Intel® Open Network Platform Server Release 1.5 Release Notes

SDN/NFV Solutions with Intel® Open Network Platform Server

Document Revision 1.1
November 2015

Revision History

Date	Revision	Comments
November 11, 2015	1.1	Added clarifications to HW specification table
September 30, 2015	1.0	Initial Release for version 1.5 of Intel® Open Network Platform Server

Contents

Figures

Tables

1.0 Introduction

This document describes release 1.5 of the Intel® Open Network Platform Server (Intel® ONP Server). The Intel® ONP Server reference architecture provides engineering guidance and ecosystem support to enable widespread adoption of Software Defined Networking (SDN) and Network Functions Virtualization (NFV) solutions across telecommunications, cloud, and enterprise sectors. The deployment of this reference architecture is done mainly using DevStack. DevStack does not make the deployment production-ready, but it does afford developers a good option to experiment with Intel's software and hardware stack.

The Intel® ONP Server defines the integration of hardware and software components, providing a framework to deliver the many benefits of Intel architecture to SDN and NFV. The reference architecture is based on a standard high-volume server (SHVS) and an Intel ONP Server open-source software stack. The software stack itself is built on open-source software created by open-standard communities like Open vSwitch (OVS), Data Plane Development Kit (DPDK), OpenDaylight (ODL), OpenStack, and Kernel-based Virtual Machine (KVM). Intel is working closely with these communities and contributing to the evolution of their open-standard projects.

Intel® ONP Server Release 1.5 focuses on performance optimizations in OVS with DPDK, including changes in the infrastructure providing the forwarding paths—to increase and optimize performance for all OVS use cases. Added capabilities—including VXLAN tunnel configurations, deployments by means of Neutron and Open Virtual Network, and flow characteristics—support OVS performance analysis and optimization to ensure that OVS with DPDK performs well across key usage models. Another release 1.5 improvement—better integration with OpenStack—enhances network management and simplifies deployment. (Benchmarking for the virtual Enterprise-Customer-Premises Equipment [vE-CPE] use case will be available soon in the *Intel® ONP Server Release 1.5 Performance Test Report*.)

The Intel® ONP Server provides a SDN/NFV reference solution that defines three main network elements: networking nodes, ODL controller, and OpenStack manager. An SDN/NFV lab environment comprised of those network elements is the basis for ONP Server software integration and customer use-case validation.

This document for the Intel® ONP Server Reference Architecture highlights the main functionalities contributed by Intel to open-source community projects that are now integrated into Intel® ONP Server Release 1.5 and details major solution limitations that remain.

2.0 System Overview

2.1 Release 1.5 Highlights

Intel ONP Server Release 1.5 introduces new capabilities and some software upgrades from the previous release:

- Integration with the latest open-source software releases:
  - Upgrade to Openstack Kilo 2015.1.1 release
  - Integration with DPDK release 2.0.0
  - Upgrade to a newer OVS version that supports the following features to boost network I/O throughput:
    - Vhost-user vif driver that uses DPDK from the host to the guest operating system
    - Virtual Extensible LAN (VxLAN) with ovs-dpdk to improve scalability
  - Upgrade to the OpenDaylight Lithium SR1 release to support service function chaining (SFC) enhancements
- Updated to support the Real-Time Linux Kernel (v3.18.16-rt13) installation

2.2 Intel ONP Server Release Distribution

Intel ONP Server Release 1.5 is delivered in the form of a reference architecture guide (*Intel® ONP Server Reference Architecture Guide*). Access to the documentation is under: Intel ONP Server Release 1.5 Software.

This guide provides instructions on how to build Intel ONP Server software, set the functionality test environment and perform the tests. It also includes commit ID details for the OpenStack components.

