

Fall 2015 ONUG Meeting Highlights

Open networking continues to mature, vision expands to cover more

Executive Summary

With over 600 attendees for the semi-annual meeting of the Open Networking User Group (ONUG), the Fall 2015 event increased the volume and the emphasis on open networking solutions. The emphasis began to change as more customers began proof of concepts (POCs) and are moving closer to using open networking in their production environments. The existing networking challenges that are driving open networking remain, but the complexity of making the move to open networking is still presenting difficulty as customers now have a much clearer understanding of the tactical challenges around implementation. Overwhelmingly the conversations at this ONUG were more substantial as customers appear to be moving from the investigative phase towards beginning the evaluations and proof of concept (POC) testing. The progression reflects an evolutionary approach to open networking and not a network revolution.

State of The Industry – Nick Lippis

The opening session is designed to set the stage and provide a high level view of what to expect at the meeting. The input from previous meetings was that the audience loved the fireside chats, so instead of the standard two that were normally held, ONUG expanded to four in order to give customers more opportunity to provide input.

Users had also indicated that they wanted more ongoing communications and more year-round activities. To this end, ONUG is scheduling a set of meetups throughout the course of the year that will be held across the country. These smaller regional meetings are designed to broaden the impact of ONUG, bringing in those that may not normally be able to commit to the longer semi-annual events. Meetup events in Japan and London are also under consideration, although there is not a commitment as of yet to host them. Additionally, the decision was made to split the major semi-annual events between the two U.S. coasts; the spring event will be held at Intuit in Mountain View, CA and the fall event will remain in New York City.

At a high level, there were several key trends that are shaping how IT is changing in the face of open networking. To begin with, the consumption models are fundamentally changing, followed by an increasing demand for new skill sets as the old way of doing IT may no longer apply.

A question was posed to the audience as to whether or not SDN is part of a larger vision or framework (i.e. is it part of a software defined infrastructure or software defined data center framework?) Clearly the attendees see open networking as part of a much larger ecosystem/environment. So as ONUG moves forward, perhaps it needs to expand its vision of open networking to allow the organization to focus at a higher level. In past meetings most vendors demonstrated their proof of concepts and their products in a vacuum, but, perhaps as a nod to the fact that open networking is part of a larger

discussion, the first multi-vendor proof of concept was being demonstrated at this event. This represents a maturity of vision, an acknowledgment that there is more to open networking than just point products, and that the open networking ecosystem must work together for the sake of customers.

ONUG took the idea of hybrid cloud head-on, something that would get replayed later in the customer breakouts; hybrid cloud as a concept may seem interesting, but as a practice, customers are less likely to try to span multiple domains because security and complexity raise too many challenges.

But what does the world of open networking look like now? The ecosystem is primarily software-focused today because this is where the innovation is happening. Hardware product cycles are much longer, which is part of the reason that customers are looking towards software to solve their most pressing challenges. Software cycles can be months, weeks or even days, bringing more agility to businesses who need to move quickly. The orchestration of this ecosystem needs to focus on more than just one product; to be truly useful to customers, orchestration needs to be far more comprehensive and fit into multiple frameworks.

Nick pointed out a significant change that will need to happen as technology shifts, a trend that was pointed out later in many of the sessions. Computer Scientists will probably be running cloud infrastructure meaning that changes in skill sets will be important as businesses try to push forward. It was pointed out that customers believe that twenty-five years of CCIE won't prepare them for the future, despite how the current crop of CCIEs emphasizes their stability and safety (these CCIEs do have a vested interest in maintaining their status quo). There is also conflict with those in Dev Ops that see this all as "living in the past." But despite these attitudes, companies still have a lot of existing equipment that needs to be managed and maintained. There is a bit of truth to both sides, like the old beer commercial "it tastes great, no, it's less filling" – sometimes it is both. In order to fully understand where networking is going, one truly needs a full stack engineer who can view the whole problem. But as vendors try to address this changing market, they need to comprehend how changes impact IT, especially those with their hands on the equipment.

90% of the customers polled indicated that they see a wider architecture that is software defined. They don't want to be operating in pockets; they want to get out of multiple proprietary (or even open) stacks, instead stepping back to view the entire enterprise as a more holistic environment. When it comes to open networking development, 42% of those surveyed said that they were in development of open networking solutions with 60% in development of SD-WAN solutions for their companies. Perhaps ONUG customers are on the leading edge of these changes, but it is clear that change is coming. Network services virtualization is lagging these other categories but those layer 4-7 functions will change; after the underlying plumbing is addressed, customers will begin to focus on application delivery.

