 4th Gen Intel Xeon Scalable processor SKUs, this can give approximately 190 Gbps of DOCSIS BPI+ cryptography bandwidth for an average packet size of 1KB (as seen while running the tests described later in the paper). 4th Gen Intel Xeon Scalable processors are also Intel's most sustainable data center processors, delivering a range of features for optimizing power, while maintaining high performance¹. By making optimal use of CPU resources, the latest processors can help achieve customers' sustainability goals. The 4th Gen Intel Xeon Scalable processors have power management capabilities to enable more control and achieve greater operational savings. Through these capabilities, MSOs can make dynamic adjustments to save on electricity consumption as computing needs fluctuate.

In addition to Intel QAT as a built-in accelerator, and the power efficiency capabilities of 4th Gen Intel Xeon Scalable processors, MSOs and independent software vendors (ISVs) can improve vCMTS performance by taking advantage of the gen-on-gen CPU architecture enhancements shown in Figure 2, which include bigger cache-sizes, higher core-counts, DDR5 memory with higher speeds, and greater I/O bandwidth with PCIe Gen 5.

DOCSIS MAC functionality can be broken into four categories: downstream data plane, upstream data plane, control plane, and system management. From a network performance perspective, the most compute-intensive workload is data plane processing, and consequently, is the focus of this paper.



Figure 2. Intel® Xeon® Scalable CPU generational enhancements.

The paper focuses specifically on how to take advantage of the following advanced features in a DOCSIS MAC data plane:

- Intel QAT
- Dual AES encryption engines
- Intel[®] AES New Instructions (Intel[®] AES-NI)
- Enhanced Intel[®] Advanced Vector Extensions 512 (Intel[®] AVX-512)
- Intel[®] Vector PCLMULQDQ carry-less multiplication instruction
- User Wait power saving instructions

The paper also provides insights into implementation options and establishes an empirical performance data baseline that can be used to estimate the capability of a vCMTS platform running on industry-standard, high-volume servers based on 4^{th} Gen Intel Xeon Scalable processor architecture. Actual measurements were taken on an Intel Xeon Gold 6428N (Networking SKU) processor-based system. This CPU has 32 cores running at a 1.8 GHz clock speed and features two Intel QAT devices per CPU⁹.

The test results demonstrate how each core of an Intel Xeon 6428N processor can support the downstream channel bandwidth for close to five orthogonal frequency-division multiplexing (OFDM) channels in a pure DOCSIS 3.1 configuration for an Internet mix (IMIX) traffic blend on a fully scaled-out, software-based vCMTS system. When Intel QAT acceleration is employed, it will be shown that each core utilizing Intel QAT can easily satisfy the maximum downstream bandwidth of a pure DOCSIS 3.1 configuration, and even exceed the bandwidth of six OFDM channels. Since vCMTS workloads exhibit good scalability, a typical server blade based on dual Intel Xeon 6428N processors (e.g. with 64 processor cores in total) can achieve compelling performance density.

The performance of the Intel Xeon 6428N processor running at 1.8 GHz is also compared to a 3rd Gen Intel Xeon Scalable processor, specifically an Intel Xeon 6338N with 32 cores running at 2.2 GHz⁹, in order to demonstrate the performance and power saving benefits of the gen-on-gen architecture enhancements.

Intel® vCMTS Reference Dataplane

Intel has developed a DOCSIS MAC data plane pipeline that is compliant with DOCSIS 3.1 specifications and based on DPDK packet processing framework. It is publicly available on the Intel® vCMTS Reference Dataplane page on the Intel® Developer Zone (Intel® DevZone) website [10]. The main objective of this project is to provide a set of tools to demonstrate the vCMTS data plane packet processing performance and other relevant capabilities of Intel Xeon Scalable processor-based platforms, and to accelerate ISVs and MSOs in deploying next-generation vCMTS solutions.

Figure 3 shows the upstream and downstream packet processing pipelines implemented for the Intel vCMTS Reference Dataplane. Both pipelines are implemented as a two-stage pipeline, which perform upper MAC and lower MAC processing respectively. Typically, the downstream data plane pipeline stages are executed on separate software threads, and affinitized to hyper-thread siblings¹⁰ on a single Intel Xeon processor core.

The upstream pipeline stages, on the other hand, run as a single software thread that is affinitized to a single hyper-thread, with either another upstream pipeline thread or management/ control-plane/telemetry threads running on the sibling. The DPDK API used for each significant DOCSIS MAC data-plane packet-processing step is also shown. A detailed description of the upstream and downstream packet processing stages shown in Figure 3 is provided in Appendix A.



Figure 3. Intel[®] vCMTS Reference Dataplane.

Many key innovations and performance optimizations are supported by the Intel vCMTS Reference Dataplane, including the following:

- Optimized multi-buffer implementation of combined AES cryptographic (crypto) and CRC processing based on Intel AES-NI and Intel AVX-512 instructions.
- Optimized multi-buffer implementation of DES crypto processing based on Intel AVX-512 instructions.
- Acceleration of AES and DES DOCSIS crypto processing using Intel QAT.
- Acceleration of combined AES crypto and CRC processing using Intel QAT.
- Flexible and efficient crypto device usage via the DPDK Cryptodev Scheduler.
- Optimized CRC32 and DOCSIS header check sequence (HCS) calculation based on Intel AVX-512 instructions.
- Optimized DOCSIS Filter and QoS Classification based on Intel AVX-512 instructions via the DPDK Access Control List (ACL) API.
- DOCSIS service flow and channel scheduling implementation based on DPDK HQoS.
- Packet Streaming Protocol (PSP) fragmentation/ reassembly implementation based on the DPDK Mbuf API.
- DOCSIS MAC data plane pre-processing using the 100GbE Intel[®] Ethernet 800 Series network interface card (NIC).

