



Sponsored by Dialogic

NFV Network Functions – Key Considerations for Profitability

Executive Summary

Over the last few years, Communications Service Providers (CSPs) have shown great interest in Network Functions Virtualization (NFV). As CSPs move towards a cloud-based infrastructure model, there should be heightened awareness on how software-based network functionality is architected when it comes to migrating services involving real-time multimedia to a virtualized environment. In order to implement NFV successfully, CSPs should take into account a number of guiding principles including software modularity, virtualization technology, and automation. By properly focusing on these principles, clear advantages start to become possible such as more efficient scaling up and scaling out of resources, lower risk to the business, high service availability, and faster time-to-market. This white paper will focus on the application of specific guiding principles that help maximize the profitability of NFV in the deployment of Virtualized Network Functions (VNFs) supporting IMS-based services.

Table of Contents

Introduction	4
The Current State of NFV	4
Industry Expectations of NFV	5
Big Bang vs. Piecemeal Approach	7
Software Modularity	7
Technology Considerations	8
Software-Centric Design	8
The Role of Open Source	8
Virtualization 101	8
Programmatic Interfaces	9
Relationship with SDN	10
A Fresh Approach to High Availability	10
VNF Implementation Examples	10
Media Resource Functionality in the Cloud	11
Virtualization of the IMS: Media Gateway Control Functionality in the Cloud	12
Summary	14
About Jim Metzler	14
About Dialogic	14

Introduction

For the last few years, Network Functions Virtualization (NFV) has been one of the hottest topics in the IT industry. The great interest in NFV comes from the fact that CSPs believe that by implementing NFV, they can reduce cost, simplify their operations, and reduce the amount of time and resources it takes to deploy new services. In January 2013, that belief caused seven of the world’s largest CSPs to create an Industry Specifications Group (ISG) for Network Functions Virtualization (NFV ISG) under the auspices of the European Telecommunications Standards Institute¹ (ETSI NFV ISG).

While NFV holds great potential for CSPs, it also holds great challenges. In line with the pivot that the ETSI NFV ISG has made away from focusing on requirements and towards focusing on implementation, the goal of this white paper is to help CSPs get started with successfully implementing NFV. As part of achieving that goal, this white paper will examine some of the technical and business drivers around NFV that lead into some guiding principles that CSPs should consider especially when looking at Virtualized Network Functions (VNFs), including how to ensure that their NFV implementation is highly profitable.

The Current State of NFV

The central concept that drives NFV is that some or all of the network functions that CSPs utilize must be available as virtual appliances that can be easily provisioned, integrated, and managed. The transition from the current hardware-centric environment to a fully virtualized environment is depicted in Figure 1.