On a final note it was pointed out that culture is a huge point of contention in change. Nick used the Grace Hopper quote "The most dangerous phrase in the language is

‘we’ve always done it this way’” to exemplify this point. Clearly, for businesses to change, they need to think differently.

From previous work by the ONUG Working Groups, a new IT Service Lifecycle Management Automation Framework was identified. This is a reflection of user needs changing as companies begin to move towards open software-defined cloud infrastructure. The four areas that will be part of this new group include: Network State Collection, Traffic Monitoring/Visibility, Automated Configuration, and SDN Federation/Operability Orchestration. Additionally, a new non-technical use case has been developed to broaden the required executive skill sets and IT infrastructure models for software-defined architecture.

The opening address finished with the proposed mission statement: “ONUG aims to enable greater choice and options for IT business leaders by advocating open interoperable hardware and software-defined infrastructure solutions that span across the entire IT stack, all in the effort to create business value.” Over the upcoming year ONUG will work to fine-tune this statement in order to accurately reflect the beliefs of the group.

Morgan Stanley and Open Frameworks, Why Do We Care? – Tsvi Gal

The main customer presentation at this event came from the CTO of Morgan Stanley, a company with 71,000 physical servers, 70,000 users and 60,000 databases who opened his session by joking, “every employee has his own database.” Morgan Stanley sees themselves as a true engineering society. All of those in IT management are “player coaches” meaning that not only are they responsible for running teams, but they also need to have the technical skills to understand what their developers are working on, and lend a hand if needed. This changes the innovation curve by challenging the norms. Mr. Gal pointed out, technology has made Morgan Stanley what it is, referring to them as “a technology company that also does some banking.”

Morgan Stanley sees the network as their bloodline and they tend to build their own systems. By default, they would prefer to choose existing commercial solutions but too often find that their needs are not being met by off-the-shelf technology, forcing them to build their own. Their goal is phrased in building where you can differentiate and buying in areas where you cannot. Morgan Stanley applies the 80/20 rule to their work: 80% is the same as other industries and here they should share; 20% is specific to their own business and they use that code to help them differentiate.

As a big proponent of open source, Morgan Stanley uses plenty of open source components and contributes much of that code back to the open source community with a position that if code is not uniquely differentiated for their own business then it should then be shared. It was pointed out that open source is not open technology frameworks, there is a difference.

Mr. Gal noted that the worst vendor lock-in is when you are your own vendor and that often the sins of the past are holding companies back from moving forward.

Morgan Stanley does have challenges with finding people with the right skill sets when hiring, meaning that they need to invest in helping build those skills in their teams and not expect that they can just hire to the needs that they have.

The Great Debate – Will Software Make Everything Better?

Each year the great debate is one of the highlights of the event. Two experts square off against each other on a topic that is important to the group. This year's topic was whether or not software was better than hardware. In an event that focuses heavily on software you'd expect a much more slanted discussion, but in reality it was pretty evenly matched. Ernest Lefner (Bank of America) hosted the discussion with Dr. Douglas Comer (Purdue) and Dr. Brighten Godfrey (UIUC) taking the two different positions.

The grand challenges of software are virtualization (which needs open APIs, softening hardware), visibility (synthesizing across different systems), security (which needs to understand and move with applications) and automation (stitching together "systems of systems"). For all of these challenges, software can excel. It was pointed out that software is already succeeding. Amazon, Google, Facebook and other large scale out companies in the cloud as well as the WAN traffic engineering that Microsoft and Google are utilizing today are examples of how software is solving these problems.

On the hardware side, the argument was made that software is bug-ridden (NYSE down for 3 hours, United cancels 800 flights, etc.). Software is sold "as is" and is sold without a warranty (something that the courts have all upheld) resulting in a world where there is no liability for bad software. Each version of software is worse than the previous as software degrades with each version, becoming more complex and more bloated. Where software is "never finished," hardware, on the other hand is clearly defined and has an end date for the development. Our culture has adopted the "as is" philosophy so it was argued that maybe this is not a software issue and is instead a cultural issue.

Hardware has design rules and both agreed that we should encourage more design rules in the development of software. But hardware has longer development cycles and a physical product, forcing lag times on the development where software can be more agile and address issues with more flexibility. And when hardware fails it needs to be physically replaced, unlike software that can usually be updated, often remotely.

