Next-Generation Network Components
Santa Clara, CA
November 6, 2014

M. R. Pamidi, Ph. D.
Editor-in-Chief
IT Newswire

Glossary

EOR	End of rack
ETSI	European Telecommunications Standards Institute
ISG	Industry Specifications Group
MOR	Middle of rack
NFV	Network Function Virtualization
ONF	Open Networking Foundation
SDN	Software-Defined Network
TOR	Top of rack
VNF	Virtualized Network Function

Executive Summary
Explosive growth of Internet traffic and increasing network speeds are driving demand for flexible carrier networks over the course of the 2010s, a decade that is seeing fundamental changes to the system development model, greater growth of mobile data, ever-increasing network speeds, and network virtualization. SDN and NFV, along with the emergence of multicore integrated processor architectures and open source software and standardized hardware, are driving new models of data center network architectures, resulting in noticeable changes by 2020 when services can be deployed in minutes.

Introduction
This was a very well-done, content-rich, one-day event on two of the many hot topics *du jour*—SDN and NFV. We are often asked to distinguish between SDN and NFV.

SDN is abstraction of the control plane from the data plane:

- Abstracts and automates provisioning
- Separates data plane and control plane
- Offers increased flexibility

NFV is abstraction of network functions from dedicated hardware:

- Replaces fixed function systems with virtualized functions on common hardware

Heavy Reading Market Overview
Heavy Reading Consultant Simon Stanley opened the event with a run at rapid speed through a decade of carrier evolution. Starting with 10 GbE circa 2010, this decade has seen networks yield to 100 GbE connections over the long haul, and by 2020 400 GbE networks will be commonplace. 3G was the standard at the start of the decade for wireless; most people now have 4G, and by 2020 there will be significant deployments of 5G, Stanley said.

However, transition from 100G to 400G will have challenges:

- Electrical interfaces and optics—increasing speeds
- Modulation—increasing complexity
- Optical modules—reducing power and size
- Packet processing—reducing time per packet

Fixed network architectures were standard in 2010, with carriers requiring days or months to enable new services. By 2015, architectures will become flexible. "People are delivering new services within a few hours or maybe a few days," Stanley said. By 2020, we'll see virtual networks driven by SDN and NFV, where services can be deployed in minutes.

Stanley (wrongly) said many systems vendors were vertically integrated just four years ago. (Actually, vertical integration was in vogue in 2000 rather than 2010.) Later, systems vendors shifted to a horizontal model—no longer developing all technology internally, but now involving best-of-breed products from third parties. And by 2020 we can expect to see a virtual development model, with network operators able to integrate multiple vendor technologies through standard APIs.

<digress>
We note this is very similar to the auto industry: A hundred years ago, Ford used to own coal mines, steel mills, and glass–making plants making it one of the most highly vertically integrated manufacturers. Today, in an average car anywhere from 60% to 80% of the parts come from third parties. Auto makers have become parts integrators and warranty and service providers.
</digress>

Stanley added components are critical to successful transitions, enabling networks to be built from off-the-shelf pieces including optical, semiconductor, software, and systems. At the same time as the network architecture is changing, bandwidth requirements are increasing on both wired and wireless networks. Data demand is driven by network-intensive applications including cloud services, video/TV on demand, voice/video over IP, and always-on applications such as instant messaging and location service, Stanley added. The Cisco Visual Networking Index: Global Mobile Data Traffic Forecast Update, 2013–2018 reported mobile data accounted for 18 times the data volume in 2013 of the entire Internet in 2000, and will see 10X growth in the next five years. Stanley added it is being driven by smart devices—smartphones, tablets, and laptops.

Network virtualization, including SDN and NFV, is key to allowing vendors and carriers to keep up with demand. Standard platforms form the foundation for network hardware. "Above that you've got a virtualization layer that essentially applies virtual resources," Stanley added. Instead of accessing real compute, storage, and networking, these resources are virtualized, yielding significant flexibility.

Industry Keynote: Orange Silicon Valley
Orange Silicon Valley's Christos Kolias preached a pragmatic approach to virtualization, saying he expects very early deployments to start in the second half of 2015 and increase through the end of the decade. He expects operators to test the technology in new markets and "green projects." "I don't see us going into our existing network and ripping out our Cisco boxes and replacing them with a white box anytime soon," he said. Rather, it will be a gradual process, with operators initially dipping their toes into the virtual world with very limited deployments in the second half of 2015, Kolias suggested. "There will be more widespread deployments in the second half of this decade.

