Gi-LAN and Dynamic Service Function Chaining for Communications Service Providers

Table of Contents

Network Topology Overview: Gi-LAN

The Gi-LAN is the segment of the network for which service providers deploy IP functions between the packet gateway and the Internet. This section of the network is where Communications Service Providers (CSPs) typically innovate, differentiate, and monetize services using unique capabilities through a combination of homegrown solutions and those provided by a wide variety of suppliers.

Figure 1 provides a high-level perspective of this segment of the network.

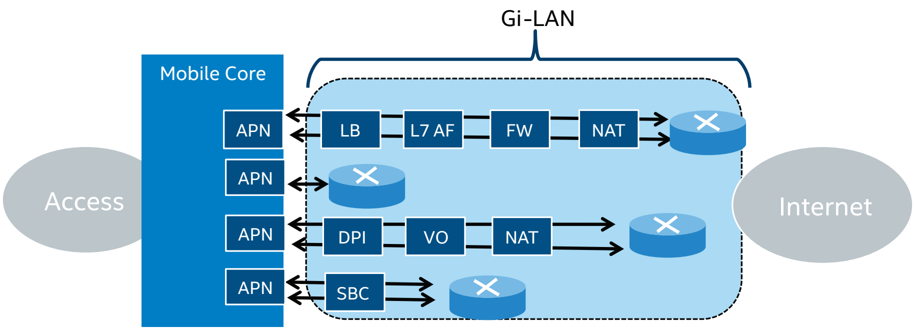

Figure 1. Gi-LAN location in the Service Provider network.

Industry Solution Architectures: Gi-LAN

Service function chaining (SFC) in coordination with software-defined networking (SDN)/Network Function Virtualization (NFV) in the Gi-LAN can be used for optimal use of data center resources. SDN-based SFC provides an approach to optimally steer traffic through Gi-LAN network functions. However, the assignment of traffic to a specific service function chain is policy-based. For mobile service providers, the implementation and coordination of SFC traffic policies will vary based on implementation and integration with a CSP's existing architecture for assigning traffic.

SDN introduces policy controls to the network that will be coordinated with the service provider's existing policy architecture. There are technical reports and studies underway in the 3rd Generation Partnership Project (3GPP*) to evaluate approaches to better integrate the SFC provisioning with the existing mobile network policy controls architecture and traffic steering policies.

Revision 1.3 May 2016

Internet Engineering Task Force (IETF) SFC Solution Architecture

While there is no global standard dictating the logic for SFC, the IETF defines an architecture and approach for SFC in packets that are classified on ingress and then processed by the specific service functions in the service chain. This architecture approach is topology agnostic, which enables it to be applied to both wireless and wireline architectures.

Figures 2-4 identify the specific architectures and the related IETF specifications that define SFC for the reference Gi-LAN:

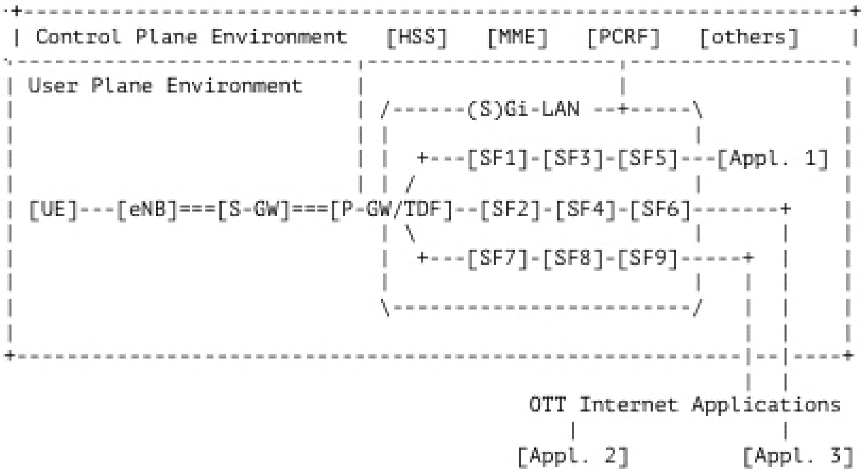

```
+------------------------------------------------------------------+
| Control Plane Environment   [HSS]   [MME]   [PCRF]   [others]    |
+---------------------------+--------------------+-----------------+
| User Plane Environment    |                    |                 |
|                           | /------(S)Gi-LAN --+-----\           |
|                           | |                        |           |
|                           | |  +---[SF1]-[SF3]-[SF5]---[Appl. 1] |
|                           | | /                      |           |
| [UE]---[eNB]===[S-GW]===[P-GW/TDF]--[SF2]-[SF4]-[SF6]-------+    |
|                             | \                      |      |    |
|                             |  +---[SF7]-[SF8]-[SF9]-----+  |    |
|                             |                        |   |  |    |
|                             \------------------------/   |  |    |
|                                                          |  |    |
+----------------------------------------------------------|--|----+
                                                           |  |
                                          OTT Internet Applications
                                             |                |
                                         [Appl. 2]         [Appl. 3]
```