Everything fails, but hardware fails less often. Software has made things more programmable, which has led to greater innovation in the market. In terms of raw product performance, hardware always trumps software when taking the standard route, but software provides more flexibility to help take a different path the answer, a path that might also be faster in the long run.

At the end of the debate, the audience was polled and was clearly split down the middle with roughly half on each side.

Fireside Chats

The four fireside chats are designed as the main mechanism for end customers to provide candid feedback without any vendors in the room. These frank discussions show the real life issues that customers are facing as they evaluate, plan, pilot and then deploy technologies. After the chats are over, there is a summary panel that reports out to the whole group, allowing the vendors to hear the customer input firsthand.

Hybrid Cloud

As the conversation began, several pointed out that the idea of “hybrid cloud” was really not what we were talking about. To the group, the traditional view of hybrid cloud is an application that spans both a public and a private cloud, but all agreed that the challenges of this kind of environment renders the idea unworkable. Instead, the more logical description for the conversation that morning should be “hybrid infrastructure” where applications can live in a public cloud OR a private cloud, but not in both. The key driver for customers was flexibility, not cost; most saw that the costs (at least in the near term) could be higher, but over time the financial benefits will start to show themselves. However, all agreed that cloud is not a compute model, cloud is a financial model with a “pay per use” function instead of the traditional IT step function. Some even viewed the cloud as “just another data center move.”

Someone posed the question: “do large organizations actually use Amazon?” The expectation was that nobody would, but there were several customers who do use Amazon, including one that is moving all of their infrastructure over (opting for an “Amazon first” strategy and defaulting back to internal services if there was no way to host on Amazon). Some had dictated that any application that was not “tier 1” must be hosted on a public provider, not in the internal data centers. In order to make this happen one customer provided free hosting to their business units for the first year through Amazon to get all of the departments on board. They still provided a monthly statement to show the departments what to plan for when the chargebacks begin in the second year.

Amazon customers were typically utilizing the direct connect method of tying into the hoster with one customer actually moving from the 1Gb to the 10Gb connection based on the amount of traffic they were now seeing. Someone paraphrased the movie quote “if you build it, they will come” with the thought that as more services are made available, users will find a way to address their needs with Amazon. The down side of this, of course, is that when resources can be added so easily, one can end up with VM sprawl (or service sprawl) making monitoring and management more difficult. The key for customers is that the data center is no longer the center of the world and they will need to work in multiple domains moving forward.

In the traditional data center, security was an important component, but with a hybrid infrastructure, security now needs to cover both the data center and the cloud, as well as the transport (data in flight as well as the pipe itself.) Most customers were in the habit of encrypting everything, both the data at rest and the data in flight.

While resources in a public cloud can be spun up in minutes, the provisioning and approval steps still add time to the overall process, so it is important to review these steps as well when factoring in the cost savings. But if customers expect to push their bad processes over onto the cloud provider, they will end up missing out on the benefits of automation and speed; businesses need to change internally with the addition of cloud if they want to take advantage of all of the savings that cloud can provide.

The other popular choice beyond Amazon was Azure. One customer was using Azure for all of their failover. While they are still using manual processes today to monitor for conditions that may require failover, they believe that they can automate this in the future. This company stated that they have stopped budgeting for servers and storage when moving to the cloud, in order to force the transition. This customer acknowledged that a key risk is that they may be running on the same cluster as X-Box. There are multiple virtualized security technologies that Azure brings, but it is incumbent upon customers to own the security and not just rely on the vendor. More importantly, on the subject of failover, there was a discussion of what would happen with another Hurricane Sandy. Would the cloud vendor have enough capacity for all of the financial customers?

A key challenge in a hybrid infrastructure is legacy applications. You can't "cloud-ify" these applications, some don't even work with existing services like load balancing. This means that customers must face the reality that even though they may want to do away with all IT infrastructure, some infrastructure must remain. For banks, much of the online banking applications are being commoditized, so this can be handled in the cloud. But some of the more proprietary applications, even though they could be moved over, may remain in the customer data centers because customers want more control over these applications in order to minimize their business risk.

One question about failover focused on the need to keep all of the failover systems online all the time versus just spinning them up when needed in the cloud. The challenge was that a spin up was ~15 minutes but if the standby site were live, this could be as little as 2 seconds. So instead, the customer focused on keeping an always-on scenario with constant replication in order to ensure systems were online and failover could happen quickly.

