EXPERT OPINION

How to unlock the business potential of NFV/SDN for CSPs

Lukasz Mendyk explains how a network as a distributed cloud can reshape itself to serve customer applications more effectively

The author, **Lukasz Mendyk**, is the OSS product manager at the Telecommunications Business Unit of Comarch

From the perspective of communications service providers (CSPs), the new technologies of NFV and SDN offer the ability to become a real cloud provider in a new sense, where a network can do much more than just provide access to data centres – it can become a cloud, serving as a platform for customer applications. It can also dynamically reshape its architecture to meet customer needs. This revolution is possible thanks to combining NFV (network functions virtualisation) and software defined networking (SDN) technologies, which means that networks can adapt by being reprogrammed. Moreover, network nodes can also become parts of distributed data centres that host not only network functions, but also applications. From the customer perspective this means that applications can be moved closer to the end user, enabling lower latency and higher speed, and leading to better customer experience.

NFV/SDN benefits

The technology also promises to open the network to innovation from the software developer ecosystem. Instead of rigid networks that are difficult to adjust to different application needs, the network becomes programmable and ready for the era of the Internet of Things (IoT), where applications can have their own virtual networks programmed.

Closer to the customer

The basic definition of NFV is to have network functions implemented on the same commodity hardware as used in data centres. If so, both network functions and customer cloud applications can be run on the same hardware. It can either mean moving network functions to data centres, or the opposite: moving applications to network edges.

When NFV is combined with its complementary technology, SDN, the cloud can become a real cloud. This means much more than running applications in the cloud from a central data centre with limited geographical distribution. The real cloud is a distributed environment where applications can be migrated as close to the customer as it takes to provide good customer experience at a reasonable cost. If network functions can run on the same hardware as end-user applications, it becomes possible to colocate applications and network functions in the same nodes, which can effectively become micro-data centres (see Figure 1).

An interesting case in point would be placing a customer application, for example a video on demand (VoD) server at the eNodeB node, which in fact means placing the video source as close as possible to the mobile user. From the perspective of traditional network topology, as defined by 3GPP, the network ▶

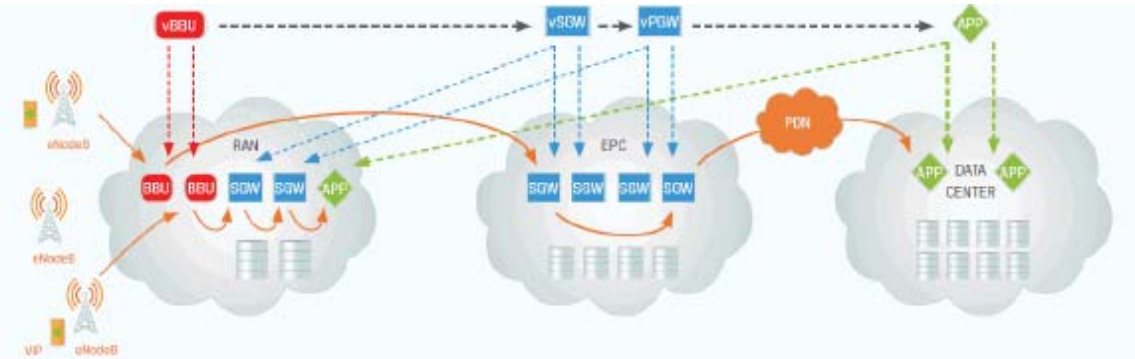

Figure 1. Extending the cloud to the network edges: RAN and core seen as micro-data centres

should at least include eNodeB, SGW (serving gateway), and PGW (packet data network gateway). Traditional network architecture assumes that there is a link between eNodeB and SGW, and it is implemented on top of the mobile backhaul, entailing many potential IP hops. The colocation of these functions in one physical node means that almost no mobile backhaul is necessary. For customers, it reduces latency, thus increasing service quality. For CSPs it decreases the burden placed on the mobile backhaul.

Control via policies: the magic answer?

Control via policies sounds like a magic answer, but this is probably easier said than done. When taking into account that the concept requires a hierarchy of policies in place, it also sounds a bit complicated. But in fact, hierarchy of policies is all about simplification. It is best to present some examples to explain this concept.

A provider of customer applications most probably shouldn't directly state, whether its application should be run at any specific network edge node. Instead, the app provider should define requirements, such as the maximum acceptable latency that can assure good customer experience, as the app provider is the one, who best understands the applications. These application requirements define the application level policy.

CSPs, who at the same time play the role of cloud providers, may treat this application policy as high-level. They may use their own policy, which additionally takes into account the network structure. For example, when many different applications are being deployed with variously defined requirements for maximum latencies, the CSP's policy may decide which application is going to be moved closer to the customer, colocated with appropriate virtual network functions, and run on hardware located at the eNodeB access node. In addition, this policy may take into account the number of customers accessing the applications from the given eNodeB (traffic statistics) to decide, which of the applications competing for resources should be moved closer to the customers.

The OSS role: Real-time OSS for NFV/SDN

As the key to benefit from the NFV/SDN potential is implementing control via multi-level hierarchical policies, the ideal solution for policy management seems to be a service catalogue where policies can be managed together with services – in fact, a policy can be treated as a service itself.

Managing definitions of policies is not enough. You need a dynamic, up-to-date view of the network and the applications to be able to execute these policies. Network inventory systems may help manage the NFV/SDN environment, as they can provide a physical infrastructure view, VNF node locations, and logical network topology, including that of micro-data centres and traditional data centres.

Service assurance solutions can provide a view of the network traffic, helping to establish the location of the congestion and suggest when it makes sense to migrate applications to actually improve customer experience. The reallocation of applications and VNFs requires programming the network. Network inventory, together with service assurance, can assist an SDN controller when making the re-routing decisions to optimise quality of service.

Reallocating application and network functions must be dynamic to provide good customer experience at reasonable costs. In addition, the need to react to changing traffic and application loads means that OSS must work in real time. Moreover, OSS needs to shift from a role of purely managing and orchestrating, towards being a true part of a dynamic network.

OSS in the cloud: the new meaning – OSS virtualisation

Real-time OSS means that performance scalability becomes critical. To transform OSS to real time can be challenging, but the solution lies in OSS virtualisation - i.e. placing OSS in the cloud. This may not be a new concept, but in this case it is more than just locating the OSS system in a data centre and having it managed by an OSS provider.

Instead, OSS may also be reallocated within the network nodes and be colocated, with VNFs and applications to assure the correct level of responsiveness. OSS must therefore be deployable both at traditional data centres but also follow network functions and applications when migrated to micro-data centres. Colocating OSS and VNF is essential when low latency is critical.

To learn more, download Comarch's free white paper:
The Business Potential of NFV/SDN for Telecoms – How a Network as a Distributed Cloud can Reshape Itself to Better Serve Customer Applications at **nfv-sdn-whitepaper.comarch.com**