WHITE PAPER

Service Providers
NFV

Harmonizing IT and Network Infrastructure to Enable the Telco Cloud

Communications Service Providers (CoSPs) have used computing servers in separate networking and IT organizations for many years. By consolidating system architectures and operational practices, they can enable a hybrid CoSP cloud

Table of Contents

Executive Overview

Communications Service Providers (CoSPs) are using virtualization to improve service agility in the network. Because Virtual Network Functions (VNFs) run on similar infrastructure to IT workloads, there is the potential for a hybrid platform to serve them both. The CoSP cloud, also known as telco cloud would enable resources to be flexibly reallocated between the network and IT applications, depending on the network traffic level.

Although the infrastructure is similar in the IT and network domains, the operational approach is different. In the network, software changes are infrequent and meticulously planned. In IT, DevOps practices are increasingly enabling more frequent updates to live systems. To enable the CoSP cloud, the network infrastructure platform needs to evolve towards using microservices in a software-defined data center, so that network services can be updated in the same way that IT systems are.

Ultimately, having greater agility in the network will enable new revenue opportunities, with the potential to more easily launch niche services, and applications for edge computing and 5G.

Background

As CoSPs prepare for 5G, they are increasingly virtualizing network functions with software to enable agile service innovation.

In a Network Functions Virtualization (NFV) strategy, the computing servers used for Network Call Control are based on the same standard hardware as used in the IT data centers for Operational Support Systems (OSS) and Business Support Systems (BSS).

As shown in Figure 1, the NFV infrastructure, as defined by the European Telecommunications Standards Institute (ETSI), exposes application environment and management interfaces which are aligned with the cloud Infrastructure-as-as-service (IaaS) model. As the IT organization adopts cloud practices, the staff will become familiar with the tools and Service Level Agreements (SLA) available. These could be reused by the IT division to provide IaaS facilities to the network division for hosting VNFs. The IT department's skills and software licenses could also be extended to support NFV requirements with marginal cost implications.

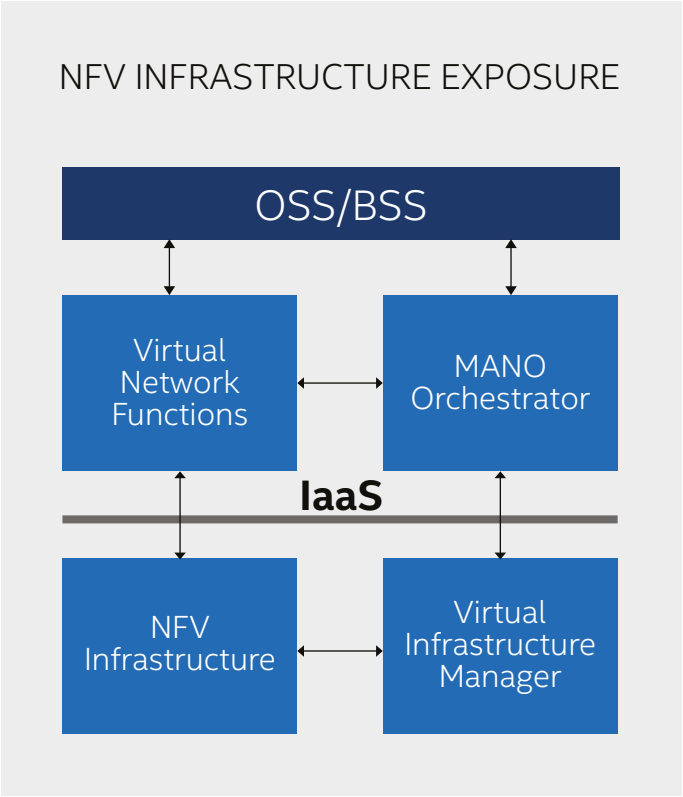

Figure 1. The NFV infrastructure and how it relates to IaaS

Intel has worked with CoSPs to define optimal server configurations to support NFV workloads in the data center. As a result of the Central Office Re-architected as a Data Center (CORD) initiative, in which Intel is a collaborator, appropriate server designs are now being manufactured using Open Compute Project modules. Further downstream in the CoSP access network, servers are also being packaged to suit street cabinets, mobile base stations, and customer premises equipment. All these network node types can present a consistent IaaS application environment to host VNFs, with workloads being placed appropriately by cloud infrastructure management tools and NFV orchestration systems.