- Configurable DOCSIS MAC data plane threading options

 dual-thread, single-thread, separate or combined
 upstream and downstream threads.
- Power-optimized C-states through the User Wait instructions, and enabled via the DPDK Power Management API [11].

The Intel vCMTS Reference Dataplane runs in either a baremetal environment or within a container-based environment with Kubernetes orchestration, as described in Appendix B.

Crypto and CRC Processing in the vCMTS Data Plane Pipeline

Encryption, specifically AES-based baseline privacy interface (BPI+) encryption, and data packet CRC generation are the two main byte-wise operations of a typical vCMTS data plane packet processing pipeline, and as a result they consume a significant portion of CPU core cycles of such a pipeline. However, by leveraging various cryptographic-related capabilities of 4th Gen Intel Xeon Scalable processors, combined with the gen-on-gen CPU architecture enhancements described earlier, the performance of this processing has been significantly improved.

Figure 4 illustrates the key crypto-related capabilities and enhancements in 4th Gen Intel Xeon Scalable processor architecture. Crypto performance was improved on 3rd Gen Intel Xeon with the addition of a second AES port for each CPU core and support for vectorized Intel AES-NI instructions. An Intel AVX-512 vectorized version of the PCLMULQDQ carryless multiplication instruction also significantly improved CRC calculation performance. All these enhancements remain in place on 4th Gen Intel Xeon Scalable processors.



Figure 4. Crypto/CRC enhancements in 4th Gen Intel® Xeon® Scalable Processor architecture.

In the case of software-based crypto and CRC processing, a key characteristic is that they have per-byte costs; hence, the cost of encryption/decryption and CRC generation for a packet is directly proportional to the size of the packet. If done as sequential steps of a vCMTS data plane pipeline, the per-byte CPU cycle cost of these operations is additive.

To reduce this per-byte cost, support was previously added to the Intel Multi-Buffer Crypto for IPSec library for combined (or stitched) crypto and CRC processing. This innovation reduces overall CPU cycle cost by performing the CRC and crypto processing of subsequent blocks of a packet in parallel, achieved by interleaving Intel AES-NI, PCLMULQDQ and other Intel AVX-512 instructions in the code. The net result is just a single byte-wise pass of each packet is required to perform both operations. Combining this with the multi-buffer capabilities of Intel Xeon Scalable processors provides highly optimized crypto and CRC software-based processing.

Intel QAT performs hardware-accelerated crypto functions and can effectively reduce the CPU cycle cost of encryption regardless of packet size. Aside from being integrated into the chipset on some 3rd Gen Intel Xeon processor SKUs, Intel QAT has typically been added to a system via a PCIe add-in card, thus consuming a valuable PCI slot on the server. The major advancement with 4th Gen Intel Xeon is that Intel QAT is now integrated into the scalable processor SKUs in the form of the latest Intel QAT Hardware version 2.0 [12]. As well as no longer requiring a PCI slot, incorporating Intel QAT into the processor package optimizes its interaction with the processor cores due to the now physical proximity.

Intel QAT can be used to offload either of the following from the IA core:

- 1. The crypto processing only, in which case the CRC must still be computed in software, albeit accelerated by the vectorized PCLMULQDQ and other Intel AVX-512 instructions.
- 2. As a recent innovation, both the crypto and CRC processing as a single combined operation.

Both Intel QAT acceleration and Intel Multi-Buffer Crypto for IPSec library acceleration are embedded in DPDK, allowing easy access to both by a vCMTS data plane pipeline implementation via the Cryptodev API. Since version 20.08 of DPDK, support for the aforementioned combined crypto and CRC processing has been available through its security API:

- For Intel Multi-Buffer Crypto for IPSec library acceleration, DPDK leverages the combined crypto and CRC processing available in the crypto library.
- For Intel QAT, a DPDK QAT Poll Mode Driver (PMD) patch file within the Intel vCMTS Reference Dataplane v22.10.0 software package [10] adds configurable support to DPDK v22.07 for the combined processing offload, with plans to upstream to a future DPDK version. In DPDK versions prior to v22.07 or if the combined offload is disabled in DPDK v22.07+, the CRC is calculated in software within the Intel QAT PMD.

An important consideration in terms of Intel QAT usage is the maximum bandwidth of the available devices. It was seen during the testing described later in the paper that each Intel QAT device on 4th Gen Intel Xeon Scalable processors has an approximate bandwidth of 47 Gbps for AES BPI+ encryption of an IMIX packet mixture with an average packet size of 1KB when accessed via the DPDK QAT PMD; with two Intel QAT devices on a 4th Gen 6428N SKU, this gives a total bandwidth of 94 Gbps.

For combined AES BPI+ encryption and CRC offload for the same IMIX packet mixture, each Intel QAT device has an approximate bandwidth of 36 Gbps, or 72 Gbps across two devices (again as seen during testing). The above bandwidth limits will determine how the overall Intel QAT bandwidth can be distributed across service groups (SGs) or vCMTS data plane pipeline instances.

Threading Model Considerations for a vCMTS Data Plane Pipeline

Intel® Hyper-Threading Technology may be used to further improve vCMTS performance. Testing done using 3rd Gen Intel Xeon Scalable processors showed that the downstream upper and lower MAC processing stages have a similar CPU cycle cost [13]. As a result, there is a benefit in deploying them separately on sibling hyper-threads (i.e., hyper-threads of the same processor core) because it enables an effectively greater number of instructions to be executed per second on a CPU core through hyper-thread time-slicing.