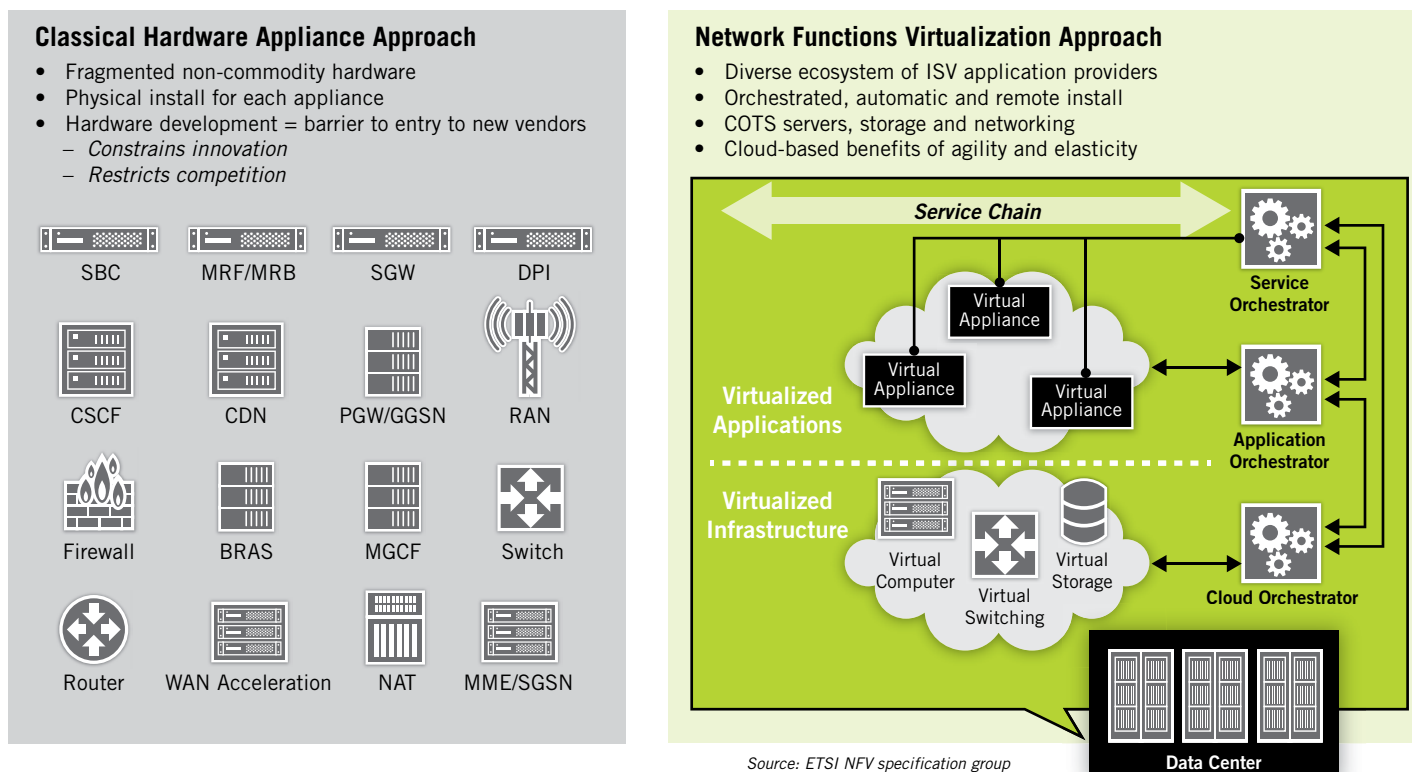


Figure 1 - The virtualization of network functions from a hardware-based approach to a cloud model

¹ <http://www.etsi.org/news-events/news/644-2013-01-isg-nfv-created>

Some of the key characteristics of the ETSI vision for NFV include ² :

- Achieving high performance virtualized network appliances which are portable between different hardware vendors and across different hypervisors or other virtualization technologies;
- Achieving co-existence with hardware based network platforms;
- Managing and orchestrating many virtual network appliances while ensuring security from attack and misconfiguration;
- Implementing automation to enable the scalability of the solutions;
- Ensuring the appropriate level of resilience to hardware and software failures.

In addition to articulating a vision and developing specifications for NFV, the ETSI ISG has placed considerable emphasis on using Proof of Concepts (PoCs) to accelerate workable implementations of the NFV specifications. These PoCs involve vendors and CSPs working jointly to put into action the technical concepts covered by the NFV ISG. The PoCs that have been put forward are varied and include demonstrations on virtualization of elements of the IP Multimedia Subsystem (IMS), the Evolved Packet Core (vEPC), the Radio Access Network (RAN), SDN and virtualized EPC gateway interaction, and stateful fault tolerance in the cloud. This healthy and diverse PoC activity for NFV principles reflects the commitment on the part of CSPs and vendors alike to take NFV from a concept to reality, and demonstrates that the efficacy and benefits of this paradigm shift are tangible.

Industry Expectations of NFV

A recent report titled *When Will NFV Cross the Chasm*³ contained the results of a survey that was given to CSPs. One of the survey questions asked the respondents to indicate the two primary factors that are driving their company's interest in NFV.

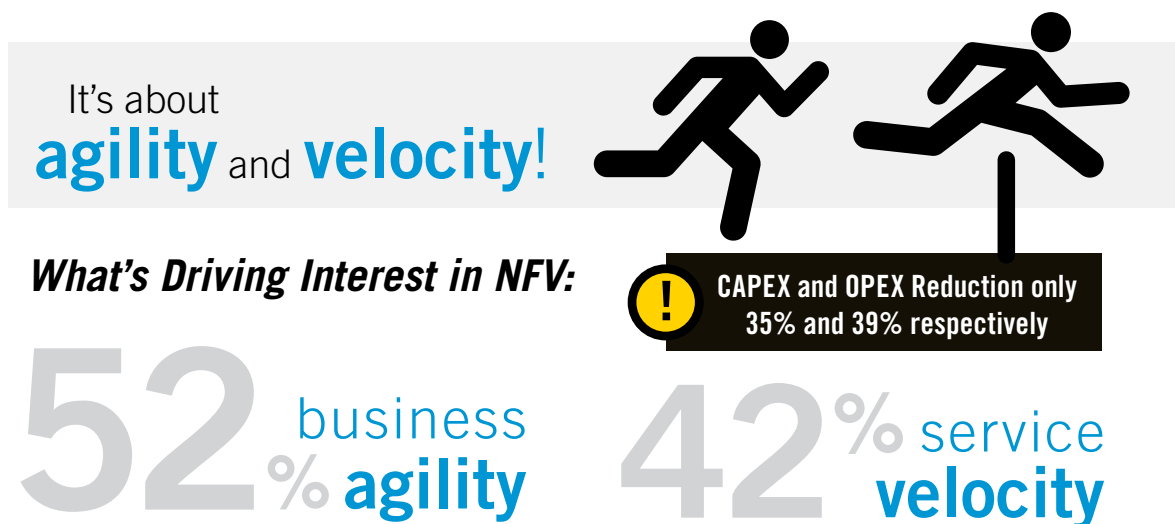


Figure 2 – For CSPs, business agility and service velocity outweigh the CAPEX and OPEX benefits of NFV

As seen in figure 2, while reducing CAPEX and OPEX is a major factor driving interest in NFV, a more significant factor is increasing business agility; e.g., reducing the time it takes to go from concept to deployed new service and enabling the company to adapt to new business or market conditions.

It is important to understand the current drivers of NFV. It is equally important to understand the current inhibitors. Towards that end, the survey respondents were asked to indicate the primary business inhibitors to their company broadly adopting NFV sometime in the next two years.

² https://portal.etsi.org/NFV/NFV_White_Paper.pdf

³ <https://www.tmforum.org/resources/research-and-analysis/virtualization-when-will-nfv-cross-the-chasm-2/>



Figure 3 – Business inhibitors to NFV deployment

Two of the leading business related inhibitors to the broad adoption of NFV are the need to reskill the employee base, and the need to make organizational changes. AT&T is an example of a carrier that is attempting to mitigate those inhibitors. In 2014, AT&T announced⁴ an online “nanodegree” program that is designed to develop a new crop of software experts. Shortly after that announcement, AT&T also announced a major change in its organizational structure. According to an article in the Wall Street Journal⁵, “The reorganization, which also includes the formation of three new business units, is designed to reduce complexity and make it easier for the carrier to offer services to customers as it transforms its hardware-focused network into a software-centric one.”

