GENBAND

WHITE PAPER

NFV, Cloud and SDN

Moving Beyond Simple Virtualization

WHITE PAPER

Table of Contents

Introduction

The evolutionary change of telecommunications networks towards all-IP mechanisms has begun to collapse the traditional boundaries between IT infrastructure and telecommunications infrastructure. The IT datacenter has for years been the location for web-based services and applications. As the telecommunications environment has moved to all-IP, more and more of the functions of the network can be available as pure software functions, deployable leveragint these previously IT centric datacenter architectures. This is certainly evidenced by the emergence of pure-IP communications networks based on IMS and SIP, transported over IP Cable, Fiber, LTE and the like. Increasing the operator's interest in this evolution are the compelling operational savings and capital expenditure efficiencies, as well as the service agility to be gained from moving historically proprietary hardware solutions to a more generic compute environment. Hence the motivation by operators to adopt all-IP networks and the introduction of cloud and Software-Defined Networking (SDN) general purpose compute datacenter architectures, including Network Functions Virtualization (NFV) specific architectures.

This paper examines the key tenets of NFV and SDN that are important for Service Providers considering implementation of these technologies, moving beyond simple virtualization of network functions to commodity hardware and embracing Elasticity, Orchestration, Scalability and Distributed Architectures for signaling and media from this foundation with the larger goals listed above as the end target.

Why NFV and SDN?

Motivations

Increasing the operator's interest in this evolution are the compelling operational savings and CAPEX efficiencies to be gained from moving historically proprietary hardware solutions to a more generic compute environment.

> "The emergence of NFV represents a game- changing shift the likes of which telecommunications has not experienced since the emergence of IP communications"

Examples include:

- Removal of the need for the operator to bear the burden of proprietary hardware platform integration from multiple vendors into their network, the associated replacement and integration costs, and doing so repetitively
- Immunity to hardware obsolescence across multiple proprietary systems
- Minimization of operational interruption from equipment upgrades
- Rapid, low-impact and low-cost Time-To-Deployment for new network functions or changes in scale across existing network functions
- The ability to focus on a single IP infrastructure based on general purpose compute platforms, enabling a highly efficient cost model where due to virtualized functions, those platforms may be used for multiple roles
- The evolution of elastic cloud based IaaS (either public, or their own private) wherein the many network resources of a communications network may be elastically scaled as-required based on actual live network use.

This represents a revolution in cost-optimization and future-proofing to the operator. When we add the follow on benefits of geo-distributed cloud/SDN implementations where dynamic change of the location of assets under software control are possible, and the geo-redundancy flexibility cloud/SDN infrastructure provides; the reasons for this technology shift become obvious and compelling for the operator.

Industry Call to Action and ETSI ISG

Seven of the world's leading telecommunications network operators initiated a new ETSI Industry Specification Group (ISG) for NFV. Over 200 other network operators, telecommunications equipment vendors, IT vendors and technology providers, have quickly joined these operators. The purpose of the NFV ISG is to define the requirements and architecture for the virtualization of network functions and to address technical challenges. The ETSI ISG has published a reference architecture framework for NFV as shown in Figure 1 on next page.

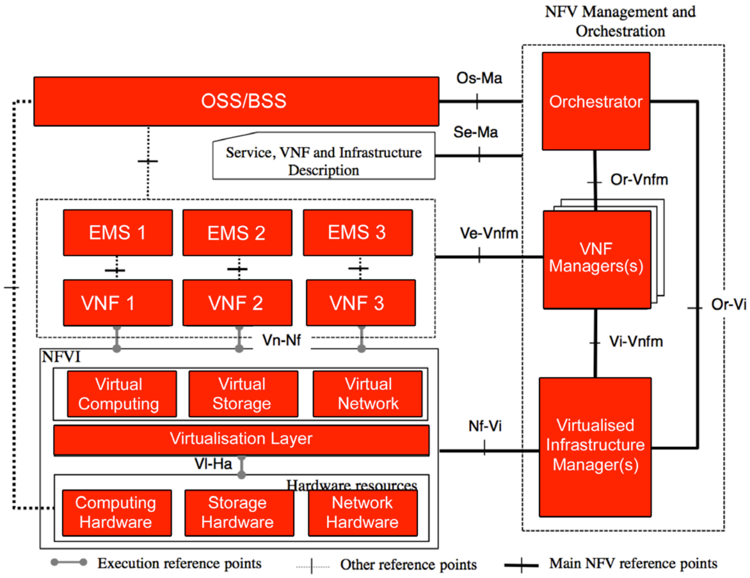

Fig. 1: ETSI NFV General Architecture

The ESTI NFV architecture creates three fundamental critical blocks:

- NFVI or NFV Infrastructure comprised of the underlying datacenter hardware (compute, networking and storage) with the associated virtualization layer, abstracting that hardware to a software controlled Cloud and SDN environment. The Virtualized Infrastructure Manager is closely correlated to this; i.e., OpenStack control functions.
- The VNF (Virtual Network Functions) and their corresponding VNF Managers. These represent the communications applications (previously deployed as custom appliance type hardware solutions) now fully virtualized, supporting elasticity, orchestration, and SDN enablement. Closely correlated with the VNFs are also the associated EMS (Element Management Systems) as well as connectivity northbound to OSS/BSS systems for provisioning, accounting and other functions.
- The new "NFV Orchestrator" enabling for the first time consideration of a standardized environment, key to enabling an operator to integrate VNFs as provided by multiple telecommunications solution providers. This layer further includes the concepts of Service Chaining, where automated deployment/configuration/management of full service-level solutions can be accomplished by orchestrating across multiple VNFs as required to create a complete service provider end user "service".

Beyond Basic Virtualization

Fundamental Components

Cloud technologies, SDN, and ETSI NFV combine various attributes to create the desired environment for flexible virtualized network functions (refer to Figure 2). Cloud technology (versus just simple virtualization of today's static appliances) enables true scalability through compute abstraction combined with support for elasticity mechanisms, and orchestration mechanisms.

SDN technologies are critical in the service provider context to enable real-time, software controlled configuration of underlying networking infrastructure in a fully automated manner to enable support of the dynamically scaling applications in the general purpose (white-box based) compute core. SDN is viewed from two important perspectives, comprised of intra-cloud SDN, and extra-cloud SDN. SDN technologies are used in both domains, enabling dynamic, general purpose compute networking flexibility within the cloud infrastructure, as well as networking flexibility outside of the datacenter in access, bearer, Wide Area Network and walled-garden service provider domains. Although this line is blurring as SDN technology evolves, these two SDN domains are generally managed somewhat independently.

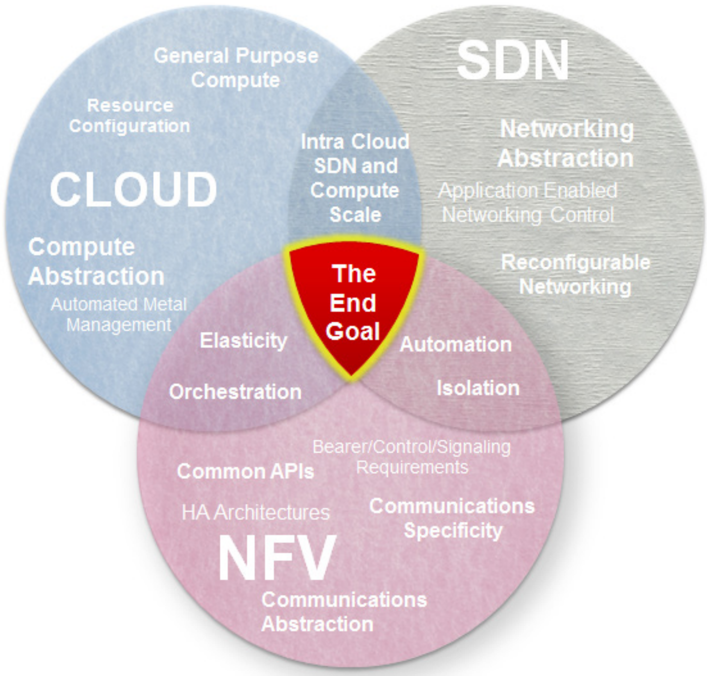

Fig. 2: NFV's Relationship to Cloud and SDN

Finally, NFV (a term associated with the ETSI ISG activity) is a result of the intersection point mentioned above, where all IP based communications and the maturation of datacenter technology intersect to create a contemporary NFV architecture which enables a new paradigm for service provider ecosystems. The emergence of NFV represents a game-changing shift the likes of which telecommunications has not experienced since the emergence of IP communications. While the ETSI NFV activity is in a nascent state, if fully realized the potential of this architecture will enable an interoperable ecosystem where a service provider can more easily absorb the many vendors providing software products into their NFV infrastructure via the standardized structure and APIs provided.

