ORACLE COMMUNICATIONS

Network Function Virtualization

The Road to NFV Success Is Paved with Intelligent Orchestration

CSPs can infuse a high degree of intelligence through business and technology rules that govern workloads and decisions based on key metrics and measurements.

ORACLE POSITION PAPER | FEBRUARY 2015

ORACLE

Introduction

The very purpose of network function virtualization (NFV) is to move beyond the limitations of proprietary hardware and platform capabilities to truly improve service agility and speed to market. However, virtualization alone will not be enough. Communications service providers (CSPs) will have to focus on the interactions that take place between virtual network functions (VNFs) and physical network functions (PNFs). These interactions must be well coordinated, or better yet orchestrated, from service design all the way through to data center operation so as to conquer the layers of complexity inherent in bridging physical and virtual environments.

This coordination will come to life through orchestration frameworks or orchestrators that best address functional requirements for operational simplicity. Simplicity is critical, no matter how complicated the networks, and no matter how many layers are added in terms of systems, functions, services, orchestration and management technologies. To achieve simplicity, an NFV implementation plan has to factor in “intelligent orchestration” based on a powerful combination of policy rules that govern network and service behaviors, as well as analytic feedback from run-time operations.

Without intelligent orchestration, CSPs will end up with “more of the same” in terms of how they order, provision, deliver, support, and bill for services. NFV implementations must instead empower CSPs to flexibly match and manage both PNFs and VNFs. CSPs must have the elasticity to piece together functionality for dynamic configuration of resources so that they can viably compete and collaborate with the likes of Facebook, Apple and Amazon. Innovative and agile, these stakeholders have upped the ante in many areas of service delivery and customer experience, and operators are poised to complement or compete, as they possess the reliable networks, the precious customer data, the brand recognition and the service management systems that others can only covet.

The missing piece has been finely controlled, quick-assembly service delivery that enables operators to experiment with new services or to test network or customer segments within “learning laboratories” in which new offers are immediately put to work. The ability to make dynamic adjustments to bandwidth and data-volume entitlements, or to gain dynamic control over quality of service (QoS) parameters would drive innovation and open the door to the types of session-based, contextually rich use cases customers really want. As important, if not more, is the ability for CSPs to gain meaningful insight into subscriber data profiles, and information about state, usage, location, entitlements and restrictions.

Experimenting with new services and getting real-time or near-real-time feedback on the performance of their businesses and networks will also help operators gauge and control the level of automation they drive into NFV implementation — for example, maintaining network operations consoles until the time is right for full automation.

Success in these areas will depend on whether NFV strategies are built around policy-driven, analytically charged management capabilities and the degree to which network, service, function and data center orchestration have been considered and invested in from the beginning.

Let the Journey Begin: Preparing for New Layers of Complexity

Though NFV's potential to revolutionize communications service delivery is undeniable, there's a "dirty little secret" often left untold in the NFV story today: just moving functions from proprietary hardware onto virtualized software will not accelerate service agility or reduce equipment and operational costs to the degree needed to justify the investment. In fact, the initial move to add virtualized functions could for the short term actually increase complexity, add points of failure and escalate costs. Bringing proven cloud computing and IT technologies into the networking domain can indeed enable a building-block approach to connecting VNFs, but managing the many virtual machines, software and processes can become quite difficult and expensive if the inherent complexities are not considered from the outset.

Nothing can immediately become "fully virtualized," nor can a single virtualization strategy be optimal for all network functions. Rather, there must be a series of evolutionary steps, starting with a top-down look at what the market wants and what a CSP desires to sell, and then what network components, functions, policies, SLAs, orders and subscriber tiers they will have to satisfy to succeed.

The first step will be an evaluation of which elements lend themselves to early implementation on industry standard, X86-based hardware and which do not. Products supporting very low-latency, high-performance data transcoding and transport will probably remain appliance-based for several years. The more likely candidates for NFV in the early phases of deployment will be the signaling-intensive functions, such as session core control plane activity and enhanced routing applications, signaling proxies, network policy elements, and routing agents.

With an eye toward rapid deployment and graceful scaling, CSPs must then discern how to best link their purpose-built appliances with these virtual machines. As they do so, they must maintain a comprehensive view across appliance-based environments and virtual network architectures, including the network services, network functions, infrastructure and applications.

There must also be an evaluation of how virtualization can impact not only core network functions, but also OSS/BSS, operations and service design and delivery processes. It will be necessary to streamline work efforts for operations teams, which will be increasingly tasked with new day-to-day operational demands as they work with infrastructure, network functions and the associated abstraction layers all the way up the stack. Failure to streamline may indeed open the door to frustration, stemming from new layers of operations processes, attrition and clashes between IT and network organizations.

As stated earlier, operational simplicity is a must, no matter how complicated networks, systems, functions, services, and management technologies get.

Achieving Operational Simplicity

For there to be elasticity in how functionality can be used in physical or virtual form, and for there to be dynamic configuration of resources as changes are needed, there must be a sophisticated level of orchestration that encompasses service design and delivery, VNF and PNF management, and of course, coordination with data center resources.