Figure 2. Typical service chaining topology (IETF).[1]

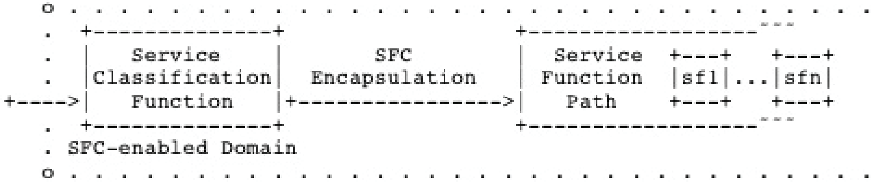

```
   o . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
   .  +--------------+                  +------------------~~~
   .  |   Service    |       SFC        |  Service  +---+   +---+
   .  |Classification|  Encapsulation   | Function  |sf1|...|sfn|
+---->|   Function   |+---------------->|   Path    +---+   +---+
   .  +--------------+                  +------------------~~~
   . SFC-enabled Domain
   o . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . .
```

Figure 3. Basic SFC architecture (IETF).[2]

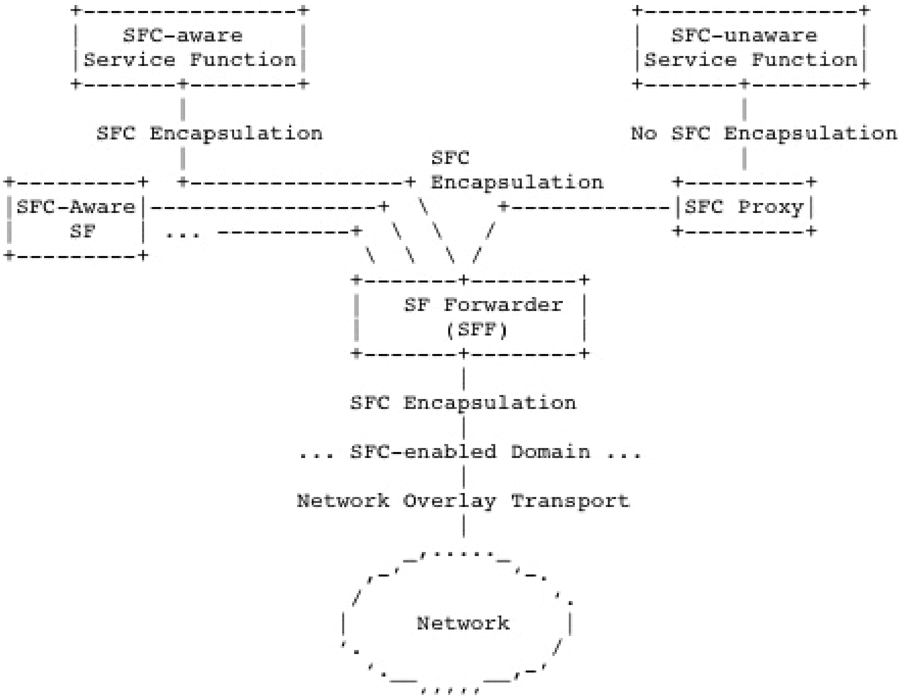

```
     +----------------+                        +----------------+
     |   SFC-aware    |                        |  SFC-unaware   |
     |Service Function|                        |Service Function|
     +-------+--------+                        +-------+--------+
             |                                         |
       SFC Encapsulation                       No SFC Encapsulation
             |                  SFC                    |
+---------+  +----------------+ Encapsulation     +---------+
|SFC-Aware|-----------------+  \     +------------|SFC Proxy|
|   SF    | ...  ----------+  \  \   /             +---------+
+---------+                \  \  \ /
                          +-------+--------+
                          |   SF Forwarder |
                          |      (SFF)     |
                          +-------+--------+
                                  |
                          SFC Encapsulation
                                  |
                      ... SFC-enabled Domain ...
                                  |
                      Network Overlay Transport
                                  |
                              _,....._
                            ,-'      '-.
                           /            '.
                          |    Network   |
                          '.            /
                            '.__    __,-'
                                '''''
```

