Network Functions Virtualisations Initiatives in Carrier Networks

: Orchid One

Table of Contents

Table of Figures

The intent of this document is to provide a look at the Software Defined Networking and Network Functions Virtualisations and its applicability in the carrier environment. These two concepts are independent of each other and each has the potential to streamline the cost to performance ratio for the service provider. Let´s look at each of the following individually and their relevance to the carrier community.

The concept of virtualisation was a response to the IT/datacenter industry where there were a large number of servers running specific applications. Imagine an Amazon or eBay or Facebook datacenters where there is a sprawl of servers running application to transact with users. These servers run in spite of there being no transaction traffic going on. This resulted in several issues:

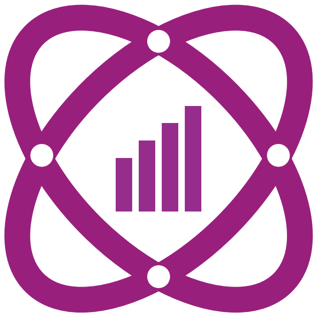

- Higher rack and real-estate foot print
- Higher power consumption
- Wastage of power

The virtualisation concept enabled the transition from one-application one-server model to many-applications running over virtualised server farm architecture. This concept of virtualisation was leveraged in a seminal paper in 2012 by a group of Tier-1 operators to investigate the application of this virtualised architecture into the CSP network, which was called Network Functions Virtualisation.

Operator networks have numerous infrastructures such as security devices (NAT/Firewalls), transactional/storage application DBs (HSS, HLR, PCRF etc.), signalling routers for dipping into these transactional DBs, edge devices such as BRAS, application infrastructures such as CDNs, DPIs etc. These infrastructures run on proprietary platforms and a major cost point for operators, both on CAPEX and OPEX.

The Goal of the operators via NFV

- Convert the above infrastructures into software applications
- Transition the proprietary hardware into a farm of X-86 processors
- Virtualise the applications in point 1 above into Virtual Network Functions (VNFs)

With the innovation in compute platforms (e.g. Intel x-86) where there is a lot more power in terms of increasing number of cores, and virtualisation, these platforms will be able to scale very efficiently to host more and more applications.

This should, in theory, reduce the complexity of the infrastructure, reduce the cost, scale efficiently and lend itself to easier OAM. However, in practice, this will be a very long process as it changes the very culture of the operations folks, requiring them to learn a new way of managing their infrastructure and induces a paradigm that demands a lot out of not only the operators, but also the standard bodies and multi-vendor co-operation to make this work.

There will be, foreseeably, several baby steps to be taken that should eventually lead to the goal. In this paper, we look at some of the challenges behind the NFV concept for service providers and role played by Cataleya in supporting the NFV initiatives of the Operators.

"There will be foreseeably, several baby steps to be taken that should eventually lead to the goal. In this paper, we look at some of the challenges behind the NFV concept for service providers and role played by Cataleya in supporting the NFV initiatives of the Operators."

Virtualisation in carrier networks

There are fundamentally two types of activities going on in the carrier networks:

- Control plane activities
- Data plane activities

1) Control Plane

The control plane activities are more oriented towards the transactional and signalling processing functions. The control plane functions provide a very ideal candidate for virtualisation. Typical control plane functions involve the following:

- Security functions - NAT/Firewalls
- Load balancers
- Transactional/storage processing - HLRs/HSS/PCRF etc.
- Signalling routers -e.g. SGSN/GGSN, DRA/DEA for Diameter etc.
- WAN accelerators
- Edge devices - BRAS etc.
- Application functions - IMS, DPI, CDN etc.

There is no media related activities in this plane and more of these applications are available in software that can run on general purpose off-the-shelf servers running Linux on X86 servers. Though scale and throughput are important in the control plane due to the sheer number of messaging transaction going on, the issue is not as critical as with media plane processing, which involve intensive processing that could impact the end subscriber's quality of experience.

2) Data plane

The data plane poses more challenges for virtualisation especially when it comes to real time applications such as voice, conversational video etc. Non real-time applications such as web browsing do not pose much of a challenge as seen with the virtualisation of the GGSN/SGSN etc.