Likewise, this orchestration activity must be rules-driven so that CSPs can infuse a high degree of intelligence through business and technology rules and bring an innate knowledge of resource availability. That intelligence will dictate how workloads are managed and how decisions are made around certain key metrics and measurements.

Finally, the hybrid physical/virtual network must become analytically charged so that CSPs can cultivate insight into network and customer behaviors that ultimately drive QoS decisions, new offers and network resource allocations.

In the area of orchestration, there are three critical areas to consider:

- Network service orchestration
- Application orchestration
- Virtual infrastructure management

All of the above come to fruition either through an orchestration framework or perhaps multiple orchestrators, depending on how organizations wish to address functional requirements.

Whether there is one premier "orchestrator" or a series of targeted orchestrators, the point is there must be "something" that addresses the following major challenges:

- Network service orchestration, which manages all capacity and virtualization coordination for both VNFs and PNFs. It must empower CSPs to respond automatically to changes in network capacity requirements and KPIs. It assumes the critical role of the VNF Manager as defined by the ETSI[i] management and orchestration (MANO) architecture. Newly provisioned devices are automatically added to device and configuration managers, and a master configuration is maintained so that operational efficiency is optimized.
- Application orchestration, which must manage the lifecycle of network functions. The VNFs can be any function, such as PCRFs, DRAs, CSCF, mobile network nodes, SBCs, or even online charging systems. They can be virtualized through software running on NFV infrastructure (NFVI), which encompasses all the hardware (computer, storage and networking resources), and virtualization or container technology used to provide the infrastructure resources wherever VNFs are deployed.
- NFV infrastructure (NFVI), the data center resources that are managed using virtual infrastructure managers (VIMs) — cloud management systems tasked with managing the virtual machines that run on hardware in the data center (e.g., Oracle OpenStack, VMWare vCloud Director). The power of the VIM to manage data center resource allocation depends on the robustness of the cloud infrastructure and its compliance with ETSI standards. It is the foundation for infrastructure management and communication with NFV analytics, which offers insight into customer actions and network resource utilization.

In the end, operators must leverage OSS assets to enable service activation. With network service orchestration, they have to simultaneously manage physical and virtual network functions, coordinating network services and integrating to an application "orchestrator" so that network services can be created and deployed quickly.

There must be an ability to manage the lifecycle of network services across different vendors' products and across diverse locations from a single application. There must also be a means of ensuring network services elastically scale and that service connectivity is automatically updated. That requires a stateful inventory that provides persistent and accurate data for all processes, not to mention service aware in its native services information model and topology. It should be dynamic and data driven with a configurable catalog of reusable models and business policies, and, it should support automated delivery through intelligent configuration of heterogenous infrastructure.

When all of these pieces of orchestration are carefully considered, CSPs can reap the full benefits of intelligent orchestration, including:

- Understanding the many ways in which virtualization affects core network functions, such as provisioning, service quality management, network performance and the overall customer experience;
- Knowing and managing how network assets are utilized and how data traverses OSS/BSS and service layer platforms — even when many service instances, physical network function instances, and VNFs come into play;
- Monitoring and acquiring vast stores of operational data that they evaluate against billions of transactions and events, as well as feeding reporting needs about network elements, servers, databases, applications, OSS/BSS, and customer support.

The ‘Intelligence’ of Intelligent Orchestration

Fueling the multiple layers of orchestration described above are two additional types of intelligence: decision intelligence, which is a set of business and technology rules that provide insight into what resources will satisfy services and business needs, and then the analytic intelligence that provides feedback on how the network performed and how customers behaved. With both, CSPs can fully understand stimulus:response intricacies and how to orchestrate communication across the virtual network environment.

