

Intel Select Solutions for NFVI v2 with Red Hat OpenStack Platform (OSP) Release 13

Communications service providers (CommSPs) are in the process of virtualizing their service platforms to gain flexibility, agility, and scale. Virtualization can be challenging and add complexity and unpredictability to network services. Intel® Select Solutions for NFVI are consistent building blocks that ease interoperability, gain stability, and speed deployment, thereby addressing the challenges faced when choosing the right infrastructure on which to implement NFV-based services.



Introduction

With the Intel Select Solutions for NFVI, Intel's ecosystem partners can deliver workload-optimized server solutions to communication service provider (CommSP) customers. Network functions virtualization infrastructure (NFVI) is the foundation of virtualized network services and has a significant impact on overall system and service performance. With a thoroughly tested hardware and software reference design, Intel Select Solutions for NFVI can help CommSPs lessen the time, effort, and expense involved with evaluating hardware and software integrations for NFV-based services and help speed deployment.

The Intel Select Solutions for NFVI v2 are based on a clustered multinode configuration. This reference design is defined in collaboration with CommSP and ecosystem partners to expose the value of an I/O balanced architecture based on a data-centric foundation of 2nd generation Intel Xeon Scalable processors, Intel Ethernet Converged Network Adapters, Intel QuickAssist Technology (Intel QAT), Intel Optane™ DC persistent memory, and Intel Solid State Drives.

The Intel Select Solutions for NFVI are developed from Intel's deep experience with industry partners including independent software vendors (ISV), operating system vendors, and original equipment and original design manufacturers (OEMs/ODMs). These partners, or CommSPs, can utilize the Intel Select Solutions for NFVI as a testing platform for determining VNF performance that is repeatable in a broad deployment and scale the configuration from testing and trial to live deployment. Because the Intel Select Solutions for NFVI platform is defined in detail, it's easier to understand and fix performance issues arising from OEM/ODM or CommSP changes to the reference design.

Intel Select Solutions for NFVI v2 contain a cloud node for network function delivery and a controller node for clustered node deployments. They include a "plus" or "base" configuration for the cloud node, and a configuration for the controller node. Together, these configurations—cloud (base or plus) and controller nodes—help realize an efficient balance of function to meet many deployment needs (see Figure 1).

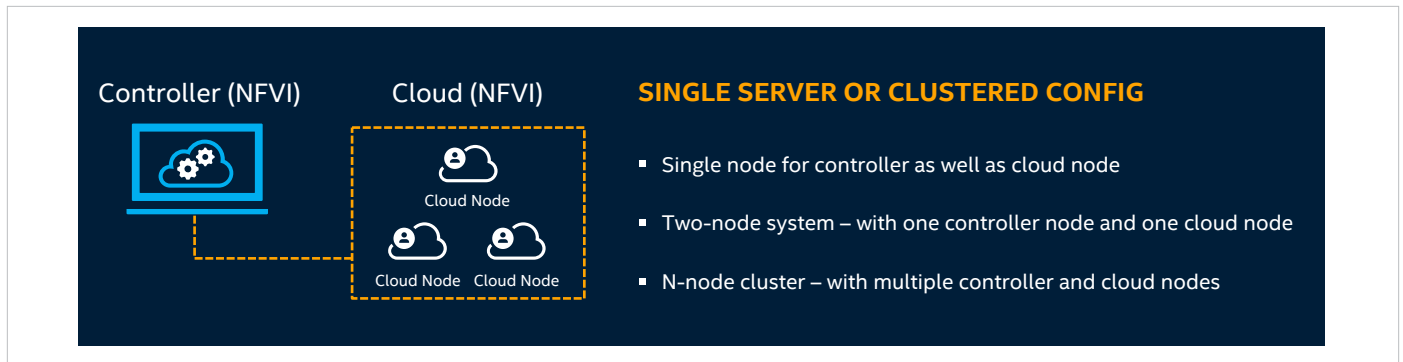


Figure 1. Intel Select Solutions for NFVI v2 clustered multinode architecture.

To fully implement the NFVI clustered multinode solution, servers based on the Intel Select Solutions for NFVI reference designs should include ingredients to implement a software defined infrastructure that resides within each cloud node instance with the controller node used for hosting the control functions in the OpenStack environment.

What Are Intel Select Solutions?

