MOOR INSIGHTS & STRATEGY

OPENSTACK NETWORK MATURITY

OPENSTACK MATURITY IS DRIVING MORE SCALABLE & FLEXIBLE NETWORKING

EXECUTIVE SUMMARY

The maturity of OpenStack is accelerating the move to cloud technology, enabling businesses to stand up their own private clouds as well as allowing telecom and cloud service providers (CSPs) to expand their public cloud offerings more easily. While networking was often cited as an OpenStack inhibitor in the past, we increasingly see improvements and alternatives that are helping reduce this barrier to entry.

Neutron—the core OpenStack networking component—has matured greatly over the past few years, enabling more deployments as it closes functional gaps. Additionally, there are now networking alternatives that can help deliver on the OpenStack promise with less of the associated networking pain of the past. Whether a telecom carrier, cloud service provider or an enterprise IT department, there is far more choice in the market today than ever before. While networking was previously seen as the "Achilles' heel" of OpenStack, it is clear that many of those challenges are now being addressed and that networking is less of a hindrance to adoption.

Additionally, Intel is delivering innovation for OpenStack that will enable enhanced networking functionality for all deployments, bringing greater visibility to the underlying hardware to drive much better cloud efficiency.

TODAY'S CUSTOMER NEEDS

Today's customers demand more agility, as deployment and provisioning times prevent them from capturing business opportunities. More often than not, legacy networking in traditional IT environments hinders agility. One financial sector customer at the most recent Open Networking User Group (ONUG) summarized, "We're taking a cloud-first approach, because we need to move faster. Virtualization allows us to deploy servers in only a few hours, but the network provisioning still takes us weeks. We're missing out on new business." In the most recent OpenStack user survey (cited throughout this paper), 79% responded that one of the top 3 reasons for choosing OpenStack was to "accelerate my organization's ability to innovate and compete by deploying applications faster." Clearly, the need to accelerate innovation is a key driver for customers.

Indicating the need for advanced networking capabilities, customers deploying OpenStack responded to the annual survey with network virtualization (through SDN and NFV) being the second most interesting technology, behind containers.

Education is a key need for those considering OpenStack, as these technologies require expertise that few have. If that need cannot be fulfilled by internal talent, then businesses will turn to partners or service providers to assist with design, deployment, and management. One of the best indicators of OpenStack adoption may not be the number of customers deploying but the number of administrators trained. The fact that we are seeing more trained administrators and more networking expertise starting to unfold is a very promising sign. As OpenStack becomes more integral to IT, in-house capabilities will increase, leading to both faster and more robust deployments.

While OpenStack is becoming a popular choice for many businesses, others have been hesitant to consider it because of their needs for greater stability, functionality, and scalability. It may be time for those who have held out to take another look, as improvements and enhancements have changed the landscape.

To understand customer needs thoroughly, it is important to view the market through different customer perspectives. There are three main groups of OpenStack customers: enterprises, telecoms, and cloud service providers. All will have similar needs for security, scalability, and high availability, but each group's ranking of these needs will vary based on their business models. All of those components are important to each of the business categories, but to differing degrees.

End Customer Needs Also Dictate Deployment Models

End customers can select from a variety of deployment models, depending on how they see themselves situated for supporting OpenStack. Few, if any, will take on the whole project themselves and build from the ground up; if they do, it is typically for development and test environments only, not for production workloads. Most often, customers will choose to host their OpenStack cloud either through a cloud service provider or with a partner who can help with the design, deployment, and even management of an OpenStack solution.

Customers typically deploy OpenStack in one of the ways shown in Figure 1. Each deployment model has its pros and cons, but there is no "right answer". The market opportunity is large enough for all models to coexist together.

FIGURE 1: OPENSTACK DEPLOYMENT MODELS

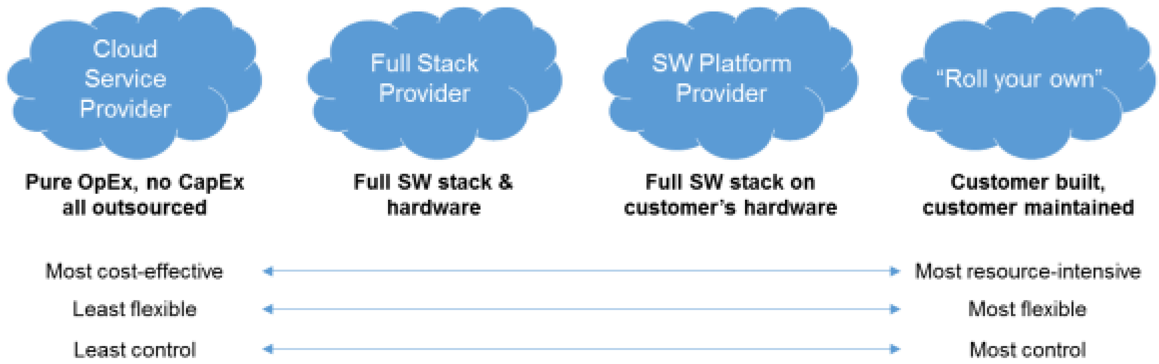

With **cloud service providers** (like Rackspace or Internap), many OpenStack network challenges are reduced through the expertise of the CSP. Hosted services from a CSP can abstract the hardware layer, turning OpenStack into a metered service.