Running the downstream data plane pipeline in a run-tocompletion model and affinitizing the upper and lower MAC stages to hyper-thread siblings can also help ensure that packet data (and other ancillary data such as cable modem data, crypto security association data, etc.) remain on the same core and therefore, potentially, the same cache for the entirety of the pipeline processing.

For DOCSIS 3.1, the upstream data plane has significantly lower bandwidth requirements than the downstream; however it can also benefit from hyper-threading by running two, singlethreaded, upstream data plane instances on sibling hyperthreads.

Figure 5 illustrates how up to 16 service groups can be deployed on a 32 core 4th Gen Intel Xeon Scalable processor, making use of hyper-threads as described above. A number of cores are reserved for the operating system (OS), telemetry and other DOCSIS MAC functions (control-plane and upstream scheduler for example); the remainder are dedicated to the upstream and downstream data plane pipelines.

orel HT2 L2 Core3 HTT HT2 L2	Core5 HTT HT2 L2 Core7 HTT HT2 L2 Core7 Core9 HTT HT2 L2 Core8 HTT HT2 L2 Core8 HTT HT2 L2 Core8 Core8 Core8 L2 Core8 Cor	CPU torel3 T HT2 L2 Corel5 Corel7 HT HT2 L2 Corel9 Corel9 Corel9 HT HT2 L2 Corel9 Corel9 Corel9 L2 Corel9 L2 Corel9 Corel9 L2 Corel9 L2 Corel9 Corel9 L2 Corel9 Corel9 Corel9 L2 Corel9 Corel9 Corel9 L2 Corel9 Corel9 Corel9 L2 Corel9 Corel9 Corel9 Corel9 L2 Corel9 Corel9 Corel9 L2 Corel9 Corel9 Corel9 Corel9 L2 Corel9 Co
		L3 Cache
ore2 1 HT2 Core4	Core6 Core8 Core10 Core12 C HT HT2 HT HT2 HT HT2 HT	corel4 Corel6 Core20 Core22 Core24 Core26 Core28 Core30 Core3 n n n n n n n n n n
L2 L2		
	OS, Orchestration, Telemetry	L2 L2 L2 L2 L2 L2 L2 L2 L2 16 x vCMTS Service-Groups on 32 CPU Cores (with Upstream Scheduling)
	L2 L2 L2 L2 OS, Orchestration, Telemetry Upstream Dataplane	L2 L2 L2 L2 L2 L2 L2 L2 16 x vCMTS Service-Groups on 32 CPU Cores (with Upstream Scheduling) • 2 Cores for OS, Telemetry, Orchestration

Figure 5. Core assignment on a 32-core 4th Gen Intel[®] Xeon[®] Scalable processor.

vCMTS Data Plane Performance Analysis

Performance benchmarks were run using the Intel vCMTS Reference Dataplane with a channel configuration that could demonstrate the maximum throughput capability for a pure DOCSIS 3.1 deployment on 4th Gen Intel Xeon Scalable processors.

For downstream traffic, the benchmarks consider three scenarios for the crypto and CRC processing:

- 1. Crypto and CRC processing performed in software, through the Intel Multi-Buffer Crypto for IPSec library, for all service groups.
- 2. Crypto processing offloaded to Intel QAT for up to 8 service groups. The CRC processing is performed in software within the DPDK QAT PMD. The 8 service groups are distributed evenly across the 2 Intel QAT devices of the Intel Xeon 6428N processor, allowing up to a maximum of 12 Gbps of crypto Intel QAT bandwidth per service group.
- 3. Crypto and CRC processing offloaded to Intel QAT for up to 6 service groups. The 6 service groups are distributed evenly across the 2 Intel QAT devices built into the Intel Xeon 6428N processor, allowing up to a maximum of 12 Gbps of crypto and CRC Intel QAT bandwidth per service group.

In the cases of scenarios 2 and 3 above, the crypto and CRC processing for any extra service groups is performed in software, again through the Intel Multi-Buffer Crypto for IPSec library.

For upstream traffic, it is assumed that CRC verification is not required. With lower throughput requirements than downstream, the upstream crypto processing is performed entirely by the Intel Multi-Buffer Crypto for IPSec library. It is worth noting that since upstream traffic generally has a higher percentage of smaller packets, there typically isn't any significant benefit in using Intel QAT. The channel configuration used is a pure DOCSIS 3.1 deployment with 6 orthogonal, frequency-division multiplexing (OFDM) channels that produces a theoretical cumulative bandwidth of 11.34 Gbps. Typically, the effective downstream bandwidth is limited to 10 Gbps for the DOCSIS 3.1 bandwidth limit per service group and this is usually enforced within the vCMTS data plane pipeline with downstream QoS rate limiting of 10 Gbps per service group. However, for the purposes of these benchmarks, this limit was increased significantly to allow the full potential of a vCMTS data plane pipeline on 4th Gen Intel Xeon Scalable processors to be demonstrated.

System Scalability and Server Sizing

For cable system architects who wish to size server requirements, scalability of the cable workload performance across CPU cores is very important. Figure 6 shows DOCSIS 3.1 (6 downstream OFDM channels with AES encryption) bidirectional throughput as it scales quite linearly with the number of service groups when passing IMIX traffic. This is achieved because the data plane workload for each service group runs independently of each other on dedicated cores or hyper-threads with optimal sharing of key CPU resources such as L3 cache. As can be seen, this linearity holds true regardless of whether the crypto and CRC processing is performed in software by the Intel Multi-Buffer Crypto for IPSec library or offloaded, either partially or fully, to Intel QAT.

Downstream traffic (IMIX) is supported on a single core per service group with two upstream traffic instances per core. For this benchmark, all processing resides on just one of the two processors of a dual processor platform. In all cases, the upstream data plane threads are handling 2 Gbps per service group.