Intel Select Solutions are pre-defined, workload-optimized solutions designed to minimize the challenges of infrastructure evaluation and deployment. Solutions are validated by OEMs/ODMs, certified by ISVs, and verified by Intel. Intel develops these solutions in extensive collaboration with hardware, software, and operating system vendor partners and with the world's leading data center and service providers. Every Intel Select Solution is a tailored combination of Intel data center compute, memory, storage, and network technologies that delivers predictable, trusted, and compelling performance.

To refer to a solution as an Intel Select Solution, a vendor must:

1. Meet the software and hardware stack requirements outlined by the solution's reference-design specifications
2. Replicate or exceed established reference-benchmark test results
3. Publish a solution brief and a detailed implementation guide to facilitate customer deployment

Solution providers can also develop their own optimizations in order to give end customers a simpler, more consistent deployment experience.

Intel Select Solutions for NFVI v2 Configurations

There are three configurations in the Intel Select Solution for NFVI v2 reference designs. There are two cloud node designs and one controller node design. Details on each of the Intel Select Solutions for NFVI v2 cloud and controller nodes are listed below.

- Intel Select Solution for NFVI v2 base configuration cloud node: This configuration is based on 2nd generation Intel Xeon Gold 6230 or Intel Xeon Gold 6230N processors and is a value/performance optimized configuration. It specifies network, storage, and add-in platform acceleration products from Intel for carrier-class use cases, which require longer life and higher reliability than many standard server use cases. The base configuration also specifies Intel QuickAssist Technology (Intel QAT) encryption/decryption and compression acceleration technology, which is required for verification, but optional for deployment.⁴
- Intel Select Solution for NFVI v2 plus configuration cloud node: This configuration is based on high-performance

2nd generation Intel Xeon Gold 6252 processors. Either of these CPUs is coupled with Intel QAT for acceleration and both components are integrated on the motherboard to deliver very high NFVI performance. It is designed to maximize virtual machine density, supporting many simultaneous applications.

- Intel Select Solution for NFVI v2 controller node: This configuration is designed to function as an OpenStack cloud controller to control various cloud nodes and features with a combination of reduced compute power with additional networking functionality. Based on the 2nd generation Intel Xeon Gold 5218 processor, the controller node can also optionally feature local application storage and hardware acceleration components.

The cloud node configurations of Intel Select Solutions for NFVI v2 feature high-performance CPUs, balanced I/O, and on-board acceleration with Intel QAT, combined with optimizations through DPDK, for exceptional throughput and low latency performance. The configurations are designed to optimize performance for virtualized packet-based network traffic workloads, and crypto/compression acceleration.

2nd Generation Intel Xeon Scalable processors:

- Offer high scalability that is cost-efficient and flexible, from the multi-cloud to the intelligent edge
- Establish a seamless performance foundation to help accelerate data's transformative impact
- Support breakthrough Intel Optane DC persistent memory technology
- Accelerate artificial-intelligence (AI) performance and help deliver AI readiness across the data center
- Provide hardware-enhanced platform protection and threat monitoring



Hardware Selections

2nd generation Intel Xeon Gold processors are specified for all models of the Intel Select Solutions for NFVI v2. Intel Xeon Gold processors offer up to 24 cores per socket and frequency up to 2.3 GHz to address a full range of value and performance needs. Intel Xeon Gold processors also feature important platform technologies that are required in Intel Select Solutions for NFVI v2:

- Intel Virtualization Technology (Intel VT) provides hardware abstraction to allow multiple workloads to co-exist and share common resources while maintaining full isolation.
- Intel Boot Guard technology is hardware-based boot integrity protection that can help prevent unauthorized software and malware takeover of boot blocks critical to a system's function.
- Intel Trusted Execution Technology (Intel TXT) is a hardware-based platform security technology that tests the authenticity of critical elements of a platform, operating system, and hypervisor against known good results.

Intel Optane DC persistent memory is designed to improve overall system performance and to lower latencies by using non-volatile media to supplement DRAM and increase the warm memory located close to the processor—reducing time-consuming disk accesses. Intel Optane DC persistent memory modules are available in per module capacities of 128 GB, 256 GB and 512 GB.

