White Paper | January 2017

radisys®

Elastic Scale Media Processing in NFV

By: Adnan Saleem, CTO - Software and Cloud Solutions, Radisys
Mohan Aravamudhan, Senior Product Manager - Virtualization and Cloud, Radisys

Introduction

Network Functions Virtualization (NFV) architectures continue to expand across the telecommunications industry due to its benefits of flexibility, scalability, and efficiencies offered by virtualized environments. Radisys released a Media Processing in NFV Architectures white paper in December 2014 that reviews the general benefits that NFV architecture offer to service and solution providers and describes the challenges posed by fully virtualized environments for real-time applications in general, and media processing applications in particular. This white paper is an evolution of that earlier work based on real-world integration challenges and experiences. To overcome these challenges, Radisys developed the MediaEngine™ Virtualized Media Resource Function (vMRF), a media processing powerhouse used for a wide range of revenue-generating interactive HD audio and HD video services. This software-only solution is highly optimized to deliver exceptional media processing performance in virtualized and cloud environments.

CONTENTS

Industry Trends

NFV quickly caught on in 2015 and 2016 with deployment announcements by global service providers that include AT&T, Verizon Wireless, Telefonica, Telstra, SK Telecom, Swisscom, Vodacom, and others. The year 2017 looks to be the year that NFV moves from proof-of-concepts and trials and into the real world with commercial production. Due to improved NFV and Software Defined Network (SDN) hardware and software, progressively more of the network infrastructure came into scope. Plus, the cost and business benefits for NFV implementation have been anticipated to grow over time due to the following key factors:

- Improved maturity and experience from NFV/SDN transformation
- Continued increase in processing and I/O capabilities of commercial (x86) servers and OS software
- Improved capabilities and performance of NFV based applications
- Broader range of traditional network elements supported as NFV based applications
- Improved efficiency of capacity utilization in a virtual environment – optimized use of available infrastructure via scalability and reuse
- Ability to consolidate (reduce the number of) hardware boxes

NFV most importantly provides the ability for operators to launch new services with much faster time to market and to respond to user expectations at faster speeds. NFV also allows operators to decouple software applications from physical location and vendor lockin.

Top NFV drivers include Deliver Agility and Flexibility, Reduce Operational Expenditures, Accelerate Time-to-Market, and Reduce Capital Expenditures.

Some additional trend statistics supported by NFV industry growth forecasts include:

- Between 2015 and 2020, the NFV market is expected to grow at a compounded annual growth rate of 42%, from $2.7 billion in 2015 to $15.5 billion in 2020
- Software is a much larger investment than the server, storage, and switch hardware, representing about $1 spent on hardware for every $4 spent on software (NFV MANO plus VNFs)
- By 2020 VNF revenue will amount to $11.4 billion worldwide
- VNFs growth at CAGR 40% from $2 billion 2015
- Market size for VNFs pertaining to the IMS core and SBC will grow from $689 million in 2015 to $2 billion in 2020 at 24% CAGR
- 7% of core IMS and SBCs sales were VNFs in 2015, heading to 57% in 2020
- VNF sales related to IMS/SBC will mount to $1.1 billion in 2020 from $117 million in 2015

In summary, industry forecasts are showing that network operators are rapidly embracing the trend to migrate networks based on physical networking equipment into flexible software-based virtualized networks. To embrace these opportunities, standards efforts around the world are in rapid development to achieve these growth and benefit expectations.

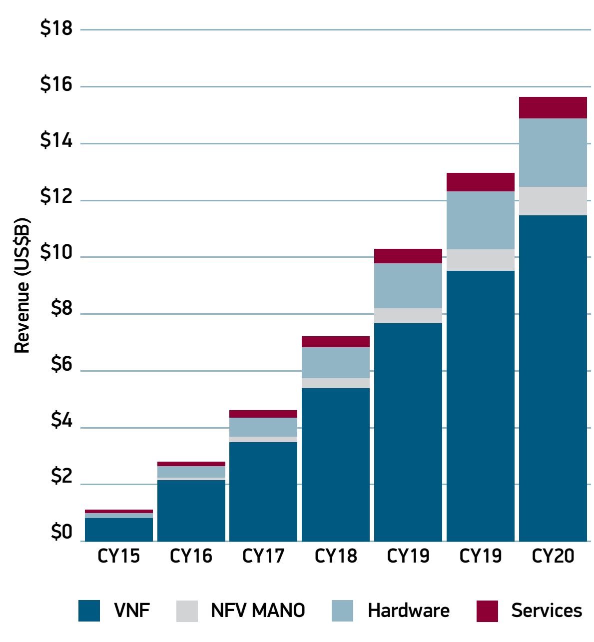

Figure 1: VNF Worldwide Market Software Forecast (breakdown by functional category)

Overview of NFV Architecture

Achieving the growth potential of virtualized media processing in a multi-vendor NFV architecture will require standardized implementation and integration interfaces amongst the NFV architectural components. The industry has largely embraced the original European Telecommunications Standard Institute Industry Specification Group (ETSI ISG) NFV architectural model (Figure 2).