White Paper | Maximizing vCMTS Data Plane Performance with 4th Gen Intel® Xeon® Scalable Processor Architecture



Figure 6. Platform throughput scalability with maximum bi-directional bandwidth per service-group.

At all service group numbers, the benefit of offloading the expensive byte-wise crypto and CRC operations to Intel QAT is clearly visible. It is worth noting that at lower service group numbers, it is preferable to offload both the crypto and CRC processing to Intel QAT. However, once 8 service groups are using Intel QAT, it is better to offload only the crypto processing to Intel QAT in terms of aggregate processor throughput. However, as will be explained in subsequent sections, there are still advantages of offloading both operations to Intel QAT at these higher service group numbers.

Note that when dimensioning a system, other processing should also be considered, which was not included in the performance data presented in this paper. For instance, the upstream scheduler, control plane, high availability (HA) standby instances, and other functionality hosted on the platform will consume additional processor cores. For the performance data presented here, CPU cores have been reserved for this processing.

Gen-on-Gen Performance Improvements

Figure 7 shows a performance comparison between Intel Xeon 6338N (3^{rd} Gen) processor and Intel Xeon 6428N (4^{th} Gen) N-SKU processor-based platforms for 16 service groups. The maximum gen-to-gen bi-directional performance gain is 32%, due primarily to CPU architecture enhancements and the addition of Intel QAT on 4^{th} Gen Intel Xeon Scalable processors. Without any Intel QAT offload, there is a 20% gain, which can be attributed solely to the general CPU architecture enhancements.



Figure 7. Platform performance comparison for 3rd Gen Intel® Xeon® Scalable and 4th Gen Intel® Xeon® Scalable processors.

While maximum performance remains crucial in all decisions regarding vCMTS deployments, performance (or throughput) per watt is becoming an increasingly important factor as the global trend to reduce power consumption across industry and society as a whole intensifies. Both the Intel Xeon 6338N (3rd Gen) processor and Intel Xeon 6428N (4th Gen) processor used in this comparison have a thermal design power (TDP) of 185W. However, the general CPU architecture enhancements and addition of Intel QAT available on the Intel Xeon 6428N results in up to a 41% improvement in the CPU's throughput per watt (0.82 Gbps/W --> 1.16 Gbps/W).

Figure 7 also shows that the actual power consumption between the Intel Xeon 6338N (174.42 W) and Intel Xeon 6428N (< 166 W) processor is lower, even when taking the increased throughput into account. This lower power envelope on the Intel Xeon 6428N processor provides more scope for other processing to be performed on the processor (or potentially turbo to be enabled on some cores) before the TDP is hit.

Further power savings (and subsequently further improvements in performance per watt) can also be achieved on 4th Gen Intel Xeon Scalable processors through new power management techniques that are applicable to data plane workloads, namely the new power-optimized C-states CO.1

and CO.2. User-space applications can directly place a core into these states using a new User Wait instruction set when traffic rates are low and, unlike the deeper C-state (C1, C1E, C6), these new C-states have negligible exit latencies.

The power savings achieved via these new power-optimized C-states can be realized with no impact on the maximum achievable throughput of the server. Support for this power management technique is integrated directly into DPDK through the Ethernet PMDs, meaning minimal code changes to a vCMTS application are required to enable these power savings.

Per Service Group Performance Comparison

Earlier results showed that, as the number of service groups scaled out towards 16, the benefit of offloading both the crypto and CRC processing to Intel QAT appeared to diminish in terms of overall system throughput. In order to see the benefit of offloading both operations to Intel QAT, the results must be analyzed a more deeply. Figure 8 shows the downstream throughput on a per service group level when 16 service groups are handling their maximum throughput, again on both Intel Xeon 6338N and Intel Xeon 6428N processors. It is important to note that each service group is also handling 2 Gbps of upstream traffic.



Figure 8. Per service group downstream performance comparison for 3rd Gen Intel[®] Xeon[®] Scalable and 4th Gen Intel[®] Xeon[®] Scalable processors.

The two leftmost results show the gen-on-gen improvement in downstream throughput, which is achieved purely through CPU architecture enhancements. At 8.7 Gbps, each of the 16 service groups can handle the downstream channel bandwidth for close to five OFDM channels in a pure DOCSIS 3.1 configuration.

The two rightmost sets of results show the additional benefit of the Intel QAT offload on 4th Gen Intel Xeon Scalable processors. Offloading either crypto-only or combined crypto-CRC processing to Intel QAT allows those service groups to easily support the maximum 10 Gbps downstream bandwidth that is specified in the DOCSIS 3.1 standard. With both crypto and CRC processing offloaded, those service groups are now capable of supporting the downstream channel bandwidth of 6 OFDM channels (11.34 Gbps), with some additional headroom (11.85 Gbps per SG). Although possibly superfluous at first sight as both Intel QAT offload scenarios are capable of handling the maximum 10 Gbps downstream bandwidth requirement, the results show that offloading the crypto and CRC operations can improve an individual service-group's performance as more processing cycles have been offloaded from the CPU core. These extra saved cycles could be used for other processing within the vCMTS pipeline or allow for some potential power savings through the power management techniques mentioned previously.

The caveat remains, however, that performing both crypto and CRC on the Intel QAT device lowers the overall bandwidth of the device and therefore, fewer service groups can take advantage of the full offload capability. Balancing the maximum service group throughput that can be achieved through using Intel QAT versus the maximum bandwidth of the available Intel QAT devices (and therefore, the number of service groups that can ultimately use Intel QAT) is an important consideration when dimensioning a vCMTS deployment on 4th Gen Intel Xeon Scalable processors.