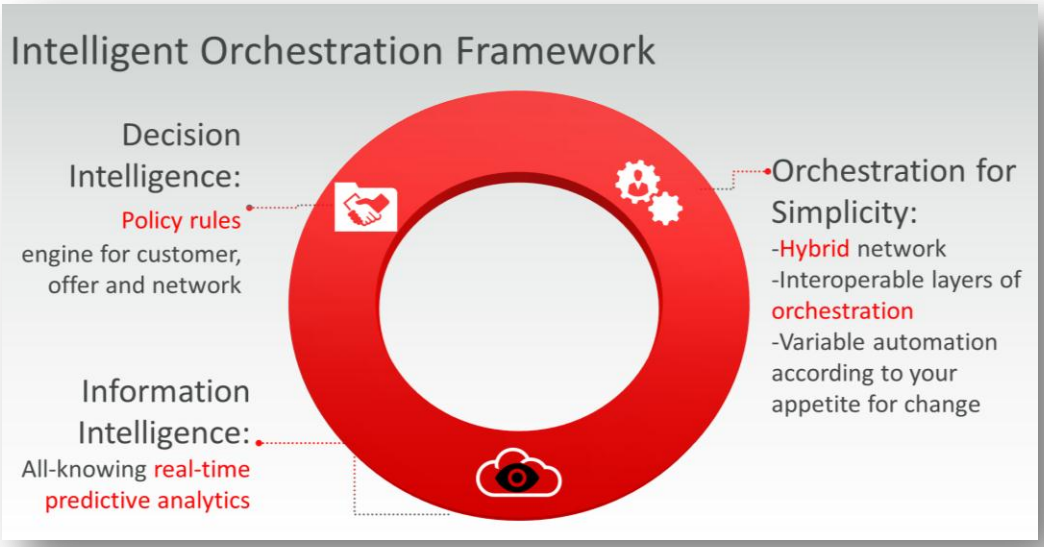

Driving this “intelligence” is a robust policy engine — one that provides a store of all the business and technology rules needed to drive revenues and ensure appropriate QoS. Policy will govern bandwidth decisions, data volume entitlement or other variables important to bandwidth shaping, traffic shaping and enforcement of QoS parameters. And policy will facilitate the creation of session-based, contextually rich use cases, based on subscriber data profile, state, usage, location, and parental control data gathered by a Subscriber Profile Repository — one that is either integrated with, or based on, existing HSS or LDAP databases. Finally, this decision intelligence ‘brain’ will possess the carrier’s rules designed to protect the NFV-driven network with geographic diversity requirements and functional anti-affinity rules. In short, it houses the full range of rules and relationships that characterize a carrier’s business.

At the same time, the policy system helps CSPs gain insight into what-if capabilities so that the potential impacts of decisions are understood *before* policy deployment. This type of informed decision making then improves revenue creation, performance and ultimately, the customer experience.

The policy engine should be fueled by relevant analytics, which filter and correlate large amounts of data so that CSPs can make faster and more informed decisions around entitlements, feature combinations, rate plans and equipment and capacity decisions.

The impact policy can have will be further optimized through trending reports and dashboards that make the most of feedback about network conditions, subscriber behaviors and policies. For example, quota tracking measures can show impact on networks and gateways, further empowering operators to report and adjust service offerings or billing accordingly, as well as helping customers adjust plans according to their usage, roaming, and budgets. These adjustments or modifications optimize network utilization and customer experience.

As valuable lessons are learned, they can be infused in a programmable way across the network, OSS, BSS and other components of the CSP ecosystem. That will in turn fuel improvements to bandwidth management, routing plans, and traffic type assignments — ultimately improving network performance and the service quality that affects overall customer experience.

Once intelligence is being cycled back into the business, operators can optimize the impact of NFV, gaining visibility and control over the compute and storage capabilities of their networks; the resources being utilized by different services; and the OSS/BSS (e.g., policy, charging, billing, etc.) leveraged for various resources, services, applications or functions.

Conclusion

The planning and components described here are all driven by what consumers and enterprises expect in a digital economy. They don't really care if networks are physical, virtual or something in between; they just want their needs to be served when they want, where they want, and with more speed, reliability and quality than ever before. They want device-driven, personalized offers and experiences that are context sensitive and persona-driven, according to who they are and what they want to accomplish at any given moment.

The finely controlled, quick-assembly service delivery possible with NFV requires more than just decoupling network functions from proprietary hardware appliances. It requires an understanding of the complexity of additional operational streams and more day-to-day operational responsibilities. It also requires a high level of expertise around network functions, as there are many ways in which virtualization can impact network functions, such as provisioning, service quality management, network performance and the overall customer experience. A high level of expertise will also be necessary in managing operations, virtual network functions, and infrastructure comprising NFV.

As NFV implementations evolve, CSPs will have to ensure they maintain visibility and control over the compute and storage capabilities of their networks; the resources being utilized by different services; and the OSS/BSS (e.g., policy, charging, billing, etc.) leveraged for various resources, services, applications or functions.

For these reasons, CSPs should consider products and new deployments built with NFV in mind; they should avoid older platforms that are retrofitted for new deployment scenarios. Additionally, they should look at technology partners who are committed to releasing network elements to VNF — whether session border controllers (SBCs), signaling routers, an EPC, or other elements. And because high-volume analytics and policy will be crucial to NFV success, they should look at suppliers that are innovating in those areas, as they will understand the interplay necessary among OSS/BSS, EMSs, VNFs, NFV instances, NFV orchestrators, and VIMs. Finally, expertise in cloud computing and data center technologies will increase in importance as a required skill set for any supply partner.

As NFV projects mature, CSPs will no doubt continuously measure the impact on SLAs and QoS, and drive what they learn back into their networks through effective policies. The goal is to continuously learn and feed that knowledge back into the network and systems managing those networks so that services improve and customer satisfaction and loyalty grow.

Network Function Virtualization

Oracle Corporation, World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065, USA

Worldwide Inquiries
Phone: +1.650.506.7000
Fax: +1.650.506.7200

ORACLE

CONNECT WITH US


blogs.oracle.com/oracle


facebook.com/oracle


twitter.com/oracle


oracle.com

Hardware and Software, Engineered to Work Together

The Road to NFV Success is Paved with Intelligent Orchestration
February 2015


Oracle is committed to developing practices and products that help protect the environment

[i] European Telecommunications Standards Institute