

Creating a Better NFV Platform: Dell, Red Hat, and Intel Foster Interoperability

Intel® Open Network Platform

“Dell, Red Hat, and Intel have worked together to deliver a truly open commercial offering. The Dell NFV Platform uses open standards, open protocols, and open source—backed by an open ecosystem—to deliver a commercial NFV platform offering today.”

- Drew Schulke, Executive Director,
Next Generation Infrastructure, Dell

The burgeoning industry acceptance of software-defined networking (SDN) and Network Functions Virtualization (NFV) has led to numerous pilot projects and the development of demonstration platforms. As these development projects move forward toward in-network deployment, communications service providers modernizing their infrastructures have gravitated toward virtualized compute platforms based on Intel® architecture. By being able to deploy NFV solutions on standard high-volume servers (SHVS), communications service providers can reduce service-delivery costs, improve agility, and enhance flexibility across the infrastructure.

Even as the momentum of NFV and SDN adoption accelerates, a key challenge remains. Integrating components from numerous open-source projects and conforming to current open standards programs adds complexity to planning and deployment. The Intel® Open Network Platform (Intel® ONP) reference architecture minimizes the confusion and lowers the risks by providing a blueprint for evaluating, designing, and deploying interoperable, open SDN and NFV solutions.

Working collaboratively, Dell, Red Hat, and Intel have developed a commercially viable NFV platform, described in this white paper, configurable to a range of CSP requirements. The Dell* NFV Platform is an open networking platform—built using industry standard Dell servers, storage, networking, and software—engineered for optimal interoperability and manageability through a broad set of open interfaces. As a part of this offering, Red Hat and Dell have worked together closely to build a fully integrated OpenStack* platform for deploying enterprise-grade private clouds.

Based on the experience gained during the development process, this paper explains the process by which NFV developers can effectively construct a commercially viable European Telecommunications Standards Institute (ETSI) NFV framework through open-source community projects and open-source ingredients.

Table of Contents

Dell and Red Hat Build a Commercial NFV Infrastructure 2

 The Dell NFV Platform 2

 OpenStack Cloud-Computing Platform 3

 Collaborative Work: Dell, Red Hat, and Intel 3

 Dell Server Setup and Intel® ONP 3

 The Dell NFV Platform for Telecommunications Environments 4

Test Setup and Results 4

 DATS Performance Testing Overview 5

 Test Descriptions 6

 Packet Forwarding without Modifications (L2 Switching) 6

 Packet Forwarding with Packet Content Modifications (L3 Routing) 7

 Adding and Removing MPLS Tags 7

 5-Tuple-Based Lookup and Decision 7

 ACL 7

 Buffering 7

 Test Results 7

ETSI NFV Framework and the Origins of Intel ONP 7

What is the Intel Open Network Platform Reference Architecture? 8

Reducing the Complexity of NFV Integration 8

Summary 9

Dell and Red Hat Build a Commercial NFV Infrastructure

With a long-standing partnership that has spanned more than 15 years, Dell and Red Hat are well equipped to co-engineer an NFV platform capable of supporting commercial deployments for telecommunications service providers. Both companies see a strong opportunity in the carrier space and mutual benefits by combining their expertise and technology to deliver value in this market. Using the Intel ONP reference architecture as a starting point, Dell and Red Hat have constructed a highly-scalable NFV platform based on high-performing servers, commercially available components, and Red Hat Enterprise Linux* OpenStack Platform. The testing that is presented in this paper reflects the work in progress at this publication date and the test results with the configured components listed.

The Dell NFV Platform

The Dell NFV Platform features the latest technologies from Dell combined with software from open-source ecosystem partners to form a fully converged, virtualized infrastructure designed to execute a wide variety of virtual network functions (VNFs), as shown in Figure 1. The platform also includes foundational software and open interfaces for management and orchestration (MANO), simplifying operation and enabling easier integration. From a deployment perspective, the Dell NFV Platform can be dimensioned and equipped for applications in a carrier environment, anywhere at any scale.

