

Deliver Scalable VoLTE and Agile Rich Communications Services

With Metaswitch Virtualized Multimedia Network Functions, driven by Intel® technologies and deployed in Hewlett Packard Enterprise's OpenNFV* infrastructure.



“Telecommunications executives need to evolve their business decision making to embrace the possibilities of cloud and software-based services.”

John Healy, General Manager
SDN Division, Intel®

Executive Overview

While holding the promise of HD quality, enriched services, and the benefits of a common packet switched access infrastructure, the adoption of infrastructure-based voice over LTE (VoLTE), along with rich communications such as video and real-time messaging, has stalled over the last few years with concerns around quality, security, and cost. During that time, the value of the phone number as a globally significant user ID and key services such as SMS and MMS have been relegated by phone manufacturers, over the top (OTT) services, and even the subscribers themselves as “technologies of last resort.” Finally, the proliferation of 4G, together with the emergence of network functions virtualization (NFV), powered by Intel®, Hewlett Packard Enterprise* (HPE), and Metaswitch*, is allowing network operators to take charge of the consumer's communications experience, once more.

Realizing the Promise of IMS

At the heart of this resurgence is the IP Multimedia Subsystem (IMS). Conceived of by the 3GPP 15 years ago, continuous innovation has ensured that the suite of deployment specifications is still relevant, both today and for future multimedia service deployments.

The IMS architectural framework was pioneering in many ways: Laying the groundwork for the migration from circuit switch to packet switch telephony infrastructures, IMS set out to ensure that historical mistakes were not replicated in next-generation consumer voice offerings. In contrast to the vertically integrated, single-vendor, proprietary stovepipes that came before, the IMS specifies a simple, three-tier, horizontally integrated model, with a clear, well-defined separation between the media, control, and application layers.

Embracing an IMS philosophy eliminates vendor lock-in while enabling network operators to deploy best-of-breed or the most cost-effective telephony network functions and scale them completely independently of each other. Employing IMS also eliminated the dependency on a single service delivery platform, allowing carriers to deploy new applications from independent software vendors (ISVs) without being beholden to the development roadmap of a single vendor.

Until recently, mobile voice network transformation from circuit switch TDM to packet-switched VoIP and IMS adoption itself was stifled not only due to the slow migration from

Deliver Secure, Scalable VoLTE and Agile Rich Communications Services

3G to 4G radio access network (RAN) technologies but also because IMS functions remained pricey components implemented on expensive platforms, in static environments, and available only from the largest network equipment vendors. This dramatically limited the implementation and application flexibility promised of the architecture, while making wide-scale deployments cost-prohibitive for network operators. These dynamics changed with the introduction of Project Clearwater* vIMS and the advent of NFV.

Project Clearwater: IMS in the Cloud

With the initial codebase contributed to the open source community by Metaswitch* in 2013, Project Clearwater* was not only at the forefront of the virtualization revolution, but it also fundamentally redefined how communications network infrastructure should be acquired and deployed. Clearwater follows IMS architectural principles and supports all the key standardized interfaces expected of an IMS Core network (see Figure 1). Unlike traditional implementations of IMS, however, Clearwater was designed from the ground up for the cloud.

By incorporating design patterns and open source software components that have been proven in many hyperscale global web applications, Clearwater achieves an unprecedented combination of massive scalability and exceptional cost-effectiveness. Clearwater was architected for highly distributed cloud infrastructures, with critical innovations unique to asynchronous communications applications and the resiliency demanded of carrier applications but not always afforded by the individual, underlying data center components.

As community-driven open source software, Clearwater completely

eliminates any individual vendor lock-in of the IMS Core itself along with its auxiliary components and associated applications. Clearwater also removes the cost associated with implementing IMS infrastructures. Data centers, built on commercial off-the-shelf hardware and employing Intel® Architecture, can serve subscribers for only a few pennies a year, affording network operators the ability to deliver innovative new services without the need to pass additional infrastructure costs onto their subscribers.

Simple Stand-Up and Scale-Out or Fast Fail

However, infrastructure built using web design methodologies and explicitly architected for virtualization, like the Clearwater vIMS* Core, is just the first step to enabling carriers to operate like nimble web companies. Network operators must have the ability to turn up new functions rapidly, beta test in live environments, and with real subscribers then turn live and dynamically scale out as usage increases. Likewise, in the event

new service offerings do not prove popular, they must have the ability to be quickly decommissioned with little to no cost and without impacting key customer applications. This is where NFV infrastructure (NFVI), such as that outlined in the HPE OpenNFV* Reference Architecture, along with HPE's Helion* Platform and the entire HPE NFV System Family excels.