A parallel was drawn between cloud infrastructure and the "sharing economy"; "aren't we just sharing compute resources like sharing a car in Uber?" But as financial models are changing with respect to hardware acquisition, as costs come down, the movement to the cloud slows, especially on more predictive workloads where one can forecast need. Customers will need a mix of traditional IT, private cloud and public cloud; one customer said that they will monitor usage over the next 5 years to figure out where things are headed because it is too difficult to make a full commitment to the cloud yet.

Legal and security considerations drove a good part of the conversation as some had requirements preventing their data from leaving the country. One customer complained that a legal agreement with Azure took 2 months, while others in the room thought this was a fast turnaround time. All security needs to be driven by the customer, not the cloud provider; the customer needs to supply the tools if they expect to pass an audit.

The final discussion was around moving between clouds. Most customers would like the flexibility to move between clouds but recognize that there is some degree of lock-in for clouds. Building a physical layer edge to the cloud so that one could remove vendor specificity would be the “crown jewel,” but most were not sure this could ever happen. If applications do move, the workflows need to stay the same as the locations change, essentially using templates to automate the work for developers and abstracting the complexity.

OpenStack in the Enterprise

OpenStack is a cultural shift for IT engineers and this movement will take time for many of them to get their heads around the fundamental changes that it brings. The technology has delivered a lot of promise to the industry as a more standardized platform for deployment. However, while there is a lot of promise, some of the usage models work better than others. Customers were in agreement that they need to figure out how to make OpenStack work in their environments because the alternative (or potential lack of an alternative) could make the situation worse. One should not expect to make a wholesale change, but instead design for scale and then execute in chunks, as it makes sense for the business. This will allow better predictability and easier management.

One of the challenges with OpenStack is the support of legacy applications; while some vendors can provide that support, most are not able to support legacy applications easily. Some were recommending that instead of having to migrate the VM, they should investigate simply rebuilding the VM on the cloud instead, then reinstalling the application. The real savings from OpenStack comes from not having to maintain multiple environments, so moving these applications to a full OpenStack environment is preferred – if they can get full application support.

There was a concern that OpenStack simply moved customers from the old lock-in to the new lock-in, as the proprietary hardware lock is replaced by a proprietary software lock. In order to mitigate this, the workflow needs to be instrumented for flexibility. Because this will require effort on the back end for customers, most are just dipping their toes in the water for now. Most believe that it will take 6-8 years for fault tolerance at the application level, which means that for the time being customers rely on redundancy in the underlying hardware instead of looking for resiliency at the application level.

Most of the customers have begun playing with OpenStack in the Infrastructure as a Service (IaaS) realm and moving small numbers of less critical applications over initially until they get a better feel for how things will work. At that point they can begin moving more applications over and moving up the stack in terms of criticality. Because of executive pressure (typically the CEO or the board) there is more of a push for OpenStack (and cloud in general) but often those desires do not comprehend the full complexity of the request. The phrase “you can’t hold up your fist to stop the wind” made it clear, the edicts are being handed down to IT so there is a requirement that these projects move forward. Over time the customers hope to move up the stack to

Platform as a Service (PaaS) but that will take time, and a lot of successful learnings from the early experiments.

The TCO analysis is important for any project, but in the world of OpenStack it is more critical because of the solutions being apples and oranges. In comparing OpenStack there are a lot of sub components that are critical to determining the total 5-year TCO. Storage and networking can vary greatly, shifting that cost. NSX, for instance, is one of the most expensive networking options to begin with, but as you get to 20-30,000 VMs the costs start to flatten out. The cost exercises are important because they help IT narrow in on where the focus needs to be. One customer pointed out that vSphere has diminishing returns and that often the applications need to change; they say that there is a “VMware tax” but not as much resource support to justify it.

Some of the pressure that IT is feeling is coming from competitors like Amazon. The services they provide are competing with IT, causing friction that pushes IT closer to OpenStack in order to close the gap. Just like Linux, OpenStack didn’t happen until companies made it happen. For most dev and test environments, there are the easy places to get started but the real savings and efficiency don’t really happen until it moves into production environments. IaaS is great, but the real goal should be getting to PaaS, which is where efficiencies really come out. Be mindful of security as this is something that always manages to get pushed back.