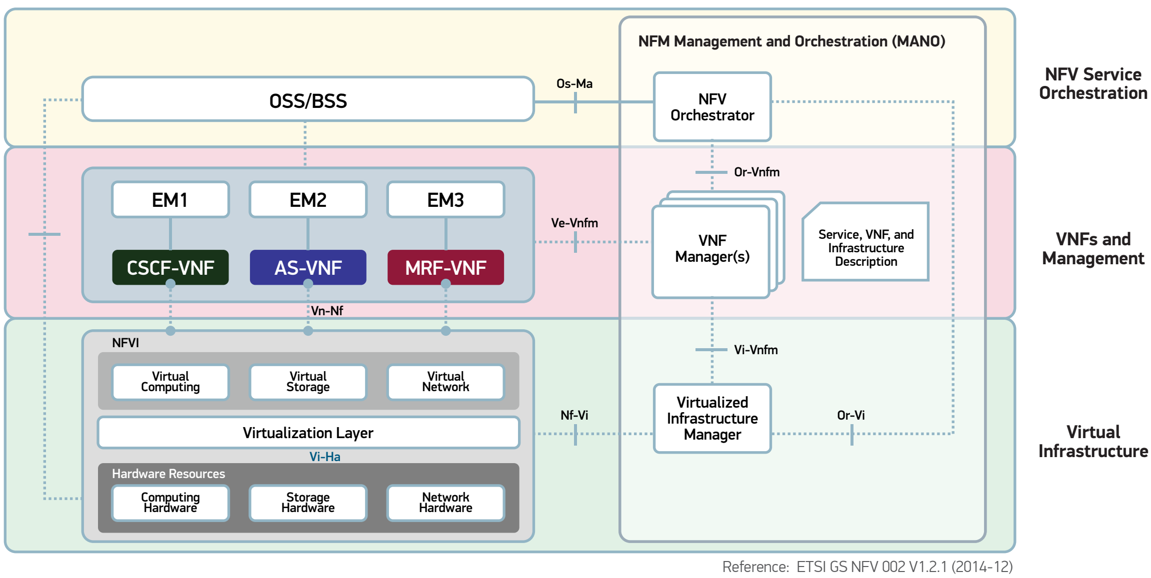

Figure 2: ETSI ISG NFV Architectural Model with IMS VNFs

Key components of the ETSI ISG NFV architectural model relative to elastic scale media processing include:

- **NFV Management and Orchestration (MANO) -** provides the overall the management and orchestration of all resources in the cloud data center, which includes compute, network, storage, and virtualized machine resources. The main focus of MANO is to allow flexible on-boarding and sidestep the chaos that can be associated with rapid spin up of network components.
- **Virtual Network Functions (VNF) Manager -** VNF Managers are critical for scaling, changing operations, adding new resources, and communicating the states of VNFs to other functional blocks in the NFV MANO architecture.
- **VNF Element Management Systems (EMS) –** The VNF EMS is responsible for the management of VNF operation, in the same way as physical network elements are managed by their respective EMS to support, for example, fault and performance management. The EMS may manage the virtual instances via VNF specific interfaces. One EMS may manage one VNF or multiple VNF instances and may report VNF performance and state to the associated VNF-Manager.
- **VNF Instances –** The Radisys MediaEngine vMRF is designed to perform as a VNF in an ETSI NFV architecture, shown in red in Figure 2. In an IP Multimedia Subsystem (IMS) architecture, an MRF is controlled either by an IMS Call State Control Function (CSCF) – green VNF in Figure 2, or a Telecom Application Server (AS – blue VNF in Figure 2), using SIP commands, or XML-based control languages like Media Server Markup Language (MSML) or VoiceXML. The MediaEngine vMRF is designed to support a OneMRF strategy, where one media processing platform can support many real-time communication applications, as shown Figure 3.