The business inhibitors are only part of the situation, and so the survey respondents were also asked to indicate the primary technological inhibitors to their company broadly adopting NFV sometime in the next two years.

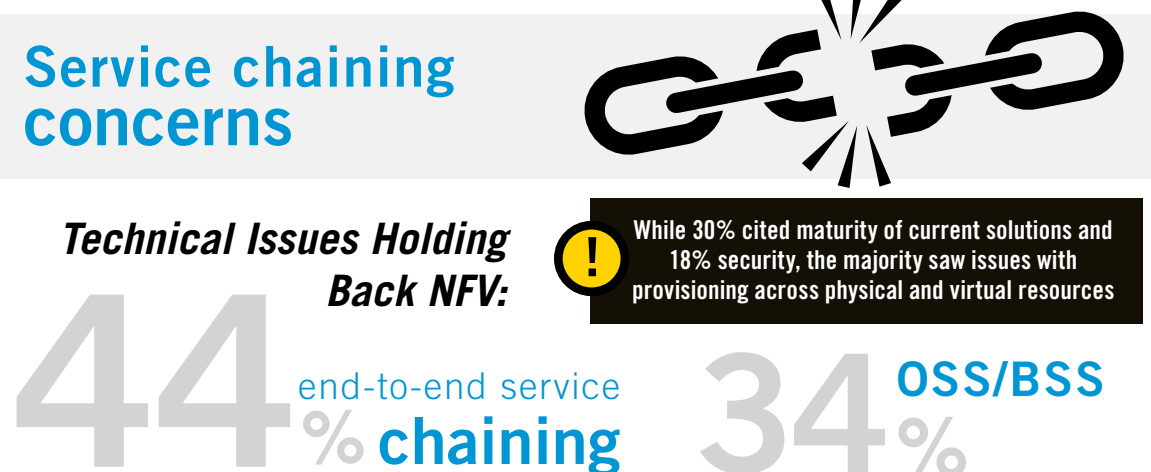


Figure 4 – Technical issues with NFV

⁴ <http://www.fiercewireless.com/tech/story/atts-nanodegree-gets-workers-ready-sdn-nfv-future/2014-06-18>

⁵ <http://blogs.wsj.com/cio/2014/09/09/att-names-new-cio-amid-it-reorganization/>

From a technical perspective, some respondents indicated that the immaturity of the current product offerings is currently inhibiting the adoption of NFV, but that will diminish over time as vendors bring out successive iterations of products. However, the concerns over end-to-end service provisioning and the need for a new generation of agile OSS/BSSs is something that needs to be addressed by the industry in the near term. One organization that is focused on those issues is the TM Forum, which recently published a blueprint for end-to-end management ⁶.

Big Bang vs. Piecemeal Approach

One of the conclusions drawn in *When Will NFV Cross the Chasm*, is that if the definition of NFV crossing the chasm means that all of the data plane packet processing and control plane functionality that is used by CSPs have been virtualized and all the previously identified characteristics of the ETSI vision such as management, orchestration, and automation are not fully in place, then NFV will not cross the chasm for the foreseeable future. However, if the focus is on a few use cases, such as IMS or vEPC ⁷, with less than the full ETSI vision for management, orchestration, and automation, then those use cases will likely cross the chasm to mainstream adoption in a year or so.

One of the key architectural questions facing CSPs then is whether to take a tactical approach (a.k.a., a piecemeal approach) or an architectural approach (a.k.a., a big bang approach) to NFV. Companies that take a piecemeal approach typically focus on one, or at most a small number of use cases for which they see clear business value and then typically participate in PoCs to demonstrate the viability of the possible solutions.

In contrast, when a company takes a big bang approach to NFV, they decide on an architecture, and possibly on some key enabling technologies and products, that the architecture will utilize in order to support any and all NFV use cases.

A CSP with a highly centralized, effective management structure would have an easier time with a big bang approach than would a CSP that was highly distributed, but the more volatile the set of enabling technologies is, the more difficult it would be to be successful with a big bang approach. However, approaching NFV in a piecemeal approach may run the risk of rolling out services with reduced synergies in terms of the platforms used, how they're managed and how they are secured.

Key Architectural Principals to Guide NFV Deployment

When CSPs implement NFV, part of their focus should be on increasing the profitability of their existing services that are delivered in a traditional manner. However, a part of their focus should also be on adopting new, profitable business models for the deployment of new services. An example of such an approach is the everything-as-a-service (XaaS) model that is utilized by cloud service providers. A CSP, or cloud service operator could, for example, offer as a service virtualized control and data plane functionality in the EPC or the IMS core in a way that incorporates some of the characteristics of the ETSI vision of NFV.

Whether it is to support existing services or new services, in order to maximize the profitability associated with implementing NFV, CSPs must minimize the cost of the implementation and the associated risk. They must also implement NFV in a way that enables them to minimize the time to market for the deployment of new services.

