

Realising the Benefits of Network Functions Virtualisation in Telecoms Networks

Network operators profit from adding computing power to the RAN and core network.



Executive Summary

Network Functions Virtualisation (NFV) is a pervasive theme within the telecoms industry, and over the past year, it has become a catalyst for major transformational change in the network. Many service providers, telecom equipment manufacturers (TEMs), and an emerging NFV ecosystem have announced proof of concepts, live trials, and commercial products embracing the NFV vision. In addition, forums such as the ETSI NFV Industry Specification Group and the Open Networking Forum have provided the necessary structures to support the initiative and its continuing momentum and evolution.

Generally speaking, as more network functions are migrated to standard IT, high-volume server environments, the stronger the business case becomes for widely deploying cloud and IT data centre concepts in the telco network. From a service provider perspective, a major driver for NFV is OpEx savings; but additionally, NFV provides a platform for rapid service deployment and monetization, which is potentially even more important as it directly addresses growth models and revenue streams. For the TEM community, NFV facilitates the development of telecom solutions with application scalability in the cloud, optimized performance, software reuse, and increased use of open source software.

Moreover, NFV is now becoming a key criterion in the vendor selection process. Service providers who have seen the benefits of NFV in concrete terms are requiring vendors to act as quickly as possible. Naturally, TEMs also see opportunities to win market share in this disruptive market environment. Opportunities exist for those willing to embrace the NFV dynamic, and the associated technologies and architectures necessary to make it a reality.

However, many network functions have extreme characteristics that must be addressed by a highly-robust and scalable software and hardware architecture. Operators expect both telco-grade quality and equipment performance to improve or at least remain the same in an NFV environment. Consequently, TEMs must achieve high performance when running their applications on standard server hardware in the cloud. The LTE evolved packet core (EPC), and especially the gateway functionality, is a good example of an application requiring not only high-performance user plane, but also extended performance scalability in both signalling and control plane. The full benefit of NFV will ultimately be realised when mature reference architectures are available that facilitate easy deployment of disparate network workloads; deliver high performance, quality, and resilience; and can be intelligently managed to allow rapid service invocation and delivery.

This paper is a joint effort by Intel, Qosmos*, and Tieto* to describe how standard IT, high-volume servers, and commercial and open source software can be used in a reference architecture vision to meet the challenges of network evolution, and thereby unleash the full potential of NFV.




Table of Contents

Executive Summary	1
Reference Architecture Description	2
Network Functions Virtualisation Infrastructure	2
Virtualised Network Functions	3
NFV Management and Orchestration	3
Network Operation (OSS/BSS)	3
Cloud Platform	3
SDN Networking	3
Application Cloudification Service Solutions	3
Virtual Network Function as a Service (VNFaaS)	4
DPI Overview	4
Business Benefits	4
Challenges to Overcome	4
Addressing the Challenges	5
Virtualisation of the Mobile Core Network (EPC) and IMS	7
Business Benefits	7
Challenges to Overcome	7
Addressing the Challenges	7
Virtualisation of the Mobile Base Station (in the RAN)	8
Business Benefits	8
Challenges to Overcome	9
Addressing the Challenges	9
Summary and Conclusions	10

Reference Architecture Description

The reference architecture vision created by Intel, Qosmos, and Tieto builds upon previous work that Intel and Tieto developed and made available to the NFV community beginning in 2012.¹ The latest evolution of the architecture is aligned to ETSI NFV ISG principles and has been extended to address intelligent deep packet inspection (DPI), traffic shaping, service invocation, and diameter signalling controller virtual functions.

- The reference architecture is based on an Intel® architecture platform, which includes the Intel® Data Plane Development Kit (Intel® DPDK) and the Intel DPDK Accelerated Open vSwitch*.
- Commercial grade components, including Qosmos virtual DPI function, Tieto IP (TIP) Stack, Tieto's virtualised Diameter Signaling Controller (DSC), and Wind River* Open Virtualisation Profile (OVP), have been combined with example applications for virtual EPC and radio access (i.e., virtual base station) and integrated to showcase an end-to-end proof-of-concept (PoC).

The high level business and technical objectives for the reference architecture (Figure 1) and the associated PoC are as follows:

- 1. Faster innovation** – Expand choices beyond traditional manufacturers and processor architectures to include open-standard, open-platform ingredients, solutions, and services.
- 2. Save costs** – Decrease OpEx and CapEx by reducing dependency on proprietary hardware and providing virtualised network functions on standard, high-volume hardware, as well as through of automated and simplified operation of load-based resource allocation, fault avoidance, and recovery.
- 3. Increase service revenue** – Offer software solutions and services more quickly and easily via a flexible architecture that is well-suited for new service creation and monetization.