MediaEngine vMRF is designed for VNF scaling and redundancy flexibility. In an NFV architecture, if you had N applications hosted on N virtualized application server VNFs, you might instantiate quantity=N MediaEngine VNFs. As traffic grows, the NFV architecture could "turn up" additional MediaEngine VNFs to meet traffic demand, and then "turn down" the VNFs during low demand freeing up NFV resources for other applications and VNFs. Finally, the MediaEngine vMRF can also be configured in redundant configurations, in N+1 or 1+1 high availability configurations, controlled by the MANO.

VNF instances always have a VNF Manager

(VNFM). A VNF Manager can have one or multiple VNFs. ETSI defines 2 types. Generic-VNFM and Specific-VNFM. MediaEngine vMRF does not have its own S-VNFM. Instead, we work with other vendor's G-VNFM.

Figure 3: OneMRF for many services

Radisys continues to be a leading system integrator in the telecommunications industry and collaborates with Tier 1 operators and vendors helping them solve the complex challenges related to real-time high capacity media processing. Radisys understands the different approaches of on-boarding and scaling media processing VNFs in an ETSI NFV architecture, and is in a unique position because of their relationships with their service provider customers.

Onboarding and Scaling MediaEngine vMRF Resources

In the previous section, we outlined the key components of the ETSI NFV architecture relevant to elastic scale media processing. In this section, we will walk through an example of how NFV components work together to deliver onboarding and scaling of media processing VNFs. Scale In/Scale Out is defined as the ability for a VNFM to add or remove VNF instances, based on the KPIs (Key Performance Indicators) reported by the VNFs. Auto scaling, or the ability to dynamically scale in/out, is a key requirement towards achieving the benefits of an NFV architecture.

The following is a five step example of MRF onboarding and lifecycle management in NFV and MANO:

1. **Onboarding vMRF VNF –** This is the first important step in introducing a new media processing VNF type (such as the MediaEngine vMRF) into an NFV cloud architecture. Radisys vMRF is packaged and deployed using Open Virtual Format (OVF), a standard file container for packaging the vMRF application with virtual machine and guest OS for cloud deployment. Onboarding also involves providing resource descriptor and configuration parameters to the MANO. This is achieved through TOSCA and HEAT templates, to add a MediaEngine VNF type into the MANO's VNF Catalog. Once this step is completed, then the MANO is now ready to start to instantiate media processing VNFs.

2. **Orchestrator Creates Service –** The MANO makes decision to deploy a new media processing VNF. It searches the VNF catalog and finds that MediaEngine vMRF is registered. So it sends commands to VNF Manager to instantiate a vMRF instance, plus updates the NFV instances table.

3. **VNF-M Instantiate MRF1 –** The VNF Manager instantiates a MediaEngine vMRF instance into production on the NFV infrastructure. We'll call this first instance MRF1.

4. **VNF Instance Reports Threshold Alert to VNF-M –** Instance MRF1 continuously reports Key Performance Indicators (KPIs) to the VNF Manager. When KPIs exceed a threshold condition, such as CPU utilization beyond a certain desired level, then the VNF-M determines that additional media processing VNFs are required.

5. **VNF-M Responds by adding MRF 2+3 –** The VNFM responds to the threshold alert by instantiating two additional MediaEngine VNF instances on the NFV infrastructure. The VNFM will also update the NFV instances table, so the MANO knows that three media processing VNFs are now available for media processing capacity. The MANO can then share the load from media processing requests (from the CSCF or AS VNFs) across the three instances.