“Dell acknowledges that an open ecosystem is the new standard. Legacy systems used by service providers are evolving into virtual, cloud-based network infrastructures. It is not happening overnight. There is an enormous amount of validation and performance testing work that needs to be done to make sure that these highly virtualized architectures simply just work.”

- Joe Barzycki, Senior Solutions Architect, Dell




Figure 1. Components of the Dell NFV Platform.

Benefits of the Dell NFV Platform include:

- An open development environment broadly available to the NFV developer community
- The capability to scale up, down, or out to address a wide range of deployment scenarios and conditions
- Support for a broad selection of open-source ingredients and third-party commercial components for flexibility and choice

Using a common infrastructure, the Dell NFV Platform supports a diverse range of workloads with widely varying service-level agreement (SLA) requirements. Through flexible selections of MANO and VNF options and other validated components, solutions based on the Dell NFV Platform can be architected to meet the performance, reliability, and workload demands of telecommunications, cloud service, and Internet of Things (IoT) environments.

Together, Dell and Red Hat have co-engineered and delivered hardened, commercial NFV solutions for the telecommunications market segment, based on RHEL OpenStack. These solutions provide a flexible, open-standards approach to virtualize and transform telecommunications infrastructures.

OpenStack Cloud-Computing Platform

As NFV gains momentum throughout the telecommunications industry, merging the latest SDN technology with cost-effective virtualized compute platforms based on Intel architecture hardware, Dell and Red Hat have strengthened the commercial offerings in this sector, focusing on carrier-grade implementations. Enhancements to Red Hat Enterprise Linux and Red Hat Enterprise Linux OpenStack Platform give telecommunications companies a proven path to make the transition

from proprietary and UNIX* platforms. OpenStack lets users create VMs on available hosts and connect networks in accordance with whatever plan IT has in effect. From within OpenStack, the administrator can configure all of the vendor's virtual application resources.

OpenStack has advanced to enterprise-ready capabilities over a succession of releases. Gartner recently made the statement, "Juno is OpenStack 1.0 for the enterprise." The collaborative work of Red Hat and Dell is directed toward ensuring that the RHEL OpenStack implementation for the Dell NFV Platform effectively meets carrier-grade requirements and can accommodate components from a wide variety of commercial third-party vendors.

Collaborative Work: Dell, Red Hat, and Intel

The three-way engagement includes joint engineering work to ensure that testing and optimization of components on the NFV platform can be fully evaluated. For NFV to be successfully deployed, the performance gap between a virtualized network function and a physical network function must be narrowed. Each component and virtual appliance must be validated and tested to ensure that the entire architectural framework operates successfully. Dell and Red Hat have consolidated the findings of their engineering work together in the [Dell Red Hat OpenStack Cloud Solution Reference Architecture Guide – Version 4.0](#) and continue to build on these reference architecture recommendations for the Dell NFV Platform.

The testing and validation work is being performed at Intel while Dell labs assist other partners looking to accelerate the deployment of virtualized network functions. The goal is to assure service

providers that these systems are not only open, but also carrier-grade, resilient, and highly scalable. Another intent is to create a platform that will be serviceable and adaptable to future requirements.

Collaborative work also extends to the standards bodies and consortiums that are shaping the NFVI ecosystem. Dell, Red Hat, and Intel are active contributors and participants in the work of OPNFV, ETSI, and numerous open-source community projects. With a rich partner ecosystem to draw from—as well as contribute to—the work on establishing a solid commercial NFV infrastructure (NFVI) continues to advance at a rapid pace.

Dell Server Setup and Intel® ONP

A number of performance advantages are conferred by key Intel ONP initiatives. Among them, these three key components are driving performance in the new systems as the architectures get built.