4. Maintain performance and resilience

– Deliver high-performance packet processing and maintain system-level resilience and quality.

5. End-to-end orchestration and automation

– Demonstrate new services and software solutions that can be quickly and effectively deployed using automation (OpenStack* Heat), and orchestration based on open interfaces and OpenStack components, where functions are managed at the system level.

The architecture consists of the following subsystems:

Network Functions Virtualisation Infrastructure

The NFV Infrastructure subsystem provides the basis for a flexible, high-performance, and highly-resilient virtualised telecoms platform.

It includes standard IT, high-volume hardware based on the latest Intel® architecture processors and Intel network interface cards (NICs). The associated PoC uses the HP* ProLiant* DL380p Gen8 Server based on Intel® Xeon® processor E5-2690 (2.9 GHz) and Intel® 82599 10 Gigabit Ethernet Controllers. In addition, there are several switches, including a high-capacity, OpenFlow*-enabled Intel® Ethernet Switch FM6764 (10 GbE/40 GbE L2/L3/L4) and several 1 GbE switches that are used for management interfaces.

The software execution environment is a Wind River* Linux* distribution with Wind River Open Virtualisation Profile providing virtualisation support, which includes KVM (Kernel-based Virtual Machine) and OpenStack, as well as real-time extensions for virtualisation.

OpenFlow* provides SDN support by configuring an optimized version of Open vSwitch that takes advantage of SR-IOV (Single Root I/O Virtualisation) and the Intel DPDK² to accelerate packet processing performance in a virtualised environment.

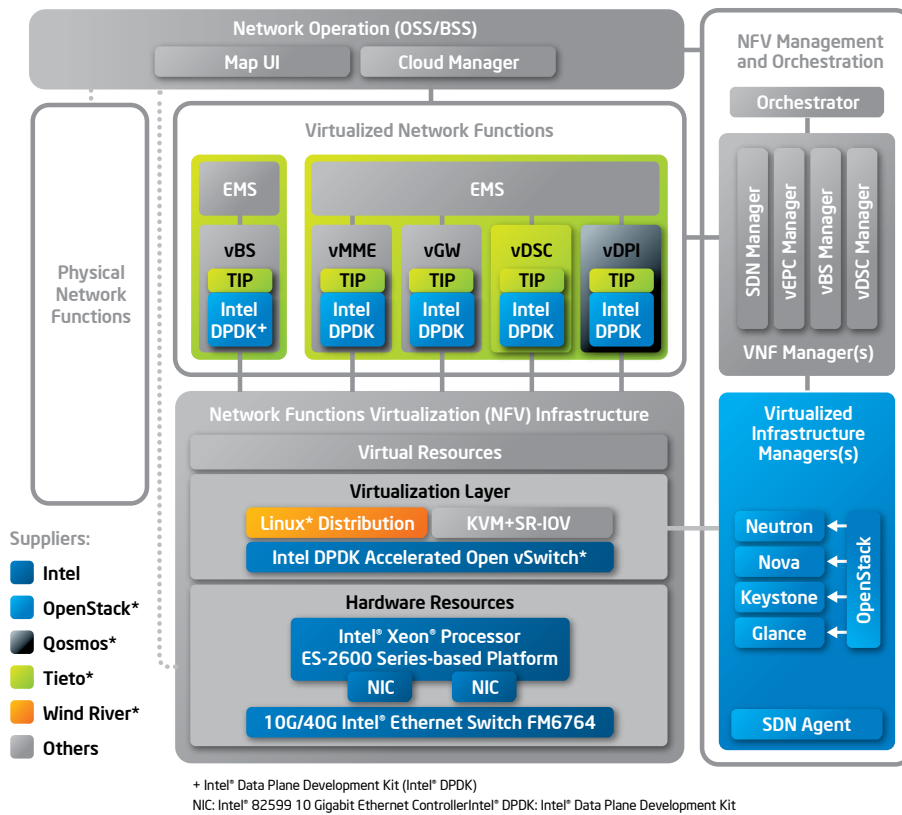


Figure 1. Architecture Description

Virtualised Network Functions

Virtual network functions (VNFs) are deployed on the NFV Infrastructure. The PoC implements reference software for LTE eNodeB and EPC (MME, SGW, and PDN GW), along with Tieto’s Diameter Signalling Controller (DSC), which is deployed as a VNF. To demonstrate intelligent DPI content inspection, Qosmos technology is used, and the reference architecture allows it to be deployed either within a VNF or as a standalone virtual networking function component (VNFC). The associated Element Management System (EMS) for each VNF is integrated within the VNF subsystem, which monitors the operational condition of the VNFs as part of the overall Telecommunications Management Network (TMN).