2.3 Reference Architecture Environment

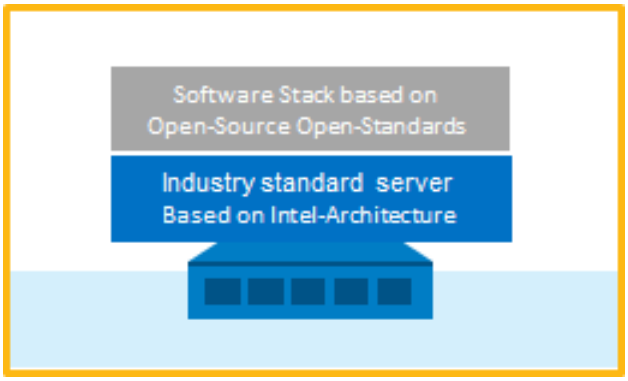

Figure 2–1. Intel ONP Server Node View

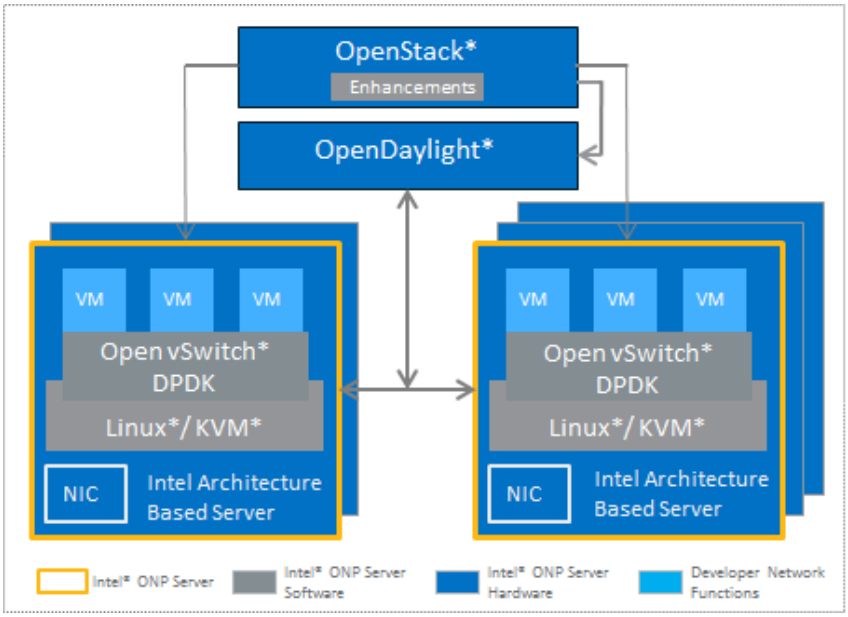

Figure 2–2. Intel ONP Server Test and Validation Environment

2.4 Network Elements

2.4.1 Compute Node Software

Table 2–1 shows the compute node software components used during the testing of VNF components — vIPS, vBNG, NUMA placement, and SR-IOV passthrough for OpenStack.

Table 2–1. Compute Node Software

Software	Notes
Fedora 21 x86_64 (Server Version)	Underlying system-level OS, based on the 3.18.8-201.fc21. x86_64 kernel.
Linux* Real-Time Kernel	Real-Time Linux Kernel 3.18.16-rt13: Provides a baseline to give real-time aspects to the compute nodes.
DPDK 2.0.0	Software libraries are used to dramatically accelerate packet processing, increasing throughput and scalability.
Open vSwitch with: – Kernel datapath: v2.3.2 – DPDK datapath: v2.4.90	Includes support for OVS with DPDK-netdev v2.4.90 commit ID: 1e77bbe565bbf5ae7f4c47f481a4097d666d3d68
libvirt-1.2.13.1.fc21.x86_64	Toolkit and API are used by QEMU-KVM to manage virtual machines. OpenStack (Nova) also uses it to manage the compute resources of the host.
QEMU-KVM version: 2.3.0.5.fc21.x86_64	Open-source machine emulator and virtualizer: Includes KVM that is used to enable hardware accelerations in Intel® platforms.

2.4.2 Controller Node Software

Table 2–2 shows the controller node software components that were used during the testing of VNF components — vIPS, vBNG, and NUMA placement, and SR-IOV passthrough for OpenStack.

Table 2–2. Controller Node Software

Software	Notes
Fedora 21 x86_64 (Server Version)	Underlying system OS: Upgraded to the 3.18.8-201.fc21. x86_ 64 kernel
OpenStack* Kilo 2015.1.1	OpenStack-related tools for building and managing clouds. Includes the DevStack shell script for automating development-environment builds.
Open vSwitch 2.3.2	Includes support for OVS
OpenDaylight Lithium SR1	Use ODL as the OpenStack network management provider through the ML2 plug-in.