It is also worth reiterating at this point that the benchmarks presented here were achieved with the DOCSIS 10 Gbps downstream bandwidth limit significantly increased to allow the full potential of 4th Gen Intel Xeon Scalable processor cores be demonstrated. Maintaining the 10 Gbps limit could potentially allow more service groups to make use of the available Intel QAT bandwidth.

Compelling vCMTS Performance and Power Savings on 4th Gen Intel Xeon Scalable Processors

More network operators are adopting software-based, virtualized solutions running on industry-standard, high-volume Intel Xeon Scalable processor-based servers to increase agility, flexibility, and cost competitiveness, while superseding the performance of custom-built proprietary solutions. Utilizing the 4th Gen Intel Xeon Scalable processor architecture offers an even higher level of performance, scalability, and powerefficiency than previous processor generations.

This paper has demonstrated how the scalability of 4th Gen Intel Xeon Scalable processor-based platforms can be applied to future OFDM-focused deployments, supporting almost five OFDM channels per CPU core in a fully software-based DOCSIS MAC pipeline. Introducing crypto and CRC offload to the Intel QAT accelerator allows the per service group downstream throughput to easily meet the 10 Gbps of bandwidth specified in DOCSIS 3.1.

The paper highlights the new features and advancements in the 4th Gen Intel Xeon Scalable processor-based platforms that can improve vCMTS data plane performance with even less power consumption than its predecessor. Results in this paper show that the 4th Gen Intel Xeon Scalable processor improves vCMTS data plane performance by up to 32% and performance per watt by up to 41%, when compared to similar N-SKUs of the 3rd Gen Intel Xeon Scalable processor product line.

With this data, network architects can better assess the benefits and costs of implementing a vCMTS on industry-standard, high-volume servers based on 4^{th} Gen Intel Xeon Scalable processors.

Learn More

For more information on networking solutions using Intel technologies, please visit networkbuilders.intel.com.



Appendix A: Intel vCMTS Reference Dataplane Packet-processing Pipeline Stages

The following describes the packet-processing stages of the Intel vCMTS Reference Dataplane Upstream and Downstream pipelines as shown in Figure 3.

Downstream Data Plane Pipeline Stages

1. Receive IP Frames

Using the DPDK Ethdev API, IP packet bursts are received via the DPDK Poll Mode Driver (PMD) from the Rx queue of a NIC virtual function (VF) port. These packets are read by a data plane software thread that begins vCMTS downstream packet processing. Packets are steered to a service group specific VF-based on destination MAC address.

2. Cable Modem Lookup and Subscriber Management

The DPDK Hash API is used to do a bulk lookup (i.e., with multiple packets) based on the destination IP address of the received frames to retrieve cable modem records containing MAC address, DOCSIS filter, DOCSIS classifier, service flow queue and security info. The Destination MAC address of the Ethernet frame is also updated to a cable modem specific address in this stage. The number of active subscriber IP addresses is checked against the DOCSIS limit (tracked by a destination IP address list per cable modem).

3. DOCSIS Filtering

The DPDK Access Control List (ACL) API is used to apply an ordered list of DOCSIS filter rules to Ethernet frames. DOCSIS filter rule configuration is described in Appendix C.

4. DOCSIS Classification

The DPDK ACL library is used to apply an ordered list of rules to classify Ethernet frames for enqueuing to cable modem service-flow scheduler queues. DOCSIS service-flow scheduler rule configuration is described in Appendix C.

5. DOCSIS QoS – Service Flow and Channel Access Scheduling

The DPDK hierarchical QoS (HQoS) scheduler API is used to apply rate-shaping, congestion control, and weighted-roundrobin (WRR) scheduling to cable modem service flow queues. The DPDK Scheduler API has also been adapted to perform channel access scheduling on data packets after service-flow scheduling. Channel access scheduling is optimized by performing it in an earlier pipeline stage than is typically done in other implementations. This scheduling stage takes into account the DEPI and DOCSIS encapsulation overhead added later in the pipeline.

6. Lower MAC Interface

A DPDK ring is used to transfer packets between upper MAC and lower MAC processing. This allows upper and lower MAC processing to be executed on separate threads.

7. DOCSIS Framing

DOCSIS MAC headers are generated, including DOCSIS header check sequence (HCS), for prepending to packets. The DPDK CRC API is used to generate the DOCSIS HCS.

Intel AVX-512 instructions are used for optimum performance on $4^{\rm th}$ Gen Intel Xeon Scalable processor platforms.

8. IP Frame CRC Generation and DOCSIS BPI+ Encryption

The 32-bit Ethernet cyclic redundancy code (CRC) of the packet is generated and DOCSIS BPI+ encryption is then applied. These two stages are performed using DPDK combined crypto-CRC processing. Intel AES-NI, Intel AVX-512 and other vectorized instructions, and Intel QAT are used for optimum performance on 4th Gen Intel Xeon Scalable processor platforms.

9. DEPI Encapsulation

DEPI encapsulation is performed based on the DOCSIS 3.1 specification. Frames are converted to Packet Streaming Protocol (PSP) segments, concatenated using DPDK Mbuf chaining, and encapsulated into L2TP frames of maximum transmission unit size. PSP segments are fragmented across DEPI frames, so all transmitted frames are of maximum transmission unit (MTU) size in order to ensure maximum utilization of the R-PHY link.

10. Transmit DEPI/L2TP Frames

Using the DPDK Ethdev API, bursts of DEPI/L2TP frames are transmitted via the DPDK PMD to the NIC VF Tx queue of the associated service group.

Upstream Data Plane Pipeline Stages

1. Receive UEPI/L2TP Frames

Using the DPDK Ethdev API, bursts of L2TP/IP frames containing UEPI encapsulated DOCSIS streams are received via the DPDK PMD from the Rx queue of a NIC VF port. These frames are read by a data plane software thread that begins vCMTS upstream packet processing. Frames are steered to the service group specific VF based on destination MAC address.