NFV Management and Orchestration

This subsystem manages overall VNF deployment and lifecycle, and it also contains SDN controller functionality

that manages flow control to enable intelligent networking. While the reference architecture is orchestrator-agnostic, the PoC’s orchestration solution is provided by OpenStack* and its subcomponents, including networking (Neutron), compute (Nova), and identity services (Keystone). The network management solution interfaces to Openstack (Heat) for automation and deployment.

Network Operation (OSS/BSS)

Interworking with network operation requirements and legacy OSS/BSS³ poses one of the biggest challenges for fast and effective deployment of NFV applications in the network. The network operation subsystem comprehends the necessity for standards-based integration of OSS/BSS infrastructure as well as for overall management of the telecoms cloud. To address these needs, the reference architecture and PoC implementation support a Netconf interface for OSS

and a Cloud Infrastructure Management Interface (CIMI). The intent is to demonstrate that a new VNF could be deployed and managed from an end-to-end perspective within the telecoms cloud, while also presenting necessary fault, configuration, accounting, performance, and security (FCAPS) information to the OSS/BSS complex.

Cloud Platform

The Cloud Platform reference architecture can be used to accelerate time to market and deliver telco grade quality to NFV deployments. It is based on OpenStack elements that Tieto integrated and performance optimized, and includes telco-grade supervision, statistics, diagnostics, fault, and performance management capabilities.

SDN Networking

The integrated SDN controller is compatible with the OpenStack Neutron and supports SDN networking and legacy network management system (NMS) integration, and provides supervision, statistics, and performance management. The OpenFlow protocol is used to communicate between the SDN controller and the Open vSwitch, which is managed by the OpenFlow controller. Efficiency is achieved via integration with the orchestrator.

Application Cloudification Service Solutions

Tieto development services for application cloudification are available to help companies accelerate time to market and off-load their own R&D. The service decouples software from hardware or decouples functions from the application, enabling them to become cloud applications and functions. The service optimizes the cost of cloud transition by taking into account the integration of the legacy equipment and the need for high performance.

ETSI NFV Use Cases

While the reference architecture is a generic, virtualised infrastructure designed to support different kinds of telecom workloads and applications, the associated PoC explicitly addresses the radio access network (RAN) and the LTE EPC. In addition, DPI is a key ingredient technology that is supported in both the architecture and the PoC. To demonstrate how the reference architecture supports key telecoms workloads, the following sections describe three **ETSI NFV use cases**,⁴ detailing the challenges to be overcome and how the Intel/Qosmos/Tieto reference architecture addresses them. These ETSI NFV use cases are:

- Virtual Network Function as a Service (VNFAaaS)
- Virtualisation of the Mobile Core Network (EPC) and IMS
- Virtualisation of the Mobile Base Station (in the RAN)

Virtual Network Function as a Service (VNFAaaS)

Many service providers offer cloud computing services in addition to network services. Supporting both types of services with common NFV infrastructure requires resource pooling mechanisms that allocate physical compute, network, and storage resources to network applications and VNFs.⁵ Once this is achieved, it will be easier to deploy individual instances of virtualised network functions and offer them as services.

The reference design demonstrates this capability in its implementation of the ETSI VNFAaaS use case by dividing the VNFs into several instances that are deployable as-a-Service.

DPI Overview

DPI enables many functions in the core network, including quality of service (QoS) management through traffic shaping, throttling, and gating. In this case, network traffic, especially obfuscated

and cyphered protocols, is inspected with fine granularity, which is possible with the DPI capabilities provided by Qosmos virtual DPI module. Likewise, telcos are creating tiered data plans by using DPI technology to turn the core network into an application-aware system.

In NFV environments, alongside the QoS enablement, information generated by DPI can also enable telcos to optimize costs through improved equipment utilization while delivering on service level agreements (SLAs) to cloud consumers. The Qosmos DPI function will present extracted information to the orchestrator/analytics system, which triggers events such as scaling virtual machine (VM) instances up or out. Virtualisation technology offers the ability to relocate a virtual machine from one host to another without shutting it down, thus giving the opportunity to dynamically optimize the placement with minimal impact on performance. Nevertheless, it is critical to be able to express application requirements related to VM placement and server state in order to define the VM placement constraints required to model a viable configuration. For instance, identifying which services are used and whether they are stateless or stateful is necessary to ensure good performance.

The reference design provides guidance on how to deploy such a VNFAaaS, while in parallel, optimizing the network topology with a network-wide, VNF-based virtual base station and EPC solution.