- **High-performance processors.** The Intel® Xeon® processor E5 v3 is the latest generation of processors designed for industry standard high volume servers.
- **High-performance Ethernet controllers.** Network adapters based on the Intel® 82599EB 10 Gigabit Ethernet Controller and the Intel® Ethernet Controller 40 Gigabit XL710, are critical to scale up the performance for higher NFV throughput.
- **Combination of powerful platform and a tuned software stack.** Optimizations available through the combination of reference architecture processors, network adapters, and the Intel ONP software stack ingredients help deliver a tuned, efficient platform suitable for scalable deployments.

Optimizing NFVI requires tying the software closely to the hardware. At the infrastructure level, foundational open community software, particularly the data plane development kit (DPDK), improves the performance of virtualized network functions. Dell, Red Hat, Intel, and a wide range of NFV partners are striving to give service providers assurance that they will get sufficient network function performance in a virtualized environment. For multi-vendor NFV solutions, service providers expect proven interoperability. Dell's collaborative work with Intel on DPDK enhancements represents a step forward on both performance and interoperability. And, improved NFV platform interoperability is delivered by the co-engineering work with Red Hat.

Dell NFV continues to contribute towards the adoption of foundational enhancement technologies, such as DPDK, to improve NFV solutions. In product lab environments across the globe, Dell is proving the effectiveness of DPDK-enhanced VNFs running on its platforms. Dell NFV partners experience a significant performance improvement when switching to DPDK-enhanced VNFs over standard virtual switch technology.

Opportunities for additional optimization exist; the test environment that is described in this paper provides an ideal framework for experimenting with components and performing trials.

The Dell NFV Platform for Telecommunications Environments

The Dell NFV Platform components are designed to meet the tolerances demanded by cloud service providers and communications service providers. The Dell NFV Platform is capable of being deployed in harsh environments with suboptimal conditions that exceed the capabilities of standard or generic enterprise servers, storage devices,

and network devices. Dell pioneered fresh-air cooling on its servers, and these servers can operate efficiently in non-data center conditions. Airflow and fan systems are designed with the same precision and expertise as is applied to building aircraft engines. Embedded monitoring tools throughout each hardware component of the NFV system relay the system status on a regular basis. Dell enhances the value of the platform components with guarantees and support typically not available for generic servers. Configured servers, storage devices, and networking devices in the NFV environment will consume less power, provide more density and higher availability, and will operate under extreme conditions—innovations not generally associated with a generic “white box” server.

Test Setup and Results

Turning the flexible, commercial Dell NFV Platform into a real evaluation and development target is important for developers of SDN and NFV. The Intel and Dell engineering teams collaborated on building an NFVI environment from commercial off-the-shelf (COTS) components: the Dell PowerEdge* R730 and Red Hat

Enterprise Linux Server. By creating this NFVI environment, a system was set up to test packet processing throughput using well-known routing techniques.

This section describes the setup and system configuration that was used throughout the test sequences. Testing took advantage of the Dataplane Automated Testing System (DATS), a flexible DPDK application, which consists of a set of Python scripts. PROX can both generate traffic and serve as a simulator for various workloads—from simple packet forwarding to complex BNG + QoS tests. DATS source code is available through 01.org.

Two separate environments were used for testing:

- **DATS test environment.** Tests from DATS v002 were run on one System Under Test (SUT) configuration, varying the SUT workload.
- **Intel ONP test environment.** Intel ONP runs the L3/L2 forwarding workload, while varying the SUT configuration.

Figure 2 illustrates the key differences between the Intel ONP and DATS test environments.




Figure 2. Comparison of Intel® Open Network Platform and DATS testing.

The following pages provide excerpts from the test results. An effort was made to perform testing that replicated conditions likely to be encountered in real-world deployments—to more closely capture the kinds of data results associated with actual commercial operations. Table 1 lists the hardware and software components used in the DATS test environment. Table 2 shows the virtualized SUT Components for the Intel ONP Environment.

