

## HIGHLIGHTS OF THE FALL 2014 ONUG MEETING

### EXECUTIVE SUMMARY

On October 28<sup>th</sup> and 29<sup>th</sup>, a host of IT technology and business professionals joined together in New York City at the offices of Credit Suisse to discuss their challenges, collaborate towards a common set of needs that could be shared with the industry and share best practices with regards to overcoming the limitations of today's networking. With a strong showing from the financial community, along with key representatives from other sectors like pharmaceuticals, retail and transportation, the meeting helped to cement the future of open networking.

### OPEN NETWORKING IS PICKING UP MOMENTUM

The trend towards open networking has been gaining steam for some time and as customers continue to demand more from their vendors, they have the opportunity now to help institute change in the industry. At this fall's meeting, the ONUG journey continued to gain momentum with reporting on the progress from the three key working groups. Clearly ONUG is having impact as vendors were there in force, taking in that feedback.

The world of open networking encompasses three key areas: Software Defined Networking (SDN) white box networking and the automation of DevOps/NetOps. The biggest underlying change that happens during this the shift to open networking is the expansion from a fully physical environment to a combined physical / virtual network consisting of both underlays (the physical) and overlays (the virtual). In the future software is running the network versus hardware and this brings an important opportunity for customers to influence the market.

How is this shift to open networking being accepted by users? In polling by ONUG, it was determined that 75% were either reviewing or in pilot with the remaining quarter already in a limited deployment. What used to be a year-long process for investigation has shrunk to only 6 months, showing that additional focus and resources are probably being applied to help accelerate the move to open networking. The biggest business drivers for this move are the need to keep pace with highly virtualized or cloud data centers. These virtualized environments shine a light on the need for a total orchestration stack. In one of the discussions a customer pointed out that it takes them

minutes to provision a VM, but close to 150 days to bring it online because of the manual networking procedures. In the past, when it took weeks to get a server ordered, received and installed, this gap was not as noticeable, but in a heavily virtualized world, the difference is startling.

Productivity is an important consideration because 80% of a network administrator's time is spent in change management and network management – basically “keeping the lights on” for the network. Because of the manual processes that most physical networks require (i.e. the underlay), most engineers can manage on average 100-200 devices. On the virtual side, when we look at the number of VMs (i.e. the overlay) that the typical engineer handles, that number is closer to 500-1000. This additional efficiency is why customers are interested in virtualizing their networks.

In looking at the wave of open networking that is readying to break, we see IT actively looking to move their networks from the proof of concepts to production deployments in 2015 with overlays and software defined WANs (SD-WAN) being the first areas for expansion. Network services virtualization (not to be confused with network function virtualization) will happen a bit later with proof of concept happening in 2015 and production deployments following.

The representatives on the user side should be primarily considering products from vendors that support their use cases when making buying decisions. Clearly the more users hold out and demand products that meet their needs, the more likely vendors are going to build around the needs that ONUG is highlighting. These use cases represent the collective input from the ONUG Community and as such reflect cross-community requirements. Companies that are represented by the board members are in the process of planning and budgeting for next-generation open network build-outs, which are reflected by the use cases defined by the ONUG board. The ONUG Board recommends and encourages all ONUG community members to review the ONUG Working Group Use Cases, modify them for unique requirements, engage the vendor community to develop multivendor interoperable demonstrations of their solutions to the use cases and then appropriate budget accordingly. ONUG Board members are seeking the vendor community to demonstrate their ONUG Working Group Use Case(s) solutions both in private and at ONUG. The ONUG Board recommends that the ONUG community do the same.

