Unleashing Network Potential through SDN on Virtual Networking Lab Environments

Pilar Somohano
Oscar Pache
Juniper Networks

Brian Mutcherson
Comcast

Boon Thau Loo
University of Pennsylvania

Abstract

Network operators and MSOs are faced with competitive challenges that require them to optimize their operations and maximize profitability. A two-pronged approach that considers both cost reductions (via improved efficiency) and revenue growth (via the introduction of new services) is necessary to meet these challenges.

Software defined networking (SDN) and network functions virtualization (NFV) are two technologies that promise to help operators meet both of these objectives. However both SDN and NFV represent a (significant) shift from the way networks are architected today. How can operators confidently evaluate and design their networks to leverage these technologies—and others—without risk to their production networks? The answer is again provided by virtualization: specifically, a virtual networking environment where SDN, NFV and other technologies can be tested and evaluated with a "virtual copy" of the production network.

This paper describes the combined application of SDN, NFV and a virtual networking environment, and how this combination can help operators create and deliver new services while optimizing network operations. The paper will also explore the differences and similarities between virtualization in the network (NFV) and virtualization of the network (virtual networking lab environments).

http://blog.cimicorp.com/?p=1100
SDN was effectively a subset of NFV, focusing on virtualizing and hosting the adaptive behavior of forwarding devices
Both of these things end up in the same place, which is that for Juniper "Software Defined Networking" is about virtualizing "layer 7"
The thing that separates it from being an SDN model is that "layer 7" reference made by Bob Muglia in his presentation. The real goal of SDN, what I've called "purist" SDN, is to centralize the control function. That would explicitly include the replacement of adaptive discovery with central control. Google's SDN application is a poster child for that. Juniper isn't talking about that at all, and in fact they are implying that the control processes remain largely distributed.

http://www.networkworld.com/news/2013/011513-juniper-muglia-265842.html

We think that BGP [Border Gateway Protocol] will be an incredibly important protocol. We will see some extensions to BGP to enable distribution of state. We will also embrace VXLAN [Virtual Extensible LAN] and NVGRE [Network Virtualization using Generic Routing Encapsulation] as data path protocols for Layer 2. I think MPLS will remain an important data path protocol. Frankly, I think OpenFlow has been overhyped in terms of its relative role in all of this. Some of the other protocols, in terms of management, have yet to full emerge.

Introduction

Service providers are faced with competitive challenges that require them to optimize their operations and maximize profitability. These providers require a two-pronged approach to meet the challenge. They need to reduce costs by improving efficiency, and they need to maximize revenue streams by offering differentiated services to win on the market place and afford expanded profitability.

Some of the most promising technologies that will help achieving these goals are:

- Software Defined Networking / Network Function Virtualization
- Virtual Networking Environments and Virtual Replication

Undoubtedly, the flexibility to enable new services and increasingly smarter implementations of the control plane enabled by SDN and NFV open a rich potential of opportunities and efficiencies. To fully unleash this potential, it is necessary to pair this intelligence with adequate tools to design and model the operation of these new applications before they get into production.

This paper describes the use of production network replicas on virtual networking environments to boost the potential of software defined networking to create new services and improve network efficiency.

Development of new applications in a Software Defined Networking environment

The enabling principle of Software Defined Networking (SDN) is a clear separation of networking functional blocks to create a modular approach to service development, network evolution and dynamic management through balanced centralized/distributed deployment of these functional elements.

The table below shows the four functional blocks that constitute the SDN.

Network Management	Unified Network Operation Smart discovery and orchestration
Services Plane	Accelerate creation of differentiated services Layer 7 applications supporting multi-source platforms
Control Plane	Enable adaptive networking applications Protocol management and dynamic routing
Forwarding Functions	Maximize bitrate forwarding performance Minimize power consumption and footprint

Table 1 – Key functional blocks of a Software Defined Network

Each of these blocks has clearly established functions, and therefore, demands different capabilities. The forwarding tier, shown on the table above, is usually based on hardware platforms and its evolution is focused on the exponential capacity growth experienced by service providers. The services and control planes, which are the focus of this paper, are the enablers of new services and network optimization policies and are the most dynamic part of the carrier's day to day operations. Finally, the management system is the centralized brain that coordinates the functional blocks, orchestrating the centralized and distributed components of the network functionality. This system provides access to the information contained on each block of the network, and this functionality is essential for a complete understanding of a particular network setup.