2. Validate Frame and Strip IP Headers

The L2TP/IP frame is validated, and IP headers are stripped.

3. UEPI Decapsulation

UEPI decapsulation is performed based on the DOCSIS 3.1 specification. UEPI/PSP sequence numbers are verified to be in order.

4. Service ID Lookup and DOCSIS Segment Reassembly

UEPI PSP header, data and trailer segments are traversed, and the data segments are reassembled into DOCSIS stream segments. The DPDK hash API is used to perform lookups based on service ID values to retrieve cable modem info.

5. DOCSIS Frame Extraction and Upstream Scheduling

DOCSIS frames are extracted from DOCSIS stream segments, including reassembly of fragmented frames using the DPDK Mbuf API. Any bandwidth requests extracted from the UEPIs are forwarded to an upstream scheduler at this point.

6. DOCSIS Frame HCS Verification

Header check sequence (HCS) verification is performed for the extracted DOCSIS frames using the DPDK CRC API.

7. DOCSIS BPI+ Decryption and IP Frame CRC Verification

DOCSIS BPI+ decryption is applied to DOCSIS frames for AES or DES encrypted frames and the 32-bit Ethernet cyclic redundancy code (CRC) of the resulting Ethernet packet is verified. These two stages are performed using the DPDK combined crypto-CRC processing. Intel AES-NI, Intel AVX-512 and other vectorized instructions are used for optimum performance on 4th Gen Intel Xeon Scalable processor platforms. Note that CRC verification is generally not required for upstream encapsulated packets, so it is disabled by default.

8. Transmit IP Frames

Using the DPDK Ethdev API, bursts of IP frames are transmitted via the DPDK PMD to the NIC VF Tx queue of the associated service group.

Appendix B: Performance Test Environment

The performance test environment for the Intel vCMTS Reference Dataplane consists of a vCMTS node and a software-based traffic-generator node, as shown in Figure 9.

The vCMTS node is based on a server blade with dual 4^{th} Gen Intel Xeon Scalable processors, with 100GbE Intel Ethernet 800 Series NICs and on-CPU Intel QAT devices.

The traffic-generator node is based on a server blade with dual 3rd or 4th Gen Intel Xeon Scalable processors⁹, again with 100GbE Intel Ethernet 800 Series NICs.

Servers based on other Intel Xeon Scalable processors or Intel Xeon D processors with different core counts are also supported. Different types of Intel Ethernet NICs may also be used.

On the vCMTS node, multiple vCMTS data plane instances run DPDK-based DOCSIS MAC upstream and downstream data plane processing pipelines for individual cable service groups. On the traffic-generator node, DPDK Pktgen-based traffic tester instances simulate traffic into corresponding vCMTS data plane instances.

During the benchmarking tests, all the vCMTS data plane instances, each representing an individual service group, handle the same traffic rates at the same time. The tests DO NOT consider a scenario where different vCMTS data plane instances handle different traffic rates.

For the purposes of this paper, the Intel vCMTS Reference Dataplane was deployed on a bare-metal Linux stack. However, it can also be deployed in a Kubernetes-orchestrated environment, based on either the Bare Metal Reference Architecture (BMRA) for containers from Intel or on Red Hat OpenShift Container Platform.



Figure 9. Performance test environment used to measure vCMTS performance.

Appendix C: Test Environment Configuration Information and Relevant Variables

CM Lookup and Subscriber Management	300 subscribers per service-group, 4 IP addresses per subscriber
DOCSIS Filtering	22 filter groups, 22 filter rules per group (18 IPv4 rules, 4 IPv6 rules) 1 filter group per cable modem (i.e. same filter group for all CPE types) 10% matched, 90% unmatched (default action – permit)
DOCSIS Classification	16 rules per subscriber 10% matched – enqueue to one of 3 service-flow queues 90% unmatched – enqueue to default service-flow queue
Downstream Service-Flow Scheduling	8 service-flow queues per subscriber (4 active)
Downstream Channel Scheduling	6 x OFDM (1.89 Gbps) channels, 2 x channel-bonding groups NOTE: channel-bonding groups are distributed evenly across cable modems
Upstream Bandwidth Scheduling	Upstream scheduler not used Upstream bandwidth pre-allocated in grants of 2KB per service ID. Bandwidth grants balanced evenly across 300 cable modems
Ethernet CRC	Downstream: 100% CRC re-generation Upstream: 0% CRC verification NOTE: CRC relates to inner frames
Encryption	100% AES, 0% DES
Packet IMIX Distribution	Upstream 65% : 70B, 18% : 256, 17% : 1280B Downstream 15% : 84B, 10% : 256B, 75% : 1280B
Core Configuration	2us1t_1ds2t: 2 single-thread upstream pipelines per core, 1 dual-threaded downstream pipeline per core, stats/management thread on OS cores
Downstream NIC RXQ Size	512