Business Benefits

The ability to deliver DPI as a VNF(C) will speed its deployment in the network, resulting in quicker TCO savings, faster service monetization, and a significant improvement over existing core networks that only implement shallow packet inspection (SPI).

In the case of the reference architecture, DPI is deployed as a separate VNF, enabling significant OpEx cost reductions associated with PDN-GW protocol

signature upgrades. This is because PDN-GW will require fewer upgrades, no longer needing bi-weekly protocol signature updates or the like. In addition, the PDN-GW function will have performance benefits without the burden of DPI processing that could be considerable when used intensively to generate more services revenue from mobile tiered pricings, application charging, etc.

Another way DPI can help monetize the network is by helping to comply with specific regulations, public advocacy requests, or government laws by detecting questionable applications and web sites, thus enabling the PDN-GW to take appropriate action. In such cases, the service providers may be able to upsell mobile services and increase average revenue per user (ARPU).

The integration of DPI can also help telcos curtail increasing CapEx and OpEx costs by enabling them to automate their network scalability based on traffic patterns and telemetry services. This supports the industry movement driven by the ETSI NFV ISG to dynamically scale up and out both the signalling (MME) and user plane (SGW, PGW). This is providing interesting new possibilities to lower their OpEx costs while controlling their CapEx thanks to the relatively low cost of the underlying commercial off-the-shelf (COTS) resources.

Challenges to Overcome

VNF deployment: Separating DPI and gateway functions, and managing DPI within a SDN framework.

DPI is playing an important role in new standards under consideration, such as the Traffic Detection Function (TDF) in 3GPP, and separating DPI from the gateway function will provide telcos with much more flexibility going forward.

Network elasticity: Scaling VNF resources.

In order to maximize the elasticity of network resources, it is important to be able to scale network resources assigned to a virtualised network function both up and down.

Efficient dynamic scaling of virtual resources is essential for achieving the elasticity required from a virtualised telecom network. Scalability will also enable the flexible allocation of different network functions based on actual traffic patterns, thereby reducing the need to overprovision network infrastructure.

Performance bottlenecks: Achieving performance on standard server hardware and minimizing the overhead inherent in virtualised systems.

A challenge for this use case is to dynamically scale the DPI instances in relation to the virtual PDN-GW without dropping packets and ensuring existing DPI instances continue to properly classify all the traffic. One hurdle is overcoming the associated performance bottlenecks, especially those related to inter-VM communications.

To enable high performance transport, Tieto IP stack (TIP) is used together with the Intel DPDK.

Addressing the Challenges

VNF deployment

Figure 2 depicts the use case where the virtual DPI sends some monitoring information to the NFVI/SDN managers to optimize resource deployment.

In Figure 2a, the virtual DPI function provides metadata information to the management system in the SDN orchestrator:

1. DPI is deployed onto the NFV Infrastructure (Tieto IP Stack + Intel DPDK + Intel DPDK Accelerated OpenvSwitch), and it monitors the traffic using a promiscuous mode (SPAN).
2. The management system receives monitoring information computed by the vDPI, and correlates it with information received from the OSS and CDR/IPDRs.
3. The management system processes the data.




Figure 2a. Extracted Information Triggers a Network Event

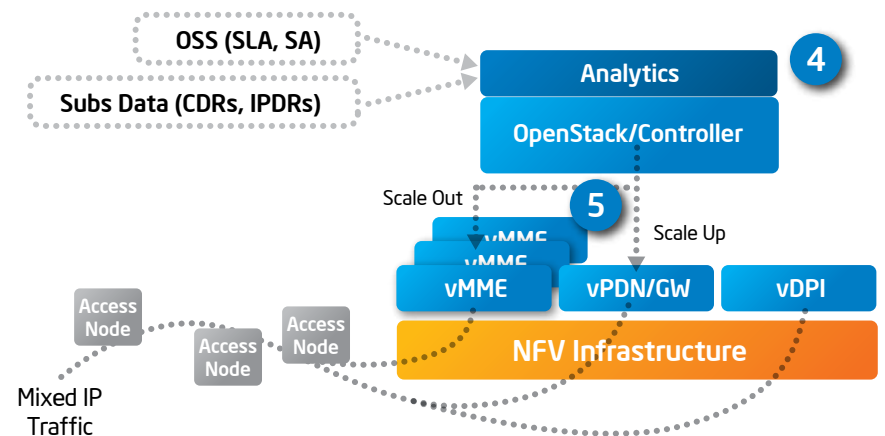


Figure 2b. Extracted Information Triggers a Network Event

In Figure 2b, the data processing triggers a network event:

4. The management system processes all the data, and based on the results, adapts the network topology and instances.
5. Based on the result, the management system triggers an update in the access node network paths and scales up and out the VMs accordingly.