While ETSI NFV matures, in parallel many of the goals articulated by that industry activity (virtualization of network functions as 'VNF's, elasticity in those VNFs, orchestration, and software defined networking infrastructure) are realizable utilizing technology available today. For example, various vendors are realizing the goals articulated above via OpenStack and other cloud infrastructures, combined with various SDN controllers (Neutron/Quantum, OpenFlow) and see the extensibility to other SDN controllers such as OpenDaylight and operator-specific mechanisms as aligned with this approach.

Embracing NFV

In order for the NFV paradigm shift to become successful, fully embracing cloud and SDN technologies is required. For example, simple virtualization (migration to a VM) of a discreet network function is not sufficient to gain the CAPEX/OPEX, operational, and technology benefits of a NFV transition. Embracing cloud and SDN requires wholesale support of elasticity in VNFs, open API support, software automated orchestration at multiple levels in the architecture, and extensively embracing SDN APIs to enable these new approaches. In many cases, this requires re-architecture of network functions to fully realize these benefits.

Elasticity, Orchestration and Scalability

Elasticity and Orchestration are key technologies for enabling the service provider to realize the CAPEX and OPEX savings promised by the NFV architecture. Elasticity fundamentally transforms the use of datacenter infrastructure. Previously, communications applications were deployed on dedicated hardware, and network engineering required worst-case "peak busy hour" sizing of this equipment to ensure the service can be provided during all network conditions. However, not every application is 100% utilized at every time in the day; the dynamic use of the applications varies by time – nor is every application busy at the same exact time.

Enter elasticity and horizontal scaling. Elasticity embraces the concept of dynamically allocating resources on an as-needed basis. When resources are no longer required, they are released back to the datacenter for other applications to enjoy. At higher application scale count, this statistically enables the datacenter operator to purchase, in aggregate, significantly less infrastructure equipment (compute, network, and storage) due to elasticity enabling a 'sharing' of the common infrastructure.

Orchestration is the act of automated control of this new paradigm. Obviously, if the operator were forced to manually scale during peak network demands, including configuration of network connectivity, provisioning, storage (etc), the demands created by elasticity would outweigh the benefits (the CAPEX savings would collapse under the higher OPEX costs of running the new architecture).

Orchestration then provides a fully software based, automated control of product deployment ("orchestrated deployment") as well as dynamic management of running VNFs that scale dynamically. Orchestration and VNF Management work through open Cloud and SDN APIs to enact the dynamic rendering and destruction of virtual machines and network infrastructure to create near hands-off operation of this elastic infrastructure.

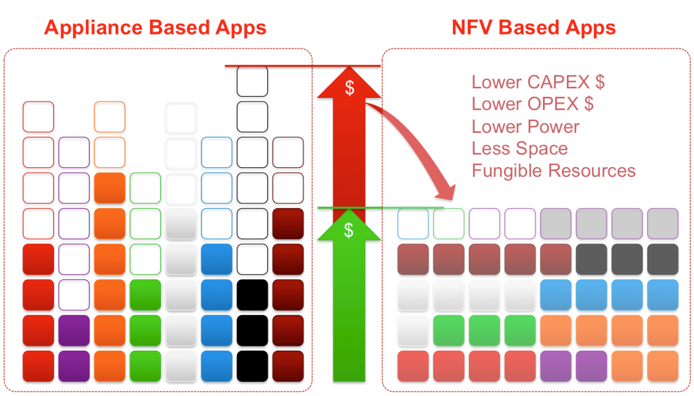

Fig. 3: Benefits of Elasticity in NFV Solutions

Finally, the combination of elasticity and orchestration create the concept of automatic scaling. The previous need to 'truck roll' new appliances to add new network capacity to any service is eliminated. The system as a whole automatically scales, within the bounds established by the cloud infrastructure management (in OpenStack, Nova Cloud Controller). Further, as detailed above, this automatic scaling can scale down equally as well, releasing resources to the other applications to continuously use the minimum infrastructure necessary to deliver the application.