Appendix D: System Configuration

vCMTS Server – based on 4th Gen Intel® Xeon® Scalable Processor

Hardware	
Platform	Intel® Customer Reference Board (Archer City)
CPU	4 th Gen Intel® Xeon® Gold 6428N Processor, 1.8 Ghz, 32 cores Uncore Frequency: 1.6 GHz Microcode: 0x2B000111 NOTE: Single CPU (of dual CPU package) used for all vCMTS performance benchmarks
Memory	16 x 32GB DDR5 4800 MT/s [4000 MT/s]
Hard Drive	1x223.6G Intel® SSD SC2KB240G8
Network Interface Cards	$3 ext{ x Intel}^{\circ}$ Ethernet Network Adapter E810-2CQDA2 (per CPU, 6 x 100 Gbps ports)
Crypto Acceleration	2 x onboard Intel® QuickAssist Technology (per CPU)
Software	
Host OS	Ubuntu 22.04, Linux Kernel v5.15.x
Kernel Options	default_hugepagesz=1G hugepagesz=1G hugepages=128 intel_iommu=on iommu=pt intel_pstate=disable isolcpus=2-31,34-63,66-95,98-127 rcu_nocbs=2-31,34-63,66-95,98-127 nohz_full=2-31,34-63,66-95,98-127 nr_cpus=128 panic=30 nmi_watchdog=0 audit=0 nosoftlockup hpet=disable mce=off tsc=reliable numa_balancing=disable memory_corruption_check=0 workqueue.power_efficient=false module_blacklist=ast modprobe.blacklist=ice,qat_4xxx,intel_qat init_on_alloc=0
Data Plane Development Kit (DPDK)	DPDK v22.07
Intel® Multi-Buffer Crypto for IPSec	intel-ipsec-mb v1.3
Intel [®] QAT Driver/Firmware	QAT20.L.0.9.6-00024
vCMTS	Intel® vCMTS Reference Dataplane v22.10.0

Tested by Intel on March 10, 2023

Traffic-Generator Server	
Hardware	
Platform	Supermicro [®] X12DPG-QT6
CPU	Intel® Xeon® Gold 5320, 2.2 GHz, 26 Cores
Memory	16 x 16GB DDR4 3200 MT/s [2933 MT/s]
Hard Drive	1x223.6GKINGSTON_SA400S37240G
Network Interface Cards	3 x Intel® Ethernet Network Adapter E810-2CQDA2 (per CPU, 6 x 100 Gbps ports)
Software	
Host OS	Ubuntu 22.04, Linux Kernel v5.15.x
Kernel Options	default_hugepagesz=1G hugepagesz=1G hugepages=64 intel_iommu=on iommu=pt isolcpus=2-25,54-77,28-51,80-103 rcu_nocbs=2-25,54-77,28-51,80-103 nohz_full=2-25,54-77,28-51,80-103 nr_cpus=104 panic=30 nmi_watchdog=0 audit=0 nosoftlockup hpet=disable mce=off tsc=reliable numa_balancing=disable memory_corruption_check=0
Data Plane Development Kit (DPDK)	DPDK v20.08
Traffic Generator	DPDK Pktgen v19.10.0

The following vCMTS server configuration was used for the performance comparison.

vCMTS Server – based on 3rd Gen Intel® Xeon® Scalable Processor

Hardware	
Platform	Supermicro® X12DPG-QT6
CPU	3 rd Gen Intel® Xeon® Gold 6338N Processor, 2.2 GHz, 32 Cores Uncore Frequency: 1.6 GHz Microcode: 0xD000375 NOTE: Single CPU (of dual CPU package) used for all vCMTS performance benchmarks
Memory	16 x 32GB DDR4 3200 MT/s [2666 MT/s]
Hard Drive	1x 223.6G Intel® SSD SC2KB240G8
Network Interface Cards	3 x Intel® Ethernet Network Adapter E810-2CQDA2 (per CPU, 4 x 100 Gbps ports)
Software	
Host OS	Ubuntu 22.04, Linux Kernel v5.15.x
Kernel Options	default_hugepagesz=1G hugepagesz=1G hugepages=128 intel_iommu=on iommu=pt intel_pstate=disable isolcpus=2-31,34-63,66-95,98-127 rcu_nocbs=2-31,34-63,66-95,98-127 nohz_full=2-31,34-63,66-95,98-127 nr_cpus=128 panic=30 nmi_watchdog=0 audit=0 nosoftlockup hpet=disable mce=off tsc=reliable numa_balancing=disable memory_corruption_check=0 workqueue.power_efficient=false module_blacklist=ast modprobe.blacklist=ice init_on_alloc=0
Data Plane Development Kit (DPDK)	DPDK v22.07
Intel® Multi-Buffer Crypto for IPSec	intel-ipsec-mb v1.3
vCMTS	Intel® vCMTS Reference Dataplane v22.10.0

Tested by Intel on May 15, 2023

Appendix E: Acronyms and Definitions

Acronym or Term	Definition
ACL	Access Control List
AES	Advanced Encryption Standard
AES-NI	Advanced Encryption Standard New Instructions
API	Application Programming Interface
AVX-512	Advanced Vector Extensions 512
BMRA	Bare Metal Reference Architecture
BPI	Baseline Privacy Interface
СМ	Cable Modem
CMTS	Cable Modem Termination System
CPU	Central Processing Unit
CRC	Cyclic Redundancy Code
DAA	Distributed Access Architecture
DEPI	Downstream External PHY Interface
DES	Data Encryption Standard
DOCSIS	Data over Cable Service Interface Specification
DPDK	Data Plane Development Kit
GHz	Gigahertz
Gbps	Gigabits Per Second
НА	High Availability
HFC	Hybrid Fiber-Coaxial
HCS	Header Check Sequence
HVAC	Heating, Ventilation and Air-Conditioning
IA	Intel Architecture
Intel® DevZone	Intel® Developer Zone
Intel® QAT	Intel® QuickAssist Technology
IMIX	Internet Mix
ISV	Independent Software Vendor
КВ	Kilobytes
L2TP	Layer Two Tunnelling Protocol
MAC	Media Access Control
ME	Micro-Engine
MSO	Multiple System Operator