2.4.3 Reference Platform Hardware

Haswell-based Platforms

The ONP Server uses the latest and greatest Haswell-based platform. More hardware details are provided in Table 2–3.

Table 2–3. Hardware Ingredients (Code-named Wildcat Pass)

Item	Description	Notes
Platform	Intel® Server Board S2600WTT	Intel® Xeon® processor-based DP server (Formerly Wildcat Pass) 120 GB SSD 2.5in SATA 6GB/s Intel Wolfsville SSDSC2BB120G4. Supports SR-IOV
Processors	Intel® Dual Xeon® Processor E5-2697 V3 2.6 GHz, 35 MB, 145 W, 14 cores	(Formerly Haswell) 14 cores, 2.60 GHz, 145 W, 35 MB total cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133, 28 hyper-threaded cores per CPU for 56 total cores
	Intel® Dual Xeon® Processor Series E5-2699 v3 2.33 GHz, 45 MB, 145 W, 18 cores	(Formerly Haswell) 18 cores, 2.33 GHz, 145 W, 45 MB total cache per processor, 9.6 GT/s QPI, DDR4-1600/1866/2133, 36 hyper-threaded cores per CPU for 72 total cores
Memory	8 GB DDR4 RDIMM Crucial CT8G4RFS423	64 GB RAM (8 x 8 GB)
NICs	Intel® Ethernet Converged Network Adapter X710-DA4	Intel Ethernet Controller XL710-AM1 (Formerly Fortville) 4 x 10 GbE ports Firmware version f4.33 a1.2 n04.42 Tested with Intel® FTLX8571D3BCV-IT and AFBR-703sDZ-IN2 transceivers
	Intel® Ethernet Converged Network Adapter XL710-QDA2	Intel Ethernet Controller XL710-AM2 (Formerly Fortville) 2 x 40 GbE ports Firmware version f4.33 a1.2 n04.42 Tested with Intel® E40QSFPSR transceiver
BIOS	SE5C610.86B.01.01.0008.031920151331 Release Date: 03/19/2015	Intel® Virtualization Technology for Direct I/O (Intel® VT-d) enabled only for SR-IOV PCI passthrough tests, hyper-threading enabled.

3.0 Functionality Highlights and Limitations

3.1 New Functionalities Contributed by Intel to Open-Source Projects in ONP 1.5

Table 3–1 summarizes new software features, bug fixes, and performance improvements that Intel has contributed to open-source projects included in ONP 1.5.

Table 3–1. New Features/Bug Fixes Included in ONP 1.5

No.	Feature Name	Category	Ingredient	Description	IA Value	Commit ID if Upstreamed, Link in Posted as Patch
1.	Vhost-user vif driver	Feature	Openstack Nova	A vif driver to connect VMs to ovs-dpdk via vhost-user ports	An accelerated interface from the host into the guest operating system using DPDK. Enables high throughput network IO into the guest.	https://review.openstack.org/149309 - Libvirt: Support for generic vhostuser vif. https://review.openstack.org/149310 - Libvirt: Support ovs plug in vhostuser vif
2.	ovs-dpdk MD and Agent	Feature	OpenStack/ Networking-ovs-dpdk	A Neutron mechanism driver and agent that supports ovs-dpdk and vhostuser	An accelerated interface from the host into the guest operating system using DPDK. Enables high throughput network IO into the guest.	https://review.openstack.org/143154 - Add Networking-OVS-DPDK project to StackForge https://review.openstack.org/152103 - exporting agent console script entry point https://review.openstack.org/150812 - implements ovs dpdk mech driver https://review.openstack.org/157827 - MD should put whole socket name in vif_details https://review.openstack.org/160746 - tox: Integrate 'pretty_tox.sh' https://review.openstack.org/162088 - Networking OVS-DPDK plugin decomposition https://review.openstack.org/162617 - Added config option for toggling hugepage allocation https://review.openstack.org/166818 - extends ovs init logging

Table 3–1. (cont'd)