Intel Solid State Drive Data Center (Intel SSD DC) Family includes the Intel SSD DC P4500 Series, Intel SSD DC P4501 Series, Intel SSD DC P4600 Series, and the Intel Optane DC SSD P4800X Series. The Intel SSD DC P4500 Series, Intel SSD DC P4501 Series, and Intel SSD DC P4600 Series are based on Intel 3D NAND Technology and designed for cloud infrastructures, offering outstanding quality, reliability, and advanced manageability and serviceability to minimize service disruptions. The Intel Optane DC SSD P4800X Series accelerates applications for fast caching and fast storage to increase scale per server and reduce transaction costs for latency sensitive workloads.

Intel QuickAssist Technology (Intel QAT) is a hardware acceleration technology that provides lookaside cryptographic and compression/decompression co-processing services for the host. Intel QAT is built into certain Intel Xeon Scalable processor platform chipsets and can be added to any server via standard PCI add-in card options.

The 25Gb **Intel Ethernet 700 Series Network Adapters** accelerate the performance of Intel Select Solutions for NFVI. The Intel Ethernet 700 Series delivers validated performance ready to meet high-quality thresholds for data resiliency and service reliability with broad interoperability.¹ All Intel Ethernet products are backed by worldwide pre- and post-sales support and offer a limited lifetime warranty.

Verified Performance Through Benchmark Testing

Intel and its solution definition partners have set minimum performance benchmarks for the Intel Select Solutions for NFVI to ensure that the design and software configuration deliver expected performance. Performance testing is specified for the following areas:

The performance baseline benchmark shows that applications running on Intel Select Solutions for NFVI will offer performance targets for latency, memory bandwidth, and jitter. These metrics will demonstrate that the BIOS, advanced technology configuration, and software stack are configured according to the reference design.

Data Plane Development Kit (DPDK) improves packet processing efficiency and is a core platform technology component of the Intel Select Solutions for NFVI for all configurations. Compliant platforms must implement DPDK software and meet benchmark performance metrics.

Intel QAT is essential for high performance encryption/decryption acceleration for the SSL layer 7 application in public key exchange as well as bulk encryption and decryption applications. This benchmark is designed to demonstrate how Intel QAT can efficiently address the growing needs for compression and critical encryption operations that are a growing portion of network traffic.

KEY KPI TARGETS		BASE CONFIGURATION ²	PLUS CONFIGURATION ³
Cryptographic and Compression Operations Performance with Intel QAT ⁴	Compression (compress and verify) throughput ⁵	24 Gb/s	54 Gb/s
	Encryption throughput ⁶	40 Gb/s	100 Gb/s
	RSA throughput ⁷	100 K Ops/s	100 K Ops/s
OpenSSL Performance (PKE operation with 1 thread)	Software 1 thread	1 K signs/s	1 K signs/s
	Utilizing Intel QAT	100 K signs/s	100 K signs/s
Open SSL Performance (AES128-CBC-HMAC-SHA1 operation with 8 threads with 16 KB packet size)	Utilizing Intel QAT	40 Gb/s	100 Gb/s
Packet Processing Performance using DPDK L3fwd RFC2544 zero packet loss test Physical Device or Physical Function (PF) Pass-Through or Single Root IO Virtualization (SR-IOV)	2 x Dual Port 10 GbE	90% line rate with packet size 256B	N/A
	2 x 25GbE Dual Port Intel Ethernet Controller XXV710	90% line rate with packet size 256 B	90% line rate with packet size 256 B
	2 x 40GbE Dual Port Intel Ethernet Controller XL710	N/A	60% line rate with packet size 256 B ⁸
NGINX Stack with OpenSSL 1.1.0 (PKE operations)	HTTPS CPS throughput (software 16 threads)	12,000 CPS	N/A
	HTTPS CPS throughput (utilizing Intel QAT engine)	75,000 CPS	80,000 CPS
Cyclictest	NFVI system latency (wake-up time of threads running in a VM)	< 10 μs average latency (1 core) < 15 μs average latency (8 cores)	< 10 μs average latency (1 core) < 15 μs average latency (8 cores)
HAProxy SSL Layer 7 Performance	Software (taskset to 1 core with Intel Hyper-Threading Technology)	100+ CPS	100+ CPS
	Intel QAT (taskset to 1 core with Intel Hyper-Threading Technology)	> 300% of the software-only performance	> 300% of the software-only performance
Business Value of the Plus Configuration			2.25 times increase in Intel QAT compression throughput, 2.5 times increase in Intel QAT encryption throughput

Table 1. Select benchmark testing specifications for Intel Select Solutions for NFVI v2.