The end result is a DPI VNF(C) becomes a service in the NFV infrastructure that provides network traffic information to the orchestrator.

Network Elasticity

The reference design employs cloud orchestration to automate the management of virtual resources using a high-availability framework, abstracted service availability, automatic scaling of network functions, and programmable APIs such as OpenFlow. Orchestration is implemented according to the NFV framework, which requires support for dynamic and automatic resource allocation.

The solution provides the means to satisfy service availability requirements by efficiently scaling virtual network functions – up and down – based on the load on the system. When a predefined

threshold is reached, automatic scaling is initiated. Additional CPU cores and VMs are commissioned as traffic volume increases, and are later scaled back as traffic decreases.

▪ **Cloud management**

The reference architecture is built on the Tieto Cloud platform reference design, which:

- Enables faster time-to-market (TTM) of cloud-based solutions by accelerating product development.
- Delivers solution accelerators, including the means to manage any OpenStack-based cloud.
- Solves telecom-specific problems, like automated deployment, monitoring, and recovery of very complex applications.

▪ **VM auto-scaling**

This feature automatically scales VM resources according to system load, creating new VM instances (similar to the existing ones) when the load exceeds a predefined limit, and retiring VMs when the load decreases substantially.

▪ **Load balancing and control signaling**

Many 3GPP and IETF defined control signaling protocols, such as Diameter, are by nature point to point protocols. The dynamic use of resources and the ability to scale/downscale VNFs must be performed without impact on control signaling between EPC/IMS nodes, as in scaling the HSS shall not require reconfiguration changes of the MME. Some 3GPP interfaces (e.g., Diameter Gx and Gy) have user sessions, which are not only identified by the point to point connection, but also with logical information in the Diameter protocol. When scaling a VNF, it must be assured that user sessions are also maintained.

Diameter-enabled VNF nodes require Diameter Signaling Controllers (DSC) that take maximal advantage of

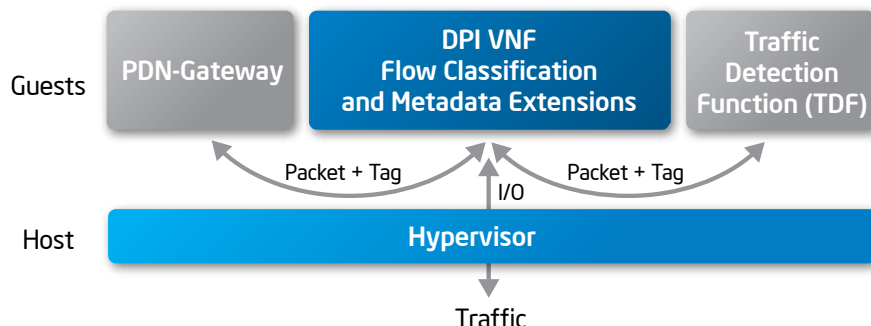


Figure 3. DPI Implemented as a Virtual Network Function

protocol capabilities to help optimize the network, perform overload control and load balancing, and implement topology hiding to protect VNFs, thus allowing VNFs to scale in a controlled way.

In the reference architecture, load balancing is explored from two different angles: a virtualised Tieto Diameter Signaling Controller ensures controlled Diameter signaling distribution in the cloud environment while Tieto IP stack provides high-performance IP load balancing.

▪ **Virtual DPI scaling**

The reference architecture decouples the virtual PDN-GW and virtual DPI, which simplifies the mechanism for scaling DPI instances. As explained previously, operators offering tiered pricing, freemium plans, and customized Apps may rely on the DPI processing to classify and recognize specific applications used by mobile plan subscribers. The ability to scale DPI capacity without impacting the virtual PDN-GW can lower security risks while simplifying the networking management of the virtual PDN-GW layer.

Each virtual DPI instance may be configured to extract metadata from a subset of a wide range of applications. For instance, operators wishing to classify 100 specific apps at 40 Gbps and within a particular latency window will need the orchestrator to provide the corresponding VM with enough virtual CPUs and RAM to meet performance requirements.

Figure 3 shows a telecom infrastructure platform with DPI implemented as a VNF component seamlessly interfaced with a virtual PDN-Gateway and a traffic detection function (TDF). Operators can scale DPI capacity using orchestration to launch and retire virtual DPI instances, as needed, to support business offerings like tiered mobile plans, application charging, freemium and premium plans, etc.