"To me, NFV is all about how you can manage and orchestrate all these new virtualized appliances," he says. This means shifting network functions to the software domain and running that code on "commoditized hardware." He also suggested that NFV applications will initially be used to virtualize elements of operators' network ahead of similar but different SDN technology. "NFV is a little bit more mature than SDN," Kolias said. SDN has further to go on the standardization path than NFV, Kolias suggested. "From our perspective as a telco, we'd like to see a standardized northbound interface," he said of ongoing work on SDN standardization.

He also noted the difference between NFV and SDN is that SDN evolved out of academia and is driven by the needs of the data center players, while NFV development has been largely driven by the service provider sector. Kolias is a founding member of the ETSI NFV ISG and is also responsible for the group's liaison with the ONF. The ISG was founded in 2013 and is expected to release documents based around phase 1 of its NFV work soon. "There's going to be another release coming up in December [of this year] or January of next year," says Kolias. The second two-year phase of the work focusing on "interoperability" and "standardizing" interfaces will begin in February 2015. Kolias warned about vendors touting their wares as NFV-ready or compliant now. "I'm not sure anyone can claim they are NFV compliant," he says. "We don't even have a program to become that."

How big are the SDN and NFV markets? It depends on whom you ask.

SDN is expected to be $1.2 billion by 2018 by one analyst; $3.7 billion by 2018 by another analyst; and a whopping $35 billion by a venture firm funding three startups in this area, this is self-fulfilling prophecy.

NFV is expected to reach $203 million in 2014 and $1.3 billion in 2019, for a CAGR of 46%, according to Research and Markets; $2 billion by 2018, according to Dell'Oro; and $5 billion by 2018 by other sources.

Kolias also summarized SDN and NFV's drivers, challenegs, and applications.

SDN

Drivers

- Network virtualization
- Multi-tenancy
- Automation (programmability)
- Elasticity, agility, simplicity
- Policy control (security, QoS, etc.)
- Analytics, Monitoring
- Lower TCO

Challenges

- Lack of single SDN controller
- SDN controller federation
- Lack of mature protocols, multiple SDN or SDN-like technologies
- Migration paths
- Northbound interface
- Interoperability, compatibility, and portability

Applications

- Data centers
- Enterprise (e. g., wireless)
- Carriers/ISPs
- Cloud

NFV

Drivers

- Faster innovation and new services
- Elasticity and flexibility
- Ease of deployment
- Better DR/BC
- Homogeneous environment
- Lower CapEx/OpEx

Challenges

- Integration
- Abundance of orchestration platforms
- Expensive software licensing
- Interoperability

- OSS/BSS alignment
- Organizational

Kolias also emphasized the need for open networking open source software; open design hardware; open standards; open interfaces, APIs, plugins, and SDKs; and an Open Community, not controlled by a single vendor. Some operators say, "80% of our software is open source, the remaining 20% is our valued-add."

To this end, he noted the September 2014 launch of the Open Platform for NFV, a collaborative, vendor/carrier community supported by the Linux Foundation to accelerate the adoption and evolution of NFV. With 40 members, it is focusing on NFV infrastructure and Virtualized Infrastructure Management.

Finally, do SDN and NFV compete? Not really, they are complementary.

- SDN can play a key role in the orchestration of the infrastructure (physical, virtual)
  - Provisioning and configuration of VNFs
  - Allocate and manage resources (e. g., bandwidth)
  - VM mobility
  - Automation and programmability
  - Security and policy control
  - Centralized network control
- Service composition (NFV Forwarding Graphs)
  - Directing traffic flows to VNFs
  - Traffic flow characterization, especially for mobile, E2E scenarios
- Ad hoc, on-demand, secure virtual tenant networks
- An SDN controller could have VNF or part of service orchestration

NFV and SDN together can create great value. A fine example is Google's Andromeda which one could label SDNFV because it combines SDN and NFV functionalities. It's an enabler to Google Cloud Platform, powering Google Compute Engine and providing an Open Network API. It also offers cloud load balancing, switching, routing, security, ACLs, firewalls, is built to scale, and delivers close-to-bare-metal performance.