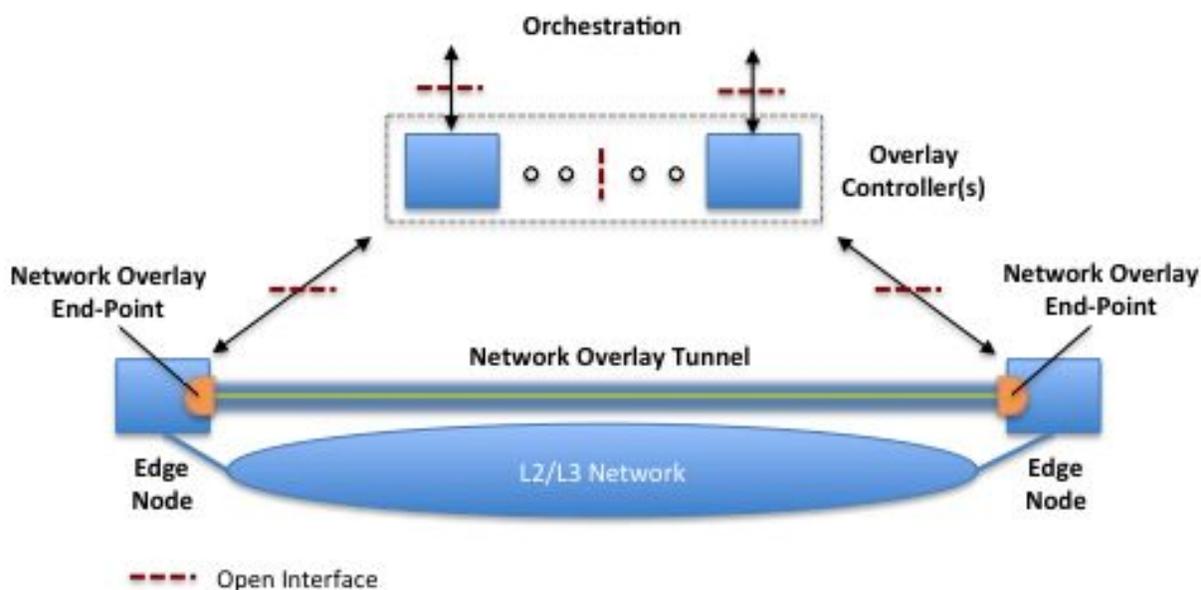
Being involved with ONUG is the best way for users to make sure that their specific needs become part of future uses cases, ultimately becoming the features of future products.

There were 3 key use cases the organization was working towards definition; this critical customer input is important to help guide vendors towards the right solutions. Voting at the meeting helped to identify the next 3 use cases that will tackled by ONUG. As ONUG moves into 2015 we'll see the following use cases being added to their work:

- Traffic Monitoring and Visibility
- Network State Collection, Correlation and Analytics
- Common Management Tools cross Network, Storage and Compute

## VIRTUAL NETWORKING / OVERLAYS

The first working group to report was the group focusing on overlays, including network virtualization.



Virtual overlays are networking solutions based on one or more packet encapsulation techniques designed to forward end-to-end service traffic, independent of underlying transport network technology and architecture. Initially, architectures like MPLS layer-3 and layer 2 VPNs (Virtual Private Networks), were introduced in the late 90's to try to address the issues that customers had. Following that, a variety of both open and proprietary solutions have been introduced to try to tackle customer challenges. This

group limited its scope to data center deployments of IP-based networks. While similar techniques for network overlays can be applied at the branch level, at the WAN level or the campus level, this group is focusing only on the data center. The same applies for network overlays for layer-2 networks, requirements for these types of network deployment scenarios are not being covered by this working group at this point.

The discussion within the meeting was lively, including the following discussion areas:

- Chief concerns began with how these overlays would impact the data plane, control plane and management plane primarily. The working group wanted to be able to comprehend the different environments and the communications – virtual to virtual, virtual to physical and physical to physical.
- Most customers seemed to adhere to the idea of a distributed control plane, with a centralized management plane. The management plane was more a function of their own management and support staffs vs. an overriding philosophy.
- In developing an overlay, a critical element is enabling the troubleshooting and management tools to span both the physical and virtual network. If the abstraction of an overlay results in disconnected layers, then as problems crop up, debugging the issues will take significantly longer or require more resources.
- The use of overlays helps enable companies to deploy applications faster and treat networking as more of a fungible, flexible resource pool. With more layer 2 flexibility, a business can put its resources anywhere. The downside to this, however, is that now a business will need to more carefully manage security and services aspect as these domains can potentially expand greatly as layer 2 flattens out.
- Today's switches are so similar to servers that some are viewing their switches as servers, with performance and scale being the key attributes to consider. Deploying at scale will be important, and customers need to understand how they can achieve convergence at scale. Their goal is to allow the servers to be as simple as possible, allowing the switches to handle all of the networking tasks.
- The starting place for overlays is breaking down layer 2 barriers, but once you do this the next question becomes how do you manage and secure the layer 4-7 services?
- Policy adherence is critical, not only being able to apply these policies but also help understand why things are breaking. There was a good discussion on the need to optimize policies across multiple levels of infrastructure.
- Will there be a VXLAN standard? Is it a standard just because Broadcom put it in their silicon (Trident)? While the Broadcom chipset is used in the majority of

network switches on the market, it is not in every switch, so claiming this as a standard is still far from guaranteed.