Performance bottlenecks

One way to avoid performance bottlenecks due to delays caused by inter-VM communications is to deploy DPI in front of incoming traffic (i.e., fast path) as an active node with an integrated tagging mechanism.⁶ Once the traffic/flows are classified, the DPI VM tags the packet headers to convey the protocol and application-ID to the other nodes, like the PDN-gateway. This approach eliminates the need for VM to VM (networking node to DPI VM) communications that can create delay and require significant processing and bandwidth. When required, the orchestration stack will scale up or out the DPI VM instance independently of the other networking VMs.

Another advantage of this approach is it allows flow management to be integrated at the DPI VM level, thus allowing the other networking nodes (PDN-GW, TDF, etc.) to be stateless, which simplifies their design.

The Tieto IP (TIP) Stack offers best-in-class transport and packet processing solution designed to ensure high-performance networks and reduce time

to market for equipment manufacturers. Well-suited for both slow and fast path designs, the solution includes a Linux user space IP stack that is fully-featured and telecom-grade, and is integrated with the Intel DPDK. Applications can use either the defacto socket API or an optimized zero-copy API for improved performance.

Virtualisation of the Mobile Core Network (EPC) and IMS

For this use case, the ETSI NFV specification provides a description and a list of high-level challenges related to the virtualisation of the mobile packet core (i.e., the core network of an LTE system) and IP Multimedia Subsystems (IMS), an architectural framework for delivering IP multimedia services.

Evolved Packet Core (EPC) Overview

The reference architecture demonstrates how mobile core network functions can be virtualised on standard Intel architecture servers using open standards while achieving carrier-grade service quality and capabilities. Network functions (MME, P-GW, S-GW, and DSC) are deployed on the NFV infrastructure (Figure 4). The main solution ingredients addressing the key challenges are described in more detail in the following.

Business Benefits

Virtualisation of the mobile core network and IMS enables several advantages for service providers. Among the most important are significant TCO savings,

efficient allocation of network resources, and increased speed of innovation via efficient deployment of new services. When increasing capacity or deploying new services, operators can add software-based services to an existing virtualised environment, instead of having to install a new network node. This will save time and money. Virtualisation also creates OpEx reduction opportunities, such as lowering power consumption during off-peak times using power management features to power down unneeded portions of the hardware platform.

Challenges to Overcome

Service awareness: Allocating resources to network functions based on the services they support.

Virtual resource monitoring and orchestration, along with service awareness, are essential for implementing elasticity effectively.

Service availability: Achieving the same level of service availability for an end-to-end virtualised mobile core network as in non-virtualised networks, but with reduced cost.

Due to the nature of telecom networks, service availability will be a key issue for a virtualised mobile core network. Since virtualisation usually leads to a performance trade-off, equipment developers must take measures to optimize data plane processing in order to satisfy carrier-grade bandwidth and latency requirements. Similarly, sufficient

control plane performance is needed to deliver an acceptable level of service and ensure the availability of regulatory services, such as emergency calls.

Addressing the Challenges

Service awareness

The reference architecture analyzes packet flows and facilitates intelligent traffic steering and service invocation by integrating DPI content inspection software into the virtualised infrastructure. The DPI engine, which scales cost-effectively across large deployments, is a powerful tool for creating service awareness and ensuring service availability. The reference architecture also demonstrates how DPI content inspection enables performance optimization in the core network by making it service aware.

Service availability

The use of standard server hardware and virtualisation technologies may introduce performance penalties (i.e., interrupt servicing and context switching) that in turn could affect the level of service availability. Therefore, maximizing performance and minimizing latency are critical success factors, which can be addressed with available technologies such as intelligent load balancing, DPI, and high-speed packet processing. In the case of the reference architecture, user plane performance in virtual LTE gateways and eNodeBs is improved using the Tieto high-performance IP Stack, SR-IOV, and the Intel DPDK.

Failure scenarios and high availability

A telecom cloud has more stringent service availability requirements than an IT cloud, which can be addressed by virtualised infrastructure capabilities, such as placing VMs in high-availability configurations (2N or N+1) and pinning VMs to specific CPU cores. The reference architecture was also designed to avoid single points of failure, supported by the OpenStack-based cloud infrastructure.




Figure 4. Virtualised Base Station and Evolved Packet Core (vEPC)

▪ Carrier-grade VM

Virtual machine (VM) failover allows a system to recover from a poorly functioning VM by moving its workload to a standby VM, which contains a clean software image and is ready to be activated. This feature takes advantage of load balancing mechanisms for diverting workloads, but it also ensures the data stays current for both the active VM and its standby VM.

Failover solutions must also identify VM state changes and trigger policy actions in the system. A change in VM state could be spontaneous, as when a VM is killed, or planned, as when the cloud administrator manually changes it.