In order to achieve these goals, CSPs should adhere to the following guiding principles:

Software Modularity

As noted, one of the components of a successful implementation of NFV is confirming that the NFV functionality has been implemented in a way to maximize profitability. One technique that is associated with maximizing profitability is modular programming. The phrase *modular programming* refers to a software design technique that emphasizes taking a piece of software-based functionality and decomposing it into independent modules, such that each module contains everything necessary to provide one component of the desired functionality. In addition,

⁶ <https://www.tmforum.org/zoom/zoom-blueprint/>

⁷ [http://www.lightreading.com/carrier-sdn/nfv-\(network-functions-virtualization\)/the-rise-of-virtual-epc/a/d-id/708394](http://www.lightreading.com/carrier-sdn/nfv-(network-functions-virtualization)/the-rise-of-virtual-epc/a/d-id/708394)

VNF modules can be deployed geographically closer to the target user community to reduce latency and improve the quality of services delivered to end users.

Technology Considerations

The adoption of NFV is still in its early stages and these early stages are characterized by rapidly changing technologies. For example, the initial discussion of NFV focused on the use of CloudStack, which have now shifted to OpenStack as a replacement for the cloud computing platform of choice. In addition, most of the discussion of VNFs to date has them running in virtual machines (VMs). However, there is beginning to be discussion about VMs being replaced by containers⁸ or unikernels⁹, which can reduce the amount of VNF overhead or a component of a VNF (VNFC), and thus accelerate scaling up and scaling out the applicable functionality.

While a VM is certainly one approach to implementing a VNF, deploying an application in a VM doesn't necessarily encourage software vendors to adopt a modular software architecture. A downside of a VM-based approach is that because a VM contains a full server hardware stack, instantiating a new VM can be relatively time consuming which negatively impacts both high availability (HA) and the time it takes to scale functionality to meet peak loads. In addition, while it is challenging to take an application running in a VM and modularize it to run in containers or unikernels can provide more flexibility since they are backwards compatible with VM-based deployments.

Software-Centric Design

As the industry makes the shift from a hardware-centric environment to a virtualized environment as depicted in Figure 1, some vendors will offer VNFs that are based on merely porting the code from their hardware-based appliance over to a VM. While that is an expedient approach, it will lead to significant performance problems since that code was designed to leverage the underlying capabilities of specialized hardware. In order to achieve maximum performance, CSPs should focus their attention on VNFs that were designed to run effectively in a software-centric environment.

The Role of Open Source

There are many groups that are shaping the evolution of NFV. That includes groups such as ETSI and the TM Forum, which define use cases, drive PoCs and work to resolve issues related to key tasks such as end-to-end provisioning. Another group driving the evolution are Standards Developing Organizations (SDOs) such as the IETF and the Alliance for Telecommunications Industry Solutions (ATIS). A third group driving the evolution of NFV is the open source community. For example, in September 2014 the Linux Foundation announced the founding of the Open Platform for NFV (OPNFV) Project¹⁰. As part of the announcement the Linux Foundation declared that OPNFV will establish a carrier-grade, integrated, open source reference platform that industry peers will build together to advance the evolution of NFV and ensure consistency, performance, and interoperability among multiple open source components. The Foundation also stated that because multiple open source NFV building blocks already exist, OPNFV will work with upstream projects to coordinate continuous integration and testing while filling development gaps. One of the goals of working with upstream projects is to ensure that it is easy to load a VNF and have it run correctly regardless of the underlying physical infrastructure.

It is unclear how the relationship between the SDOs and the open source community will develop. One option is that after a group such as the OPNFV has made progress on creating an open source reference platform for NFV, that one or more SDOs will establish working groups to

Virtualization 101

Virtual Machines (VMs) are an abstraction of physical hardware in which each VM has a full server hardware stack from virtualized BIOS to virtualized network adapters, storage, and CPU. VMs also have their own instance of whatever operating system is in play.

Containers don't virtualize the entire server hardware stack. Instead, only the application and its dependencies – not the entire operating system – are placed in a virtual container. While they share a single OS, they are compartmentalized in all other respects so that a container has no interaction with any other container.

Unikernels are specialized operating system kernels that act as individual software components. A full application or appliance consists of a set of running unikernels working together as a distributed system.

⁸ <http://searchsqlserver.techtarget.com/definition/container>

⁹ <https://queue.acm.org/detail.cfm?id=2566628>

¹⁰ <http://www.linuxfoundation.org/news-media/announcements/2014/09/telecom-industry-and-vendors-unite-build-common-open-platform>

create standards for some of the key tasks that are part of the reference platform. However, since SDO working groups have historically taken years to create new standards, another option is that whatever functionality is part of the reference platform will become de facto standards. Regardless, much collaboration is taking place between the various SDO groups involved to address gaps, conflicts, and overlaps in the developing standards. This should help expedite a more widespread adoption of NFV. Thus, it is important that VNF vendors integrate support for Open Source in architecting and implementing NFV applications.

Programmatic Interfaces

Using the data center as a model, there is a strong movement away from a static environment, and towards a software defined data center (SDDC). In a static environment, infrastructure functionality is provided on a piece of dedicated hardware and the interface into the hardware is manually accessed. In a SDDC:

- Computing, storage, and networking are virtualized;
- All resources are pooled;
- Programmatic interfaces exist into data center resources;
- Automated management delivers a framework for policy-based management of data center application and services.

