VIRTUAL ENVIRONMENT FOR NETWORKING TESTING AND DESIGN

Judy Beningson, Colby Barth, Brendan Hayes
Juniper Networks, Inc.

Abstract

This paper describes the use of virtual environments for the testing, design and modeling of networks. This paper will also explain the architecture and technology behind these virtual networking environments, and will highlight two real world use cases. The paper will also cover the benefits and limitations for cloud-based network modeling and testing to help operators determine the best uses.

INTRODUCTION

Operators who own and run IP transport networks understand that testing new protocols, design changes and/or modeling service introductions can be challenging. Most operators have access to a test lab for such purposes, but these labs have limitations in terms of scale and flexibility. Even the largest test labs do not approximate the size of an actual production network; smaller operators' labs may be non-existent or so small that any realistic control plane scalability testing is simply not feasible.

Due to size, budget availability and space limitations of current physical test labs, it can be difficult to test or design for the same level of scale as an operational network. Additional challenges result from the requirement for physical "racking and stacking". To test different topologies or configurations typically means making changes to physical connections and systems, which can be time-consuming and in some cases can have an impact on the number of test iterations.

Physical labs are also costly to both acquire and maintain. There is typically some level of capital outlay required for new projects, and once equipment is purchased, there are recurring costs associated with power, space, cooling and maintenance.

While physical labs are absolutely a critical part of any operator's test and design toolkit, because of the aforementioned limitations in terms of scalability, flexibility and costs many have considered the possibility of moving some testing and design exercises into the software realm. In fact, there exists several commercial and open-source software-based network simulation tools (e.g., GNS3, Olive), but these introduce another set of challenges and limitations. Generally these solutions are not officially supported by the major network equipment manufacturers, so features, protocol behavior and capabilities vary between what is available in software and what one will see on an actual network. For example, some of the router simulation software options lack forwarding capabilities. Other offline modeling tools can show results that diverge from actual world behavior. While these software solutions certainly have their place, to be able to test and design with confidence, one needs to conduct tests with the actual code that will run in your physical network.

To help fill the gap between physical test labs (realistic but limited scale and flexibility) and traditional software simulation solutions (flexible but limited realism), networking equipment vendors such as Juniper Networks are now offering cloud-based services that enable operators to create and run networks in a virtual environment. These environments enable users to create and operate virtual networks consisting of fully functioning router/switch "stacks" of network equipment operating systems. Some solutions also include virtual machines of the test equipment you would expect to see in a physical lab.

These cloud-based environments have the benefit of using virtual resources—so they are immensely flexible and scalable—and are also fully supported by network equipment vendors. This latter point ensures feature parity across multiple versions of router OSes and protocol consistency across both the virtual environment and physical gear.

Use Case	Virtual environment solution
Scalability	✓
Protocol interop	✓
OSS/BSS integration	✓
What-if scenarios	✓
Alternate Network architectures	✓
Training/Education	✓
Hardware testing	✗
Forwarding performance	✗

Table 1: Virtual Testing Environment use-cases

Within a virtual environment, operators can essentially replicate their production network and conduct test and design exercises with a level of scale and realism not otherwise possible, along with many other use cases. Refer to table 1. However, because it is a virtual environment, some tests are simply not possible. In this paper, we will outline the technology behind these virtual environments; examine some real-world use case examples; and discuss the benefits and limitations of such solutions.

VIRTUAL ENVIRONMENT

The network virtualization environment used for the tests described in this paper is a Juniper solution (marketed under the name Junosphere), and it is essentially used to create networks in virtual, rather than physical space. These virtual networks can be used for design, test and training exercises without the need for physical gear while providing a true instance of a router operating system (in this case, Junos) along with an emulated data-plane.