With granular control capabilities, the NFVI requirements, base configurations, and the relationships between individual virtualized network functions (VNFs) and components are defined within the NFV management and orchestration (MANO) layer (See Figure 2). This allows for practically one-click deployments of even the most complex VNFs, like those comprising a vIMS Core. New elements and applications, such as an innovative rich communications service, can be easily added to an existing vIMS Core or easily spun up within a completely standalone vIMS instance. If successful, the service can be expanded dynamically, with the individual functions themselves alerting the

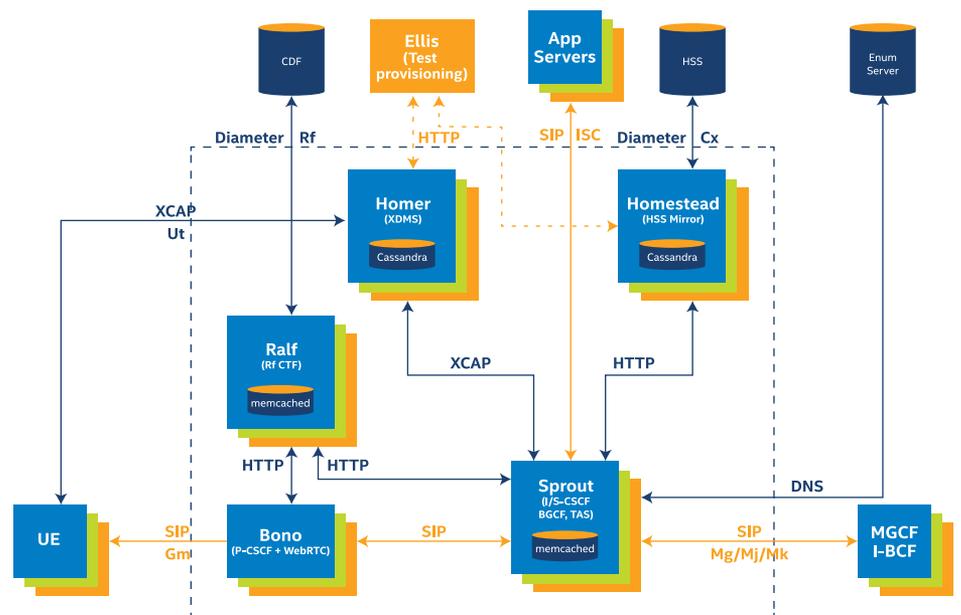


Figure 1. Project Clearwater Architecture and Components

Deliver Secure, Scalable VoLTE and Agile Rich Communications Services

NFV MANO later to deploy additional capacity, where required.

The IMS Core has long been considered an excellent candidate for NFV, and Project Clearwater, itself, was the earliest example and still remains one of the most widely deployed independent vIMS VNFs, supporting fixed line and mobile VoIP, VoLTE, voice over Wi-Fi* (VoWiFi), and rich communications suite (RCS) services around the globe. With wide scale adoption of packet-based services, however, comes the need to protect and secure VoIP infrastructure—a critical function not addressed in the IMS Core specifications.

Fortify Your Edge and Protect Your Core

Residing at both the network access layer and at the interconnect points between network operators, the Session Border Controller (SBC) is at the heart of not only VoIP security, but it also guarantees interoperability between potentially disparate signaling and media versions or types. The location of an SBC within the network architecture and its functionality also provides a perfect point from which to perform network analysis. While, like the IMS Core, the SBC's role is invaluable to the smooth operation of a multimedia communications infrastructure, with each end-to-end VoLTE call traversing an SBC element at least twice, the function must be deployed in the most cost-effective and scalable manner possible.

Historically, these SBCs combined signaling and media functions within a single proprietary hardware platform with many specialized chipsets. Not only did this design characteristic prevent SBC virtualization, but also the integrated nature of the platform went against IMS philosophies of separated control and data planes. This resulted in a solution where these very distinct functional elements could not be

placed in their most optimal locations—edge for media and core for signaling—or scaled independently. This resulted in an overly expensive solution with large amounts of stranded (unused/ unusable) capacity around a network.

Recognizing that advancements in the Intel architecture, along with our own code-level innovations, enabled Metaswitch to break with historical precedence, the Perimeta* SBC was architected from inception as a pure software solution deployable on standardized server hardware platforms and therefore perfect for cloud NFV evolution. Moreover, Perimeta decoupled media and signaling, allowing for totally independent scaling and optimal placement of these functions within emerging networks.