Voice applications are impacted by network impairments. There are too many types of voice end points in the market today (fixed line, mobile, WebRTC etc.) and transcoding becomes a very essential function. The transcoding process is very CPU intensive and running transcoding on general purpose computing platforms poses challenges from a cost to performance perspective. This is one of the major challenges for the Orchid One platform, when it comes to handling the transcoding function.

The other challenge with NFV is that the hypervisor tends to add more processing delay in the packet processing and this will impair the subscriber's quality of experience. This area is being investigated for the performance metrics, to estimate the additional latency that will be induced by this layer and its impact on the overall voice quality.

Orchid One platform

The requirement behind the Orchid One platform is to have a sizeable scale to cater to the IPX markets. The platform handles 100K concurrent sessions with media and capable of delivering 15K transcoding sessions. On top of this, there is the requirement to do real time analytics, end-to-end QoS/QoE etc., that needs a lot of computing resources for data collection and analysis. This is the reason why the hardware based Network Processor Unit (NPU) was chosen.

The following diagram shows the high level architecture of the Orchid One platform.

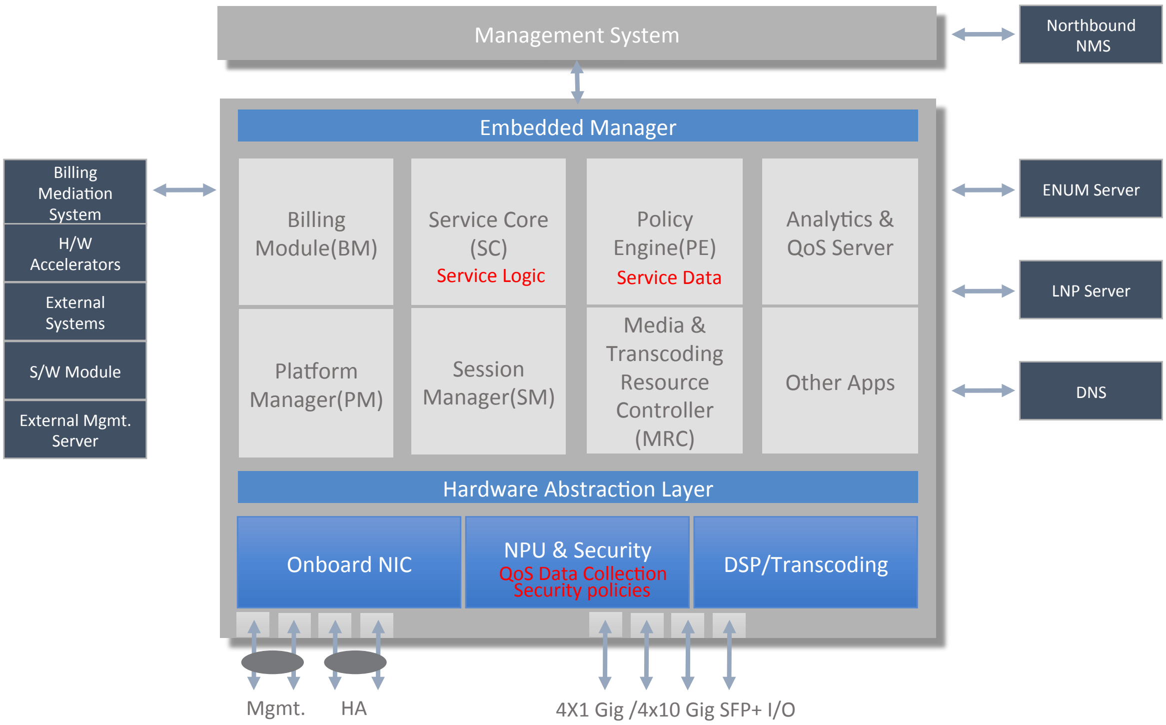

Figure1 High level architecture of the Orchid One platform

The NPU packs an abundant amount of processing power needed for packet manipulation, processing and QoS data collection. There is the Intel's DPDK (data plane development kit), a set of libraries/drivers for fast packet processing that runs on general purpose Intel platforms. The drawback is that the DPDK is optimised for data packets and not for voice RTP. RTP packets are very small (minimum 64 bytes) compared with the 1500 byte data packets and the DPDK is not optimised to handle the smaller RTP packets. This is one of the areas of innovation that the industry is looking at and when these developments are in place, we should be able to see DPDK handle significantly more voice traffic.