Table 1. System Under Test (SUT) Hardware and Software Components for DATS Test Cases



SUT HARDWARE COMPONENTS	
Platform	Dell PowerEdge* R730
Processor	Intel® Xeon® processor E5-2680 v3 (30M cache, 2.50 GHz)
Number of cores	24
RAM	192 GB
Data Plane Development Kit (DPDK) ports	4x Intel® Ethernet Controller X710 for 10GbE SFP+
SUT SOFTWARE COMPONENTS	
BIOS version	1.2.10
BIOS release date	03/09/2015
OS	Red Hat Enterprise Linux* Server 7.1
Kernel	3.10.0-229el7.x86_64
PROX version	V0.21
DPDK version	V2.1.0
Hugepages – 1 GB	16

Table 2. Virtualized System Under Test (SUT) Components for the Intel® ONP Environment

VIRTUAL MACHINE SUT HARDWARE COMPONENTS	
Platform	Standard PC (i440FX + PIIX, 1996)
Processor	Intel® Xeon® processor E5-2680 v3 (30M cache, 2.50 GHz)
Number of cores	8
RAM	8 GB
Data Plane Development Kit (DPDK) ports	4x Virtio network devices
VIRTUAL MACHINE SUT SOFTWARE COMPONENTS	
BIOS version	rel-1.8.1-0-g4adadb-20150316_085822-nilsson.home.kraxel.org
BIOS release date	04/01/2014
OS	Red Hat Enterprise Linux* Server 7.1
Kernel	3.10.0-229el7.x86_64
PROX version	V0.21
DPDK version	V2.1.0
Hugepages – 1 GB	5

DATS Performance Testing Overview

The DATS performance testing includes these individual tests:

- L2 forwarding without modifying packets
- L3 routing with packet modifications
- Add/Remove MPLS tags
- 5-Tuple-based lookup and decision
- Access Control List (ACL)
- Buffering packets in memory

Each test progressively stresses the packet processing on the server platform, with more complex algorithms being added on each new test. The high-level sequence is as follows:

- 1) Tester runs PROX to generate traffic.
- 2) PROX is also run to measure statistics.
- 3) The SUT runs PROX to simulate a workload.
- 4) DATS carries out these operations:
 - a. Performs test cases.
 - b. Starts and stops the SUT workload.
 - c. Starts and stops traffic generation.
 - d. Collects statistics.
 - e. Generates reports.

Figure 3 shows the progression of testing.




Figure 3. DATS test progression.

Table 3. Executed Tests

TEST NAME	KEY PERFORMANCE INDICATOR (KPI)	KPI LOAD
Port forwarding without touching packets	16.74 Mpps	64-byte packets, 4 streams
Port forwarding with touching packets	17.08 Mpps	64-byte packets, 4 streams
Adding and removing MPLS tags	14.07 Mpps	68-byte packets, 4 streams
5-Tuple-based lookup and decision	8.37 Mpps	64-byte packets, 4 streams
ACL (access control list)	11.38 Mpps	64-byte packets, 4 streams
Buffering	5.73 Mpps	64-byte packets, 1 stream, 125 ms buffer

Tester and SUT machines were connected back-to-back by four 10Gb links. The management network was set up using 1Gb network to avoid interfering with the tests.

DATS and PROX packages can be obtained from <https://01.org/intel-data-plane-performance-demonstrators>

- PROX: <https://01.org/sites/default/files/downloads/intel-data-plane-performance-demonstrators/dppd-prox-v021.zip>
- DATS: <https://01.org/sites/default/files/downloads/intel-data-plane-performance-demonstrators/dppd-dats-v002.tar.gz>

Test Descriptions

Table 3 shows the tests that were executed, which are explained in more detail in the pages that follow. The L3 routing throughput values of the key performance indicators (KPI) are shown in million packets per second (Mpps). The tolerated packet loss for the tests was 0.001 percent.

Packet Forwarding without Modifications (L2 Switching)

The workload takes packets from one port and forwards them unmodified to another port. This use case is not representing any actual workload, but essentially provides a baseline and a check of the test environment. The KPI represents the number of packets per second for 64-byte packets with an accepted minimal packet loss.

Packet Forwarding with Packet Content Modifications (L3 Routing)

The workload takes packets from one port, updates source and destination MACs, and forwards them to another port.