Intel also publishes the Intel® Select Solutions for NFVI performance benchmark for a commercial off-the-shelf (COTS) server optimized for NFV workloads. Commercial server products are then validated against that benchmark, so that CoSPs can procure hardware from the competitive market with confidence that it will perform to the required level.

Ultimately, IT and network workloads could be consolidated into one CoSP cloud platform. This would enable resources to be flexibly reallocated between the network and the IT systems, depending on the time of day and network traffic level.

Despite similarities at the infrastructure level, the operational approach for the IT and network resources can be very different. The mission-critical nature of the network means any changes to the platform (such as major software upgrades) are planned extensively and undertaken infrequently, with extensive support from Network Equipment Providers (NEPs). For IT systems, there is more involvement of systems integrators and in-house developers.

To gain the best business benefit from NFV it is essential to blend the IT and network approaches to enable frequent incremental software changes in the network, as it adapts to become part of the CoSP cloud. Evolving the NFV technologies to align with IT technologies will involve several stages, as described in this paper. Ultimately, the use of microservices for NFV will enable new services to be delivered with lower risk and greater agility, more frequently. This will help to harmonize the operating culture of the network and IT domains, smoothing the transition to a hybrid CoSP cloud.

Evolution of the Network Infrastructure Platform

The evolution of network functions, from classic appliances to microservices, is expected to go through a number of stages, as shown in Figure 2.

Stage 1 - Classic Appliance

The classic appliance that has made up the network infrastructure for many years houses a network function inside a dedicated device, with the hardware and software supplied by the same NEP as an integrated unit. Because resources cannot be flexibly reallocated as they can with virtualization, the appliances must often be overprovisioned to meet peak capacity needs and to ensure continuity. Historically, a network appliance would be complemented by a second unit to provide backup resilience in the event of a device failure, regardless of the actual level of appliance utilization.

Additionally, CoSPs cannot benefit from the economies that come from buying standard hardware in a competitive market to run their network functions.

With NFV, one standard server can support several functions simultaneously, and one machine can provide backup for a range of different workloads.

Stage 2 - Software Appliance

Step one in the network evolution is to introduce software appliances, which are software versions of legacy network devices. These are hosted on virtualized COTS server platforms and can be based on Intel® Xeon® processors. This step in the evolution provides a first opportunity for network staff to become familiar with the virtualization