Also, in order to get the 15K transcoding sessions, the hardware based transcoding chips are used. Intel architectures today would not prove cost effective to achieve this scale of transcoding needed. There could be future innovation from Intel that results in less generalised, more specialised architectures that is capable of better performance for a narrower range of problems such as transcoding.

> This is one of the areas of innovation that the industry is looking at and when these developments are in place, we should be able to see data plane development kit (DPDK) handle significantly more voice traffic.

Steps to Orchid One Virtualisation

The first step towards virtualisation is to create a software version of the Orchid One platform. This will entail using the DPDK. The software version of the SBC will provide a lower scale per instance and these numbers are being verified. Once the software version is created, the next step is to bring in the virtualisation and verify the scale numbers against this.

The following diagram shows the software version of the SBC in a non-virtualised environment.

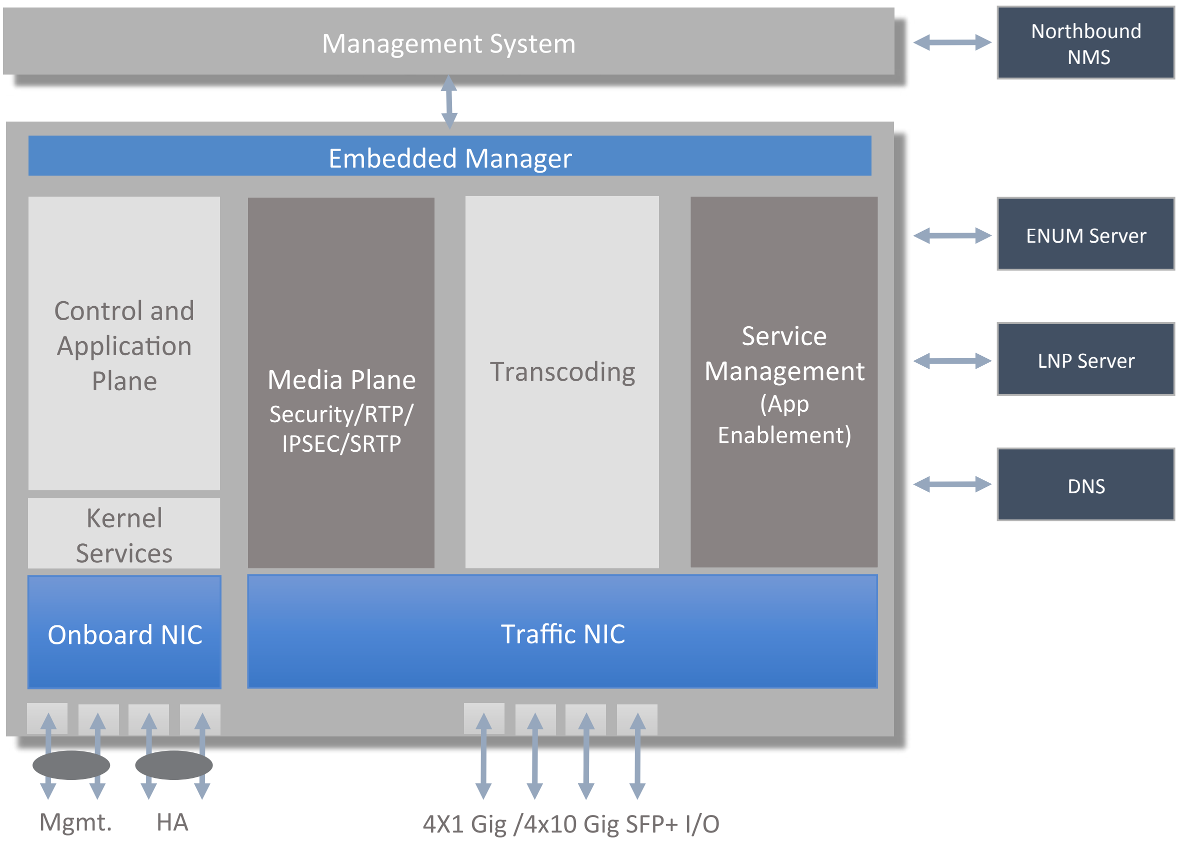

Figure2 Orchid One in a Software version : Non-Virtualised

The above diagram shows the NPU and DSP cards replaced with the Intel traffic NIC that will handle the traffic and transcoding.