Adding and Removing MPLS Tags

The workload takes packets from one port, adds an MPLS tag, and forwards them to another port. While forwarding packets in the other direction, MPLS tags are removed. The KPI is the number of packets per second for 68-byte packets with an accepted minimal packet loss.

5-Tuple-Based Lookup and Decision

A lookup based on a sufficiently long key extracted from the packet can be used to formulate different resource-demanding use cases, such as:

- Load distribution
- Symmetric load distribution
- Routing
- Policing

ACL

The test measures how well the SUT can exploit structures in the list of ACL rules. The ACL rules are matched against a 7-tuple of the input packet: the regular 5-tuple and two VLAN tags.

The rules in the rule set allow the packet to be forwarded, and the rule set contains a default match all. The rule set has a moderate number of rules with moderate similarity between the rules and the fraction of rules that were used.

Buffering

This test measures the impact of packets being buffered (staying in memory for an extend period of time). For buffering, 125 milliseconds of packets at 10Gb link speed requires 4 GB of RAM available for the application on the SUT. Packets are forwarded without being modified. The test runs only on the first port of the SUT.

Test Results

Although packet sizes from 64-byte to 1518-byte sizes were measured, for simplicity only three different packet sizes are presented—64, 256, 1518—to show a range of workloads situations on a network, shown in Figure 4. The smallest packets represent the worst-case scenario for both the data plane and the workload.

The longest packets are often used as the optimal test conditions. For the largest packets, the system’s throughput was nearly at 100 percent of the line rate for most workloads.




Figure 4. Test results for different network workloads.

ETSI NFV Framework and the Origins of Intel ONP

The Dell NFV Platform is based on two industry initiatives: the ETSI NFV architectural framework and Intel ONP, creating a path to commercialization as shown in Figure 5. In 2013, the ETSI through an Industry Specification Group (ISG) created an architectural framework for the industry to implement NFV technology. The NFV Architectural Framework identifies functional blocks and the main interfaces between the functions. Over the last two years many ETSI NFV ISG members have contributed proof-of-concept (PoC) projects, API proposals, and code to open-source community projects. Many of these PoCs have validated and refined the framework, so that the entire industry can develop and deploy commercial solutions.




Figure 5. Progression toward commercialization of Network Functions Virtualization.

The ETSI NFV Framework served as inspiration for Intel ONP. More information about the relationship between the ETSI NFV Architectural Framework and Intel ONP can be found in these documents:

- “Intel ONP Server Reference Architecture, ETSI NFV Framework”¹
- “NFV Architectural Framework; ETSI GS NFV 002 v1.1.1”²

The ETSI NFV Architectural Framework has helped accelerate the design and development of commercial NFV solutions, establishing a technical basis for supporting virtual network operations.

What is the Intel ONP Reference Architecture?

Developed to streamline the design, deployment, and manageability of open solutions within SDN and NFV environments, the Intel ONP reference architecture provides a blueprint for enabling commercial adoption of validated hardware and software in key industry sectors. Telecommunications carrier networks, enterprise networks, and cloud data centers can more easily build solutions using an open-source

software stack running on standard high-volume servers, as prescribed by the ONP reference architecture. The ecosystem that has grown around the tenets of the Intel ONP reference architecture is supported by collaboration, contributions, and industry participation. Participants include industry consortiums, telecommunications and cloud providers, Intel® Network Builders, and leading companies involved in open-source projects.

Despite the many technology advances in virtualized networking that improve interoperability and simplify integration of components, setting up systems is still a complex endeavor that requires substantial experience and expertise.

Reducing the Complexity of NFV Integration

Integrating and testing multiple open-source projects—each with its own release cycles, distinct features, and licensing models—poses unique challenges. Which versions of the open-source community projects do you select? How do you determine whether a test suite exploits the full range of capabilities of the open-

source community projects, while allowing for performance innovations on the industry standard high volume servers? And, once you have defined the integration stack, how do you implement performance tests?