MTU	Maximum Transmission Unit
NIC	Network Interface Card
OFDM	Orthogonal Frequency-Division Multiplexing
OS	Operating System
PCI	Peripheral Component Interconnect
PCIe	Peripheral Component Interconnect Express
PMD	Poll Mode Driver
PSP	Packet Streaming Protocol
R-PHY	Remote-PHY
SG	Service Group
SR-IOV	Single-Root I/O Virtualization
SW	Software
TDP	Thermal Design Power
UEPI	Upstream External PHY Interface
UGC	User Generated Content
vCore	Virtualized Core
vCMTS	Virtualized Cable Modem Termination System
VF	Virtual Function
VNF	Virtual Network Function
VPN	Virtual Private Network
WRR	Weighted Round Robin

References

- [1] C. Cullen, "Global Internet Phenomena COVID-19 Spotlight: VPN's Rise During Work From Home," [Online]. Available: https://www.sandvine.com/blog/global-internet-phenomena-covid-19-spotlight-vpns-rise-during-work-from-home.
- [2] "The Global Internet Phenomena Report COVID-19 Spotlight," [Online]. Available: https://www.sandvine.com/covid-19-trends.
- [3] "Measuring Fixed Broadband Twelfth Report," [Online]. Available: https://www.fcc.gov/reports-research/reports/ measuring-broadband-america/measuring-fixed-broadband-twelfth-report.
- [4] "H.R.3684 Infrastructure Investment and Jobs Act," [Online]. Available: https://www.congress.gov/bill/117th-congress/ house-bill/3684.
- [5] "Broadband Equity, Access, and Deployment (BEAD) Program," [Online]. Available: https://broadbandusa.ntia.doc.gov/taxonomy/term/158/broadband-equity-access-and-deployment-bead-program.
- [6] "CableLabs," [Online]. Available: https://www.cablelabs.com/.
- [7] T. Muders, R. Elftmann, T. Nguyen and E. Heaton, "Vodafone Network Evolution Paves the Wat for Energy Savings," [Online]. Available: https://www.vodafone.com/news/technology/vodafone-intel-paper-lower-network-energy-billsbetter-environmental-footprint.
- [8] "Data Plane Development Kit (DPDK)," [Online]. Available: https://www.dpdk.org/.
- [9] "Intel® Multi-Buffer Crypto for IPsec Library," [Online]. Available: https://github.com/intel/intel-ipsec-mb.

White Paper | Maximizing vCMTS Data Plane Performance with 4th Gen Intel® Xeon® Scalable Processor Architecture

- [10] "Intel® vCMTS Reference Dataplane," [Online]. Available: https://www.intel.com/content/www/us/en/developer/topic-technology/open/vcmts-reference-dataplane/overview.html.
- [11] "DPDK Power Management Library," [Online]. Available: https://doc.dpdk.org/guides/prog_guide/power_man.html.
- [12] "Intel[®] QuickAssist Technology Documentation Hardware Version 2.0," [Online]. Available: https://intel.github.io/ quickassist/.
- [13] B. Ryan, M. O'Hanlon, D. Coyle, R. Sexton and S. Ravisundar, "Maximizing vCMTS Data Plane Performance with 3rd Gen Intel® Xeon® Scalable Processor Architecture," [Online]. Available: https://networkbuilders.intel.com/solutionslibrary/ maximizing-vcmts-data-plane-performance-with-3rd-gen-intel-xeon-scalable-processor-architecture.

Footnotes

- 1. More details at intel.com/processorclaims: 4th Gen Intel[®] Xeon[®] Scalable processors. Results may vary.
- 2. 60% more L2 cache based on 1.25 MB per core on 3rd Gen Intel® Xeon® Scalable processor and 2MB per core on 4th Gen Intel® Xeon® Scalable processor.
- 3. 50% more CPU cores based on max core-count of 40 cores per CPU on 3rd Gen Intel® Xeon® Scalable processor and max core-count of 60 cores per CPU on 4th Gen Intel® Xeon® Scalable processor.
- 4. 50% more Memory B/W based on 8 memory channels per CPU up to 3200 MT/s on 3rd Gen Intel[®] Xeon[®] Scalable processor and 8 channels up to 4800 MT/s per CPU on 4th Gen Intel[®] Xeon[®] Scalable processor.
- 5. 25% more L3 cache per core based on 1.5MB per core on 3rd Gen Intel[®] Xeon[®] Scalable processor and 1.875MB per core on 4th Gen Intel[®] Xeon[®] Scalable processor.
- 6. 25% more PCI Lanes based on 4 x 16 Lanes per CPU on 3rd Gen Intel® Xeon® Scalable processor and 5 x 16 Lanes per CPU on 4th Gen Intel® Xeon® Scalable processor.
- 7. 2 x PCI B/W based on PCIe Gen4 on 3rd Gen Intel[®] Xeon[®] Scalable processor and PCIe Gen5 on 4th Gen Intel[®] Xeon[®] Scalable processor.
- 8. Up to 4 x Intel® QAT devices on Platinum SKUs of 4th Gen Intel® Xeon® Scalable processor.
- 9. Performance measured using the test environment, scenario and system configuration described in Appendix B, C and D. Results based on a different test environment, scenario or system configuration may differ. Note that there will be a margin of error due to the action of taking performance measurements. Results shown are for a reference implementation of a vCMTS data plane and not a production system. These numbers should be treated strictly as a reference only.
- 10. Hyper-threaded siblings are hardware threads of execution contained within the same physical CPU core and which share the same set of core resources. For data-plane cores on the Intel® vCMTS Reference Dataplane system, each hyper-thread runs its own data plane software thread.

intel

Notices & Disclaimers

Performance varies by use, configuration, and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

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

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.

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.

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 nonexclusive, 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.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0723/BB/DJA/HO9/PDF OPPlease Recycle 355428-001US