Figure 4. SFC architecture components.

3GPP Policy Solution Architecture

Mobile service provider networks follow comprehensive 3GPP policy and charging standards to architect charging controls and policy enforcement within the network. These standards define the manner in which the policy functions are deployed within the CSP's network. Figure 5 shows an example of the logical policy architecture for a mobile network as defined by 3GPP TS 23.203.[3]

It is important to note that the above diagram is a logical architecture and it is common for vendor provided solutions to collapse these functions. As described in the Intel Gi-LAN Solution Brief, the manner in which these functions are collapsed, or provided in a standalone mode, impact how they scale.

3GPP is in the process of defining the integration of 3GPP policy standards with external policy standards, such as for SDN controllers. As an example, 3GPP TR 23.718 defines a new interface (St) between the Policy and Charging Rules Function (PCRF) and a new Service Chain Traffic Controller Function (SCTCF). This interface, among other capabilities, allows the PCRF to interface to the SFC controller functions to provide traffic description filters that enable more comprehensive and coordinated implementation service function chains in the Gi-LAN.

The 3GPP reference architecture diagrams in Figures 6 and 7 highlight the proposed integration of these different domains to enable an integrated policy architecture for Gi-LAN.

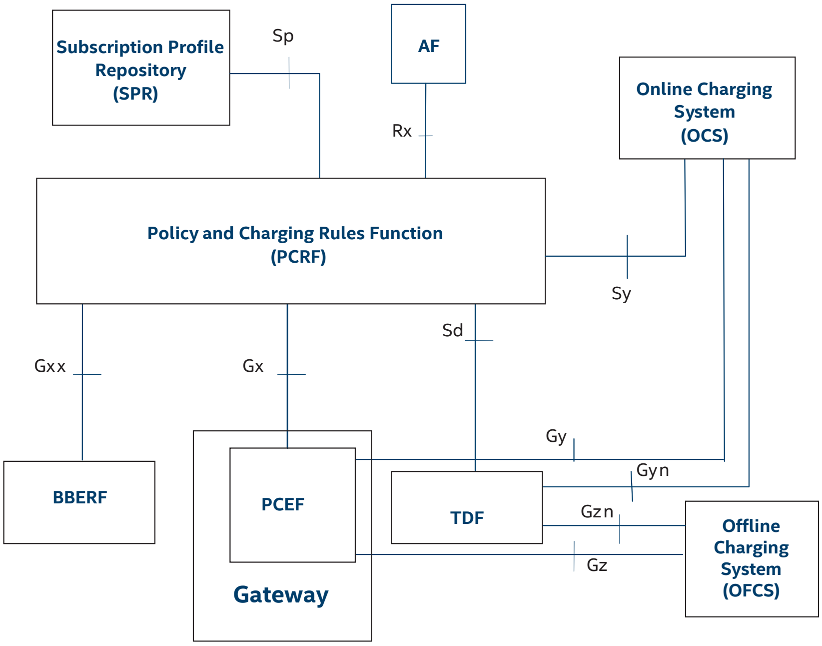

Figure 5. Overall Proof Carrying Code (PCC) logical architecture (non-roaming).[4]