The key components of a virtualized networking system are:

- A secure, multi-tenant Data Center, optimized for high-speed networking between servers and network-attached storage
- A virtual machine (VM) management

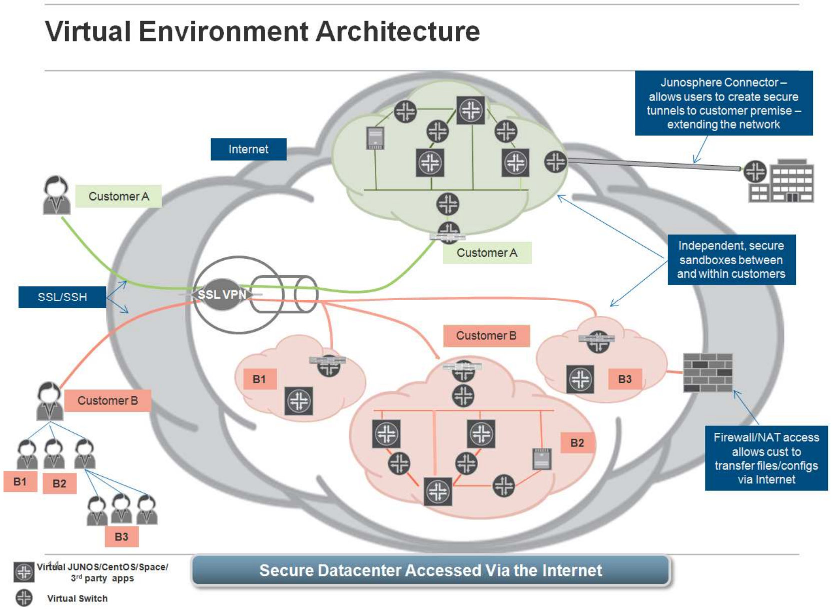
Virtual Environment Architecture
Junosphere Connector – allows users to create secure tunnels to customer premise – extending the network
Internet
Customer A
Customer A
Independent, secure sandboxes between and within customers
SSL/SSH
SSL VPN
Customer B
Customer B
B1
B2
B3
B1
B2
B3
Firewall/NAT access allows cust to transfer files/configs via Internet
Virtual JUNOS/CentOS/Space/ 3rd party apps
Virtual Switch
Secure Datacenter Accessed Via the Internet

layer customized for creation of network topologies

- A series of VM images, that users can load on demand
- A graphical user interface which allows users to save and store custom topologies as well as control permissions and access to the service

Each of these components is covered in more detail below.

Data Center

Because the virtual environment will be used to create and operate networks, the demands on it are quite different from most cloud environments or services, which traditionally are priced and offered based on compute power and/or storage. It would be very difficult to simulate a network in these environments, so it was necessary to build out an entirely new, next generation, cloud Data Center for the foundation of this virtual networking environment. The data center is a combination of Intel-based servers and network-attached-storage, with all Ethernet ports connected together via Juniper EX Ethernet switches. DC file upload and download protection is provided via high-end firewalls, and end-user topology access is secured via the SSL VPN gateway software. The cloud is located in a high-availability colocation facility that provides rack space, cooling, redundant power and high-speed, redundant Internet access. DC uptime is designed to be 24x7, 365 days per year, with service maintenance windows roughly occurring monthly. Finally, a publicly accessible URL completes the access.

Virtual Machine Manager

The real brain of the solution is the Virtual Machine Manager (VMM) software that handles the virtual machine creation and deletion as well as the unique job of VM inter-working. A purpose-built cloud for this virtual networking environment was required because we are building customer-specified networks of VMs, and not just leasing workload CPU cycles and/or access to storage.

The VMM used is a Juniper-developed KVM/QEMU-based solution that provides the ability to scale according to the size of the computing platform, offering support of complex network topologies as well as hosting a mixture of Junos, Unix and other 3rd-party VMs. VMM takes in via its API an execution script that, in conjunction with the Virtual Distributed Ethernet (VDE) switches, provide emulated Ethernet segments to which virtual machines are able to interconnect. VMs within a user’s space are able to communicate over these emulated segments, the interfaces operating in the same way that a Layer2/Layer3 interfaces on a regular physical device would. VMM, thus, creates a “VMM topology” per customer which is a unique instantiation of the VDE Switch process, the number of VMs, and the type of VMs. The spaces are “secure”; VMs from User A are unable to communicate with those of User B.