From Moore's Law to Multicore and Beyond

The continued increase in processor speeds and the number of individual cores per CPU, plus some kernel-level packet processing smarts, enables Perimeta to out-perform proprietary hardware-based SBCs, right out-of-the-box, while leveraging the cost

benefits afforded by commercial off-the-shelf Intel architecture-based server hardware. Perimeta employs other standard Intel® processor features to eliminate the need for the custom silicon found in classic SBCs. This includes the Intel® Advanced Encryption Standard instruction set for cryptography features and a core's large Layer 3 cache memory to hold the IP address lookup tables typically stored in a Ternary Content Addressable Memory (TCAM) Application-Specific Integrated Circuit (ASIC). When deployed on Intel platforms, Perimeta can secure a network operator's VoIP and VoLTE infrastructure from the most devastating DDoS attacks while servicing 100 percent of calls.

When virtualizing the SBC function, Metaswitch again looked to Intel's innovations to ensure the Perimeta vSBC could benefit from the flexibility of NFV without suffering from any detrimental performance degradation (See Figure 3). Dramatically reducing CPU interrupts and the number of times packets originating at the Network Interface Controller (NIC) must be copied before they can be processed by the SBC application, Intel architecture data plane acceleration

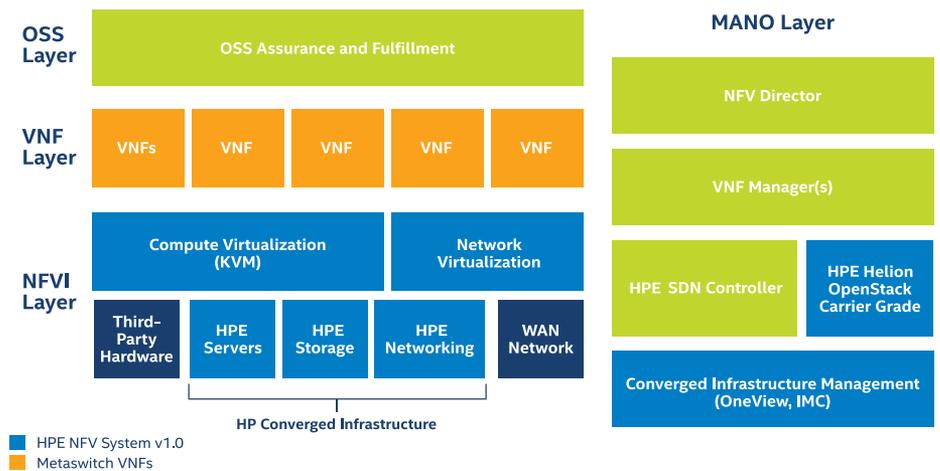


Figure 2. Metaswitch VNFs within the HPE OpenNFV Reference Architecture

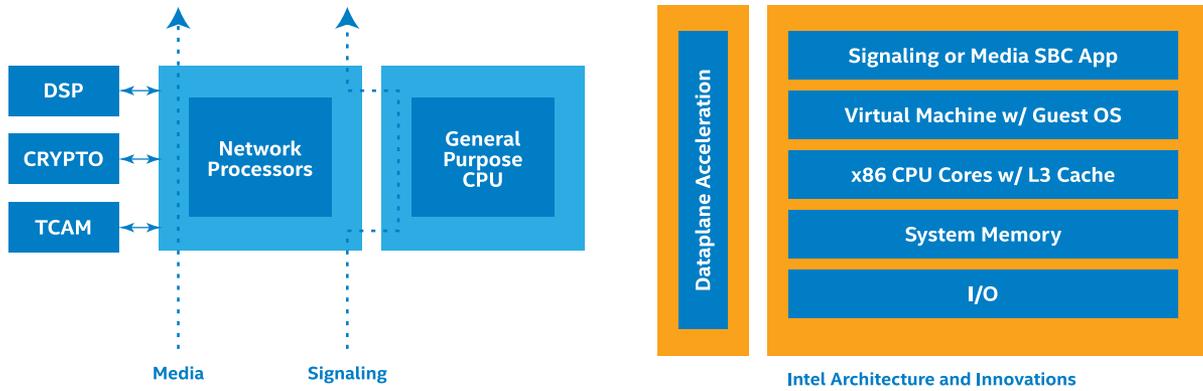


Figure 3. The Proprietary Hardware-base SBC Architecture and the Metaswitch PerimetavSBC Architecture

techniques, such as Single Root I/O Virtualization (SR-IOV) and the open source Data Plane Development Kit* (DPDK) community, driven by Intel, are important components in virtualization and NFV. With accelerated vSwitch implementations, employing modified DPDK implementations, the Perimeta SBC can be deployed in a fully orchestrated NFV environment on a software-defined network (SDN) foundation with only a minimal reduction in overall performance, which is more than compensated by the dynamic flexibility and extensibility afforded by NFV and SDN.