No.	Feature Name	Category	Ingredient	Description	IA Value	Commit ID if Upstreamed, Link in Posted as Patch
3.	ovs-dpdk vxlan support	Feature	OpenStack/ Networking-ovs-dpdk	Support VXLAN with ovs-dpdk	Improves scalability associated with large cloud computing deployments.	https://review.openstack.org/175999 – enables ovs native tunneling support
4.	OVSDB Plugin	Feature	ODL	Network virtualization using Openstack, ODL Netvirt and OVS (kernel datapath). The main differences with Lithium are the use of the new refactored OVSDB southbound plugin, better DVR capabilities, more robust and stable, and better parity with Neutron features.	Improves orchestration of IA servers running OVS when controlled by ODL as the SDN controller.	Lithium release
5.	SFC	Feature	ODL	SFC and Group-Based Policy Integration Phase 1. This level of integration allows for a service chain to be specified as an 'action' of a group-based policy.	Demonstrates the integration of two ODL projects that provides more options for orchestration and control of IA servers running OVS with SFC capabilities.	Lithium release
6.	SFC	Feature	ODL	Enhancements for SFC. Implemented APIs for monitoring service function information and capabilities via NETCONF. This allows, e.g., monitoring of a service function's CPU use to support the load-balancing service function selection algorithm.	Enhances orchestration and control options for SFC.	Lithium release
7.	SFC	Feature	ODL	Adds service function selection algorithms to ODL (e.g., random, round-robin, load balancing, and shortest path). This provides different methods for selecting a service function instance from a pool of available instances when rendering a service path.	Enhances orchestration and control options for SFC.	Lithium release

Table 3–1. (cont'd)

No.	Feature Name	Category	Ingredient	Description	IA Value	Commit ID if Upstreamed, Link in Posted as Patch
8.	SFC	Feature	ODL	Allows selection of service function selection algorithm in the ODL Web GUI	Enhances orchestration and control of SFC	Lithium release
9.	SFC	Feature	ODL	Implemented unit and integration tests for the SFC project.	Improves the reliability and stability of the SFC functionality in the ODL controller.	Lithium release
10.	Various perf improve-ments	Perfor-mance Improve-ment	vSwitch	• 115f248163e1651ef932e1210cc4e49e3ebf61c5 miniflow: Fix miniflow push of L4 port numbers. • 9154f798ef0011ea9d1d7fb1dc91b51b60da82d3 netdev-dpdk: Use default NIC configuration. • fc82e877efc03400e65b44588fb40eb507a98bf4 dpif-netdev: Increase the number of EMC entries • 95e9881f843896751a76481cfe7869e2c0c1270b netdev-dpdk: Add vhost enqueue retries. • 4345e1b5bf563ebfd7a7dcf489eac0fdf68135cf netdev-dpdk: Change phy rx burst size. Various (upstreamed): • 7dd671f08e03ba2d133389aa25680baffd2ce0d5 dpif-netdev: log port/core affinity Various (not upstreamed): • dpif-netdev: Make EMC Size Configurable (Rejected by OVS community)	Improve the performance of OVS on Intel Archi-tecture	N/A

Table 3–1. (cont'd)

No.	Feature Name	Category	Ingredient	Description	IA Value	Commit ID if Upstreamed, Link in Posted as Patch
11.	Various Bugfixes	Bug Fix	vSwitch	Various (ONP 1.5): • 30f4d875f53c2e3817628faba20112eaea3e7715 Documentation: DPDK IVSHMEM VM Communications: • 3bcc10c0701c241ef62bdb32c5d21c060ad7590b dpif-netdev: Fix non-pmd thread queue id. • 1b99bb055218e56603cff764df6dd2f1d166a48d netdev-dpdk: Reset RSS hash on transmit 3088fab7cb5c af2308dfda3f2d713c87ac4f35df INSTALL.DPDK.md: Update ivshmem page size restrictions: • dbde55e7fa21881af18a48502c91168be269482a INSTALL.DPDK.md: Fix whitespace: • 418d2485aa029015389768bbc2f66a50c6281880 travis: fix errors • 618f44f7a406d6c3e90110420a7fd183d40f1bff netdev-dpdk: Put cuse thread into quiescent state. • c876a4bb9bcc881befac5647a1e311b5af09d549 netdev: Fix user space tunneling for set_tunnel action. • 58be9c9fd732b5bdd3d4c2e9b8cc2313f570094d auto-make.mk: Improve schema checksum error message. Various (upstreamed): • 84072381c60d112c49ecbb634898069d682e23cb docs: Fix alignment for diagram in native-tunneling.md. Various (not upstreamed): • netdev-dpdk: Silence DPDK Wbad-function-cast warning (Rejected by OVS community) • Reported by Intel or discovered in ONP as part of integration. • dc6ba5dc8b996210f85d4bc6884af469a5244b96 netdev-dpdk: Do not flush tx queue which is shared among CPUs because it is always flushed.	Enhance the usability of OVS with DPDK by removing potentially harmful bugs	N/A