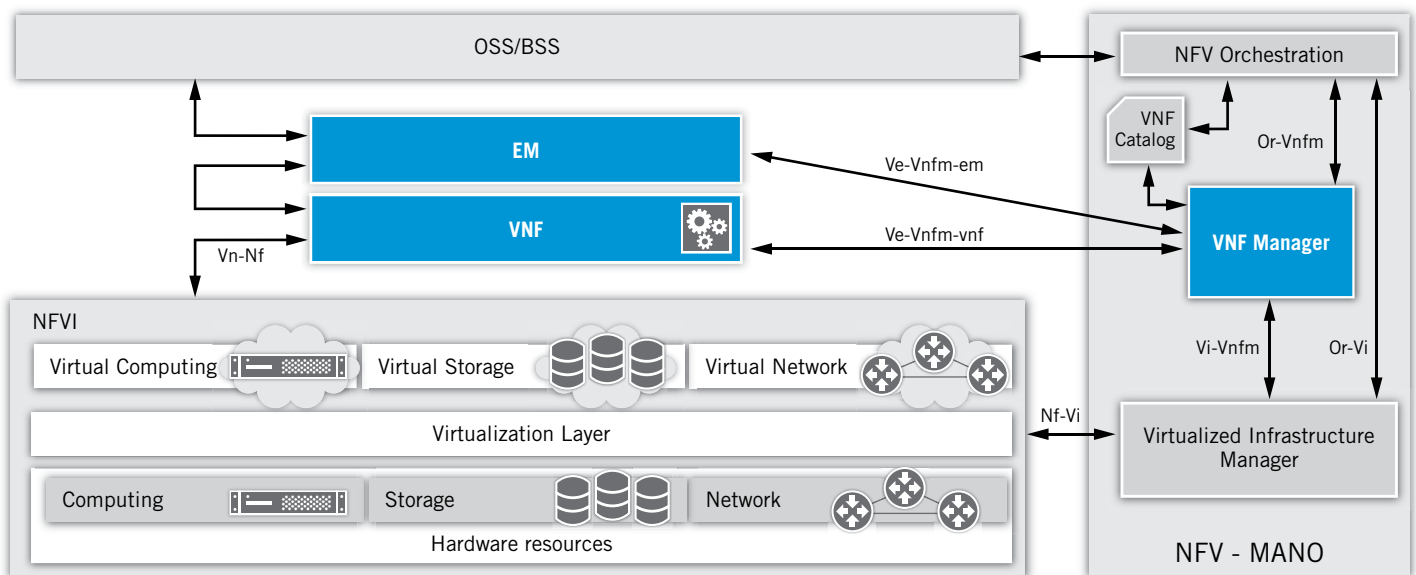


Figure 5 – One important role of the VNF Manager is to take key performance indicators from the VNF and EM to determine whether to scale up or down capacity to meet increased or decreased demand for application resources

For VNFs, automation, scalability, and programmability are not “nice to have” concepts when it comes to maximizing the profitability of NFV, but rather “must have” goals. VNFs should provide sufficient intelligence to assist in the process of automating the rapid scaling out and scaling up of capacity in response to traffic loads in an NFV environment. Automation, and programmability are also important components to high availability, as well as efficient use of NFVI resources. Thus, one should carefully look for VNF vendors that incorporate the capability to automate many of the tasks of installing and configuring virtualized functionality to help realize the benefits of a cloud-based architecture.

Relationship with SDN

In March 2014 the Open Networking Foundation (ONF) released a document entitled the *OpenFlow-enabled SDN and NFV Solution Brief*¹¹. That document discussed how OpenFlow-enabled SDN can meet the need for automated, open, and programmable network connectivity to support some of the ETSI-defined use cases such as *Network Functions Virtualization Infrastructure as a Service and Virtual Network Function Forwarding Graph*.

In a recent white paper¹² ETSI expressed their belief that NFV and SDN are highly complementary efforts. The ETSI view is that both efforts are seeking to leverage virtualization and software-based architectures to make network infrastructures more cost-effective and more agile by integrating the ability to accommodate the dynamic nature of the workflows demanded by applications and end users. While NFV can be implemented using a non-SDN infrastructure, the ETSI vision is that NFV and SDN will increasingly be intertwined into a broad, unified, software-based networking paradigm based on the ability to abstract and programmatically control network resources dynamically and automatically.

To exemplify the potential interaction between SDN and NFV, consider a situation where a load balancing service is implemented as a VNF. If demand for load balancing capacity increases, a network orchestration layer can rapidly spin up new load balancing instances and also adjust the network switching infrastructure to accommodate the changed traffic patterns. In turn, the load balancing VNF entity can interact with the SDN controller to assess network performance and capacity and use this additional information to balance traffic better, or even to request provisioning of additional VNF resources.

A Fresh Approach to High Availability

CSPs have historically designed their systems for High Availability (HA), with the typical goal being five 9s of availability. This means that a system is available 99.999% of the time, or conversely, that a system is unavailable for roughly 5 minutes or less a year. This level of availability is achieved through a variety of techniques involving geographically diverse transport, hardened physical facilities, and system and component level redundancy.

The implementation of NFV enables CSPs to rethink HA. NFV enables CSPs to shift the focus on HA away from redundant systems of highly reliable physical components to a focus on highly available services that are supported in part by the inherent capabilities of the underlying NFV infrastructure (NFVI) layer. With NFV, when there is a failure, the impacted traffic will be re-directed to a new instance or a load-shared instance of that application either in the same data center or across disparate data centers. The potential benefit here is that CSPs would have less CAPEX tied to idle capacity for the purpose of meeting high availability objectives.