- Standards around encapsulation are good, but vendors need to stay aligned with standards, deliver incremental technology that addresses near-term pain points while staying focused on the long term strategy (and obviously that strategy should encompass open standards.)
- In questioning whether there would be multiple controllers managing multiple overlays or one controller to rule them all, most believe that there will be multiple usage models and both scenarios will exist. A strong point regarding controllers was “how large do you want your failure domain to be?” Essentially the idea of a single controller is challenged by the thought of a single point of failure.

A report that outlines the results from the virtual overlays working group [can be found here](#).

## SD WAN WORKING GROUP

This was one of the earliest customer needs and became one of the primary use cases identified for the original ONUG. The original definition was not identified early on as “software defined WAN” but as the group began to dig into the issues, the use case for Software Defined Data Center (SDDC) very quickly coalesced around that definition, one that the industry has also identified and begun to develop products around. The team looked not only at the problems in today’s enterprise WANs, but also looked at the architectural models that were prevalent in most enterprises. From there, they were able to focus on the desired features and functionality that vendors need to be delivering in their products. Finally, they focused on the implications of SD-WAN from a services, support, tools and delivery process.

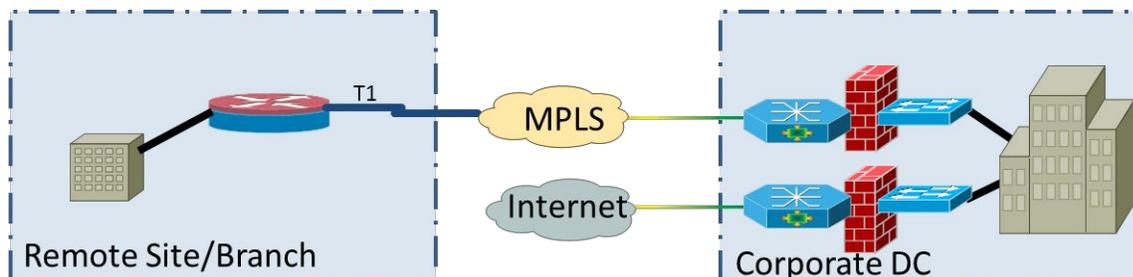


Figure 2 A traditional WAN connection to a branch office and the Internet

The discussion at the meeting revolved around several key topics and points:

- The first question asked what percentage of the overall needs are being met by the latest offerings. The panel believed that a surprising 90-95% of the customer requirements are being met by the vendors today, a testament to the influence that ONUG has had on the industry.
- It was identified that the top 3-4 use cases are about using the WAN transport in an agnostic manner, allowing the customers to not only connect to a variety of services but also a variety of providers.
- Networks are in the business of delivering applications with a security, compliance and SLA. Richer applications, including video, are driving more need for better WAN optimization. Customers want more performance and they are more skeptical of vendors claiming savings.
- A lively discussion in the fireside chats revolved around vendors' claims of 10X savings, most believe the savings could be 40-70% but it depends on how recently one had negotiated contracts with carriers. The thought that you have to "spend some to save some" was there, but the savings could be realized very quickly in reduced carrier costs. Reducing the time to deploy and security were both other sources for savings.
- The formation of ROIs needs to be done on a per-customer basis because, while hardware and carrier costs may look similar, there are too many variables in operational costs on the customer side to adequately develop a "one size fits all" ROI pattern.
- There was a discussion of solutions being carrier agnostic. If customers have more control and choice they might not need the same level of relationship with carriers. Solutions should compete OEM to OEM and the customers are not interested in WANs being provided as a service, they prefer to consume the hardware and software in order to maintain control and flexibility.
- It was pointed out that complexity does not go away with outsourcing of WAN services, you are just pushing the accountability off onto someone else.
- For future deployments, customers should architect the WAN for the cloud, allowing security to extend from the data center to the cloud. Thinking only of data center to data center is far too limited.
- Developing the operational model is critical; the technology is here, but the tooling infrastructure and visibility to the underlay are still lacking.

A report that outlines the results from the SD-WAN working group [can be found here](#).