Table 3–1. (cont'd)

No.	Feature Name	Category	Ingredient	Description	IA Value	Commit ID if Upstreamed, Link in Posted as Patch
12.	Enable Missing Statistics	Feature	vSwitch	Some standard OVS statistics are missing from DPDK ports. This patch enables them.	Improve the usability of OVS with DPDK by ensuring all relevant statistics are available to the user.	http://openvswitch.org/pipermail/dev/2015-July/057013.html
13.	DPDK 2.0 support	Feature	vSwitch	This patch ports the OVS datapath to use DPDK 2.0	Update OVS to use the latest version of DPDK. This allows users to take advantage of the latest DDPK features and performance improvements.	543342a41cbceffaac30ace2c66b6e489eb359c8
14.	Bonding	Feature	vSwitch	Unlike system interfaces, DPDK-enabled interfaces must have their interface type explicitly set when used to create bonded ports. This patch updates documentation to indicate how to create a bonded DPDK port.	Allows bonding of DPDK ports using standard OVS interfaces.	77c180cea6f646229c88e47be632ca7c29248c0a
15.	VFIO	Feature	vSwitch	Since DPDK 1.7, VFIO is supported in place of UIO. This allows a user to avoid having to insert a non-standard kernel module.	Allows a user to use DPDK ports without having to insert a non-standard Linux kernel module.	491c2ea3230f53ecbe65a556a0a1cc68647d7b99

3.2 Release 1.5 Limitations

Table 3–2 summarizes the limitations for release 1.5.

Table 3–2. Release 1.5 Limitations

No.	Open Software	Limitation
1.	OpenStack	After the stack compilation is successful on the controller and compute nodes, the user logs in to the OpenStack UI to spawn the VM and an error displays: `Failed to connect to server (code 1006).` The workaround is to flush the iptables on the compute node with the following command as stack user: `sudo iptables -F.`
2.	ODL	On the ODL controller node, the native Java version from the yum repository of Fedora 21 is v1.8. ODL Lithium, the current released version used in this release, however, only supports Java v1.7. It is necessary, therefore, to manually install Java 1.7 on Fedora 21 system and make it the default Java version. Refer to *Intel® Open Network Platform Server Reference Architecture (Release 1.5)*, section 6.2.1, for details of the installation procedure.
3.	ODL	On the ODL controller node a bug exists in the */opt/stack/networking-odl/devstack/plugin.py* file that tries to install openjava-1.7 from the Fedora 21 repository, but it does not exist. Refer to *Intel® Open Network Platform Server Reference Architecture (Release 1.5)*, section 6.2.2, for details of the workaround.
4.	ODL	A bug exists in the */opt/stack/networking-odl/setup.py* file that it requires Python module pbr>=1.3 for stable/kilo branch, whereas other OpenStack components are using pbr=0.11.0. This conflict causes a stacking failure. The workaround is to change "pbr>=1.3" to "pbr" in */opt/stack/networking-odl/setup.py.*
5.	ODL	On the ODL controller, after it is up and running, ODL does not add the manager IP address to the OpenStack bridges br-int and br-ex. A workaround is to add this address to the bridges. Refer to *Intel® Open Network Platform Server Reference Architecture (Release 1.5)*, section 6.2.3, for details of the workaround.
6.	ODL	On the ODL controller, the VXLAN port of the remote peer is not created in bridge br-int. Refer to *Intel® Open Network Platform Server Reference Architecture (Release 1.5)*, section 6.2.3, for details of the workaround.
7.	ODL	The default OpenStack tenant network (private network) does not pass traffic. Users are advised to create their own tenant network in order to create a fully functional instance (VM). Refer to *Intel® Open Network Platform Server Reference Architecture (Release 1.5)*, section 6.1.1.2, for how to create a tenant network.