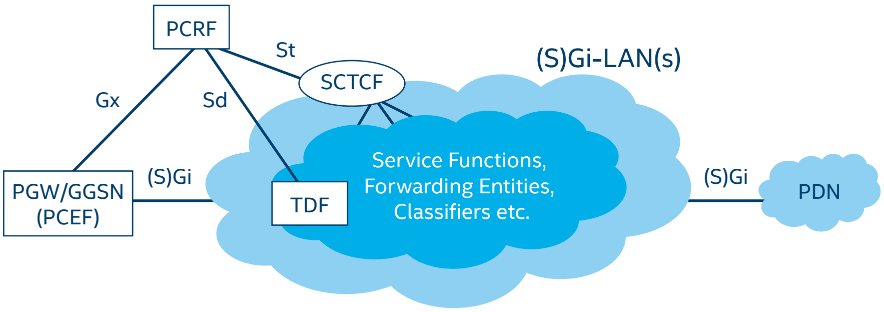

Figure 6. Reference architecture for Service Steering Policy Interface St (3GPP).[5]

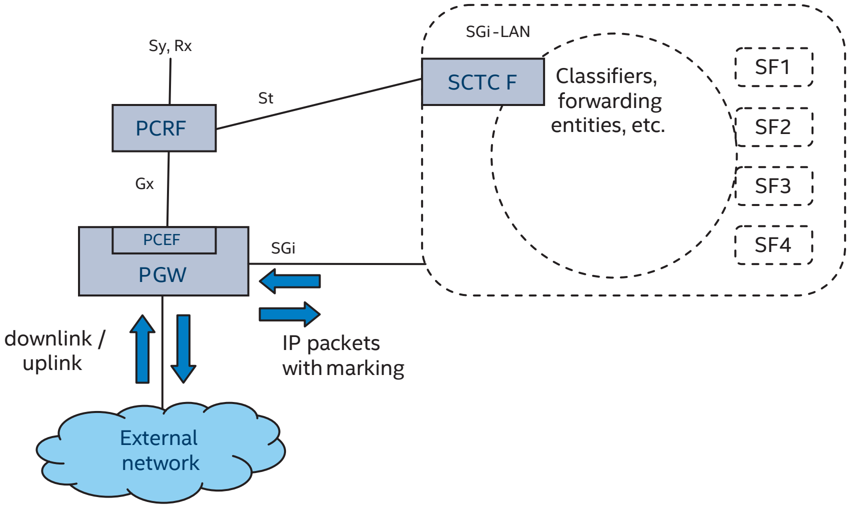

Figure 7. Example of traffic steering with packet marking in Policy Charging and Enforcement Function.[6]

Platform Reference Architectures

Table 1 shows an example of a hardware bill of materials for a Gi-LAN reference architecture.

Additional technologies, including Enhanced Platform Awareness (EPA),[7] Intel® Resource Director Technology (RDT),[8] Intel® QuickAssist Technology (Intel® QAT),[9] Intel® Trusted Execution Technology (Intel® TXT), [10] and Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI),[11] among others, are all relevant to the Gi-LAN reference architecture. Please note that not all processors support all relevant technology.

For the latest Intel® Open Network Platform (Intel® ONP) reference architectures, please refer to 01.org.[12]

Table 1. Example of a hardware bill of materials.

Processor	Dual Intel® Xeon® processor E5-2680 v3 and v4	
Platform	Intel® Server Board S2600WT 3 TB HDD Seagate Constellation* CS 2x Intel® Solid State Drive Data Center S3700 Series/200 GB	
Memory	256 GB Memory	
NICs	2x Intel® Ethernet Converged Network Adaptor X520-DA2	Niantic 2-port 10 Gb. Fortville 25GE (xl710) is also targeted for Gi-LAN. For the latest on controllers see Intel® Ethernet Controller
Add-In	Intel® QuickAssist Adapter 8950	PCIe* Gen 3 Refer to Drivers for Intel® QuickAssist Technology for the latest drivers and patches

Table 2. Links to specific capabilities.