VNF Implementation Examples

Moving network functions to the cloud is not a trivial task. Applications that involve both control plane and data plane processing require stringent response performance that puts demands on latency characteristics of virtualized environments. NFV performance will be put to the test even more with services contemplated by 5G, which will involve a network infrastructure with always-on coverage characteristics, and end-to-end latency that gives the user a real-time experience. Innovative applications delivered over 5G mobile networks - for example, autonomous driving, tactile Internet, virtual and augmented reality, multi-person video conferencing, real-time gaming and remote control - will require fast processing by all layers of the NFV to ensure fast download rates and low latency.

IMS is already starting to play an important role in the delivery of real-time multimedia services and it involves both control (or signaling) plane and media plane processing. We'll look at two core IMS functions; the Media Resource Function (MRF) and the Media Gateway Control Function (MGCF), which have historically been implemented using proprietary hardware platforms. We'll discuss how these functions can be deployed in cloud environments in a manner that adopts some of the guiding principles for NFV profitability that were presented earlier.

¹¹ <https://www.opennetworking.org/images/stories/downloads/sdn-resources/solution-briefs/sb-sdn-nfv-solution.pdf>

¹² http://portal.etsi.org/NFV/NFV_White_Paper3.pdf

Media Resource Functionality in the Cloud

The MRF provides signaling and multimedia processing for IMS/VoLTE networks. An MRF is a key functional entity in IMS networks and includes the following functionality:

- Audio/video conferencing;
- Voice mail;
- Transcoding, transrating, and transizing of audio and video streams;
- Tone detection;
- Playing announcements;
- Prompts for Interactive Voice Response (IVR) systems.

In an IMS session, the media is played, recorded, or bridged by a media resource application that is controlled by SIP or MGCP signaling protocols and is initiated by an application server as a call progresses through the system.

The MRF handles data plane traffic performing computationally intensive processing on both real-time and time-shifted audio and video. The stringent performance requirements imposed by real-time applications make implementation of scalable MRF VNFs challenging. By following the architectural principles laid out earlier, these challenges can be overcome. Figure 6 shows how the MRF could be implemented in a modular fashion as a composite VNF. Using this approach, the MRF is comprised of a number of components, including an RTP proxy, a load balancer, and individual media resource functions.

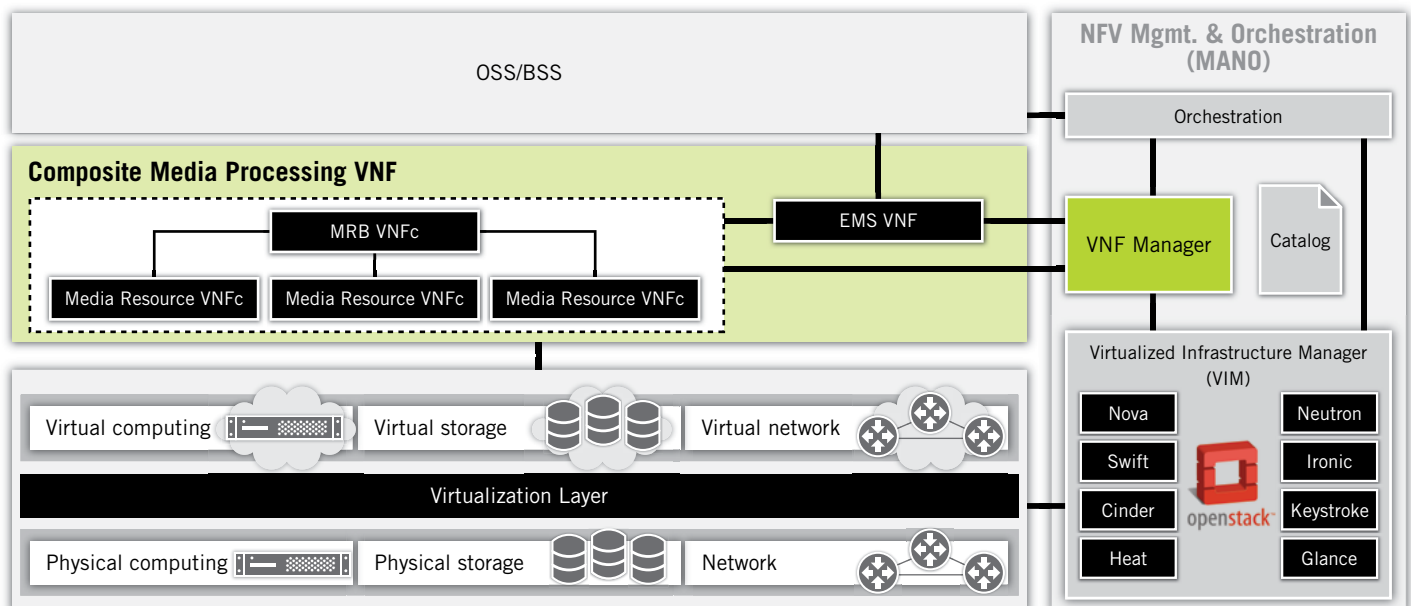


Figure 6 – Composite Media Processing VNF

For the sake of example, assume that additional MRF conferencing resources are needed due to increased demand. If the MRF is a single monolithic piece of code, then the only way to support this requirement would be to create another instance of a MRF VNF. However, this approach could result in the following:

- Additional time to load the OS and associated application;
- Additional CPU and storage resources being required.