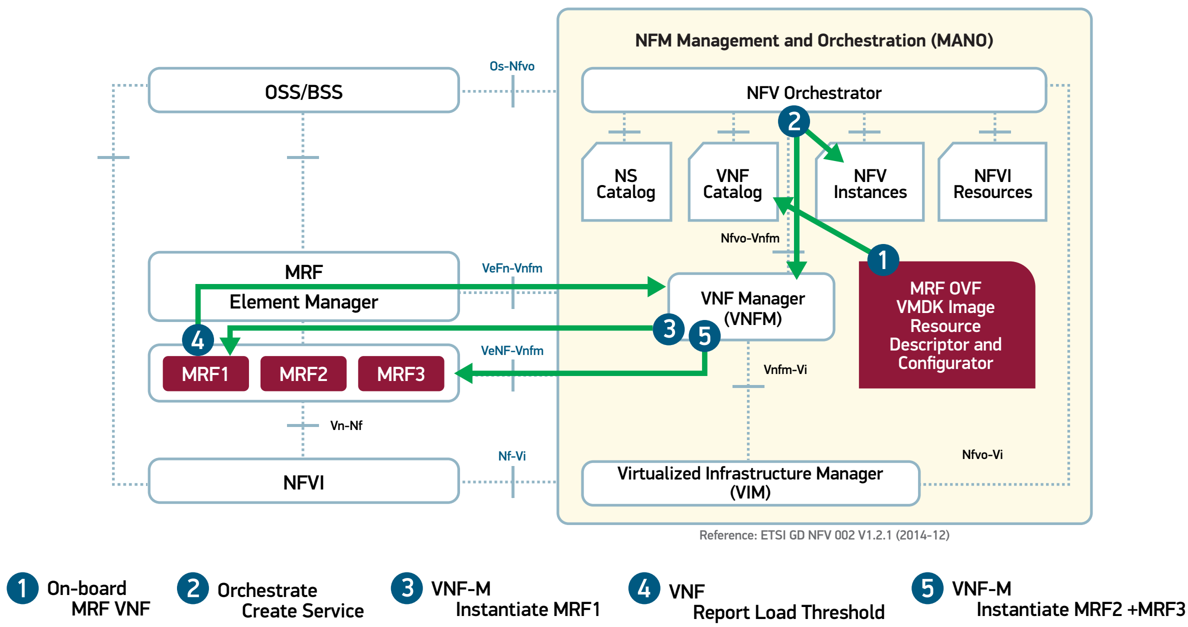

Figure 4: MRF Onboarding and Lifecycle in NFV/MANO

The five step example above walked through scaling up and dynamic auto-scaling of media processing VNFs due to increased demand. Scaling down follows the same process in reverse. As demand upon a media processing VNF declines, an "idle" threshold alert would be reported to the EMS. The EMS reports the idle alert to the VNFM, which then "scales down" the number of VNF instances to better match the current traffic load.

Implementation Considerations

While the architectural principles and interfaces provide a foundation for elastic scaling of media processing resources in multi-vendor NFV implementations, the devil is in the details. This section outlines some of the challenges in integrating virtualized media processing in multi-vendor NFV environments, and how Radisys MediaEngine vMRF products, engineered for operation in a multi-vendor NFV architecture, address the challenges.

Not Everyone's MANO is the Same

For all the work to develop MANO, there is a surprising lack of consistency in how "orchestration" is defined by vendors. Furthermore, there is a major disconnect between the formal notion of orchestration for NFV and what operators actually need to meet their virtualization goals. That disconnect must be addressed or it will threaten the progress and adoption rates of NFV deployment.

Standards are still evolving and every company's VNF Manager is different and all VNF instances are assumed to have an associated VNF Manager. A VNF Manager may be assigned the management of a single VNF instance or multiple VNF instances. The managed VNFs can be of the same or different types. Finally, VNF manager functions are assumed to be generic and can be applied to any VNF.

Auto deployment VNFs, on the other hand, can be instantiated using cloud application lifecycle management techniques that automate deployment, reducing operational cost and time-to-launch. The NFV architecture taps into the agile application lifecycle management opportunity that cloud delivers to achieve operators' service agility goals. However, VNFs need further configuration at runtime to fulfill customer-specific services. A cloud deployment template can help to turn up a new VNF instance so that it is operationally ready, but such a template does not have an API that can configure the VNF dynamically, on demand. This is where runtime configuration approaches are needed to complete the fulfillment process. NETCONF/YANG is an example of one such approach that delivers a programmatic configuration experience that can be harnessed to support the automated lifecycle management goals of NFV.

To overcome the variations in early MANO implementations, the Radisys MediaEngine vMRF includes certain flexibilities and features. As a standards-defined IMS core network element, this solution integrates well within OpenStack-defined management and orchestration capabilities. The NFV Infrastructure (NFVI) Manager manages the complete suite of network elements comprising this infrastructure, including compute, network I/O, and storage. The VNF Managers, as well as the

ZTE Standardizes IP Service Delivery Solutions on Radisys' Virtualized MRF

Mobile operators are deploying NFV to bring the benefits of the cloud to their networks, such as scalability, increased efficiencies and reduced costs. By leveraging Radisys' common media processing platform and TOSCA templates for all of its cloud-based offerings, ZTE is helping its customers accelerate the deployment of new interactive HD audio and HD video services, while reducing upfront capital expenditures and ongoing operating expenses.