Once the virtualisation is in place, it would be possible to run multiple instances of the SBCs on a hypervisor as shown in the diagram below. The service management module provides the service chaining logic via the "service building blocks". This provides the flexibility to create new services by chaining virtual network functions. The service management module may be virtualised on its own thereby providing the service logic feature to all the SBC instances.

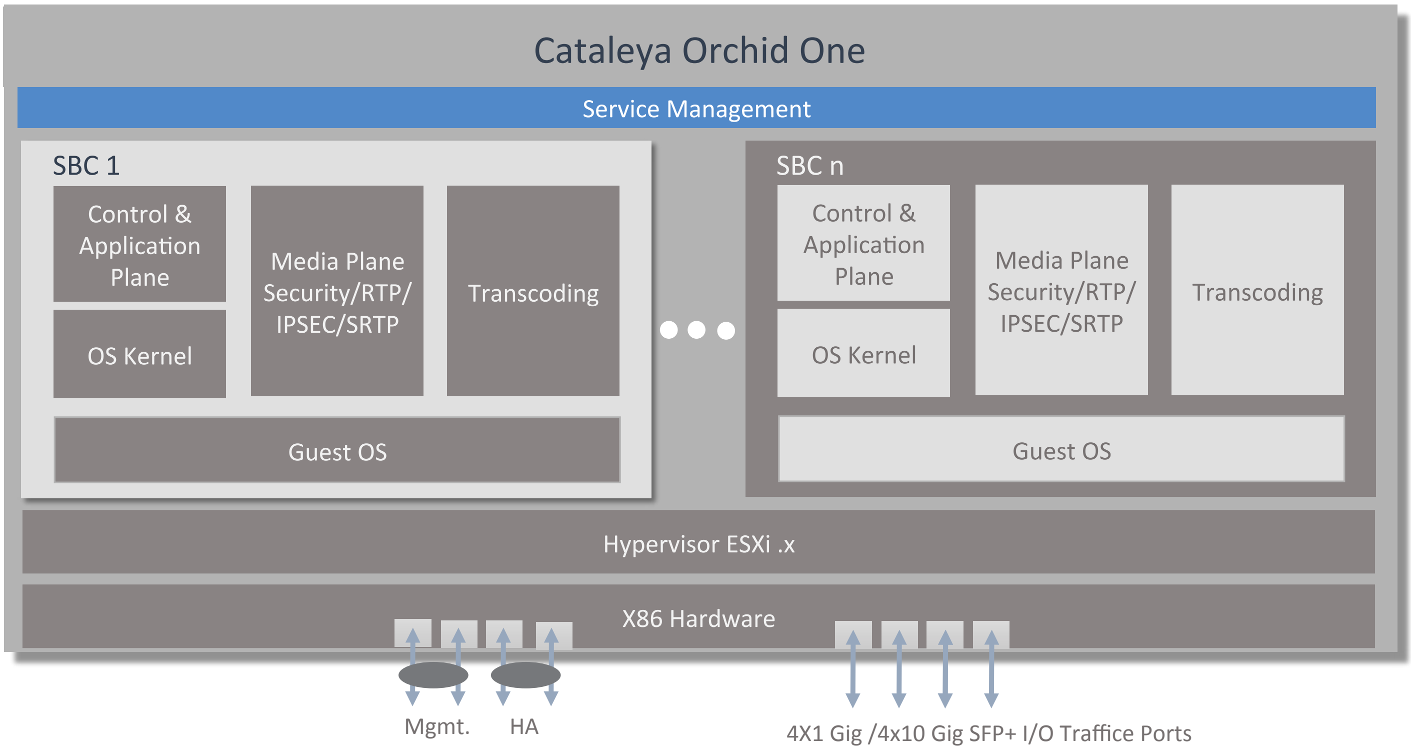

Figure3 Orchid One instances in a virtualized environment

NFV Service chain

It is perceivable that operators deploying NFV will look at implementing one or more virtualised network functions (VNF). There is significantly more advantage for the operator in deploying several virtualised network functions and provide the ability to service chain these VNFs to provide more complex revenue generating services. This is where we are focusing on the service management capability that can be virtualised and interconnect multiple virtualised services in a controlled, secure, reliable and scalable manner.