With the MRF designed in a modular fashion with programmatic interfaces it is possible to add only additional media mixing and bridging resources to respond to increases in user demand. In practice, this process can be automated with analytics from the MRF VNF. For example, on-demand MRF VNF instantiation can occur based on monitoring and reporting of key performance metrics such as capacity, media, system resources or other KPIs collected by the MRF VNF and supplied to VNF management and orchestration functions. These metrics, especially metrics on RTP performance, collected in real-time as the MRF processes various media requests are invaluable in the analysis and assessment of end-to-end quality of service (QoS) and quality of experience (QoE) measurement. Based on this data, necessary VNF instances can be spun up when and where they're needed to reduce delay and improve overall application performance.

An important point is that in moving IMS/VoLTE applications to the cloud requires special consideration. When it comes to media plane processing, applying the guiding principles when porting network functions to a virtualized environment can help achieve a CSP's NFV goals including profitability.

Virtualization of the IMS: Media Gateway Control Functionality in the Cloud

The IMS core in VoLTE applications contains multiple elements involved in the control and routing of voice and video sessions. CSPs are already working on moving this functionality to cloud environments and in doing so are opening up opportunities to further decompose the capabilities found in more monolithic functions such as call session control, media gateway control, breakout gateway control, and interconnect border gateway control. One of the many advantages of implementing a modular software architecture is that there may be times when, in order to increase performance, a CSP wants to move a VNF or VNFC closer to the user. This is easy to do with a VNF and very difficult to do with a Physical Network Function (PNF).

Similar to the MRF, the MGCF is an example of a capability, which is best implemented in a modular fashion with programmatic interfaces. In the IMS, the MGCF provides key functionality that supports routing and interworking of media and signaling between the core IMS/VoLTE network and TDM networks. Some of the functions of the MGCF include:

- Signaling control of calls from the IMS to the PSTN;
- Media channel control of associated IMS Media Gateways;
- Determination of next hop for inbound calls from the PSTN;
- Signaling protocol conversion between TDM, SIP-I and the IMS.

As part of implementing the MGCF in a modular fashion, the MGCF can be decomposed into several core-processing components, for example:

- Service Logic Execution (SLE): Service creation module used to developed enhanced services
- Route Policy Engine (RPE): Route policy and session treatment module;
- Element Management (EM): Systems monitoring, alarming, and diagnostics;
- IP Call Element (ICE): Management of incoming and outgoing voice IP calls;
- SS7 Call Element (SCE): Management of incoming and outgoing SS7/SIGTRAN calls.

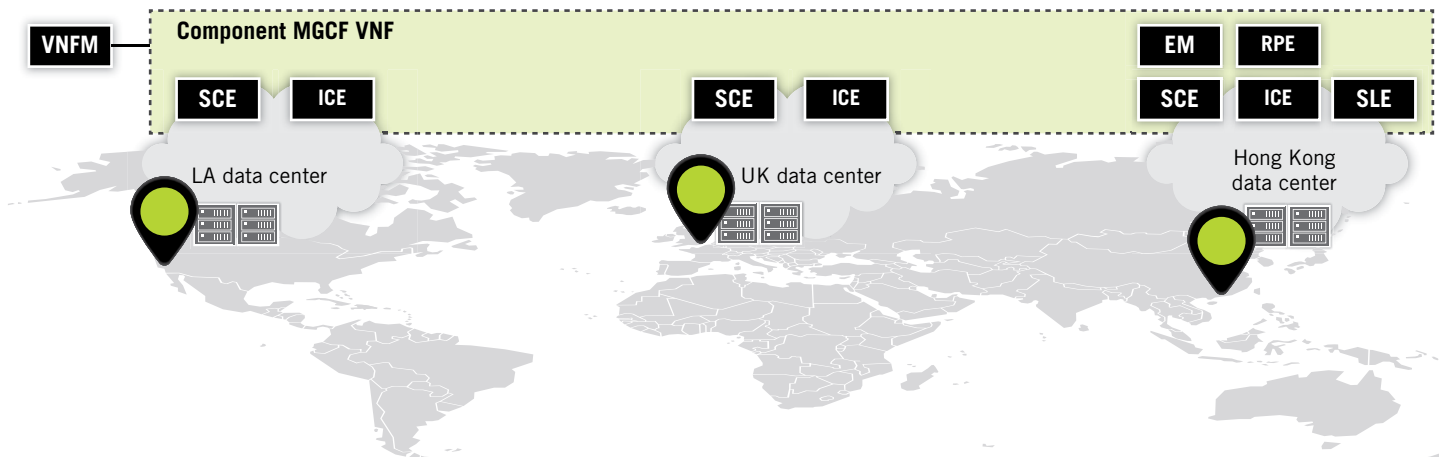


Figure 7 – An MGCF VNF with decomposed functionality can spin up VNFC instances of specific resources where and when they are needed to provide more efficient performance

These modules work together in a distributed fashion in order to process, modify, interwork, route, and charge sessions and calls that traverse between the IMS and the PSTN.

In contrast to a modular software architecture that was designed to run in a virtualized environment, a monolithic software architecture ported to a virtualized environment will have to scale the entire MGCF or spawn a new instance of the entire application if any component of the application runs out of resources. A modular approach will most likely scale differently, and if the VNF is not decomposed properly, it can make inefficient use of the underlying virtual and physical resources. This could lead to delays in scaling up to meet bursty customer demand.

The preceding discussion focused on how to implement the MRF in a modular fashion with programmatic interfaces. If the MGCF is designed in a similar fashion, it's possible to cost efficiently scale the MGCF by only adding the resources that are needed to support an increase in demand. For example, when faced with an increase in VoIP calls, it is highly efficient to only add an instance of the VNFC that is reaching exhaust (e.g., the IP Call Element) where and when it is needed.