The SDN architecture opens the opportunity to develop and combine multiple “best-in-breed” applications at each functional layer; and most importantly, it facilitates the creation of unique services that can be implemented faster than ever to accelerate revenue generation and market differentiation.

In this fast-evolving environment of new services and optimized performance, it is essential to have a scalable environment where the service and control applications can be tested efficiently, and as fast as the new services and policies are to be introduced on the production network. This early test permits a fast and clean implementation reducing the risk of carrying unnoticed bugs into the production environment.

Scalability test of service and control applications in a Virtual Network environment

When a new service is developed, or new control policies are planned, or when the network grows in capacity or reach, it is necessary to test these changes before they are activated on the production network. The traditional test process uses physical equipment in a lab to verify the new applications and control configurations. This approach is limited in scale, while the inherent richness of SDN applications conveys new complexities that should be addressed in an adequate test environment.

Figure 1 shows a new approach that combines the intelligence available on the network management system and the flexibility of a virtual networking environment to create a virtual replica of the production environment to execute realistic “what-if” scenarios of new service applications and dynamic network control algorithms to perfect them before they enter production.

Using this method, the node state information is used to define the properties of nodes in the virtual environment which are logically identical to the physical nodes on the production network. Because this method is automated, the operators can create, instantly, multiple virtual replicas. With this, multiple test scenarios can be conducted simultaneously efficiently and cost-effectively.

Since the network replica on the virtual environment has the same logical properties of the production network and it is scalable, it allows for a realistic test of the new applications so potential problems are detected and corrected early, before impacting the production environment. This approach reduces guesswork that results in operational problems such as introducing errors on the production environment, and then, attempting to debug them in that live network.

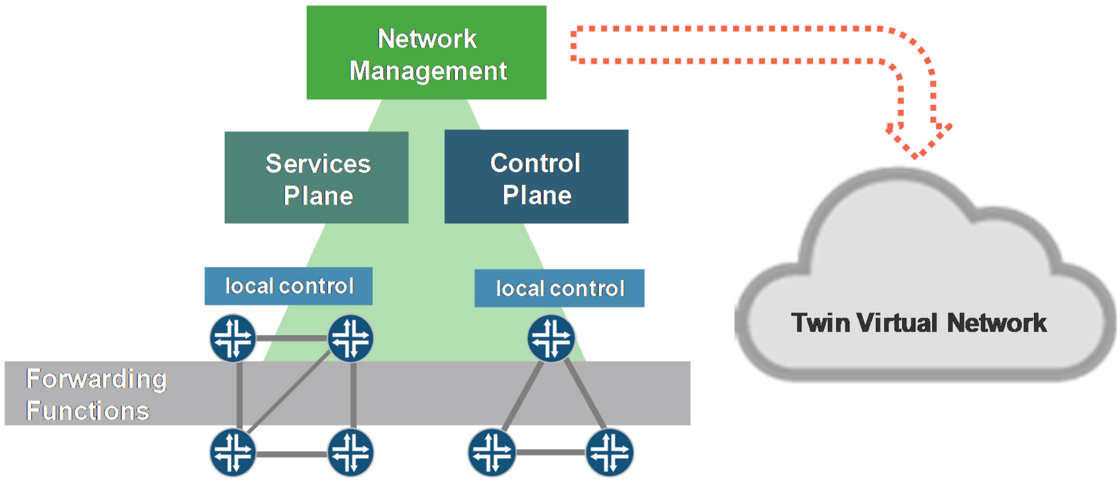

Figure 1 – Virtual replication of a production environment for pre-deployment test

Application Cases

Case I: Introduction of new protocols – BGP-TE

BGP-TE (Traffic Engineering) allows to collect Link State and TE information from networks and to share them with external components using the BGP routing protocol [RFC4271]. It extends BGP with new Network Layer Reachability Information (NLRI) encoding format as documented in the IETF draft (reference 1).

To make better decisions, those external components, like Path Computation Element (PCE) or ALTO Server, require TE information visibility outside one IGP area or Autonomous System boundary. Both applications need to gather information about the topologies and capabilities of the network in order to be able to fulfill their function. This is carried out by BGP-TE with the new protocol extensions.