Full stack providers (like Ericsson or Huawei) have the expertise that helps them deliver a complete stack from the hardware all the way up through the software layer. However, while this turnkey solution eases the deployment and operational challenges of OpenStack, it also limits choice for customers.

Software platform providers (like Red Hat, Mirantis, or Canonical) allow OpenStack customers to leverage the provider's expertise working through networking challenges. This approach offers a more limited set of options than if the customer were to "roll their own". Unlike a full stack provider who will also deliver the hardware, a software platform provider will usually recommend reference platforms that are part of a "tested and proven" solution. This class of provider focuses on helping customers characterize their workloads to boost scalability. They have also worked out and documented many of the networking best practices.

Only a small portion of the market (for now) will choose a completely self-built **"roll your own"** OpenStack environment. For those customers, networking is where administrators focus most of their effort. Luckily, there are replacements for Neutron from companies like Big Switch Networks, Cisco Systems, Juniper Networks, Midokura, Plumgrid, and others that can help overcome complexity, visibility, and scalability challenges. Additionally, companies like Wind River (an Intel company) provide carrier-grade performance and high-availability capabilities that augment Neutron.

Regardless of how a customer decides to approach OpenStack, Neutron has evolved rapidly and significantly, thus easing concerns and improving adoption.

NEUTRON CONTINUES TO IMPROVE

While 90% of OpenStack customers are using Neutron in one way or another (according to the user survey), there is still room to improve Neutron's functionality. One OpenStack survey respondent stated, "I hear from too many operators that there are simply too many things to configure to know which one to tweak when things don't work well." This is also true of Neutron, where the complexity—with so many configuration options to choose from—makes it difficult for users to deploy, configure, and manage, even those who are experienced networking professionals.

Most survey respondents said they needed to reduce the complexity of Neutron, with better documentation as a way to reduce confusion. Additionally, the need for better integration with compute (Nova) and the PaaS layer were critical. One of the best ways to alleviate these problems is to choose tested and proven configurations from a partner who has already put in the time and effort to validate, document, and automate these best practices. This level of integration is often more valuable than custom code, as it delivers a more solid foundation for deployment.

Despite past challenges, Neutron continues to mature and add capability. In the most recent OpenStack user survey, distributed virtual routing and virtual router redundancy for high availability (through VRRP) both ranked above 40% as actively used, interested in, or planned for use. This indicates that users are finding value in some of the advanced features beyond just load balancing and DNS. Open vSwitch remains the leading switching technology for OpenStack, and in the latest user survey, the number of Open vSwitch users in production jumped to 68% from just 13% only one year ago, signaling more maturity and momentum for the project.

The OpenStack roadmap was recently presented to analysts at the OpenStack Summit in Austin, TX (April 2016). In relation to Neutron networking, the technical areas outlined in Table 1 are being addressed in Mitaka (the newest release) and Newton (the next release). Additionally, the community is already starting to consider blueprints for the Octa release which is further out on the horizon.

TABLE 1: OPENSTACK ROADMAP RELEASES

	Mitaka	Newton	Octa
Elements	20 specs, 22 blueprints	9 blueprints	TBD blueprints
Functions	External DNS Resolution Tenant Delete L2 API Extensions Neutron Flavor Framework Add Availability Zone BGP Dynamic Routing LBaaS L7 Rules Network IP Usage API RBAC QoS	FWaaS API 2.0 Multiple L3 Backends DHCP Options per Subnet Keystone v3 VM without IP Address	

OPEN SOURCE ALTERNATIVES FOR OPENSTACK NETWORKING

While much of the public cloud has been built on open source today, public cloud service providers are sometimes seen as less giving when it comes to contributing code back to the broader community. The OpenStack community, on the other hand, is seen as having a strong push to bring code and enhancements back to the project, making it a more flexible and collaborative environment for developers and customers.

Many open source community projects are focused on delivering functionality back into OpenStack, and some of the key pioneering functions have come directly from the community. Open source networking has been a strong focus with many products and code being contributed to help expand functionality, especially in the realm of software defined networking (SDN). One of the key drivers for SDN is creating a networking fabric that allows for more agility in an OpenStack deployment, extending the core Neutron capability. This fabric is more flexible and scalable and enables service chaining, which allows a single network connection to have a suite of servers that each have different characteristics (such as firewall, load balancing, routing, NAT, etc.)