There was also discussion about the idea of microservices and containers. Early tests did not meet one customer’s needs, but that was mainly due to not configuring properly; after figuring that out, they were able to support 7000 applications. Using options like CloudFoundry and OpenShift allows for a lightweight choice that gives customers enough of what they need without all of the OS/Hypervisor overhead and challenges. Warden and Garden, along with Docker containers are bringing customers more choice and more options, allowing them to create a marketplace within IT that gives end users more choice and flexibility. This is all part of the preparation for the day when IT may no longer have data centers to manage.

One customer pointed out that “I will never be as good as Amazon on infrastructure, that is their business, and I am only a banker.”

Open Cloud Infrastructure Security

Every conversation about the cloud will start with security and if it does not, it will get there quickly. This group discussed a variety of security issues, the largest of which were configuration management and insider threats.

Configuration management is so critical because vulnerabilities often arise through improperly configured devices that leave openings for an intruder. Especially in cloud/multi-tenant environments where, while there is an implied separation between the different domains, improper configuration could allow one environment to provide an intruder access into another. Insider threats are equally difficult to defend against because unlike intruders from outside of the company, those that work inside the

company do require some degree of access in order to do their jobs; misconfiguration here can lead to an insider being somewhere that they shouldn't be.

Some believed that in many cases the move to an open cloud could actually enhance security. This is because in the world of open systems there are more eyes on the code and issues that pop up become visible quickly (often publicly) and can be fixed just as fast. Whereas in a world of closed proprietary code (or clouds), any vulnerability that arises can be conveniently pushed to the back until the vendor is ready to deal with it, if it is dealt with at all. Often the desire to keep issues out of the public eye can have a detrimental impact, as fixing them is essentially "admitting guilt."

Tools are emerging to provide security in the cloud as well as provide traffic visibility across the fabric. As the split between the physical underlay and the virtual overlay becomes more pronounced, the criticality of these tools will increase. More work is required to open up the security tools and adjust to the new architectures on the horizon. With new cloud services like containerization and micro-segmentation, security vulnerabilities continue to increase.

Enterprises do know that they need to change the way that they provide IT services to their end users. Cloud adoption is on the rise, but the customers mentioned that the move to the cloud is being slowed by hesitancy that is coming from data and application security concerns. Many of these customers still do not trust the public cloud providers with their data, leaving customers in a position of relying more on private cloud services, hosted internally, in order to provide a better level of security. But because these new services are also new to most of the organizations, security expectations may be different as those charged with securing the enterprise not only have to learn new tools and infrastructures, but also new threats and vulnerabilities.

Customers need to take ownership of security in the cloud and not leave securing the resources to their cloud providers. As pointed out in the other fireside chats, it is not likely that a customer could pass a security audit for applications that are hosted in a public cloud if they relied only on the security included from their cloud provider. The security practices and protocols (as well as tools) that are protecting an enterprise within the data center should be extended to the cloud in order to enable consistent policies and intents to be delivered across both domains.

The Edward Snowden events reminded customers that data security is a serious issue, causing some customers to pull back from their cloud plans and reassess their data strategy. The 512-bit AES security hack was a very important event, but not enough people are taking that seriously and considering the implications.

The customers indicated that companies should start with the assumption that they have already been compromised and then work their way backwards in the planning process. However, cloud vendors make this difficult as they are not doing an effective job in the integration and scaling of security into a complete package. Many of these customers are writing their own code to address the interoperability challenges instead of either waiting for or relying on the cloud vendors to deliver that functionality.

For smaller organizations there is a greater chance that the customer will purchase off-the-shelf solutions exclusively; in these instances, the cloud vendor is more apt to provide a better security solution because they own the whole stack. But for most larger organizations the off-the-shelf solutions rarely meet their needs, driving them to custom solutions that require the customer to write their own custom code and build their own security stacks in order to support the needs of their business.

Many of the “pure play” security products are not open, making them hard to integrate into the cloud. The customers discussed the fact that open source detractors are struggling for relevance in this new cloud world. Vendors that do not support open source are fighting against becoming obsolete by working to control the standards, thereby controlling the coders. This has created a shift in the industry where security appliance vendors are beginning to “white box” their solutions to allow for the integration of customized IP. Enterprises are allocating more resources and capital on hardening their operating systems, kernels, applications, and data. These customers believe that there is much less focus on protecting the network than there needs to be.