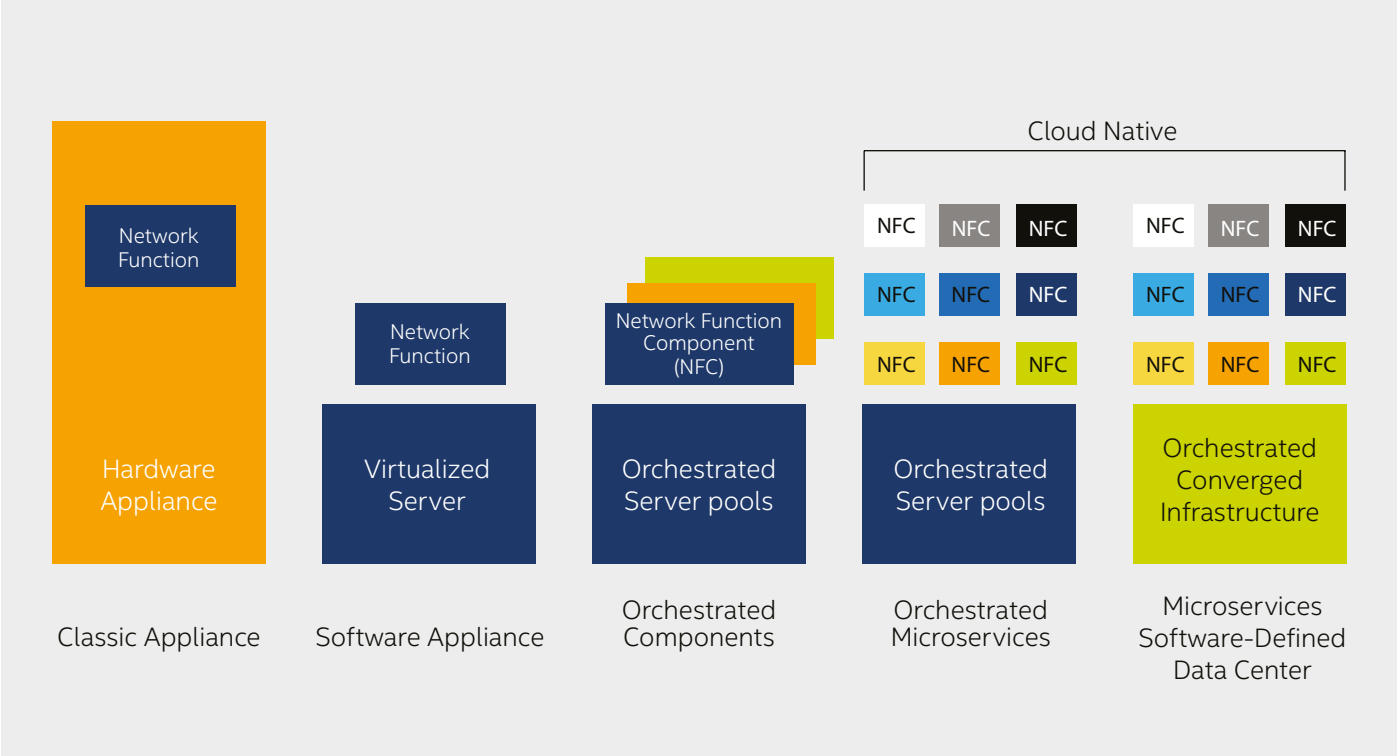

Figure 2. How network functions are expected to evolve

layer. Competitive procurement of infrastructure hardware becomes possible, while the operational staff can keep their familiar network management interface. Multiple workloads can share one server platform, making it easier to extend the functionality of the network node. The software appliance model simplifies field logistics with standard hardware deployment in the network and the downloading of software to remote locations.

For wireline CoSPs a popular use case for the software appliance is the Universal CPE (uCPE) server for Managed Enterprise Network Services in enterprise branch site locations. In this case, additional functions need to be deployed to handle Software Defined WAN (SD-WAN) connections that support resilient network connections to corporate data centers and external cloud services.

For wireless CoSPs, virtualized versions of the 3GPP IP Multimedia Subsystem (IMS) call servers needed for Voice over LTE (VoLTE) have been an attractive starting point to gain experience before moving on to re-engineer the Gi Services LAN in the Mobile Wireless Packet Core. It is expected to be much simpler to operate the Services LAN with software compared to using hardware appliances for individual network service functions such as deep packet inspection (DPI) traffic management, content processors, and firewalls. Compared to forecasting the future requirements for each hardware function, using software network functions reduces infrastructure planning significantly, with COTS servers available at short notice.

Additionally, use of Internet Engineering Task Force (IETF) RFC 8300 Network Service Headers enables simple traffic steering through a series of network functions to create a GiLAN Service. This removes the need for dedicated Ethernet switching and load balancing hardware between each function in the service chain. Network staff can change the service by simply updating the header properties, and without any need to reconfigure the underlying data center VLANs.

Stage 3 - Orchestrated Components

The next stage of infrastructure innovation uses orchestration to vary the IaaS resources used to suit the level of traffic and minimize overall power consumption. A good example is the Voice over IP Session Border Controller (VoIP SBC), which has two main components: the signaling controller and the data plane gateway. In a virtualized model, the controller and the gateway can be built as separate VNFs, and become independently scalable depending on the traffic needs, unlike integrated appliances. These VNFs would be hosted on pools of virtual machines, under the control of an NFV orchestrator. Software instances can be created or removed based on feedback received from the Element Management System (EMS) associated with the SBC.

Stage 4 - Orchestrated Microservices

Further alignment between network and IT practices is achieved when the network functions are disaggregated

to become multiple cloud-native microservices, which can simplify scalability and maintenance.

Classic network data plane functions are generally bidirectional proxies, where the "go" and "return" traffic flows pass through a single software instance. This makes it hard to scale dynamically, presenting challenges in the design of the load balancing of the processing clusters.

By comparison, cloud-native microservices are transaction based, where a response is produced when a message is received. A network function designed for cloud working can be scaled easily with additional servers brought into service to meet peak traffic demands. With cloud-native network functions, separate web services can be established to handle data packets in each direction. The real-time network state can be shared between the web services using modern in-memory database techniques.