Board Design Configurations

The Intel Select Solution for NFVI v2 reference designs are anchored on a balanced I/O for network and storage media. All workloads benefit from a deterministic latency and performance for their resources, including NFVI workloads when they are deployed within the same non-uniform memory access (NUMA) node. The table below identifies the high-level hardware elements for the controller node and each of the cloud nodes.

INGREDIENT	CONTROLLER NODE REQUIREMENT	CLOUD NODE BASE REQUIREMENT	CLOUD NODE PLUS REQUIREMENT
PROCESSOR	2 x Intel Xeon Gold 5218 processor (2.1 GHz, 16 cores), or Intel Xeon Gold 5218N processor (2.3 GHz, 16 cores), or Intel Xeon Gold 5218R (2.1 GHz, 20 cores), or higher number SKU	2 x Intel Xeon Gold 6230 processor (2.1 GHz, 20 cores), or Intel Xeon Gold 6230N processor (2.3 GHz, 20 cores), or Intel Xeon Gold 6230R (2.1 GHz, 26 cores), or higher number SKU	2 x Intel Xeon Gold 6252 processor (2.1 GHz, 24 cores), or Intel Xeon Gold 6238R (2.2 GHz, 28 cores), or higher number SKU
MEMORY	Option 1: 192 GB DRAM Option 2: 192 GB DRAM plus 512 GB Intel Optane DC persistent memory ⁹	Option 1: 384 GB DRAM Option 2: 192 GB DRAM plus 512 GB Intel Optane DC persistent memory ⁹	Option 1: 384 GB DRAM Option 2: 192 GB DRAM plus 1.0 or 1.5 TB Intel Optane DC persistent memory ⁹
NIC	2 x 25GbE Dual Port Intel Ethernet Controller XXV710 or 2 x 10GbE Dual Port Intel Ethernet Controller X710, or 2 x 10GbE Dual Port Intel Ethernet Converged Network Adapter X520	2 x 25GbE Dual Port Intel Ethernet Controller XXV710, or 2 x 10GbE Dual Port Intel Ethernet Controller X710, or 2 x 10GbE Dual Port Intel Ethernet Converged Network Adapter X520	2 x 40GbE Dual Port Intel Ethernet Controller XL710, or 2 x 25GbE Dual Port Intel Ethernet Controller XXV710
LAN ON MOTHERBOARD (LOM)	10 Gbps or 25 Gbps port for pre-boot execution environment (PXE) and operation, administration and management (OAM), and 1 Gbps or 10 Gbps port for management NIC	10 Gbps or 25 Gbps port for PXE and OAM, and 1 Gbps or 10 Gbps port for management NIC	10 Gbps or 25 Gbps port for PXE and OAM, and 1 Gbps or 10 Gbps port for management NIC
INTEL QAT	Intel QuickAssist Adapter 8970 (PCIe) add in card (AIC), ⁹ or Equivalent third-party Intel C620 series chipset Intel QAT enabled PCIe AIC with minimum 8 or 16 lanes of PCIe connectivity, ⁹ or Multiple Intel QAT AIC Adapters installed in the platform ⁹	Intel QuickAssist Adapter 8970 (PCIe) AIC, ⁴ or Equivalent third-party Intel C620 series chipset Intel QAT enabled PCIe AIC with minimum 8 or 16 lanes of PCIe connectivity, ⁴ or Multiple Intel QAT AIC Adapters installed in the platform ⁴	Intel C620 series chipset family integrated on base board Intel C627/C628 chipset, integrated symmetrically attached to each CPU or minimum 16 PCIe lane connectivity to one CPU
STORAGE	2 x Intel SATA Solid State Drive @ 480 GB or equivalent boot drive, and 2 x Intel SSD Data Center P4510 Series for NVMe @ 2.0 TB or equivalent drive, symmetrically attached to both CPU sockets ⁹	2 x Intel SATA Solid State Drive @ 480 GB or equivalent boot drive, and 2 x Intel SSD Data Center P4510 Series for NVMe @ 2.0 TB or equivalent drive, symmetrically attached to both CPU sockets	2 x Intel SATA Solid State Drive @ 480 GB or equivalent boot drive, and 4 x Intel SSD Data Center P4510 Series for NVMe @ 2.0 TB or equivalent drive, symmetrically attached to both CPU sockets

Table 2. Example hardware configurations for Intel Select Solutions for NFVI v2.