Intel® Resource Director Technology	http://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html
Intel® QuickAssist Technology	http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/communications-quick-assist-paper.pdf https://01.org/packet-processing/intel®-quickassist-technology-drivers-and-patches
Intel® Trusted Execution Technology	http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/malware-reduction-general-technology.html http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/trusted-execution-technology-security-paper.html http://www.intel.com/content/dam/www/public/us/en/documents/guides/intel-txt-software-development-guide.pdf
Intel® Advanced Encryption Standards New Instructions	https://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/aes-ipsec-performance-linux-paper.pdf
Enhanced Platform Awareness	https://software.intel.com/sites/default/files/managed/8e/63/OpenStack_Enhanced_Platform_Awareness.pdf https://networkbuilders.intel.com/docs/openStack_Kilo_wp_v2.pdf
Open vSwitch*	https://networkbuilders.intel.com/docs/open-vswitch-enables-sdn-and-nfv-transformation-paper.pdf
Data Plane Development Kit	http://www.intel.com/content/www/us/en/intelligent-systems/intel-technology/dpdk-packet-processing-ia-overview-presentation.html https://networkbuilders.intel.com/docs/aug_17/Future_Enhancements_to_DPDK_Framework.pdf
Hardware Offload	http://www.intel.com/content/www/us/en/ethernet-products/controllers/overview.html

vGi-LAN Solution Deployment Considerations

Performance

For all virtual applications, adequate performance on Intel® architecture-based standard high-volume servers (SHVS) is necessary to justify the business case to transition from physical- to virtual-based network functions. Therefore, the ability to run more virtualized applications on the same physical hardware enhances the business case for SHVS.

However, the wide-ranging functionality of the disparate Gi-LAN Virtualized Network Functions (VNF) requires the SHVS to deliver the capabilities to enable VNFs to scale across multiple scenarios. Some examples of the vGi-LAN application requirements include throughput, latency, connections per second, statefulness, transcoding, and file servicing.

NFV and standard high volume servers in the Gi-LAN provide increased flexibility on how these logical functions can be bundled or decomposed into discrete functions or deployed as micro-services of several discrete logical functions. With all of these functions, the multi-dimensional nature of the Gi-LAN (i.e. data-plane, control plane, encryption, and statefulness) will impact the performance and how the functions scale.

To meet these requirements and to assure optimal utilization of data center resources, the environment should include the latest Intel architecture-based chipset and the capabilities identified in the previous section (for example, pQoS, Intel QAT, Intel AES-NI, Intel TXT, and EPA). This applies not only to the capabilities available in the infrastructure but also to the specific capabilities leveraged by the virtual applications.

Lab trials and proofs-of-concept (PoCs) continue to demonstrate and allow for the evaluation of different paths to reach optimal performance for virtualized functions on SHVS. Standard kernel-based Open vSwitch* (OVS) is not sufficient to meet the performance requirements to prove the business case for SHVS. While single root I/O virtualization (SR-IOV) enables better performance than standard OVS, it has certain dependencies (for example, hardware and specific drivers) that create an operational challenge. In addition, its architecture does not ultimately deliver the benefits of virtualization. However, results from recent PoCs and lab trials of Data Plane Development Kit (DPDK)-enabled OVS (user space OVS) have delivered performance results very near that of SR-IOV.

With this in mind, the recommended reference architecture to support the business case for the SHVS replacement of dedicated physical appliances must deliver an environment that supports SR-IOV and DPDK-enabled OVS. This approach enables CSPs to get to market immediately on SHVS using DPDK-enabled OVS or SR-IOV, if required. The transition to DPDK-enabled OVS will support the realization of the benefits of virtualization, such as optimization and specific feature sets.

Operational

As identified in the previous section, the successful implementation of a virtualized Gi-LAN solution requires coordination of policy across 3GPP and SDN domains. However, as part of the decision process to enable the policy framework, Gi-LAN solutions also require the virtualization and integration of hooks across all layers of the software stack. As shown in the figure below, horizontal solutions must link capabilities across all the disparate layers of the NFV/SDN stack to actually deploy, monitor, and bill for services. Optimally, vGi-LAN service realization requires an appropriate environment, with resource models coordinated with subscription plans and entitlements in the policy framework.