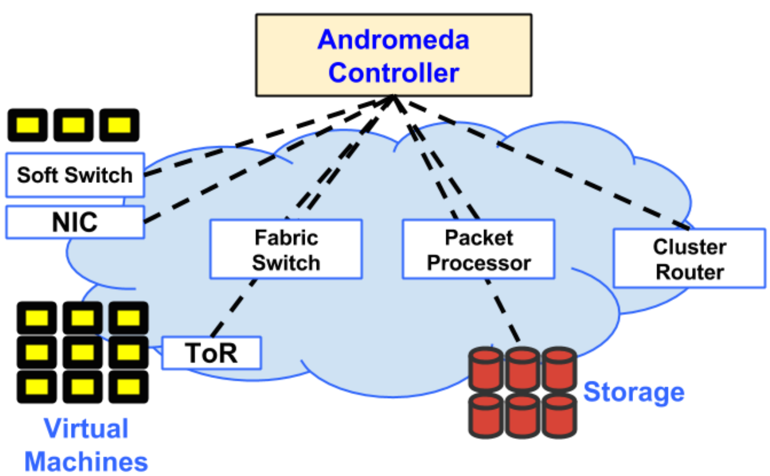

Source: Google

Data Center Interconnects
Server performance and data center network traffic are growing dramatically. Data centers are scaling to hyper, cloud-scale architectures with huge workloads and 100 GbE networks. While mega-scale data centers are being built by the likes of Amazon, Facebook, Google, and Microsoft often with white boxes, there is also a trend towards deploying smaller, localized data centers, especially to meet many governments' (especially in

the EU) restrictions on moving personal data from out of their geographical boundaries. So, while some vendors talk about easily moving workloads from one cloud to another with their cloud fabric, no one is addressing the cost, task, and legality of moving petabytes of data to different clouds. Mega-scale data centers also tend to be more homogeneous for easier systems management. For instance, Facebook reportedly has "hundreds of thousands of machines", over 80% are from Quanta, and claims every operations staffer at one of Facebook's data centers can manage at least 20,000 servers, with some handling as many as 26,000.

There will be major transition from 10GbE to 100GbE in the next year. Reducing component count is a major goal. Optical interfaces on the backplane are still expensive. Data center interconnects will be copper in the foreseeable future, fibre is still too expensive. Copper will dominate TOR, but EOR and MOR may move to fibre. Low latency and high bandwidth continue to be major drivers. The biggest data center challenges in the data center are supply chain management with optical components; power, scale, and workload management; reducing cost and designing software to make the DC fault-tolerant; and flexibility of network infrastructure. Power reduction is becoming imperative with the Kyoto Protocol.

Solutions are being developed to support low-cost 100G optical links up to 2km and optical connections within systems. 28 Gb/s serial links over backplanes and to optical modules are also becoming available and the next generation currently in development will support 56 Gb/s.

SDN adoption is driving next-gen data centers architecture from 1st gen integrated router to 2nd gen split control/data plane to 3rd gen separate control/data plane.

Another factor affecting next-gen data centers is the evolution of integrated multicore processors combining high-speed networking I/O, packet processing acceleration of up to 100 Gb/s, and fourth-generation 28 nm or smaller technology. The leading vendors are:

- PowerPC, Freescale, Avago
- MIPS (Imagination Technologies), Broadcom, Cavium
- ARM, AMD, Applied Micro, Avago, Broadcom, Cavium, Freescale
- X86, AMD, Intel (Atom, Xeon Processor D Family)

Also, chip vendors are putting more and more functionality in to the chip itself, thus blurring the line between a server and a switch. By 2022 we can expect 1,000 64-bit CPU cores with 100X throughput.

Wireless Infrastructure: Developments for LTE Advanced

LTE Advanced features such as carrier aggregation and DPI-based policy management are quickly becoming essential for new wireless deployments. At the same time, carriers are shifting towards heterogeneous networks (HetNet) with a mix of small and macro base stations to provide the coverage and bandwidth that is needed to meet subscriber expectations. The availability of scalable and highly integrated components that support LTE Advanced features is a key issue for carriers and mobile network equipment manufacturers.

SDN & NFV: Assembling the Best Components

SDN and NFV are changing the way service providers plan, deploy, and manage networks in datacenters, wireless infrastructure, and optical networks. These new networks are built using virtual systems that communicate through open interfaces such as OpenFlow and OpenStack and are being implemented on a wide range of hardware platforms, including standard servers, blade servers, ATCA, and rackmount appliances.

Conclusions

Overall, this was an extremely informative conference with a lot of content compressed into a one-day event. SDN and NFV will definitely alter the data center landscape, but not anytime soon. Customers won't rip and replace 25+ years of investments they have made on infrastructure on 'proprietary' networking gear. SDN was initially touted as a "Cisco killer," but Cisco ain't going to yield its turf and has come back with vengeance and a plethora of offerings—ACI, APIC, Cisco onePK, Cisco ONE—making this a landscape worth watching over the coming years.