Each virtualised application typically has a template that defines which state changes are relevant for that application and the policy actions to be taken for each defined state change. Currently, the policy actions supported by the reference architecture are restart the virtual machine; apply the scalability policy to start a new VM for up-scaling; and stop a VM for downscaling. The reference architecture supports several failure scenarios, including VM, host, and network failure. It is also Service Availability Forum compliant (see www.saforum.org) such that middleware (e.g., OpenSAF) can run within the VM to protect the application within the VM.

Virtualisation of the Mobile Base Station (in the RAN)

Containing a large number of nodes, the radio access network (RAN) constitutes a major share of the TCO and energy consumption in mobile networks. In large mobile operators' networks, multiple RAN nodes from various vendors are usually operated with different mobile network systems (e.g., 3G, LTE, and WiMAX*) in the same area. It is possible to consolidate these platforms into a virtualised base station (BS) using IT virtualisation technologies.⁵




Figure 5. Pool of Virtual Base Stations and Remote Radio Heads

Today, the NFV vision is being applied to the RAN, yielding benefits from virtualisation; however, there are some challenges that still need to be addressed to enable a virtualised RAN.

Business Benefits

A node in the RAN traditionally serves a limited geographical area and is typically provisioned for the maximum usage expected in that area, even though the average load is usually far less. As a result, RAN nodes are often underutilized, and today a RAN node cannot share its resources with other nodes.

The pooling of base station nodes (Figure 5) over a large geographical area allows network operators to more precisely dimension capacity and avoid bottlenecks in an area served by multiple nodes, thus improving equipment utilization by sharing resources in the pool.

Base station (BS) virtualisation can achieve resource sharing among multiple logical RAN nodes from different systems, dynamically allocating the resource as well as reducing power consumption. BS is a generic term referring to 2G BS, 3G Node B, and 4G eNodeB.⁵

Virtualisation of a base station node allows functionality to run on standard servers, storages, and switches, enabling resource sharing with other network functions, including third-party network functions. Virtualisation creates a competitive environment for vendors supplying innovative network functions that unlock the traditional proprietary boundaries of RAN nodes.

In some cases, it may be possible to co-locate virtual base stations and virtualised evolved packet cores (vEPCs) in the data centre, enabling operations, administration, and management (OAM) savings. The feasibility of this deployment approach may hinge on front-haul network lengths and the associated propagation delays.

A virtualised server platform close to the edge can host services, such as content caching. Not only does this provide an opportunity for network operators to deploy revenue-generating services at the access edge, users will also benefit from improved performance.

Challenges to Overcome

Challenges that need to be addressed include:

Real-time performance: Upholding the strict requirements of delay and jitter.

Virtualisation of the base station introduces new sources of delay and jitter that must be minimized in order to maintain adequate service in the radio access network. These sources of delay and jitter need to be included in the total delay budget, from the base station to the radio interface, together with other sources not related to virtualisation, such as front-haul delays. Note that the front-haul challenges related to centralized deployment are not addressed in this paper.

Network elasticity: Optimizing the utilization of network resources by scaling the capacity of the virtualised base station in run-time.

The virtualised base station must scale to handle the network load and respond to changing demands of service and capacity. At the same time, it needs to minimize the power consumption of the radio access network. Even when the load is low and the virtualised base station is consolidated to a few virtual machines, capacity needs must be met to ensure service fulfilment.

Network management: Decoupling management of infrastructure and network services in a multi-vendor environment with both physical network functions and virtualised network functions with geographical and topology limitations.

The base station includes both virtualised and physical components, such as physical remote radio units, which are tightly interconnected, limiting the flexibility of deployment. The base station is part of a dynamic network in a constantly-changing environment with VM instances that come and go based on load. Hardware

and software is decoupled, which further introduces challenges in management of the network. In addition, the network operation needs to be able to handle the coexistence of legacy base stations with virtualised base stations that must be combined with the large number of existing nodes in the RAN.

Addressing the Challenges

Real-time performance

It is possible to achieve near-real-time performance in virtualised environments when several main issues are addressed. Foremost, it is necessary to minimize the interrupt latency and the overhead associated with virtualised, standard servers. A major source of performance loss is from VM enters and exits, which typically occur when the hypervisor must service an interrupt or handle a special event. These transitions are expensive operations because execution contexts must be saved and retrieved, and during this time the VM is stalled.

The overhead from VM enters and exits can become substantial since it is not uncommon for I/O-intensive applications, such as base stations, to have hundreds or thousands of interrupts arriving in a second. Similarly, a KVM guest may need to take thousands of VM exits per second because of the internal timer interrupt. These constant disruptions cannot be tolerated with communications applications because of the resulting degradation in performance, latency, and determinism.