"To enable our mobile operator customers' transition to an NFV environment, ZTE is migrating its IMS solutions and VoLTE core network service offerings to a cloud architecture, driving an important requirement for virtualized media processing. Radisys' vMRF delivers HD audio and HD video media processing and transcoding that are required by today's increasingly multimedia-intensive applications", said Yang Wei, Chief Architect of Core Network Deployment, ZTE Corporation.

ZTE is a market leader in value-added service deployment, with 20 percent worldwide market share for voicemail that includes major deployments throughout China, India, Latin America and Europe. China Mobile will leverage ZTE's IMS core and ring-back tones solutions, all powered by Radisys' virtualized MRF.

"ZTE is a long-time customer of Radisys, having deployed Radisys' hardware MRF portfolio for many years to provide the critical media processing for our value-added services, such as voicemail and ring-back tones, used by more than 600 million subscribers around the world. As we looked to transition our VAS service offerings to the cloud, we made the strategic decision to standardize across our IP service delivery solutions on Radisys' virtualized MRF. This will allow ZTE to rapidly deploy new interactive multimedia services such as VoWiFi and WebRTC, while reducing our customers' total cost of ownership", said Huang Xiaobing, Director of VAS Product Planning Department, ZTE Corporation.

Orchestrator functions, enable MRF VNFs to be managed along with other network elements in the signaling domain, such as Application Servers or Session Border Controllers.

Leading Open Source MANO initiatives in the market today include:

- ECOMP
- ETSI NFV Open Source MANO (OSM)
- Open-O
- OpenStack Tacker
- Open-Baton

All of these implementations are being open sourced. The MediaEngine vMRF has been engineered for rapid integration into all open source MANO initiatives. These initiatives will accelerate adoption of NFV deployments.

The VNF Managers run under the control of the higher-level NFV Orchestrator, which is typically responsible for controlling the higher level business service logic—for example, instantiating new vMRF instances during peak conferencing hours, followed by winding them down and releasing resources during off-hours. The Orchestration controller typically needs to be customizable to individual NFV deployments, as it must provide flexibility suited to the business logic. A deployment within a flexible, NFV-compliant management and orchestration platform such as the OpenStack framework is a cornerstone to the capabilities and performance characteristics expected from a media plane element such as the MediaEngine vMRF.

TOSCA and OpenStack Heat

Topology and Orchestration Specification for Cloud Applications (TOSCA) is one of the cloud application deployment template initiatives attempting to establish a standard in the cloud environment. The principles of using a TOSCA template for the deployment phase of the VNF lifecycle equally apply to other types of cloud application deployment templates. TOSCA templates can also be used for later lifecycle operations such as scaling, healing and software update.

Heat is the main project in the OpenStack Orchestration program and implements an orchestration engine to launch multiple composite cloud applications based on templates in the form of text files that can be treated like code. A native Heat template format is evolving, but Heat also endeavors to provide compatibility with the AWS CloudFormation template format, so that many existing CloudFormation templates can be launched on OpenStack. Heat provides both an OpenStack-native REST API and a CloudFormation-compatible Query API.

Nokia Networks Accelerates VoLTE Deployments for Communications Service Providers

Communications service providers are deploying IP Multimedia Subsystem (IMS) functionality as part of their LTE rollouts to enable VoLTE, as well as other enhanced services including video and HD voice. Nokia Networks, the world's specialist in mobile broadband and provider of the most efficient mobile networks, is leveraging Radisys' combination of hardware and cloud-ready software platforms to accelerate media processing deployments for their customers.

"Partnering with Radisys on media processing is a great example of how Nokia Networks can extend its portfolio with solutions from partner companies. In this case, the Radisys product family extends our end-to-end VoLTE solution. We're able to quickly provide operators with the right media processing solution for their VoLTE deployment, whether it's a virtualized solution or not", said Michael Clever, Senior Vice President of Core, Nokia Networks.

Recognizing the importance of helping its customers move quickly in a competitive space, Nokia Networks has chosen to standardize on Radisys' OneMRF portfolio, consisting of the Virtualized MRF (vMRF) and the high-capacity MediaEngine™ MPX-12000 Broadband MRF, for its Open Telecommunication Application Server (TAS) built according to IMS core architecture.

By standardizing on Radisys' vMRF portfolio, Nokia Networks can deploy a OneMRF strategy – one media processing platform for all real-time media processing services and applications. A OneMRF approach reduces CapEx, Total Cost of Ownership and network complexity, while enabling faster service introduction. Nokia Networks' CSP customers can choose either the vMPX or the MPX-12000, depending on their economic and technical requirements for VoLTE and VoWiFi service.