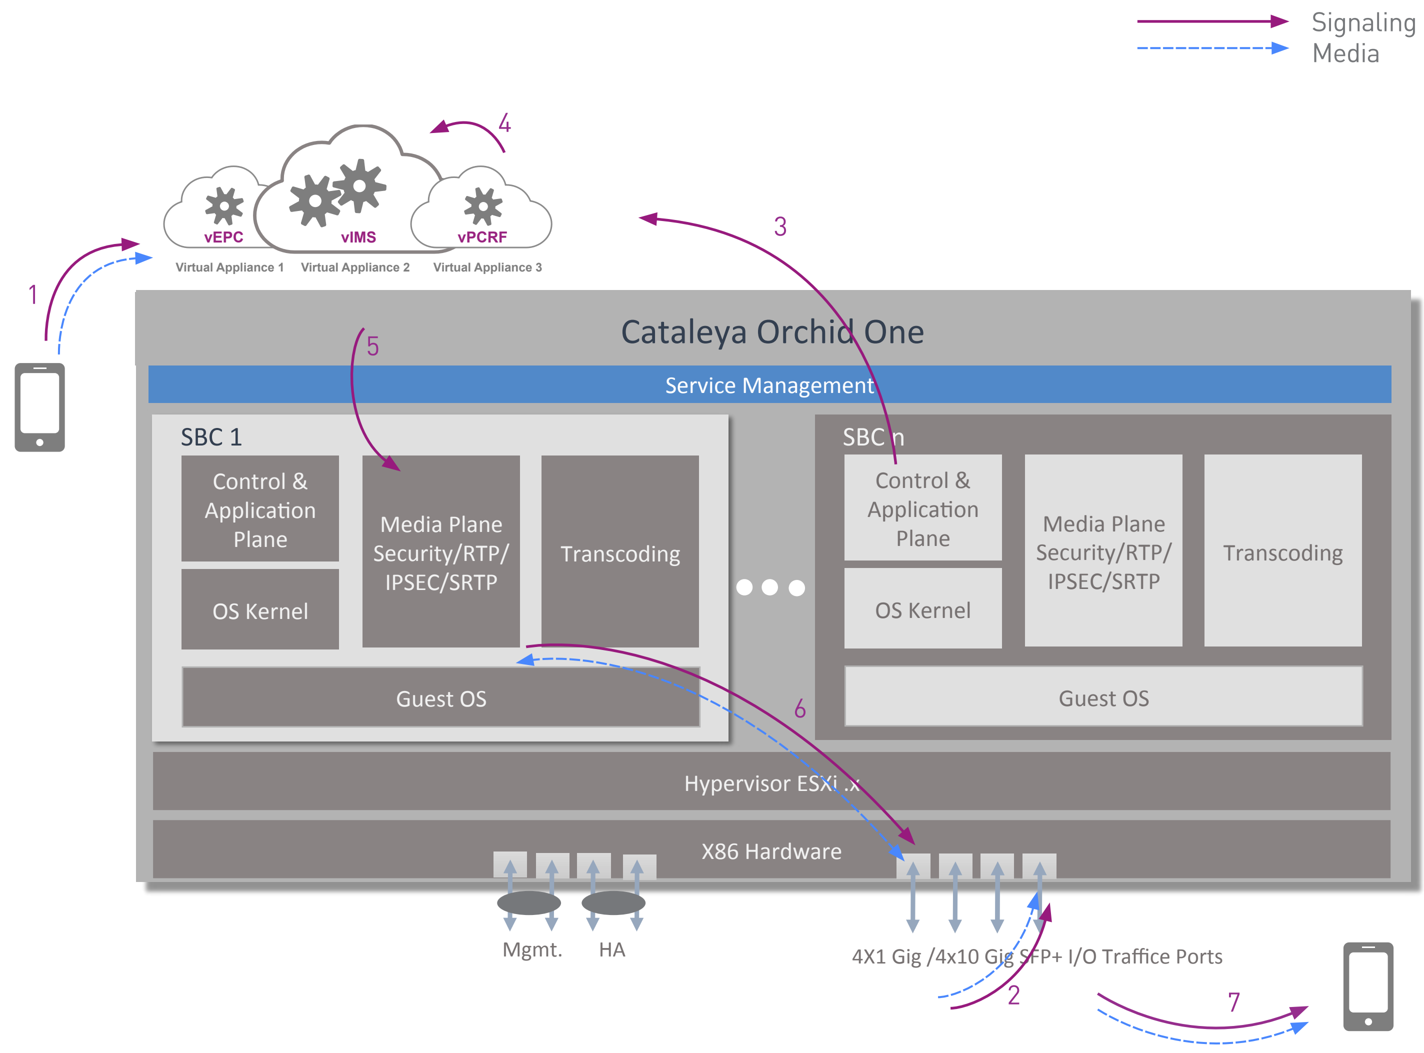

Figure4 NFV service chaining

The above diagram depicts the use of the service management of Orchid One to orchestrate the service chaining in order to deliver a service from the virtual EPC network function to a handset (Phone 2) while invoking the services of the virtual PCRF and the virtual IMS network functions. The significant advantages of virtualising the network function will lie in the ability to service chain these VNFs into meaningful revenue generating services that can be delivered reliably, securely and in a scalable manner.

VNF SERVICE CHAINING -SIGNALING

1-3. vEPC to Orchid One
4. Orchid One Service Mgmt to vPCRF
5. Orchid One Service Mgmt to vIMS
6-7. Orchid One Service Mgmt to Phone2

Challenges with NFV

It is very intuitive to see that the NFV could deliver significant cost savings to the operator, but this does present several challenges that are listed below:

- Change in the paradigm, which the operators are used to, for provisioning, upgrading and managing network resources.

- A new breed of management or OAM tools will need to come into place to operate, administer and monitor the virtual network functions (VNFs).

- Operators, standards bodies and vendors related to a particular network function need to create consensus and have an open API.

- Service chaining across VNFs means that there should be consensus among various VNF vendors to expose APIs to each other to enable creation of services.

- Migration timetables will tend to be too long and since this is a new paradigm, the solutions are bound to go thorugh several stability issues before the bugs get ironed out.

- There may not be an immediate cost saving to the operator - in fact, even an increase in cost initially - but in the long run this will justify the cost to performance ratio.

- Real time applications such as voice still have several issues that prevent them from being a perfect candidate for virtualisation and hopefully these innovations/ optimisations can come sooner than later.