Software and Firmware Stack

To ensure consistent performance, firmware, operating systems, drivers, hypervisors, BIOS, and other software components for Intel Select Solutions for NFVI as specified in the reference design must be integrated in the solution stack for proper operation, stability, and optimal performance.

Operating systems are the foundation of the software stack and this solution integrates Red Hat Enterprise Linux (RHEL) for its host operating system. The Intel Select Solutions for NFVI v2 must also be certified with Red Hat OpenStack Platform as the deployed open NFVI platform for next-generation networking services and applications.

To meet the optimized deterministic performance requirements for an Intel Select Solution for NFVI platform, the BIOS settings must be set to maximize performance using Intel Turbo Boost Technology (turbo mode).

The software solution stack of the reference design is described in Table 3. Intel continually validates and integrates new versions of this firmware and software to ensure that software updates will perform correctly when integrated into a deployed system. The components of the software stack are therefore subject to change.

	INGREDIENT	SW VERSION DETAILS		
FIRMWARE	BIOS MCU	SE5C620.86B.0D.01.0241 Release Date Nov 19'2018 0x04000010		
	Firmware for Intel Ethernet Converged Network Adapter X520	E68793-005_rev1_0		
	Intel Ethernet Network Adapter X722 Firmware	v3.3 or later		
	Intel Ethernet Controller XXV710 Firmware Intel Ethernet Controller XXV710 Dynamic Device Personalization	v6.02 or later GTPv1 or later		
	Intel SSD Data Center P4510 Series for NVMe Firmware	VDV10131		
	Intel Optane DC SSD Firmware DIMM Firmware	NVMDIMMDriver: v01.00.00.3371 NVMDIMMHii: v01.00.00.3371		
HOST	OS	Red Hat Enterprise Linux	RHEL7.6-kernel-3.10.0-957.el7.x86_64	
	Hypervisor	KVM/QEMU	2.12.0	
	Libvirt	Libvirt	4.5.0	
	Development Kit	DPDK ¹⁰	18.11	
	APPs	CollectD ¹¹	5.8.0	
	Drivers	Intel QAT	1.7-L.4.3.0-00033	
		i40e	2.3.2-k	
		ixgbe	5.1.0-k-rh7.6	
NVMe		1.0		
GUEST	APPs	DPDK	18.11	
	OS	Red Hat Enterprise Linux	7.6	
		CentOS	7.6	
		Ubuntu	18.04 LTS	
	Drivers	i40evf	3.2.2-k	
		ixgbev	4.1.0-k-rh7.6	

Table 3. The software solution stack for the Intel Select Solutions for NFVI v2 with Red Hat OpenStack Platform (OSP) Release 13.¹²

Conclusion

Intel has developed the Intel Select Solutions for NFVI reference designs with feedback and in cooperation with industry leaders that are playing a key role in the CommSP virtualization transformation. The Intel Select Solutions for NFVI integrate balanced networking and optimized

compute and storage options, and they provide acceleration for demanding encryption and compression workloads. In addition, the solution is based on well-known BIOS configurations, current instances of the middleware from host OS providers, and driver combinations that are regularly integrated and performance-tested to assure optimal performance.

Intel Select Solutions for NFVI provide Intel's OEM/ODM partners with a reference design that meets or exceeds the production traffic needs of CommSPs in a cloud network world. In addition, Intel Select Solutions for NFVI provide ISV partners with a reference design up through the virtualization layer that can be used as a baseline for structural validation when establishing performance metrics. And finally for the CommSP, the Intel Select Solutions for NFVI provide a vertical platform that can serve as the foundation for a network cloud design—removing many person years of research

and analysis. Intel and its leading partners in the ecosystem have invested significant time up front to enable CommSPs to be more confident in the performance of their virtualized solutions and to drive a much more rapid deployment schedule.

For more information on the Intel Select Solutions for NFVI, visit the Intel Select Solutions page on the Intel Network Builders website: <http://builders.intel.com/intelselectsolutions>.