Radisys' vMRF has been certified by Nokia Networks in their NFV cloud architecture, including the use and integration of OpenStack, giving CSPs assurance that their virtualized media processing solution has been pre-tested for flawless operation and integrated into Nokia Networks' OpenTAS solution for a complete end-to-end VoLTE solution.

TOSCA is the open cloud standard, that is enabling an ecosystem that is supported and being adapted by a larger number of industry leaders. TOSCA defines the interoperable description of applications; enabling portability and automated management across cloud providers regardless of underlying platform or infrastructure. TOSCA addresses critical cloud challenges of agility, speed and is agnostic to the underlying cloud infrastructure. The OpenStack Newton Release adopts TOSCA NFV related templates and conversion of TOSCA to Heat. The OpenStack Tacker project addresses the NFVO and VNF Manager. Tacker provides a VNF Catalog to on-board VNF Descriptors were using OASIS TOSCA NFV standards and provides APIs for Life-Cycle Management of VNFs along with capabilities like VNF monitoring, auto-scaling and self-healing.

VNF Key Performance Indicators

As shown in the walk-through section above, the ability of VNFs to report threshold alerts and key performance indicators (KPIs) are critical to achieving dynamic scaling of VNF resources, and the business benefits of the NFV architecture. VNFs deliver the actual network functions that create value; however, they are not autonomous. They require an EMS and VNF Manager. An example of the importance of a VNF Manager is key performance indicator (KPI) monitoring for scaling operations. Ultimately, the VNF Manager maintains the virtualized resources that support the VNF functionality without interfering with the logical functions performed by the VNFs. The services provided by the VNF Manager can be employed by authenticated and properly authorized NFV management and orchestration functions, for example functions that manage network services.

The MediaEngine vMRF implements the subsequent KPIs for VNF Management. The MediaEngine vMRF supports VNF management by receiving KPI query request from VNF manager via REST interface and responding with the corresponding real-time performance data, for example CPU, memory, DSP, Ethernet bandwidth and port license usage. The KPIs are reported at short intervals to ensure that almost instantaneous scalability is achieved. The MediaEngine vMRF VNF Manager interface is based on OpenStack Tracker Project Rest API standards.

Real-Time Performance

Real-time media processing is a challenging application. vMRF products are engineered with careful consideration to minimize processing latency, packetization intervals, and jitter. Engineering vMRF products to run on “bare metal” with real-time performance on dedicated COTS Intel servers is already a tough engineering challenge. Re-engineering a software-based MRF to run as a VNF over a hypervisor in an NFV architecture introduces even more challenges, as hypervisors have a tendency to introduce additional negative performance impacts.

Radisys’ MediaEngine vMRF has been carefully engineered to achieve more than 95% of the performance and capacity available in bare metal configurations while running in a virtualized environment – a significant achievement compared to alternative virtualized media processing solution available today, and one that maximizes return on hardware investment and total cost of ownership.

Real-Time Scaling Implications

NFV needs to be massively scalable to support large numbers of data centers and millions of subscribers. In addition, the scalability of resources will be far more dynamic in a virtualized environment. The major advantage here is elasticity as VNFs can be created, adjusted, and destroyed in real time, and on demand. The MediaEngine vMRF is engineered such that when network triggers are reached, capacity can be dynamically added or removed from the overall network such that capacity and performance can constantly change to reflect the current demand.

Elasticity requirements for license management and billing systems

Elasticity requires automatic scaling via orchestrator and its integration to License Management and billing systems. Radisys has implemented a license server to support elasticity and ease of implementation. The software MRF VNF is responsible of managing licenses and synchronization between different MRF VNFs using the remote license server. The Radisys License Management system provides the foundation to implement features that prevent unauthorized and over usage.

Redundancy and Service Continuity

Resiliency of the network service in a fault tolerant VNF deployment model with fully stateful VNF virtual machine redundancy, means keeping a duplicate copy of full virtual machine state and implicitly all session state. Therefore, if a failure occurs the standby VNF Component will have a replica of the full virtual machine state and will continue providing the service as if a failure never happened. The network access redundancy requires dual access to the service provider's network for site redundancy, such as primary and secondary, to use in the event of an unplanned network outage. The enterprise network must be able to route traffic toward the alternate site in the event of failure. The hardware redundancy is key to maintain service continuity by providing backup hardware systems with machine state preservation to take control in the event of system failure.