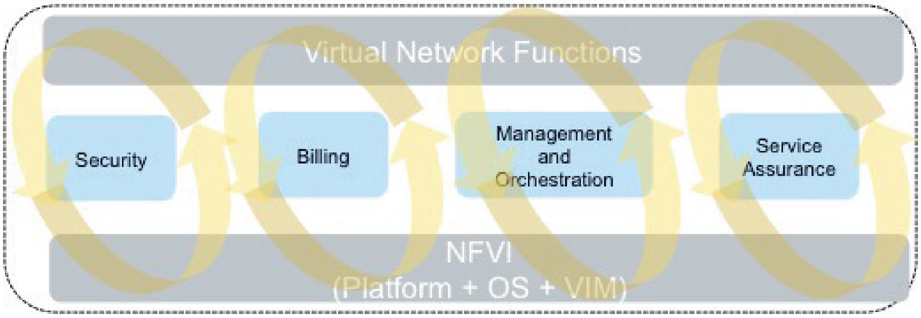

Figure 7. Horizontal solution linkage

How, where, what, when, and how much?

There are important questions that must be addressed to effectively design, integrate, and operate a virtualized Gi-LAN (for example, how, where, what, when, and how much). The ability to deploy VNFs successfully on SHVS requires an understanding of the environment's capabilities and location and must be coordinated with the application's requirements. SHVS provides significant flexibility and workload placement from edge to core, driving the requirements for the operational tools to ensure the environment is trusted, secure, and has the appropriate resources for the task.

As an example, Intel is driving industry initiatives and making contributions, such as OpenStack[13] EPA, to ensure that Intel architecture-based platform capabilities are exposed appropriately to the CSP's management and orchestration implementations and enable optimal utilization of data center resources.

Network service life cycle management, monitoring, security, and billing models are all examples of the horizontal solution capabilities required for success and scaling in a virtualized environment. Intel is working closely with CSPs to enable and, in some instances, drive industry efforts to accelerate maturation of these horizontal capabilities. Future versions of this document will provide more detail and examples of those efforts.

As described in Intel's Service Provider Network Maturity Model,[14] the initial deployments of network virtualized functions will rely on single-vendor implementations for end-to-end vertical solutions that include the necessary horizontal capabilities. As the horizontal capabilities mature and service providers transition from single-vendor implementations, the business case for creating commercial business services from the virtualization of network functions will continue to improve and will also help drive new revenue opportunities from the new offerings.

Next steps

- To learn more about Intel's technology for NFV, attend the courses available in the Intel Network Builders University at https://networkbuilders.intel.com/university.
- To learn more about Intel Network Builder partners for vGiLAN and other NFV products, visit https://networkbuilders.intel.com/solutionscatalog.
- To build a testbed using the Intel ONP Reference Architecture, download the documentation at https://01.org/packet-processing/intel%C2%AE-onp.
- To get the best security in your NFV systems, specify Intel Cloud Integrity Technology in your infrastructure and VNF procurements.
- To get the highest performance from your NFV systems, specify compatibility with the Data Plane Development Kit in your Infrastructure and VNF procurements.
- To get the highest return on investment from your NFV systems, specify use of Enhanced Platform Awareness in your Orchestration, Infrastructure and VNF procurements.

[1] https://datatracker.ietf.org/doc/draft-ietf-sfc-use-case-mobility/

[2] http://www.rfc-editor.org/info/rfc7665

[3] http://www.3gpp.org/DynaReport/23203.htm

[4] http://www.3gpp.org/DynaReport/23203.htm

[5] http://www.3gpp.org/DynaReport/23203.htm

[6] http://www.3gpp.org/DynaReport/23718.htm

[7] https://software.intel.com/en-us/articles/openstack-enhanced-platform-awareness

[8] http://www.intel.com/content/www/us/en/architecture-and-technology/resource-director-technology.html

[9] https://01.org/packet-processing/intel%C2%AE-quickassist-technology-drivers-and-patches

[10] http://www.intel.com/content/www/us/en/architecture-and-technology/trusted-execution-technology/malware-reduction-general-technology.html

[11] https://software.intel.com/en-us/articles/intel-advanced-encryption-standard-instructions-aes-ni

[12] https://download.01.org/packet-processing

[13] www.openstack.org

[14] http://www.intel.com/content/www/us/en/communications/service-provider-network-maturity-paper.html

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.

Copyright © 2016 Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Printed in USA 0516/JL/PDF ♻ Please Recycle 334453-001US