easily extended to deliver IR-94 video and comprehensive RCS services. Even WebRTC* interfaces and applications can be supported with the addition of a gateway such as that integrated within Project Clearwater. Moreover, as the unlicensed Wi-Fi spectrum is increasingly employed to provide offload or roaming services, Clearwater and Perimeta can enable IMS over Wi-Fi, or VoWiFi, which is fully compliant with the GSMA IR.51 specifications (See Figure 4). This entire infrastructure features comprehensive integrated

Advancing VoLTE Deployments with Metaswitch, HPE and Intel

Adopting NFV is critical to delivering a commercial VoLTE and multimedia rich communications suite services offering, which is both flexible and scalable while meeting the pricing pressures faced by network operators around the globe. Anchored by a Clearwater vIMS Core and the Perimeta vSBC, carriers can offer a complete IR-92 VoLTE service with feature parity aligned to the legacy circuit switched network and Enhanced Single Radio Voice Call Continuity (eSRVCC) for seamless fallback when roaming across 3G infrastructures. This VoLTE offering is

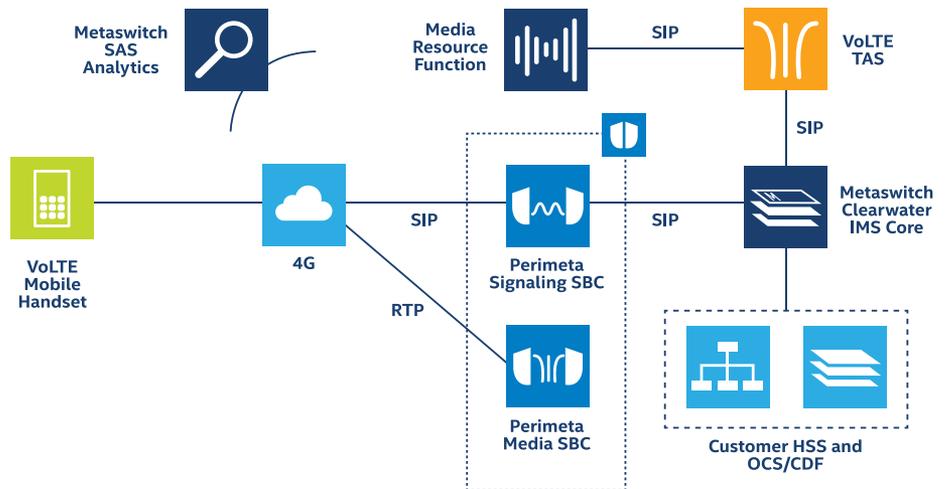


Figure 4. Metaswitch VNFs within a VoLTE Architecture

Test Results for Metaswitch Perimeta* SBC and HPE OpenNFV* Reference Architecture

The Metaswitch Perimeta vSBC was successfully tested on HPE Helion*, demonstrating elastic scalability through HPE's OpenNFV platform, along with complete management and orchestration using HPE's NFV Director. During testing, The Metaswitch Perimeta vSBC exhibited:

- Onboarding multi-VM pattern VNF for basic compatibility testing with Helion KVM/OpenStack* virtualization
- Verification of network paths, traffic routing, and basic functionality
- Verification of redundancy, failure, and recovery scenarios
- Verification of basic call flows

Clearwater vIMS and the Metaswitch Service Assurance Server have also been tested in this environment.

HPE OpenNFV* Labs

In HPE OpenNFV Labs, we and our partners test applications to make sure they run as expected on our reference architecture. We provide a testing center where partners can test a set of critical applications simultaneously to make sure they are ready to be deployed in carrier networks on a shared NFV infrastructure. In addition, HPE has dedicated lab facilities where we can stage NFV proof-of-concept projects with partners. Carriers can also conduct NFV proof-of-concept

and feasibility tests for new NFV applications on our HPE NFV reference architecture. Our labs are located in:

- Houston, Texas
- Ft. Collins, Colorado'
- Grenoble, France
- Tel Aviv, Israel
- Seoul, South Korea

About Hewlett Packard Enterprise

With its unparalleled experience in IT and a long standing telecommunications expertise, Hewlett Packard Enterprise is a trusted partner to Communications Service Providers as they embark on a journey to the Telco Cloud. Recognizing that the benefits of NFV cannot be fully realized without open solutions and a robust ecosystem of partners, HPE launched its OpenNFV Partner Program for network equipment providers, independent software vendors, and system integrators. The OpenNFV Partner program is a very important part to create a rich, vibrant, and open ecosystem of VNFs. The goal for HPE's OpenNFV program is to create a platform on which CSPs have the freedom to choose applications from their vendor of choice.