8.	OpenStack	A fresh installation on the controller running `prepare_stack.sh` may fail with an error message pointing to `Neutron did not start`. This issue was discovered recently and so far has not been root-caused. A workaround is available by performing a restack, running the following commands as a stack user:
		`cd /home/stack/devstack`
		`./unstack.sh`
		`./stack.sh`

9.	OpenStack	The recent release of Python module sqlalchemy-migrate (v0.10.0) picked by DevStack has a dependency of pbr v1.6, which is in conflict with pbr v0.11.0 used by other OpenStack components (e.g., Keystone, Nova, Neutron, etc). This conflict causes a stacking failure. The workaround is to use an older version of sqlalchemy-migrate (to 0.9.7) by changing the requirement in */opt/stack/requirements/global-requirements.txt* and */opt/stack/keystone/requirements.txt*, and then restack. See no. 8 above for how to perform a restack.

Table 3–2. (cont'd)

No.	Open Software	Limitation

10.	OpenStack	A fresh installation on the controller and compute running `prepare_stack.sh` may fail with an error due to a version conflict related to the Python module Babel. DevStack now downloads Babel verison 2.0, while the requirements are for Babel 1.3. The workaround is to uninstall Babel, reinstall the compatible version, and then restack:
		`sudo pip uninstall Babel`
		`sudo pip install Babel==1.3`
		See no. 8 above for how to perform a restack.

11.	OpenStack	The installation on the controller and compute nodes may fail with an error due to a version conflict related to Python module oslo.utils. DevStack now downloads oslo-utils version 1.4.1, which breaks the version dependency of the other Python module(s), e.g., oslo.log. The workaround is to uninstall oslo.utils, reinstall the compatible version, and then restack:
		`sudo pip uninstall oslo.utils`
		`sudo pip install oslo.utils==1.4.0`
		See no. 8 above for how to perform a restack.
		If the physical ports and IPs are not configured to the bridges that are created by DevStack, follow the example given in step 9 of section 5.2.2.2 in the Intel® Open Network Platform Server Reference Architecture (Release 1.5).

4.0 Node Software Installation Instructions

This is a high-level description of the installation process. Read the README file described below for more details:

Instructions:

1. Download the Intel ONP Server, Release 1.5 Scripts tarball from:

   https://01.org/packet-processing/intel®-onp-servers

   The bundle contains the files shown in the table below.

Files	Description	Notes
README	A script to do the initial system configuration and pull the right updates from the Fedora repos	—
prepare_system.sh	A script to do the initial system configuration and pull the right updates from the Fedora repos	—
prepare_stack.sh	A wrapper to DevStack scripts	—
onps_config	The main configuration file with the necessary setup parameters	—
onps_config_stand_alone_compute	A version of onps_config that sets up a standalone compute node without any controllers involved	—
local.conf-controller	A DevStack template configuration file	Do not modify.
local.conf-compute	A DevStack template configuration file	Do not modify.
onps_commit_ids	Commit IDs or tags for the OpenStack components	Do not modify.
settings.xml	A template file for the settings of the Maven configuration used by ODL	Do not modify.

2. Plan ahead to decide what hosts you are going to use and what interfaces of your hosts will belong to management and/or the data plane network.
3. Install a fresh Fedora 21 installation on each host, and install the tarball in step 1 on each host.
4. Untar the tarball and confirm the files described above are present.
5. Manually edit the onps_config configuration file on each host based on the planned role of that host (controller/node), type of interfaces, type of vSwitch desired, etc.
6. Execute `prepare_system.sh` and reboot. The script will parse `onps_config` and prepare the system accordingly.
7. Run `yum update` to update the system, and then reboot.
8. Execute `prepare_stack.sh`. The script will kick off the DevStack installation.

After successfully executiing this procedure in all the hosts, a fully deployed set of compute and controller nodes should be available with the version information detailed in section 2.4.

Note: These sets of scripts and configuration files are for evaluation and facilitation only.

Legal Information

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

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

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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.

The products described in this document 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. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase. For more complete information about performance and benchmark results, visit http://www.intel.com/performance.

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

Intel does not control or audit third-party web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.