Virtual Machine Images

During the instantiation of the VM by the VMM software, a personality (image file) is loaded onto the VM. This personality decides the operation of the VM. Within the virtual environment discussed in this paper, the available image files included:

- VJX1000 – a virtual version of a Juniper router/switch – that supports current releases of the Junos operating system. It is a “real” operating system, with an emulated forwarding plane capable of supporting all routing (MPLS, VPLS, v4, v6, multicast) and firewalling (stateful firewall) features. The virtual machine is able to operate as a regular Juniper device, without the need for hardware to be present.

- Junos Space - a network management application platform that can be used to provision, monitor, and manage Juniper devices
- Centos – a Unix host image for customers to add custom applications or host configurations
- Partner images from leading design and test vendors such as:
  - Cariden Technologies (MATE) [1]
  - Packet Design Insight Manager [2]
  - Spirent Virtual Test Center [3]
  - Mu Dynamics Studio [4]

This paper describes specific experiences, and therefore the images above are restricted to what was available within the existing virtual environment. It is possible that virtual machine image files representing other vendors or technologies could be incorporated into a similar virtual networking environment.

User Interface

The user interacts with the virtual network via a web-based user interface (UI) that lets users access the environment from any browser-equipped laptop or tablet. The UI is an application built as a multi-tenant provisioning tool for account, capacity and library management. It provides the GUI-based control of resources, allowing users to schedule their access times, store their topology files, and build their unique networks on-demand.

IN THE WILD

As previously mentioned, a virtual environment can provide significant value when trying to evaluate new technology and/or test specific large-scale protocol scenarios for a network. A physical lab environment is essential for router/switch hardware testing and validation but in almost all cases cannot provide the topologic resources to determine how a technology or protocol with act on an actual network.

In the next two sub-sections, we will discuss two scenarios where a virtual environment is used to validate network operation in the presence of new technology. For each use-case we will briefly describe the problem and/or challenge followed by a description of how virtual networking resources were used to solve the problem.

Use-case #1: Large Scale Core Network Scaling

In this example, an operator is trying to validate several simultaneous technologies to enable a more efficient method of scaling their core network. This represented a fundamental architectural shift that required a much more detailed test environment than could be provided by a set of off-line modeling tools and a few routers in a lab. The goal was two fold:

- Introduction of a MPLS “optimized” packet forwarding paradigm through the use of BGP labeled-unicast sub-address-family [5]
- Introduction of a multi-plane core architecture and the Aggregation/Edge connectivity

The network and technology migration is illustrated in the figure below.

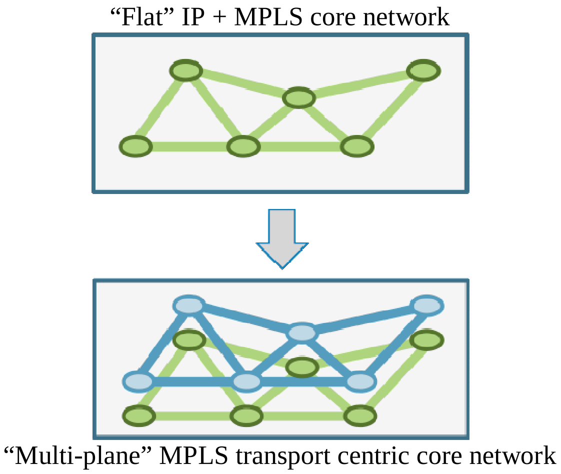
“Flat” IP + MPLS core network
“Multi-plane” MPLS transport centric core network

Figure 2: Network Architecture validation

The challenge the operator faced was how to conceptualize and visualize the target network, test the required protocol modifications, test the introduction of new protocols, and subsequently validate the forwarding properties in the network.

It was essential to be able to validate the changes on a mirror image of the current core network which consisted of a number (10’s) of PoPs geographically dispersed across the U.S. in order to ensure the correct routing policy changes, interaction of additional protocols, and validate the protocol architecture.

In addition to generally validating the modified network architecture, the operator now had a working virtual model of the target network in order to train their operations teams, practice and validate change-order methods and procedures as well a working documented target network.