## NSV WORKING GROUP

Network Services Virtualization (NSV) is not to be confused with NFV (Network Function Virtualization). NFV mainly focused on carrier-led efforts to virtualize the many appliances they need to install at the edge of a customers' data center to provide connectivity.

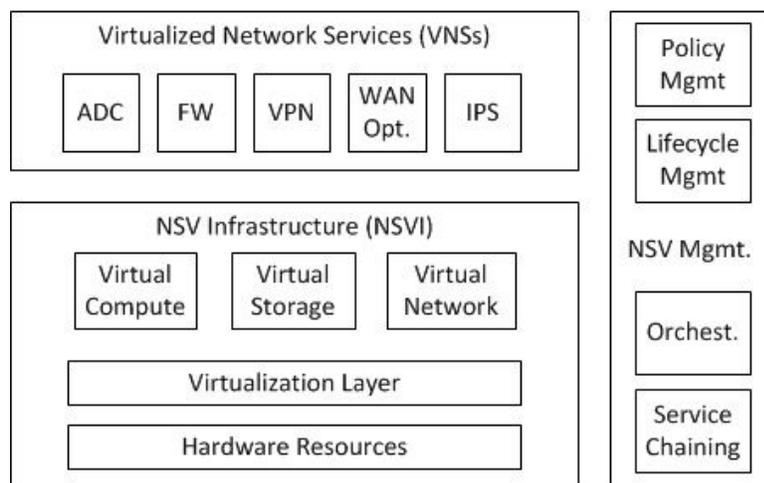


Figure 3 A high level view of the NSV framework and main components

NSV focuses on the problem with virtualizing layer 4-7 services; these are typically deployed by appliances within the data center. Customers are not happy with the appliance model because it brings too many variables and limits flexibility. Moving these from hardware-based appliances to virtualized software-based services can help drive better efficiency in managing data centers. The cost and complexity of managing a large number of layer 4-7 appliances from different vendors (each with their own tools) creates a headache for all large organizations. Appliances like server load balancers, application delivery controllers, WAN optimization, firewalls and security/intrusion detection are not optimized to scale easily today. With difficulty in provisioning and managing these devices, driving them to a software-based solution could help drive down costs and increase productivity.

The session around NSV began with clearing the air about NSV vs. NFV so that everyone was on the same page. After that, the group was able to find a lot to discuss.

- In NSV, enterprises are looking to get the same benefits that service providers are looking for (with NFV), but the key driver is flexibility in provisioning new services and servers. For carriers, reducing the support infrastructure and making it easier to deploy services remotely is key for NFV. But for NSV, speeding up provisioning time *inside* of the data center was critical.

- Security policies were important to the discussion, and the exposure from poorly provisioned (or incorrectly provisioned) security was highlighted. Customers would like to see the service chain run all the way up to layer 7 to help enforce not only application priorities but also security.
- Customers want to move away from appliances and towards services. But for now they need to slowly work into that state, adding the right set of services for the right access levels. As this happens, more global-based security policies will be important.
- The customers would like to see vendors form a group that defines a declarative language. But declarative policy cannot become a policy per vendor, which does not scale well. No longer being tied to any one vendor is critical – clean APIs and interoperability are the keys to bringing flexibility. More operability to allow a business to create best of breed. Customers don't care as much about how this is done, but that they have a way to ensure adherence to policies, better agility and lower cost.
- There was a discussion of the timing between business policy and service insertion. Customers want to know how they can insert new services into the data center without disrupting operations or overriding some of the existing policies.
- Good orchestration tools are available to spin up new VMs and services, but the security gap is too large, which creates problems. Customers want to build a security model that can cover both the data center and out to the cloud.
- It used to be the case that customers would build data centers and protect from the outside; now they build applications and need to protect at an application level, especially because most applications are having some contact with the outside world in one way or another.
- Branch offices are a real concern. Customers would like to reduce the amount of infrastructure in order to have a branch in a box. At that point they can manage all of the branch infrastructure as a unified entity, managing the state, policy, security, as well as multiple carriers.
- Customers would like to be able to start with a bare metal server, add Linux virtualization or containers and then deploy their network services on top of that stack (versus using appliances.) By utilizing an industry-standard platform that can host multiple services, they can cut down on the number of physical boxes to manage and improve flexibility. Commoditization of the underlay allows for open networking – and customers can then focus on rightsizing the applications over a shared infrastructure. By using an industry-standard server with the right

software and not needing to over provision, customers can future proof their services layer.