When a VNF such as an MGCF is virtualized, it makes it easier to establish a checkpoint on the deployment of a new version and roll back the upgrade if there is a problem. This approach to deploying new functionality reduces risk and would be extremely difficult to implement with a PNF MGCF. Risk is not only mitigated, but also being able to test new functionality in a virtualized lab environment reduces the time and the cost of testing due to the simplicity of downloading a new version of software versus acquiring and implementing new hardware. For example, if the MGCF is virtualized, it is easy to prepare new instances of it in a CSP's lab, test those instances and move them into production. This approach ultimately results in less down time, less operational cycles devoted to testing, and an increased likelihood of interoperability with the rest of the network than would be the case if the MGCF were a physical appliance. In addition to saving time and thus money, being able to test new functionality in a virtualized environment enables CSPs to roll out those new capabilities more quickly to customers.

Summary

When choosing which VNFs to acquire and implement, CSPs should consider the modularity of the underlying vendor's VNF architecture and its ability to adapt to newer virtualization technology while supporting existing approaches.

Some of the ways that a forward looking approach to implementing VNFs impacts the profitability of CSPs moving towards a cloud-based infrastructure environment include:

- **Cost Savings:** Making it possible to scale capacity by adding just the needed functionality and not all of the VNF, saves cost;
- **Faster response:** With efficient software design and use of advanced virtualization techniques such as containers, scaling the needed functionality to support peak loads gets faster;
- **Increased Network Efficiency:** Implementing a smart VNF architecture that allows for putting resources closer to the sources, and sinks of traffic can help keep latency minimized. This is important for real-time multimedia services, and for CSPs with a global footprint;
- **Lower Risk:** Testing in a virtualized environment results in an increased likelihood of a successful upgrade and better ability to rollback an upgrade if there is a problem;
- **High Availability:** Instances of needed functionality can be instantiated in less time, which leads to higher availability and faster recovery from failures;
- **Faster Time-To-Market:** The time to test new functionality is reduced which lowers the cost of testing and shortens the time needed to bring new infrastructure and services to market.

About Jim Metzler

Dr. Metzler has worked in many positions in the networking industry. This includes creating software tools to design customer networks for a leading communications service provider (CSP); being an engineering manager for high speed data services for a different CSP; being a product manager for network hardware; managing networks at two Fortune 500 companies; directing and performing market research at a major industry analyst firm; and running a consulting organization. Jim's current research interests include Application Delivery, Software Defined Networking (SDN) and Network Functions Virtualization (NFV).



Jim holds a Ph.D. in Numerical Analysis from Boston University. He has been on the faculty of several universities including Bentley University, Northeastern University and Drew University. He has co-authored a book, published in the Prentice Hall series in computer networking and distributed systems, entitled "Layer 3 Switching: A Guide for IT Professionals". He writes for numerous trade magazines and is a frequent speaker at conferences and seminars.

About Dialogic

Dialogic has a rich history of software innovation and is helping CSPs reimagine and deploy core capabilities in their networks and applications using a cloud-ready approach. Dialogic fully embraces NFV and as such, this movement towards software-based network processing has been an intrinsic part of Dialogic's DNA and is exemplified by the visionary developments in our applications and participation in public cloud POCs and industry organizations. Dialogic is primarily focused on developing virtualized network infrastructure solutions and our products include solutions like media processing, virtualized SBC, softswitch and signaling controller applications.

Dialogic is a member of OPNFV and a participant in the ETSI NFV Industry Specification Group. To find out more information about Dialogic, go to www.dialogic.com.

Dialogic®

www.dialogic.com

For a list of Dialogic locations and offices, please visit: <https://www.dialogic.com/contact.aspx>

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH PRODUCTS OF DIALOGIC CORPORATION AND ITS AFFILIATES OR SUBSIDIARIES ("DIALOGIC"). NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN A SIGNED AGREEMENT BETWEEN YOU AND DIALOGIC, DIALOGIC ASSUMES NO LIABILITY WHATSOEVER, AND DIALOGIC DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF DIALOGIC PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY INTELLECTUAL PROPERTY RIGHT OF A THIRD PARTY.

Dialogic products are not intended for use in certain safety-affecting situations. Please see <http://www.dialogic.com/company/terms-of-use.aspx> for more details. Dialogic may make changes to specifications, product descriptions, and plans at any time, without notice.

Dialogic is a registered trademark of Dialogic Corporation and its affiliates or subsidiaries. Dialogic's trademarks may be used publicly only with permission from Dialogic. Such permission may only be granted by Dialogic's legal department at 6700 de la Cote-de-Liesse Road, Suite 100, Borough of Saint-Laurent, Montreal, Quebec, Canada H4T 2B5. Any authorized use of Dialogic's trademarks will be subject to full respect of the trademark guidelines published by Dialogic from time to time and any use of Dialogic's trademarks requires proper acknowledgement.

The names of actual companies and products mentioned herein are the trademarks of their respective owners. Dialogic encourages all users of its products to procure all necessary intellectual property licenses required to implement their concepts or applications, which licenses may vary from country to country.

Any use case(s) shown and/or described herein represent one or more examples of the various ways, scenarios or environments in which Dialogic® products can be used. Such use case(s) are non-limiting and do not represent recommendations of Dialogic as to whether or how to use Dialogic products