- Service chaining implies that the relevant VNFs that can be chained must be certified in order to disallow any service chaining that may lead to infinite combinations that may not all work.

> “The significant advantages of virtualising the network function will lie in the ability to service chain these VNFs into meaningful revenue generating services that can be delivered reliably, securely and in a scalable manner.”

Application of the Orchid One virtualised solution

The Virtualised Orchid One sees several applications in the following areas:

- Virtual EPC environments - here the virtualised Orchid One can serve as a Diameter and SIP signaling router between the EPC to the IMS environment and to the PDN.
- Interconnect peering solutions - here the virtualised solution will provide scale by creating several instances of the SBCs.
- Enterprise SIP trunking solutions - here the virtual Orchid One can provide scale in the aggregated site as well as provide a cloud solution for the customer premise SBC requirements. The fact that the Orchid One can provide end-to-end Qos/QoE from the customer premise to the aggregated site to the peering side would provide significant advantages to the enterprise customers.
- Service chaining virtual network functions into delivering new services with faster time to revenue.

About Cataleya

Cataleya is a leader in IP networking innovation, with a strong track record in developing and deploying next generation carrier grade switching systems, pushing the envelope in an all IP paradigm.

Cataleya is headquartered in Singapore with its own technology development team in Silicon Valley and a wholly owned subsidiary of Epsilon Global Communications.

Cataleya is another outstanding result of Epsilon's innovation DNA and reflects a strong service provider influence in the design and functionality of its technology. A new approach to new challenges has led to a product of unparalleled performance, simplicity to operate and reduced cost of ownership.

cataleya
let your network blossom

General Contact Details:

Singapore

New Tech Park #06-02
151 Lorong Chuan
Singapore 556741
Telephone: +65 3106 4020

USA

1900 McCarthy Blvd, #412
Milpitas CA 95035

Telephone: +1 408 571 2200