- The goal is to build fewer out-of-band services and move to more in-band services with more integrated into the network stack through in-line processing. Being able to tap into flows in real time allows for more proactive monitoring and alerting, versus reporting on issues after the fact.
- More proactive monitoring and reporting should be established so that IT does not have to continually monitor, ping, and inquire on the app health; apps should be self-reporting. There should be more focus on management for the future, better contextual awareness for applications
- There needs to be a tight linkage between firewalls and overlays; today it is manual. In the future, as changes happen, flows need to reflect the change and adjust automatically.

A report that outlines the results from the NSV working group [can be found here](#).

Across the board, the input that came directly from customers is critical in helping to shape the conversations with vendors. Some of the sessions allowed the vendors to participate, but as always, customers are most open and most forthcoming with their concerns when they are in the “safer confines” of customer-only meetings.

Through the working groups ONUG will continue to drive these customer needs back to the vendor community and standards organizations, helping to refine future products and narrow the gap between demonstrated needs and delivered features.

### AN ANALYSTS' VIEWPOINT

At Moor Insights & Strategy, we participate in a variety of industry events that promise to bring together both customers and vendors to “shape the future” of an industry. Few accomplish this lofty goal; most become an opportunity for the vendors to pay lip service to the customer needs while pushing their products. Most industry consortiums focus on their industry, but do so with an inside-out view. With vendors looking to develop sustained competitive advantages and lock out the competition, true cooperation on standards does not often occur. When it does, it is often tied to standardizing technology for manufacturability or support, not the customers.

What made this meeting, and ONUG, different is that they are an end-user focused organization. This provides a much stronger base for recommendations as the end customers have the power of the purse – and more than a 100 billion dollars' worth of

power. Representing some of the largest buyers in the industry, the members of ONUG are not content to sit on the sidelines and accept the status quo.

By actively driving their needs back to the vendors they are putting their stake in the ground and making their demands known. But in speaking to the vendors this is not an adversarial situation at all. There is a genuine understanding of what is at stake and an earnest desire on the part of most of the vendors to try to meet these market needs. The products that come out of these meetings is still not known because product cycles in the enterprise space tend to be much longer than consumer products, but there is every indication that the messages will be carried back to the headquarters.

Finally, one of the gutsier moves was scheduling a debate on SDN at the end of the first day. With two strong debaters who were able to hold their ground, the event clearly was not afraid to take on the thorny questions that are on everyone's mind.

## IMPORTANT INFORMATION ABOUT THIS PAPER

### *AUTHOR*

[John Fruehe](#), Senior Analyst at [Moor Insights & Strategy](#)

### *REVIEW*

[Patrick Moorhead](#), President & Principal Analyst at [Moor Insights & Strategy](#)

### *INQUIRIES*

Please contact us here if you would like to discuss this report and Moor Insights & Strategy will promptly respond.

### *CITATIONS*

This note or paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title and "Moor Insights & Strategy". Non-press and non-analysts must receive prior written permission by Moor Insights & Strategy for any citations.

### *LICENSING*

This document, including any supporting materials, is owned by Moor Insights & Strategy. This publication may not be reproduced, distributed, or shared in any form without Moor Insights & Strategy's prior written permission.

### *DISCLOSURES*

Moor Insights & Strategy provides research, analysis, advising, and consulting to many high-tech companies mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

### *DISCLAIMER*

The information presented in this document is for informational purposes only and may contain technical inaccuracies, omissions and typographical errors. Moor Insights & Strategy disclaims all warranties as to the accuracy, completeness or adequacy of such information and shall have no liability for errors, omissions or inadequacies in such information. This document consists of the opinions of Moor Insights & Strategy and should not be construed as statements of fact. The opinions expressed herein are subject to change without notice.

Moor Insights & Strategy provides forecasts and forward-looking statements as directional indicators and not as precise predictions of future events. While our forecasts and forward-looking statements represent our current judgment on what the future holds, they are subject to risks and uncertainties that could cause actual results to differ materially. You are cautioned not to place undue reliance on these forecasts and forward-looking statements, which reflect our opinions only as of the date of publication for this document. Please keep in mind that we are not obligating ourselves to revise or publicly release the results of any revision to these forecasts and forward-looking statements in light of new information or future events.

©2014 Moor Insights & Strategy.

Company and product names are used for informational purposes only and may be trademarks of their respective owners.