The result is a highly efficient, automatic infrastructure tightly coupled to the applications hosted within it.

Horizonal vs Vertical or "Scaling Out Instead of Up"

In legacy applications, product performance and capacity scale was often directly coupled to the horsepower of individual servers. To achieve higher scale, the next generation processor was often required. This resulted in time to market issues as well as continuous hardware replacement over time. Increasing scale within a product in this manner was called 'vertical scaling'.

Horizontal scaling or 'scaling out instead of up' fundamentally changes the approach by scaling by instance count instead of per instance size. With this approach each instance (as a virtual machine in the cloud) can be much smaller in size, even utilizing sub-server sized resources, and scale is achieved by adding instances until the scaling need is met.

The end result is an approach immune to dependency on processor generational performance, and results in near unlimited scale potential in any datacenter environment without delaying deployment.

Embracing The Future

Realizing True NFV and SDN

Referring to Figure 4, a viable architecture for enabling the key tenets of NFV and SDN is shown. Note that along with the evolution of the elastic VNF applications, a common NFV Manager is created aligned with the ETSI NFV architecture, where the NFV Manager interworks with various APIs to enable orchestrated deployment, elasticity and automation for the VNFs being managed. We will use OpenStack for this example architecture, although, multiple target cloud environments are possible including VMWare, Joyent, Eucalyptus, Apache CloudStack, etc. It is important to note that these solutions must be architected to work equally well in any selected cloud infrastructure.

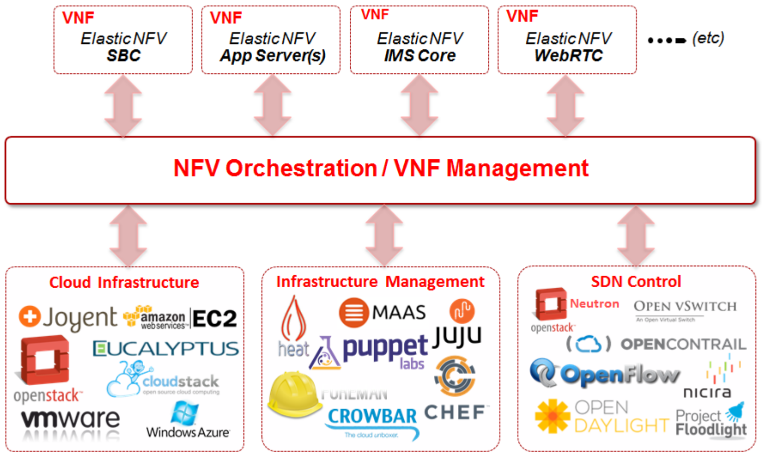

Fig. 4: Introduction of NFV Orchestration and VNF Management

In particular, note the use of MaaS (Metal as a Service) APIs for metal-aware VNF assignment, of orchestration services such as Heat, Puppet, Chef or Juju for automated deployment and management of the infrastructure, as well as APIs to the cloud infrastructure (in OpenStack; Glance for software image management, Nova for VM allocation / scaling, Neutron/ Quantum for intra-cloud SDN control) and finally extra-cloud SDN via OpenFlow or interfaces as appropriate to external SDN controllers. Note that in this context, telecom solution vendors are a supplier of VNFs and VNF Managers, which will utilize SDN infrastructure as deployed by the service provider, and typically depend on the usual datacenter equipment vendors for SDN enabled bearer plane assets such as switches and routers. Also, the VNF Managers tend to provide the role of an SDN application utilizing SDN controllers and SDN APIs.

Not shown but included are northbound REST, NETCONF, and SOAP/XML or other APIs for adaptation of the NFV Manager into a larger orchestrated ecosystem, such as one that a service provider may use to orchestrate VNFs across multiple vendors per the NFV architecture.

Continuing our OpenStack based example, we will add a scalable VNF to further illustrate the enablement of elasticity within the NFV infrastructure. Referencing the high level VNF architecture depicted in Figure 5, the architecture of a elastic and SDN enabled VNF is apparent. Note the application layer, which now scales horizontally, elastically and in real time per cloud tenets, versus vertically as typical in an appliance-based solution. VM instances of the application are dynamically rendered from snapshots under automated elastic and orchestration control of the NFV Manager, via dynamic and stateful manipulation of the Cloud and SDN APIs mentioned above. High availability is provided at the application level, as well as in the underlying cloud/SDN NFVI infrastructure.