Learn More


Intel Select Solutions web page: intel.com/selectsolutions

Intel Xeon processors: intel.com/xeonscalable

Intel Select Solutions are supported by the Intel Builders Program: <https://builders.intel.com/>

Red Hat NFV: <http://www.redhat.com/NFV>

Red Hat OpenStack Platform: <https://www.redhat.com/en/technologies/linux-platforms/openstack-platform>



¹ The Intel Ethernet 700 Series includes extensively tested network adapters, accessories (optics and cables), hardware, and software, in addition to broad operating system support. A full list of the product portfolio's solutions is available at intel.com/ethernet. Hardware and software is thoroughly validated across Intel Xeon Scalable processors and the networking ecosystem. The products are optimized for Intel architecture and a broad operating system ecosystem: Windows, Linux kernel, FreeBSD, Red Hat Enterprise Linux (RHEL), SUSE, Ubuntu, Oracle Solaris, and VMware ESXi. Supported connections and media types for the Intel Ethernet 700 Series are: direct-attach copper and fiber SR/LR (QSFP+, SFP+, SFP28, XLPP1/CR4, 25G-CA/25G-SR/25G-LR), twisted-pair copper (1000BASE-T/10GBASE-T), backplane (XLAUI/XAUI/SFI/KR/KR4/KX/SGMII). Note that Intel is the only vendor offering the QSFP+ media type. The Intel Ethernet 700 Series supported speeds include 10GbE, 25GbE, 40GbE.

² Testing conducted by Intel on Feb. 27, 2019, with the following hardware and software configurations: 1 Node, 2x Intel Xeon Gold 6230 processor; 1x Intel Server Board S2600WFT, Intel C628 series chipset; Total Memory 384 GB, 24 slots/16 GB/2666 MT/s DDR4 RDIMM; Intel Hyper-Threading Technology (Intel HT Technology): Enable; Intel Turbo Boost Technology: Enable; C-State: Enable; Storage (boot): 2X 480 GB Intel SATA SSD; Storage (capacity): 2x 2TB Intel SSD DC P4510 Series PCIe NVMe; Network devices: 2x 25GbE Dual Port Intel Ethernet Controller XXV710; Network speed: 25 GbE; ucode:0x04000010; OS/Software: Red Hat Enterprise Linux 7.6 (Kernel:RHEL7.6-kernel-3.10.0-957.el7.x86_64).

³ Testing conducted by Intel on Feb. 27, 2019, with the following hardware and software configurations: 1 Node, 2x Intel Xeon Gold 6252 processor; 1x Intel Server Board S2600WFT, Intel C628 series chipset; Total Memory 384 GB, 24 slots/16 GB/2666 MT/s DDR4 RDIMM; Intel Hyper-Threading Technology (Intel HT Technology): Enable; Intel Turbo Boost Technology: Enable; C-State: Enable; Storage (boot): 2X 480 GB Intel SATA SSD; Storage (capacity): 4x 2 TB Intel SSD DC P4510 Series PCIe NVMe; Network devices: 2x 25GbE Dual Port Intel Ethernet Controller XXV710; Network speed: 25 GbE; ucode:0x04000010; OS/Software: Red Hat Enterprise Linux 7.6 (Kernel:RHEL7.6-kernel-3.10.0-957.el7.x86_64).

⁴ Intel QuickAssist Technology is required for Intel Select Solutions technical verification. For the base configuration, it is not required to ship with the server for workloads that do not require crypto or compression acceleration.

⁵ Performance to be measured at 8 KB packet size.

⁶ Performance to be measured at 4 KB packet size.

⁷ Performance to be measured at 2 KB packet size.

⁸ The Intel Ethernet Network Adapter XL710 is a 40 GbE controller. The 2 x 40 GbE adapter using this controller is not intended to be a 2 x 40 GbE but a 1 x 40 GbE with an active back-up port. When attempting to use line-rate traffic involving both ports, the internal switch is saturated and the combined bandwidth between the two ports are limited to a total of 50 Gb/s.

⁹ Recommended.

¹⁰ DPDK available as an extra package from Red Hat Enterprise Linux

¹¹ Available as part of Red Hat OpenStack Version 13.

¹² The software versions listed are minimum requirements. It is recommended to use the latest version if available.

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.

Performance results are based on testing as of the date set forth in the Configurations 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. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com/selectsolutions.

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, and other Intel marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Other names and brands may be claimed as the property of others.