Key to the success of any new protocol is the early experience provided by diverse entities in the network operator through their feedback during the design and development phases. It is here where the usage of a virtual networking environment plays an invaluable role. We have loaded in this environment an experimental control plane release supporting BGP-TE and exposing it to the operator entities. To do so, we have built a very large topology and setup RSVP-TE LSPs both for Inter Area and Inter AS scenarios as depicted on figure 2.

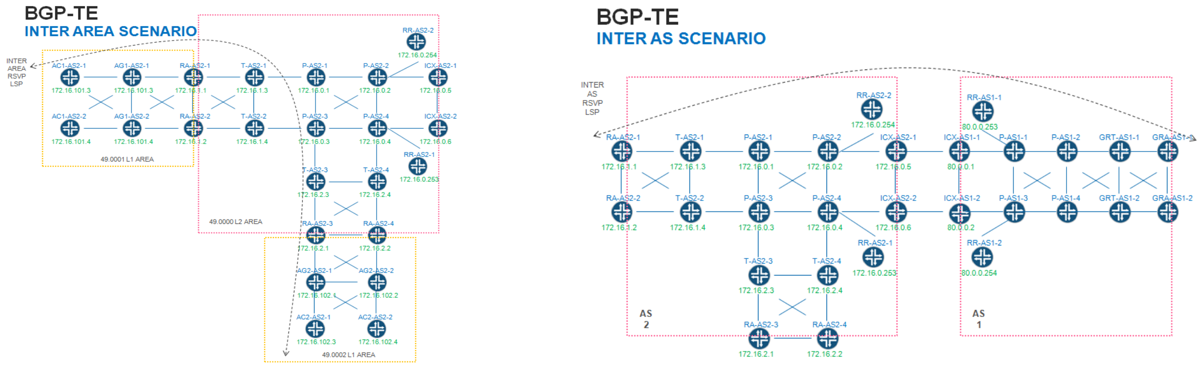

Figure 2 – Topology diagrams used on a BGP-TE evaluation

We have leveraged the capabilities of the virtual networking environment, as large virtualization scale and full ISIS and MPLS support, and we have complemented them with the new BGP-TE extensions.

We are getting very interesting feedback from different groups including Marketing, Planning, Technology and Network Operations, that will be taken into consideration to make a final evaluation about the value of a BGP-TE implementation.

Case II: Adaptive networking via Declarative Programming

At the University of Pennsylvania (UPenn), there is ongoing research work on developing new programming tools and platforms for software-defined networks. The technology, called *declarative networking [3]*, enables network routing protocols to be concisely specified in few lines of code. This approach provides ease and compactness of specification, and others additional benefits such as optimizability and the potential for safety checks.

The development of declarative networking began in 2004 with an initial goal of enabling safe and extensible routers As evidence of its widespread applicability, declarative techniques have been used in several domains including fault tolerance protocols, cloud computing, sensor networks, overlay network compositions, anonymity systems, mobile ad-hoc networks, wireless channel selection, network configuration management, and as a basis for course projects in a distributed systems class at the UPenn. Interested readers are referred to [3] for a survey of recent use cases.

The combination of SDN and virtualization has provided a compelling platform for showcasing declarative networking technologies. Currently, researchers at UPenn are leveraging this infrastructure to explore various adaptive networking techniques that can dynamically modify existing network routing protocols to meet QoS requirements, or to reroute around possible failures and network intrusion. This involves integrating the RapidNet declarative networking engine [4] within the SDN controller for enforcing performance and security goals at the control plane, and then leveraging the virtualization capabilities to evaluate the performance of these protocols at scale.

References

1) North-Bound Distribution of Link-State and TE Information using BGP
   http://tools.ietf.org/html/draft-ietf-idr-ls-distribution

2) Virtual Network Optimizing A Physical Network
   P.Somohano, B.O'Sullivan, H.Stern; M.Yip; A.Mints
   Patent application Ref. No. JNP-2039

3) Recent Advances in Declarative Networking.
   Boon Thau Loo, Harjot Gill, Changbin Liu, Yun Mao, William R. Marczak, Micah Sherr, Anduo Wang, and Wenchao Zhou.
   Fourteenth International Symposium on Practical Aspects of Declarative Languages (PADL), co-located with POPL, Jan 2012.

4) RapidNet declarative networking platform.
   http://netdb.cis.upenn.edu/rapidnet/