The goal of Intel ONP is to create a balance between demonstrating the server platform optimizations and running on the latest open-source community project releases. Performance expectations make it essential to gain access to the latest server platform and the improvements that are useful for NFV. For this reason, each release of Intel ONP includes a well-documented system and software configuration, providing information that any ecosystem vendor or service provider end user can use to confirm the interoperability of components in their labs. Another notable open industry initiative in this area—sponsored by Network Functions Virtualization Research Group, a group within the Internet Engineering Task Force (<https://irtf.org/nfvrg>), helps accelerate NFV-applied research and is mapped well to open-standards initiatives (ETSI NFV, and so on) and open-source initiatives, such as OPNFV.




Figure 6. Carrier operations before and after virtualization.

Summary

The test results provide an indication of the performance levels that can be achieved—using varying workloads to simulate conditions in a typical work environment provisioned with commercially available servers and an optimized operating system and application stack. These results are consistent with performance levels aligned with the requirements of telecommunications companies in support of a transition from a conventional platform to a virtualized environment. As shown by the testing, there is minimal performance degradation moving from L2 switching to 5-Tuple based lookup and decision testing.

Among the insights gained from the test procedures:

- The test environment should be designed using the Intel ONP reference architecture, relying on commercially available solutions for the NFV infrastructure. By aligning to the ETSI architecture framework and the Intel ONP reference architecture, the NFVI test environment can simplify the analysis of performance data across multiple instantiations.
- To more precisely characterize NFV performance under typical workload conditions, performance testing must extend beyond simply performing packet forwarding operations in a bare bones virtualized environment to include varying both the workloads and system configurations.

Future test procedures will focus on simulating a more complex workload profile that includes procedures for quality-of-service testing. Intel, Dell, and Red Hat continue to collaborate on carrier-grade NFV solutions targeting the telecommunications market sector.

The collaboration work involving Intel, Dell, and Red Hat has resulted in a platform—the Dell NFV Platform—that is based on commercially available hardware and a hardened software stack using components from Red Hat. The resulting platform meets the service requirements of a wide range of industry segments, including telecommunications firms, as confirmed by the test data.

“Both Dell and Intel are longtime Red Hat partners, and have worked with us for many years to help bring open source innovation to enterprise customers around the world. Today, Red Hat Enterprise Linux OpenStack Platform is emerging as a leading platform for CSPs who are looking to NFV to help modernize their businesses and bring new services to customers faster than ever before. By working with Dell and Intel on OpenStack and NFV initiatives, we are working together to bring carrier-grade solutions to CSPs who are looking to NFV to help modernize their businesses and bring new services to customers faster than ever before.”

- Darrel Jordan-Smith, VP of Sales, Red Hat

Learn more about the Intel ONP:

www.intel.com/ONP

See how Dell plans to turn today's infrastructure into tomorrow's growth engine with NFV:

www.dell.com/en-us/work/learn/tme-telecommunications-solutions-telecom-nfv

For more information about Dell and NFV, contact:

Ask_NFV_PLM@Dell.com

Learn more about the strategic alliance between Dell and Red Hat:

<https://www.redhat.com/en/partners/strategic-alliance/dell>

Download the Intel ONP Server Reference Architecture for NFV and SDN:

www.01.org

Learn more about Intel Network Builders:

<https://networkbuilders.intel.com/>



¹ "Intel ONP Server Reference Architecture ETSI NFV Framework." <http://images.tmcnet.com/online-communities/nfvessentials/images/NFV%20Everywhere%20Workshop%202015-SEPT-2015%20-%20Intel%20ONP%20V01.pdf>

² "Network Functions Virtualisation (NFV); Architectural Framework." http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.01.01_60/gs_NFV002v010101p.pdf

Intel technologies may require enabled hardware, specific software, or services activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer.

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. UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

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. Copies of documents, which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or by visiting Intel's Web site at www.intel.com.

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

Copyright © 2015 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and other countries.

1115/NU/MESH/PDF 333516-002US