Traditional telecom infrastructure demand extremely high reliability for each infrastructure component. OpenStack provides immense horizontal scalability but additionally raises the need for Applications to have high availability feature at the VNF also. The vMRF redundancy provides HA functionality at application layer, independent of the HA capability provided by the cloud infrastructure (e.g.: VMware supports live migration via vMotion, and OpenStack supports virtual machine evacuation to recover from nova compute node failure). The HA capability provided by the virtualized platform can allow the application to migrate to new location and preserve the existing calls in the case of HW or virtual machine failure, however any fatal error in the application can't preserve the existing calls unless redundancy support is available in the application layer. The Radisys MediaEngine vMRF addresses application level by the support both 1+1 and N+1 mode.

OPNFV Models

NFV and VNF are often used interchangeably, both focus primarily on optimization of the network services, contrary to SDN, which separates the control and forwarding plane for a centralized view of the network. VNF is designed to consolidate and deliver the networking components necessary to support a full virtualized environment. However, in an NFV environment, a VNF takes on the responsibility of handling specific network functions that run on one or more virtual machines on top of the hardware networking infrastructure, like routers or switches. Individual VNFs can be connected or combined together as building blocks to offer a full-scale networking communication service.

Open Platform for Network Functions Virtualization (OPNFV), is an open-source software project formed under the Linux Foundation. The goal of the organization is to create a new open-source NFV reference platform, covering both hardware and software. OPNFV will build this reference platform based on ETSI NFV reference architecture. Currently, every operator creates an implementation of this architecture individually, reinventing the wheel. Operators can instead leverage the OPNFV reference platform to focus their resources on new services implementation.

VNFs and PNFs

Hybrid Network Functions (HNF) for Radisys' MediaEngine consists of two deploymnet models, the MRF-VNF (Virtual) and MRF-PNF (Physical), where high density scale is required beyond what VNFs can delivery today:

- Network Service (NS) – Composed of VNFs and/or PNFs
- VNFs – Fully Virtualizable Functions, Standard COTS compute, I/O, Storage
- PNFs – Non Virtualizable Functions, High-density, Purpose Built (eg: DSPs/HW Accelerators)
- Consistent, Common, Reusable Orchestration of VNFs and PNFs

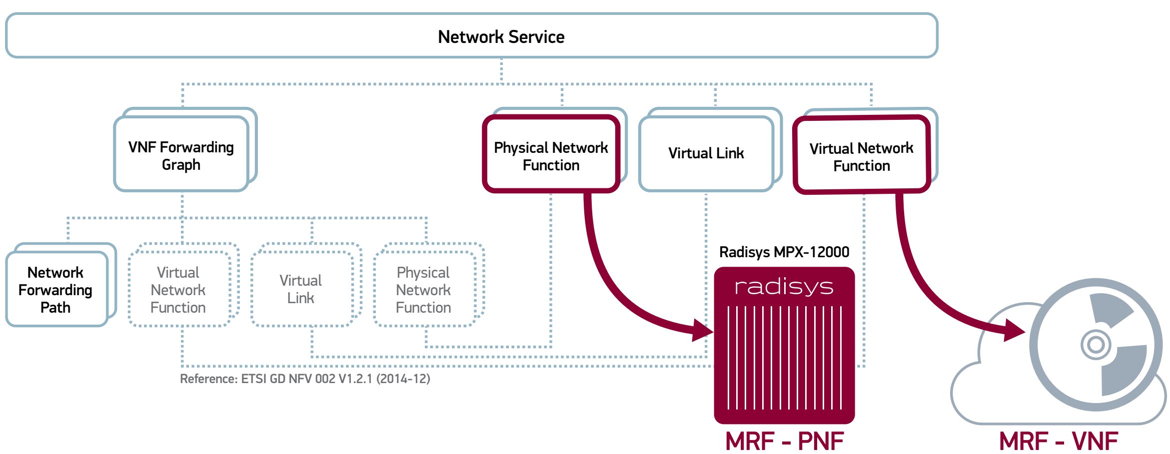

Figure 5: Radisys MediaEngine MRF-VNF (Virtual) and MRF-PNF (Physical)

NFV Architectures using Acceleration Technologies

ETSI standardization efforts are increasingly recognizing that virtualized instances of certain network functions will always be limited without some hardware acceleration. For example, real-time media transcoding is a specific use case called out in ETSI's NFV Report on Acceleration Technologies and Use Cases