As part of the OpenNFV platform, HPE provides the NFV infrastructure (NFVI) running on HPE's converged servers; the HPE Helion OpenStack Carrier Grade which provides

the virtualization layer and VIM functionality as well as the HPE NFV Director for NFVO (NFV Orchestrator) and in many cases VNFM functionality. This platform then acts as a foundational reference architecture on which VNFs (from HPE or from any third party partners) can be tested and benchmarked for performance.

OpenNFV Labs (located in the U.S., France, Israel, and South Korea) help partners accelerate their design, proof-of-concept, trial, and deployment of cloud-enabled network services. The primary goal of HPE NFV Labs is to assure CSPs that solutions being proposed to them from multiple vendors are pre-tested and integrated – thereby saving them valuable time and effort in network validation during deployment.

Hewlett Packard Enterprise is committed to developing open solutions and an open ecosystem that will accelerate the CSP transformation journey to the Telco Cloud. HPE is a key contributor to open source initiatives like OpenStack, OpenDaylight Project, OPNFV and numerous other Industry initiatives and organizations.

Deliver Secure, Scalable VoLTE and Agile Rich Communications Services

Ecosystem Partners Deliver Innovation and Performance

Intel is leading network transformation through open source, open standard technologies that form the building blocks of SDN and NFV. With HPE as a technology partner, Intel collaborates closely with Intel® Network Builders members such as Metaswitch Networks. This collaborative ecosystem enables an open, flexible environment for communication service providers (CSPs) to make the transition to NFV on the HP NFVI platform, which includes Intel® Xeon® processors, a DPDK, and the Intel® Open Network Platform (Intel® ONP) reference architecture.

The Intel, HPE, and Intel Network Builders ecosystem provides CSPs with the resources to respond quickly to market demands while accelerating the overall transformation toward NFV. HPE's NFVI platform streamlines the design, proof-of-concept, trial, and deployment of new cloud-enabled network services and innovations while lowering capital expenditures, operating expenditures, and risk. Intel is a key contributor to the HPE OpenNFV Lab testing, standards work, and the development of the NFVI platform.

Intel and select subsidiaries are also working with HPE to help CSPs deliver commercial solutions targeting European Telecommunication Standards Institute (ETSI) NFV use cases on the HPE Helion OpenStack* Carrier Grade platform. Further integrations will take place with Intel and open source communities, such as the Open Platform for NFV* (OPNFV*), when those integrations meet CSPs' performance requirements.

Intel® Open Network Platform

Intel developed an architecture platform based on the ETSI standards for network functions virtualization (NFV) in an effort to accelerate development of commercial hardware and software platforms. The Intel® Open Network Platform (Intel® ONP) provides a reference architecture that redefines network architectures by decoupling the network functions from the hardware itself. This provides the NFV infrastructure necessary to virtualize functions.

With Intel ONP, companies in telecom carrier networks, enterprise environments, and cloud data centers can more easily build solutions using an open source software stack running on commercial, off-the-shelf servers. With the Intel ONP reference architecture, solution providers can plan, evaluate, and benchmark designs in advance of NFV deployments.

More About Metaswitch Networks*

Metaswitch is powering the transition of communication networks into a cloud-based, software-centric, all-IP future.

As the world's leading network software provider, we design, develop, deliver and support commercial and open source software solutions for network operators. Our high performance software runs on commercial, off-the-shelf hardware, as appliances or in the cloud. We package this software into solutions that are redefining consumer and business communications and enabling the interconnection between diverse network services and technologies. We also apply our software development expertise to removing network virtualization complexities in the data center, with a solution that easily scales and secures workload interconnection in support of mission-critical IT and real-time communication applications.

To learn more about Project Clearwater vIMS, visit www.projectclearwater.org. For further information about Clearwater support services from Metaswitch, the Perimeta vSBC, or the Service Assurance Server, go to www.metaswitch.com.

Get more information about storage solutions powered by Intel: visit <http://hpe.com/go/nfv> or <https://networkbuilders.intel.com>



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

0116/TT/ICMCSW/PDF

Please Recycle

333824-001US