Cloud-native microservices also help with In-Service Software Upgrades (ISSU) on production networks. Having validated that a new software release functions correctly, the relatively small software image can be distributed to remote network nodes and brought into service using standard cloud services tools, such as Kubernetes* and Docker*.

Stage 5 - Microservices on Software-Defined Data Center

The last stage of NFV infrastructure innovation is to exploit software-defined data center techniques. Protocols such as the Redfish* API, created by the Distributed Management Task Force, could be used to create a virtual server to support NFV workloads, exploiting processing, storage and networking facilities available from hyper converged infrastructure.

This should allow complete consolidation of IT and networking workloads in one infrastructure platform, improving overall operational efficiency. It will also provide a platform that supports more DevOps-friendly practices in the network, enabling more flexibility and greater freedom for innovation.

Operating the Future Network

The transition to microservices presents opportunities for CoSPs to adopt DevOps practices, which are already becoming popular in the IT domain. These practices enable rapid time-to-market of new services by enabling frequent minor changes to be made to the production IT systems. This is in stark contrast to how CoSP networks have historically been managed, with infrequent major software upgrades and a heavy dependence on equipment manufacturers for support.

Under a more DevOps friendly model, the underlying infrastructure platform software used for NFV will evolve slowly. However, there will be frequent changes to the VNF software and the associated service descriptor configurations, enabling new network services to be created.

With the overall complexity of the NFV software stack, the expectation is that CoSPs will use automated testing tools to validate the functionality and performance of any modified component prior to production use. Once the software has been validated, Network Operations Center (NOC) staff will then be involved in loading it on the network. This will require new software skills in the network division to handle the expected rate of change, and more cultural alignment with IT practices.

Delivering New Services on the New Network

With NFV, a key feature is the open application environment that allows software from a wide range of independent software vendors (ISVs) to be hosted inside the network. This is a significant change from the classic scenario where network functions were provided as fixed function appliances from a limited number of NEPs.

Having re-engineered the network and the NOC, the CoSP will be able to exploit new business opportunities in three main areas.

Firstly, the reduced cost of implementing a new service in software means that niche services can become commercially viable. One-off custom solutions for major clients can be managed effectively with standard tools. Product managers can test new services in the market, with the freedom to take risks and innovate. Unsuccessful experiments can fail quickly, while successes can be more easily developed to production.

Secondly, the flexible deployment of NFV services around the network enables processing close to customers. That means application response times can be faster compared to typical use of remote over-the-top service providers. Processing at the edge of the network reduces core network traffic, such as content delivery systems for consumer and enterprise services. Edge processing also enables some new high bandwidth applications such as augmented and virtual reality systems. One example is a smart stadium trial undertaken by Intel with China Unicom, Nokia and Tencent Cloud. HD video of events was streamed to smartphones in the stadium with low latency, using local storage and avoiding mobile core traffic.

Finally, while CoSPs can introduce new services on today's wireless networks, further opportunities will become available when virtual Radio Access Network (vRAN) technologies are added with 5G architectures. Automated driving, healthcare, Internet of Things, smart city, and Industry 4.0 applications are expected to benefit from ubiquitous and high bandwidth wireless access, coupled with the ability to undertake local event processing within the network. 5G network slicing based on NFV will enable CoSPs to offer differentiated services to customers without duplicating infrastructure and systems.

Conclusion

This paper outlines how the NFV infrastructure model can be evolved towards microservices, to bring it into greater harmony with the IT systems. In the future, it may be possible to use a single hybrid infrastructure, the CoSP cloud, to serve both domains.

The new network can support new revenue opportunities for CoSPs, including in niche services, 5G and edge computing.

Find the solution that is right for your organization. Contact your Intel representative or visit **https://www.intel.com/content/www/us/en/communications/network-transformation.html**

Learn More

- A Step Towards 5G: Using MEC to Cut Live Video Latency at Concerts and Sports Arenas
- ETSI NFV Architecture Framework
- CORD Reference Implementation
- IETF Network Service Header (NSH) RFC 8300
- DMTF Redfish API
- 5G PPP use cases and performance evaluation models
- Intel® Select Solutions for NFVI
- Intel® Network Builders
- Intel® Network Builders University
- Intel® Data Center Builders
- Intel® Xeon® Scalable processors

Solution Provided By:

0618/CAT/GC/PDF 337737-001EN