> "Elasticity and Orchestration are key technologies to enabling the service provider to realize the CAPEX and OPEX savings promised by the NFV architecture."

Note also that transition to a decomposed VNF also enables particularly successful implementations where the separation of the database layer enables a unique, highly horizontally scalable architecture, effectively decoupling the cloud based data storage scaling from the signaling/application scaling, enabling each to scale independently as required. In this manner, a highly flexible and dynamically scalable VNF is created which will grow based on network load, yet release resources back to the cloud for other applications when the compute resources are not required.

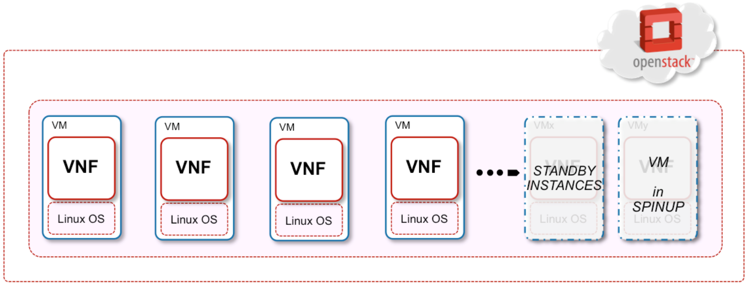

Fig. 5: Elastically Scalable, Orchestrated VNF

Separation of Signaling and Bearer Planes

Further enhancing the concept of composite VNFs, the separation of bearer plane and signaling/management VNFs/VMs into discrete instances with standard APIs between them further enhances deployment flexibility. This enables the operator to design optimized datacenters where compute resources and network infrastructure for bearer plane, and control/management plane, can be optimized in the datacenter environment. Further, this architecture enables geo-distribution of the above mentioned network functions, for ideal placement to optimize the network traffic flow.

This enables and support geo-distributed datacenters with dynamic VNF rendering to accommodate network load events and/or high availability concerns.

NFV Enabled Datacenter Architectures

While NFV may be implemented on typical datacenter infrastructure largely unchanged, an NFV datacenter designed for hosting NFV applications for service provider infrastructure may benefit from several differences as compared to a typical datacenter.

In general, bearer plane infrastructure requires orders of magnitude more bandwidth and is particularly latency sensitive. However, signaling/management applications typically can use lower bandwidth connections. Therefore, rather than burdening the entire Data Center with the requirements of bearer plane traffic infrastructure (networking, compute, etc.) the service provider may choose to optimize the datacenter design to ensure certain types of datacenter equipment are used for bearer plane, different from signaling and management. OpenStack supports this via Zone definitions and geo-separated infrastructure management, further enhanced by metal-layer management features.

> "…all IP based communications and the maturation of datacenter technology intersect to create a contemporary NFV architecture which enables a new paradigm for service provider ecosystems…"

A relevant example would be interface speeds, where 10/40/100G infrastructure is needed for bearer plane, but not for signaling. Additionally, examples also include more powerful servers (such as Intel Crystal Forest enabled servers leveraging DPDK) for efficient IP bearer plane, DSP and media interworking functions. In this manner, significant cost optimization can be obtained in the datacenter, while preventing the complexity and cost associated with equipping the entire datacenter with bearer plane capable infrastructure. A version of this is shown in Figure 6 below where servers to the left are dedicated to various applications; however, servers to the right of the diagram are dedicated to bearer plane processing functionality. These resources are then scaled independently, and may be physically decoupled into geo-separated Data Centers.

Separated datacenter architectures also enable optimizations in the network topology as required differently for applications, versus, bearer plane functions. Examples here include private/public network topologies and security isolation, internal Cloud datacenter network topologies for database access, bonded interfaces for redundancy, and metal scaling methods. Effectively, this difference in Data Center design supports the fundamental difference NFV introduces when not only signaling/management applications are introduced as NFV entities, but also enables optimum support of bearer plane NFV entities.

Selected 'white box' compute platforms for bearer plane applications may produce significantly higher performance for bearer plane functions, such as IP NAT, IP rate limiting, forwarding (etc.) exceeding 10x the performance of other white box solutions when specific compute architectures for the Cloud nodes are selected. Examples here include use of the Intel Crystal Forest server architecture where use of DPDK and SPDK hardware acceleration within that generic compute server can be leveraged for these performance increases.