Use-case #2: Protocol Scaling Characterization

In this use case, an operator wanted to very specifically characterize the memory and forwarding impact on their routing infrastructure if they enabled a new protocol extension. The protocol extension was a Border Gateway Protocol (BGP) extension called Add-path [6]. We will briefly describe BGP Add-path in the next few paragraphs before getting into the specific operator example.

BGP has implicit withdraw semantics on each of its peering sessions, where an advertisement for a given prefix replaces any previously announcement of that prefix. If the prefix completely goes away, then it’s explicitly withdrawn. BGP scaling techniques such as route-reflector and confederations are widely used in networks of all shapes and sizes. These techniques result in information hiding—for example, available backup routes are hidden. This may be good for scaling, but can problematic in other ways. BGP Add-path addresses some of these inefficiencies.

There are a number of reasons to enable BGP Add-path.

- Faster convergence, robustness and graceful shutdown schemes that require backup paths. This is because route reflectors eliminate backup paths.
- Stability and correctness schemes that require additional paths. For example fixes for MED oscillation or MED misrouting
- Multipath schemes that require multiple next hops
- And, implicit withdraw alone is potentially a problem for some types of inter-AS backup schemes

As you can see, much like the previous use-case, the operator was faced with multiple challenges:

- Would BGP Add-path provide the expected functionality?
- How would the additional BGP paths affect the routing resources of their network?

- Do they leverage the current BGP design or could further optimizations be realized?

It was essential for the operator to build a virtual representation of their current International core network to baseline BGP behavior and resource utilization. Another requirement was the need to be able to access and import, as closely as possible, their current peering locations in order to replicate the current BGP table "attributes".

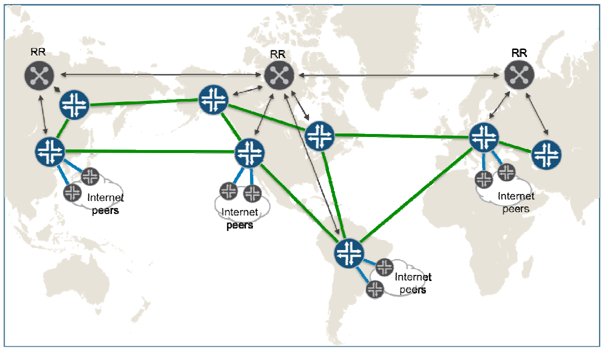

Figure 3: International Core network with regional route-reflectors (RR) for BGP scaling

The resulting virtual network representation allowed the operator to not only characterize their current design, validate BGP add-path and understand specific add-path configuration requirements but also developed multiple future architectural scenarios where indeed BGP Add-path not only delivered the required functionality but could also result in reducing the network resources required to scale BGP.

CONCLUSIONS

Virtual networking environments are a new development that leverage the technologies and concepts popular in cloud computing, and apply them in new ways to solve a fundamental problem for network operators. While virtualized environments will never be a complete replacement for hardware testing, they can provide the resources that allow operators to perform large-scale topology design or testing exercises that would not otherwise be possible. In this paper, we have outlined the technologies behind a specific virtual networking environment implementation, and several use cases, but these technologies and use cases can vary beyond what was discussed within the scope of this paper. In any form, virtual networking environments can be a powerful addition to an operator's design and testing toolkit.

FURTHER READING

QEMU/KVM references/publications

http://www.linux-kvm.org/page/Main_Page
http://wiki.qemu.org/Main_Page

Network virtualization references:

Flexible Cloud Environment for Network Studies:
http://edusigcomm.info.ucl.ac.be/Workshop2011/20110311002

BGP Route Reflection:

http://www.ietf.org/rfc/rfc2796.txt

REFERENCES

[1] http://www.cariden.com/
[2] www.spirent.com
[3] www.packetdesign.com
[4] http://www.mudynamics.com/
[5] http://tools.ietf.org/html/rfc3107
[6] http://datatracker.ietf.org/doc/draft-ietf-idr-add-paths-guidelines/