Regulation is another area that drives customers towards utilizing their own tools. Those customers who are in healthcare or financial segments can't rely on their cloud provider, they need to use their own tools in order to stay within the bounds of regulation.

Security needs to be addressed as an end to end solution, not just firewall and intrusion detection. Indemnification is a first (and important) part of any cloud security discussion. This is because hacking has an impact beyond just the obvious. As the group discussed, recent high profile hacking incidents had an impact on the business in many ways, including share price and brand recognition. Depending on the incident, the SEC may become involved if the intrusion might be viewed as something that either impacts the trading of the stock or may release sensitive information into the public domain.

Fireside Chat – SD-WAN

Software Defined Wide Area Networks (SD-WAN) are the area of Software Defined Networking (SDN) that is gaining the most traction in the market. Part of this acceptance revolves around the fact that WANs are essentially point-to-point connections, so it is much easier for customers to tackle the challenges. Changes here can be done in a more controlled manner with less complexity than addressing some other parts of the network that are more interconnected.

Customers were quick to point out that while SDN is still a nascent technology for many of them, SD-WAN is taking hold and becoming a real-world production solution. Most, at this point, are working with the technology, and in the initial data from the main session it was pointed out that roughly 60% are in the deployment stage. The usage cases for SD-WAN are real and customers are moving to this technology today because it can solve real business issues.

The WAN landscape is changing; in the past it was monopolized by a few select players, but with SD-WAN there are plenty of new companies hitting the market. While this may be seen as a positive for some customers, the uncertainty about the long-term

prospects for some of these smaller companies does slow down the decision process. Many are not comfortable making a wholesale commitment over to a vendor who might not make it or may be consumed by a larger competitor in a future acquisition.

Most of the customers see a large potential for add-on services, allowing companies to become more agile, consistent and improve upon time to market.

There is a customer that was discussed who has over 800 locations that are all being connected via an SD-WAN; this scale of change would not have been possible with traditional WAN technology.

Most of the customers are looking at multiple transports for their SD-WAN, typically broadband to go along with the standard MPLS. One of the benefits of an SD-WAN is the ability to easily set up multiple transports. This is giving customers the flexibility for driving better economics on their WAN solutions. Because providers often set up WAN connections for customers, this move to SD-WAN may bring additional work along with the additional flexibility for the customer.

An important part of making the argument in favor of SD-WAN is the ROI. The ROI is clear for customers, but that ROI does not represent a business case. Business units are far less likely to be driving the need for SD-WAN, the more probable situation is that IT is asked to either make WAN changes or lower costs; these business needs then become a part of the exploration. When comparing the Gartner numbers to the actual savings, customers aren't necessarily seeing that level of savings in most cases, but the savings are still compelling enough that it makes sense for these customers to continue investigations or move to a POC.

MPLS circuits comprise most of a customers' WAN costs (typically ~70% of the WAN costs) and the savings from moving to an SD-WAN is typically ~50%. This makes it easy for the customer to justify the decision without having to get into all of the other "softer" savings that might be harder to quantify or harder to justify to the business.

Most of the customers are conservative, getting into SD-WAN slowly in most cases. They want to "dip their toes in the water," so the change is not a wholesale change, it typically starts in small pockets and then expands over time. Additionally, most customers are not abandoning their MPLS circuits for broadband in a wholesale manner, most are overlaying a second circuit for now until they are comfortable to remove MPLS (if they remove it at all). The additional broadband connection is allowing them to step down their MPLS traffic/bandwidth, which helps reduce (but not completely remove) the costs.

A large number of the customers are in the proof of concept phase as this point or about to enter it. But with so many companies offering SD-WAN, it is hard to make the commitment based on future long-term business probabilities. Some customers voiced that they want to see others deploying successfully before they jump in. This meeting was a positive step in that direction as a number of customers have moved out of POC and into production, boosting the outlook for the technology.

Some customers are going halfway to SD-WAN, using both standard edge routers and SD-WAN routers for now. Over time they will make the determination on when to move completely to SD-WAN. Part of that determination is based less on the customers' business and more on the vendors' business. Some are seeing too many vendors in the market as a potential negative and there is a desire to have a more stable, long-term partner proven out before making the full commitment.

Another use case that is gaining traction in the SD-WAN space is in kiosks. Some customers are utilizing more kiosks for customer interaction and this is driving the need for connectivity as each kiosk needs to be treated like a branch office (security was a prime concern for these deployments.)