The rapid adoption of VoLTE, VoWiFi and WebRTC for multimedia communication services is introducing new HD quality audio and video codecs into the network. Interoperability between new services and legacy devices will require transcoding and/or transrating of the media streams exchanged between end user terminals to adapt them to the capabilities supported by each device. Software-based transcoding and transrating of modern HD audio and especially HD video codecs for real-time communications requires extensive compute processing power and can quickly bog down even the highest performance compute resources in an NFV infrastructure. Hence, NFV architectures that incorporate hardware accelerators can provide a more efficient and effective solution with respect to software-based transcoding in a fully virtualised environment.

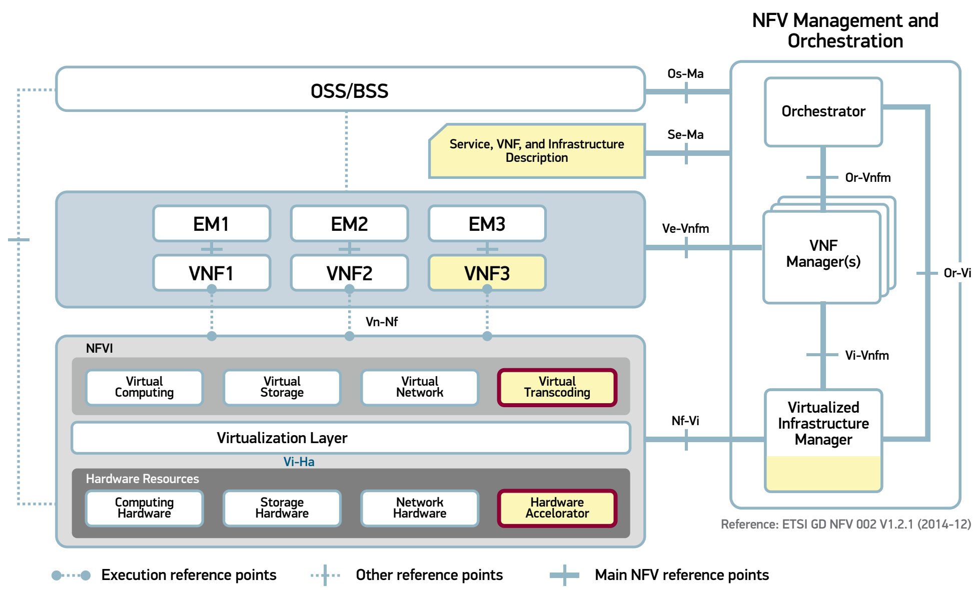

Figure 6: NFV Architecture with Hardware Accelerators for Virtualized Transcoding

Radisys is at the forefront of media processing engineering for real-time communications, and has ongoing development efforts underway to enhance the MediaEngine vMRF software offering for NFV architectures with hardware accelerators based on Digital Signal Processing (DSP) and Graphic Processor Unit (GPU) technologies.

Conclusion

Service providers are expanding their deployments and achieving success in their implementations of virtualized networking solutions. While early growth is promising, wide scale deployment will still require work to achieve operational and economic savings. Growth is predicated on adherence to NFV architecture and open interfaces. There are a number of considerations when implementing elastic scaling for media processing VNFs, including on-boarding, scaling, resilience, redundancy, lifecycle management, MANO, TOSCA, and HEAT, to name a few. Radisys' MediaEngine products provide a number of standards-based control interfaces for broad application server interoperability, delivers the richest audio and video feature set to serve hundreds of media application cases, and boasts deployments around the world. With over 25 years of experience in the telecom industry, combined with Radisys' collaboration with Tier 1 operators and vendors, Radisys can help communications service providers and their IP application developers, to deploy VNFs delivering the promise of elastic media processing scalability.

References

- 2016 Mega NFV Report, SDX Central, 2016
- NFV Hardware, Software, and Services – Annual Market Report, IHS Technology, July 2016
- Terminology for Main Concepts in NFV, ETSI GS NFV 002 V1.2.1 (2014-12)
- NFV Acceleration Technologies and Use Cases, ETSI GS NFV-IFA 001 V1.1.1 (2015-12)

Corporate Headquarters
5435 NE Dawson Creek Drive
Hillsboro, OR 97124 USA
503-615-1100 | Fax 503-615-1121
Toll-Free: 800-950-0044
www.radisys.com | info@radisys.com