The use of standard server hardware and virtualisation technologies will typically introduce performance penalties (i.e., interrupt servicing and context switching) that in turn could affect the level of service availability. The reference architecture mitigates these penalties in a number of ways, including:

- 1) The Tieto high-performance IP (TIP) stack implements SR-IOV (Single Root I/O Virtualisation) in order to eliminate the hypervisor's involvement in data movement by providing independent memory space, interrupts, and DMA streams for each VM.




Figure 6. Latency distribution of GTP-U traffic when testing 1) without power management in a purely polled setup and 2) in a setup with power management and the one-shot interrupt feature (based on an Intel® Xeon® Processor E5-2600 @ 3.4GHz).

- 2) The Intel DPDK enables Intel processor cores to process packets continuously without being interrupted by the operating system, other applications, or interrupts, which greatly increases performance and determinism.
- 3) Core pinning guarantees that a particular flow, as identified by a five-tuple (e.g., IP address, port and protocol type) or some other predetermined criteria, is always sent to the same VM for processing. This eliminates the need for sharing connection and forwarding information among VMs because each VM only needs to know about its own connections.

These approaches, along with others, help the TIP stack to process network traffic with 15 to 20 microsecond latency. Figure 6 models the round trip time (RTT) from an external virtualised application to the virtualised base station application and back.

Network elasticity

The reference architecture addresses the challenge of network elasticity in the RAN, enabling run-time scaling of the base station from a few virtual machines to utilizing the complete virtualised pool, based on the load within the covered area.

The virtualised base station's different components are deployed in virtual machines implemented with a high-availability framework and automatic scaling. It uses cloud orchestration to automate the management of virtual resources. The reference architecture

is implemented according to the NFV framework, and supports the required dynamic and automatic resource allocation.

The solution provides the means to satisfy elasticity requirements by efficiently scaling number of VMs (up/down) in the virtualised base station based on the load in the covered area. When a predefined threshold is reached, automatic scaling is initiated. Additional CPU cores and VMs are commissioned as traffic volume increases, and are later scaled back as traffic decreases. Off-loading cells are automatically commissioned as traffic volume increases in hot-spots. These off-loading cells are deactivated as traffic decreases in the hot-spot areas.

Network management

The reference architecture integrates Tieto's O&M framework for embedded O&M in network functions, VNF managers, and element management systems. The O&M framework provides a flexible solution that is easy to maintain and extend, making it possible to tailor to specific needs and requirements. The embedded O&M framework easily scales network functions to meet a wide range of needs, from systems with a small memory foot-print to high-performance EMS. The reference architecture addresses SDN/NFV concepts and interfaces, for example decoupling of the infrastructure and the network functions as well as integration with legacy systems.

Summary and Conclusions

This paper describes how the reference architecture developed by Intel, Qosmos, and Tieto facilitates the deployment of NFV in telecoms networks and creates opportunities for telecom equipment manufacturers to fulfil the vision for mobile service providers. Designed in accordance with ETSI NFV ISG principles, the reference architecture addresses key barriers to NFV adoption, and provides the flexibility and inherent intelligence needed to unleash innovation and rapid service development and monetization.

The industry is in the midst of transformational change, and industry players who are proactively embracing these changes will likely be the ones to reap the benefits in the future. A collaborative approach between service providers, network equipment providers, silicon vendors, and independent software vendors (ISVs) is vital for accelerating telecoms network evolution, where the benefits described in this paper can be realised by all.

For more information about Tieto, visit www.tieto.com/pds. For information about Tieto components used in the reference architecture, visit www.tieto.com/signaling and www.tieto.com/tip.

For more information about Intel® solutions for Telecoms, visit www.intel.com/go/commsinfrastructure.

For more information about Qosmos, visit www.qosmos.com.

¹ See previous Intel and Tieto work at www.intel.com/go/commsinfrastructure and www.tieto.com/product-development/communications-infrastructure/tieto-carrier-cloud-network-equipment-providers.

² For more detailed information about the Intel® Data Plane Development Kit (Intel® DPDK), see: www.intel.com/go/dpdk.

³ Operations Support System (OSS) and Business Support System (BSS).

⁴ See: <http://www.etsi.org/technologies-clusters/technologies/nfv/nfv-poc>.

⁵ ETSI GS NFV 001 V1.1.1 Network Functions Virtualisation (NFV): Use Cases www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf.

⁶ Several initiatives are being drafted in the IETF Service Function Chaining (SFC) working group.

Copyright © 2014 Intel Corporation. All rights reserved. Intel, the Intel logo and Intel Core are trademarks of Intel Corporation in the U.S. and other countries.

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