While Neutron delivers SDN capabilities, OpenDaylight is another open initiative that some have deployed as an alternative SDN controller for network virtualization capabilities, enabling administrators to manage both the physical and virtual infrastructure. Neutron's Modular Layer 2 can delegate network management to the OpenDaylight controller, enabling the OpenDaylight controller to directly manage Open vSwitch virtual switches on the physical hosts.

Open Virtual Network (OVN) is an extension to the Open vSwitch initiative that brings control plane functionality to the virtual switch, creating a virtual network through SDN. Although more limited in scope, it does bring some functionality for connecting groups of VMs or containers into private networks (both L2 and L3) without resorting to VLANs.

Containers, which are a complete runtime environment wrapped into a single package, are being used today by some of the largest cloud scale datacenters to help manage virtual workloads. Containers are now hitting the radar of traditional enterprises through OpenStack. While similar to virtual machines in how they abstract operating environments from the underlying hardware, the container model generally brings more ease in deployment, consumption, efficiency, and management. With containers, all of the application dependencies are bundled into the container, which is more abstracted from the physical networking than a VM would be. Container projects like Docker and Canal bring great potential for abstracting networking in order to deal with some of the limitations and challenges, while projects like Kubernetes (orchestration) and Mesosphere (distributed datacenter OS) help enhance container capabilities.

Finally, Open Platform for NFV (OPNFV) is a midstream project primarily focused at telecom carriers and cloud service providers. Network function virtualization (NFV) is a key driver for carrier clouds, as it virtualizes many of the network functions that physical appliances deliver today. OpenStack is the *de facto* virtual infrastructure manager (VIM) layer for NFV, and as long as Open Platform for NFV is focused closely on the OpenStack market, there is a great opportunity for synergy.

Together, these projects are helping round out OpenStack's networking capabilities. As OpenStack continues to mature and gain momentum, there are more options to help deliver networking functionality, not only within OpenStack but also outside the project.

Intel Innovation for OpenStack

Much of the innovation for OpenStack is happening at the network level. But traditional (and even startup) networking companies are not the only ones investing time and resources into making OpenStack networking more robust. Intel has been putting significant efforts into helping expand OpenStack's capabilities and drive a better networking experience for customers.

Enhanced Platform Awareness

Enhanced Platform Awareness (EPA) allows a more fine-grained approach to managing OpenStack by enabling workloads to be better matched to the underlying physical platform resources. Cloud computing has traditionally relied on abstracting VMs from the underlying hardware to provide portability. But it makes sense to correlate the workload with underlying hardware, especially when specialized hardware (like GPUs, IP packet handling, or FPGA capabilities) is in the underlying servers. By understanding workload requirements and platform capabilities before deployment, VMs can be spun up on the physical servers that better correlate to specific workload needs, delivering a more optimized experience. More importantly, enabling the VM to "see" the underlying hardware capabilities can mean better networking optimization. Instead of being tied to the most generic "lowest common denominator" networking, the VM can take advantage of hardware optimizations to maximize network performance and efficiency.

Extra capabilities are exposed to Nova (the OpenStack compute component), so a more optimized network environment can be established. If a virtual VPN server were being deployed, underlying capabilities like IPsec encryption / decryption as well as Trusted Execution Technology (TXT) can be offloaded to accelerators and specialized hardware on the physical host for better performance. Packet performance optimizations can also be handled for platforms that may have hardware-based packet acceleration capabilities installed. The open source DPDK libraries and drivers can then have access to the host hardware for better data plane performance. For instance, session border controllers can perform better by optimizing for the underlying cache capabilities of the platform.

Using OpenStack in the NFV Architecture

One of the most prevalent and forward-moving use cases for OpenStack is as a VIM for NFV in carriers and CSPs. In fact, OpenStack is the preferred manager for NFV deployments at carriers as they attempt to move network functions (like firewall, load balancing, packet inspection) from inflexible physical appliances to much more flexible virtual network functions (VNFs).

To truly take advantage of NFV's virtualized networking benefits, there does need to be an orchestration layer, as OpenStack does not readily provide this. The European Telecommunications Standards Institute (ETSI), who created NFV, also created OSM: Open Source MANO for NFV management and orchestration. Intel is a part of the OSM group and is working to help drive this standard into the NFV and OpenStack worlds. Visibility by the VIM layer to both the layers above and below can help collect and forward performance metrics and events to drive better overall application performance.