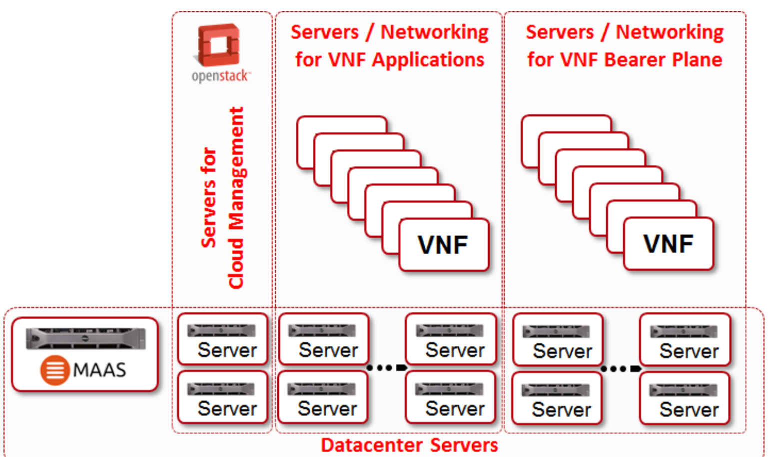

Fig. 6: NFV Optimized Datacenter

Interconnect and routing topologies to and from the datacenter to effectively support access and interconnect IP flows as well as support effective HA architectures are often different for NFVI versus typical Cloud Data Center implementations.

In summary, consideration of the type of NFV VNF that will be deployed on the Data Center Infrastructure, optimization of the selected NFV Infrastructure resources (switching/routing equipment, compute systems, topologies) should be made to ensure the Data Center Infrastructure is optimally designed for NFV.

Conclusion

The NFV evolution represents a watershed event for telecommunications service providers and their vendors. Incredible levels of OPEX and CAPEX reduction combined with dramatically faster time to market for services are enabled by a highly automated infrastructure based on a pure software based environment. As can be seen from the possibilities outlined in this paper, truly unique approaches to leveraging this technology are emerging, ranging from elastic and orchestrated applications, geo-distributed datacenter and communications solutions, and next generation operator datacenter architectures. This is a very exciting time for network operators, as they head into a next generation technology, which is proving to solve critical business issues unique to the communications service provider market.

About the Authors

Paul Miller, Jr.
Vice President, Technology and Strategy
GENBAND

As Vice President of Technology and Strategy at GENBAND, Paul Miller is responsible for the company's strategic vision and technology planning. His efforts are currently focused on GENBAND's NFV, Cloud and SDN enabled strategy and solutions. He is also responsible for developing and managing strategic relationships with key partners and customers.

Miller has over 25 years of telecommunications and advanced applications technology leadership at both large companies and successful startups. He has led the architecture and development of various switching, IMS, IP media, call control and web applications solutions employed by multiple tier-one operators worldwide. His contributions throughout his career have enabled many communications service providers worldwide to create new revenue streams, while dramatically reducing operating costs.

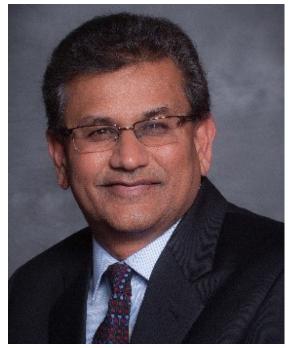

Sanjay Bhatia
Sr. Director, Solutions Marketing
GENBAND

Serving as Senior Director of Solutions Marketing for GENBAND, Sanjay Bhatia oversees the company's strategic marketing organization and alignment of GENBAND's go-to-market strategy across different market segments including NFV, Cloud, Unified Communications and Wireless. Bhatia is an accomplished telecommunications professional with over 26 years of wide-ranging global experience.

Bhatia has held a variety of senior leadership roles in Marketing, Product Line Management, and Research & Development for leading technology companies. He possesses a strong track record of successfully introducing new technologies via innovative product and solutions launches. Bhatia is an accomplished speaker who has presented at several global events.

Prior to joining GENBAND, Bhatia served in senior business unit and marketing roles at Tekelec and various R&D leadership positions at Nortel Networks. Bhatia holds bachelors and masters degrees in Electrical Engineering from Virginia Tech University.