Some carriers (like Verizon) are offering an SD-WAN solution today. There is some loss of flexibility and there may not be as great of a savings as customers may see on their own, but it is a healthy compromise that gives customers better control over their network while also providing more dynamic bandwidth.

An Analyst's Viewpoint

At our third ONUG event it is clear that there is continued traction in the market and demand for open networking solutions. Here are some of our team's observations:

- **Open networking is gaining traction** – With each meeting we check the pulse of the industry on their views of open networking; each time they warm up a bit more. The customers are firmly committed to open frameworks and open source, but when it comes to betting their network on open networking, they tend to be more risk-averse. However, as we have found in previous meetings, that comfort level continues to ratchet up and the customers are more likely to get involved.
- **SD-WAN is the clear leader in adoption** – As we have pointed out before, SD-WAN is a point to point connection (unlike the typical mesh of most networking) so it is an easy area for customers to take that first step. We see that the number of customers getting engaged in SD-WAN is definitely on the rise and the number of customers who are in POC is clearly up from previous quarters. SD-WAN may be the “gateway drug” of open networking, so vendors should be watching developments here to determine how the market will shake out. But as always, we caution that moving to more complex networking solutions will take longer in both the POC and the deployments.
- **Traditional vendors are coming around** – Distrust of the larger traditional networking vendors around open initiatives is shrinking. At this event there seemed to be more confidence that traditional vendors are “coming around” to the idea of open networking. But before any of these vendors starts to take a deep breath and relax, keep in mind that most of the customers discussed what vendors were saying versus what they were doing. Most of the customers said that the vendors seemed less dismissive of open networking, but that is hardly a “win” for either open networking or for the customers. Many said the proof is in the products, and until customers see true change there, many are still skeptical of the motives of the traditional vendors.

- **Making the business case is still difficult** – Outside of SD-WAN, it is still hard for many of the customers to make the business case for changing their networks. They all believe that open networking brings them significant benefits, but those are mainly being articulated on the IT side. What these customers need help with is building the business case on the business unit side. Business units are not coming to IT screaming for XYZ new technology, so IT is not going to go shake up the stack. It is important for these vendors to help articulate their business benefits, beyond just the technical benefits.
- **There is still a hesitancy** – IT is not sitting around pontificating about what should be, they are instead focused on projects and getting things done. Unless the demand comes from the BU, it is unlikely that things will change (see bullet above.) As a more conservative segment of IT, networking teams are hesitant to make changes that they don't need or have new technologies deployed. While there is a genuine interest in open networking technologies, the change is happening more slowly than hoped (especially as voiced by the vendors).
- **Still very focused audience** – While the financial sector tends to be the early adopters, vendors will need to see wider demand across other segments in order to solidify their commitment to the market. ONUG is helping to address this by having a west coast meeting in the spring. By moving to two locations on other sides of the country, this may help to broaden the customer base. However, as the market is already very tech-focused, it would be good to see the west coast meeting bring customers outside of the traditional technology companies.
- **Skill sets will need to change** – The biggest challenge that many customers voiced was the need for skill sets to change. This is not only true of their existing IT departments, but also of the new recruits. This says two things to the market. First, vendors need to be helping customers better align their internal resources. Professional services, training and recommendations will help with changing the mindset. There is clearly an opportunity for some thought leadership in the vendor community that can ultimately pay off down the road in customer acquisition and retention. Additionally, the educational world needs to shift towards offering better choices for both those in college curriculum and those who might be in tech-specific programs. The world keeps pumping out more CCIEs but everyone at the event sees this as a dead-end direction that in 10 years may be a drag on a career, not an advantage.
- **Too many choices** – Sometimes it feels like there are more vendors than deployed customers in some areas. In the SD-WAN space there were ~10 companies all vying for the same business. Some consolidation needs to happen, which is part of the market dynamic. In order to induce them into investigating technology, customers need to have choices (to avoid lock-in and better match their needs.) But too many choices will lengthen the evaluation process and add more uncertainty. We are definitely not recommending that companies either get out of open networking or don't enter the space, we are only pointing out that the plethora of choices leads to longer evaluation cycles and more hesitancy in making the decision. Over time the market will clearly work itself out, but for now the pendulum is clearly over to one side.

Important Information About This Brief

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

This paper was commissioned by ONUG. 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.

©2015 Moor Insights & Strategy.

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