Addressing Specific OpenStack Use Cases

Security is a key area for networking, as the two are tied together with security being one of the biggest reasons that network provisioning requires so much time. The Intel Security Controller is designed to be deployed as a virtual machine and act as a broker between the security layer and the virtual infrastructure. This layer helps drive better service assurance by leveraging the telemetry data within OpenStack, enforcing better security adherence during the orchestration and management processes.

Virtual CPU Pinning

The beauty of a cloud-based infrastructure is that it can consume underlying resources in any manner and assign those resources to virtual machines based on load, utilization, and availability. This dynamic allocation of resources allows a cloud to be elastic and expand or contract based on application need, whether it is a public cloud like Amazon Web Services or a private cloud built on OpenStack.

However, there may be times when pinning virtual CPUs to physical CPUs may make sense. Normally, a cloud platform decides how resources are allocated, but if certain applications require a specific performance level or the networking components require a certain quality of service, it may make sense to "hard code" that relationship and "pin" a virtual CPU to a particular platform's physical CPU to achieve certain performance guarantees. The biggest beneficiary of this scenario is cache management. CPU pinning can prevent "cache thrashing" where cache lines are emptied and refilled constantly based on changing virtual workloads that may be overprovisioned and battling for the underlying resources.

Intel delivered the functionality to pin CPUs in the OpenStack Kilo release. Although customers are just now learning how to take advantage of this feature, CPU pinning is one more example of the capabilities that Intel is providing to the OpenStack community to help drive better functionality, greater performance, and more manageability.

Other Initiatives

The OpenStack community is driving innovation with many upcoming projects designed to drive particular features into OpenStack. A pair of projects, Murano and Tacker, have an opportunity to make some fundamental changes in the networking layer as well.

Murano, an initiative started by Mirantis, provides an application catalog for OpenStack. Customers can simply choose an application from the catalog and have it, along with all of its dependent technology components, delivered in one shot. The process is similar to how smartphone apps are purchased from the Apple or Android stores: underlying technology configuration is handled seamlessly, and customers are abstracted from much of the underlying complexity. This complexity includes the networking aspects, because the virtual networking is handled from within the application deployment mechanism. Customers are shielded from all of the networking knobs and dials traditionally associated with deploying applications. Enterprises will be the earliest beneficiaries, as Murano will help them better standardize and deploy applications on

OpenStack. Carriers and cloud service providers may leverage Murano, but many are building out their own application catalogs for their specific customers already.

Tacker is a project for building a generic VNF manager built on the MANO Architectural Framework. Tracker will help orchestrate and manage the network services that are supposed to be delivered through the virtual network functions in an NFV environment. With the biggest initial support from telecom carriers who are using NFV and OpenStack, the Tacker initiative is being driven by Brocade, Cisco, and others to help deliver true lifecycle management. While it will bring benefits to telecoms first, cloud service providers are eyeing this space as well, because virtualized network functions are also a critical part of their infrastructure.

CALL TO ACTION

OpenStack has reached critical mass from a community and vendor perspective. Enterprise customers are now finding that it has the robustness and capability that can allow them to move beyond virtualization and deliver true elastic, public cloud-like capability to their datacenters. Telecom providers and cloud service providers are leveraging OpenStack as a way to drive a more standardized and scalable offering.

While networking had previously been the weak link in the chain, significant improvements over the last few OpenStack releases have addressed many functional gaps. Additionally, most customers are finding that the path to OpenStack clouds has been insulated from networking shortcomings by partner delivery or other technologies that helped neutralize some of Neutron's previous shortcomings. As a technology provider, Intel is committing resources to help drive stability, functionality, and scalability with significant investments into both Neutron and the overall OpenStack project.

Those who are either considering or now working with OpenStack will find the best opportunity to learn more about OpenStack by becoming a contributor to the project and participating in the frequent community discussions, including many of the meetups around the world. More information about Intel's involvement in OpenStack can be found in the Intel Cloud For All initiative in the OpenStack Innovation Center.

Customers looking for the agility that cloud environments can drive—coupled with control over their infrastructure or data—can find what they need in OpenStack and should consider OpenStack for their business. Cloud service providers and telecoms who have been searching for a more standardized product offering can now turn to OpenStack with more confidence that networking has been significantly enhanced.

IMPORTANT INFORMATION ABOUT THIS PAPER

AUTHOR

John Fruehe, Senior Analyst at Moor Insights & Strategy

REVIEWER / PUBLISHER

Patrick Moorhead, Founder, President, & Principal Analyst at Moor Insights & Strategy

EDITOR / DESIGN

Scott McCutcheon, Director of Research at Moor Insights & Strategy

INQUIRIES

Contact us if you would like to discuss this report, and Moor Insights & Strategy will respond promptly.

CITATIONS

This 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

DISCLOSURES

This paper was commissioned by Intel. 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.